



UNIVERSITÀ DEGLI STUDI
DI GENOVA

Facoltà di Scienze Matematiche, Fisiche e Naturali
Scuola di Dottorato in Fisica
XXXII Ciclo, A.A. 2016-2019

**Magnetic and Newtonian noises in
Advanced Virgo: evaluation and
mitigation strategies**

A thesis presented for the degree of
Dottore di Ricerca in Fisica

Author
Alessio Cirone

Supervisors
Dr. Andrea Chincarini
Dr. Maurizio Canepa

This work was performed with the support of the “Istituto Nazionale di Fisica Nucleare” (INFN) Doctoral Fellowship at the University of Genova, and under the supervision of Andrea Chincarini (INFN – sezione di Genova). The work was partially carried out at the European Gravitational Observatory (EGO) site (Cascina, Pisa). Additional financial support for scientific missions was received by the COST (European Cooperation in Science and Technology) funding organization, within the COST action CA17137: a network for gravitational waves, geophysics and machine learning.



Dedicated to the ones I love

Abstract

In the present study, I table the first detailed estimation of the magnetic noise contribution to the Advanced Virgo sensitivity to gravitational waves. I tackle the topic by performing experimental assessments and numerical finite element simulations, all accompanied by careful data analysis. Results suggest that the magnetic noise impact for Advanced Virgo is not dramatic, but it will eventually be a considerable issue once the detector will approach its final design. In anticipation of that, I propose a mitigation strategy based on passive magnetic field shielding. In the second part, I deal with seismic newtonian noise, focusing on two crucial aspects involving the noise cancellation pipeline. These are the choice of the subtraction filter and the optimization of the seismic sensor array placement. The former issue required the definition of a machine learning algorithm based on deep neural networks, and its fine tuning. Results give some indication of good performances compared to the standard Wiener filter approach. The problem of the sensors deployment is instead addressed with the finite element analysis of the actual Virgo infrastructure and underground soil layers surrounding the test masses.

Contents

| | |
|---|-----------|
| List of acronyms | 11 |
| Introduction | 13 |
| 1 Gravitational astronomy background | 15 |
| 1.1 Gravitational waves | 15 |
| 1.1.1 Derivation from general relativity | 15 |
| 1.1.2 Astrophysical and cosmological sources of GWs | 18 |
| 1.1.3 Data Analysis techniques | 26 |
| 1.2 The Advanced Virgo detector | 28 |
| 1.2.1 Interferometer layout | 29 |
| 1.2.2 Suspension system | 31 |
| 1.2.3 Sensitivity and horizon distance | 34 |
| 1.2.4 Noise sources | 36 |
| 1.3 Observing Runs and Detections | 48 |
| 1.4 Multimessenger Astronomy | 50 |
| 1.5 Future prospects | 51 |
| 2 Magnetic noise | 57 |
| 2.1 Global magnetic noise | 57 |
| 2.1.1 Schumann resonances | 57 |
| 2.1.2 Search for a low-noise location | 59 |
| 2.1.3 Coherence measurements | 64 |
| 2.1.4 Noise subtraction with Wiener Filter | 67 |
| 2.2 Local environmental magnetic noise | 69 |
| 2.2.1 Typical magnetic sources | 69 |
| 2.2.2 Evaluation of the vacuum chamber shielding properties | 72 |
| 2.2.3 Identification of point magnetic sources | 73 |
| 2.2.4 Far-field magnetic noise injections | 74 |
| 2.3 Magnetic coupling modeling | 78 |
| 2.3.1 Magnetic coupling to the payloads | 79 |
| 2.3.2 System modeling | 80 |
| 2.3.3 Model validation | 82 |
| 2.3.4 Magnetic strain noise calculation | 87 |
| 2.4 Magnetic noise projections | 91 |
| 2.4.1 Projection from FE analysis | 91 |
| 2.4.2 Projection from far-field injections | 92 |
| 2.4.3 Projection with coherence | 94 |
| 2.5 Proposed mitigation strategy | 95 |

| | | |
|----------|--|------------|
| 2.6 | Discussion | 97 |
| 3 | Seismic Newtonian noise | 101 |
| 3.1 | Basic principles | 101 |
| 3.1.1 | Seismic waves | 101 |
| 3.1.2 | Seismic gravity perturbations | 103 |
| 3.1.3 | The Saulson model | 104 |
| 3.2 | Newtonian noise cancellation | 107 |
| 3.3 | Cancellation system infrastructure | 109 |
| 3.3.1 | Seismic sensors | 110 |
| 3.3.2 | Sensor array configuration and first tests | 114 |
| 3.3.3 | Online and offline subtraction schemes | 116 |
| 3.3.4 | Cancellation with Wiener filter | 116 |
| 3.4 | Cancellation with Machine Learning | 117 |
| 3.4.1 | Data processing | 118 |
| 3.4.2 | Neural Network building | 121 |
| 3.4.3 | Neural Network Optimization | 125 |
| 3.4.4 | Results | 131 |
| 3.5 | Sensor array optimization | 136 |
| 3.5.1 | The Lamb's problem | 138 |
| 3.5.2 | Numerical simulations | 139 |
| 3.5.3 | Temporal FE solution | 144 |
| 3.5.4 | Probability maps with normal modes | 148 |
| 3.6 | Discussion | 153 |
| | Conclusion | 155 |
| | Appendix A Linear MISO Wiener filter | 158 |
| | Appendix B Artificial Neural Networks | 160 |
| | Appendix C Fast PEM sensor maps | 163 |
| | Bibliography | 165 |
| | Acknowledgements | 183 |

List of acronyms

| | |
|-------------|---------------------------------|
| AdL | Advanced LIGO |
| AdV | Advanced Virgo |
| ANN | Artificial Neural Network |
| BBH | Binary Black Hole |
| BH | Black Hole |
| BNS | Binary Neutron Star |
| CAD | Computer-Aided Design |
| CE | Cosmic Explorer |
| CBC | Compact Binary Coalescence |
| CMB | Cosmic Microwave Background |
| CW | Continuous Wave |
| DAQ | Data Acquisition |
| DNN | Deep Neural Network |
| DQ | Data Quality |
| EM | Electromagnetic |
| ET | Einstein Telescope |
| FE | Finite Element |
| FP | Fabry-Perot |
| FGST | Fermi Gamma-ray Space Telescope |
| GA | Genetic Algorithm |
| GR | General Relativity |
| GRB | Gamma-ray Burst |
| GW | Gravitational Wave |
| IMC | Input Mode Cleaner |

IMP Input Mirror Payload
IP Inverted Pendulum
LIGO Laser Interferometer Gravitational-Wave Observatory
ML Machine Learning
MN Magnetic Noise
MGAS Monolithic Geometric Anti Spring
MSE Mean Squared Error
NS Neutron Star
PAY Payload
PEM Physical Environment Monitor
PN Post-Newtonian
PR Power Recycling
ReLU Rectified Linear Unit
RF Radio-frequency
SA Superattenuator
SGWB Stochastic Gravitational Wave Background
SN Supernova
SNN Seismic Newtonian Noise
SNR Signal to Noise Ratio
SR Signal Recycling
TCS Thermal Compensation System
TM Test Mass

Introduction

Why do we need gravitational information to understand the universe? Essentially because the record of its history is still not complete, especially going back to the first 400,000 years after the Big Bang, when the universe was opaque to light. In addition, our knowledge of the most explosive phenomena in the universe, like compact binary coalescences and supernovae, is still weak.

If we imagine the universe as a jungle, then the electromagnetic observations are analogous to watching a video with no sound of animals or flowing water; on the contrary, gravitational waves allow us to switch on the audio. The sound analogy is not accidental and is useful for a couple of reasons. Present gravitational wave detectors share the same observing frequency band with the human hearing (tens to thousands of Hertz). If these frequencies were sounds, our ear could actually hear something like a bird chirp. In addition, the gravitational wave detectors do not have any pointing capabilities, unlike the most electromagnetic telescopes. Accordingly, the sensible observing angle is almost isotropic, just like our ears, as opposed, for instance, to our optical vision, which manifests itself through the pointing attribute of a telescope. With enough “ears” in place, we can aspire to triangulate the signal and pinpoint its source.

Virgo is a giant L-shaped detector, with 3 kilometers long arms, built to sense the infinitesimal rippling of spacetime created by a passing gravitational wave. Until a few years ago, that goal seemed like a moonshot, as it was for Albert Einstein, who thought it would never be possible to confirm that gravitational waves exist. Even the most powerful black hole merger in our own galaxy would induce a variation of the Earth’s diameter of the size of the uranium nucleus. Finally, gravitational waves were detected on September 14, 2015. From that first detection and those that followed, we have begun a journey into the darkest side of the universe. These discoveries represent a completely new way of observing the cosmos, proving that multi-messenger astronomy is possible: the study of the same astrophysical objects in two different “message carriers”, namely gravitational waves and electromagnetic waves. Virgo, with LIGO and KAGRA, is part of this multiple cosmic observatory, which now has a potential of nearly ten detections a month.

The detection capability of a gravitational wave interferometric detector lies entirely on its spectral sensitivity, which is affected by a large number of noises. Detections from compact binary coalescences can be improved by extending the sensitivity to lower frequencies. In particular, the frequency range around 10 – 50 Hz is interesting for the observation of high mass mergers and to better constrain certain signal parameters, leading to more stringent tests of general relativity. Two of the low-frequency noise contributions come from terrestrial gravity fluctuations, also known as seismic newtonian noise, and magnetic field gradients.

The latter impact the data analysis by coupling with electromagnetic actuators

of the interferometer test masses and causing their displacement. The actuators fall victim not only to the local anthropic magnetic field, but also to the natural global component, due to the Schumann resonances. Although the global signal is orders of magnitude weaker than the local one, it is expected to produce a correlated effect in multiple gravitational wave detectors. If not suppressed, the Schumann magnetic noise would spoil the coherent search for the stochastic gravitational wave background. In chapter 2, I address the two problems of identifying magnetic noise sources and assessing their influence on the Virgo spectral sensitivity.

Likewise, seismic newtonian noise is expected to limit the sensitivity of the detector in two ways. First, it contains a stationary component that, at least for the current configuration of the interferometers, is predicted to be weaker than other instrument noise. Second, it exhibits the non-stationary, transient characteristics of the seismic field. A noise cancellation system will target both the components and can be achieved by monitoring their sources with seismic surface arrays. The goal is to produce an estimate of seismic newtonian noise, which is subsequently subtracted from the interferometer data. In chapter 3, I present a filter design proposal based on machine learning and a comparative study of its performance. Finally, I face the non-trivial problem of optimizing the sensor array configurations.

Chapter 1

Gravitational astronomy background

In this chapter, I will give an overview of the physical context on which my PhD work has been based during the past three years, while in no way pretending to be exhaustive. I start from the Gravitational Waves (GW) concept, predicted by the Einstein teleinstein's theory of General Relativity (GR), moving on to the ways the Laser Interferometer Gravitational-Wave Observatory (LIGO), on September 14, 2015, got to the first direct detection of GWs, almost 100 years after the famous Einstein's discovery.

1.1 Gravitational waves

Einstein's GR theory, published for the first time in June 1916 [1], led to the prediction of the existence of GWs. In January 1918 he published another paper entitled "On gravitational waves" [2] and its content is pretty much like what one can find in the textbooks today.

As far as we know, the idea of GWs did not start with Einstein, since Clifford was the first who imagined curvature waves [3]. Einstein wondered whether GWs would ever be discovered, because of their faintness and weak interaction with matter. Even if they were detectable, he also wondered if they would ever be useful in science. Now we know that GW detections are possible and are providing a plethora of scientific results.

1.1.1 Derivation from general relativity

Einstein presented his generally covariant theory of gravitation on November 1915. His field equations are now usually written in the form

$$G_{\mu\nu} = \frac{8\pi G_N}{c^4} T_{\mu\nu} \quad (1.1)$$

with G_N the Newton's gravitational constant and c the speed of light. They describe how the energy-momentum density of matter, $T_{\mu\nu}$, generates the curvature of spacetime. Here the Einstein tensor is

$$G_{\mu\nu} := R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R, \quad (1.2)$$

with $g_{\mu\nu}$ being the metric tensor (i.e. $ds^2 = g_{\mu\nu}dx^\mu dx^\nu$), $R_{\mu\nu} := g^{\sigma\rho}R_{\sigma\mu\rho\nu}$ the Ricci curvature tensor and $R = g^{\mu\nu}R_{\mu\nu}$ the Ricci scalar curvature. The latter two are obtained by contracting the Riemann tensor:

$$R^\alpha_{\beta\mu\nu} := \partial_\mu\Gamma^\alpha_{\beta\nu} - \partial_\nu\Gamma^\alpha_{\beta\mu} + \Gamma^\alpha_{\gamma\mu}\Gamma^\gamma_{\beta\nu} - \Gamma^\alpha_{\gamma\nu}\Gamma^\gamma_{\beta\mu}, \quad (1.3)$$

with the Christoffel connection given in terms of the metric by

$$\Gamma^\alpha_{\beta\gamma} := \frac{1}{2}g^{\alpha\delta}(\partial_\beta g_{\delta\gamma} + \partial_\gamma g_{\delta\beta} - \partial_\delta g_{\beta\gamma}) \quad (1.4)$$

Now, consider the case where the general spacetime metric is the sum of the Minkowski metric of special relativity ($\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$) plus some small perturbation (weak field limit):

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad \text{where } |h_{\mu\nu}| \ll 1 \quad (1.5)$$

If we substitute this into the expression for the Riemann curvature tensor, the Christoffel symbol, the Ricci tensor and the Ricci scalar curvature and only keep terms to leading (first) order in h , then the linear order of the Einstein equation takes the form

$$\square \bar{h}_{\mu\nu} + \eta_{\mu\nu}\partial_\alpha\partial_\beta\bar{h}^{\alpha\beta} - \partial_\mu\partial_\alpha\bar{h}^\alpha_\nu - \partial_\nu\partial_\alpha\bar{h}^\alpha_\mu = -\frac{16\pi G_N}{c^4}T_{\mu\nu}^{(1)} \quad (1.6)$$

where the wave operator is defined as $\square := \eta^{\mu\nu}\partial_\mu\partial_\nu$. Here it was technically more convenient to express the field equation in terms of the ‘‘trace reversed’’ tensor

$$\bar{h}^{\alpha\beta} := h^{\alpha\beta} - \frac{1}{2}\eta^{\alpha\beta}h \quad \rightarrow \quad \bar{h} = \bar{h}^\alpha_\alpha = \eta_{\alpha\beta}\bar{h}^{\alpha\beta} = -h \quad (1.7)$$

as Einstein found in 1916. As a generally covariant theory we still have gauge freedom that we can use to simplify our expressions. Considering the infinitesimal coordinate transformation

$$x'^\mu = x^\mu + \xi^\mu, \quad \text{where } |\xi^\mu| \ll 1, \quad (1.8)$$

a convenient gauge choice for gravity is the harmonic gauge, where $g^{\alpha\beta}\Gamma^\mu_{\alpha\beta} = 0$. In this condition, equation 1.6 reduces to a wave equation:

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G_N}{c^4}T_{\mu\nu}^{(1)} \quad (1.9)$$

Let us consider the case in which we are away from the sources, so that we only need the vacuum solutions with $T_{\mu\nu} = 0$. A simple one is the plane gravitational wave:

$$h_{\mu\nu} = \epsilon_{\mu\nu}e^{ik_\alpha x^\alpha} \quad (1.10)$$

characterized by a 10-component constant polarization tensor $\epsilon_{\mu\nu}$ and a wave vector k_α , which is consistent with a wave that travels at the speed of light (by substituting in the field equation: $k_\alpha k^\alpha = 0$).

If we count the degrees of freedom of the GW, there are only 2 independent ones: all the gauge freedom is fixed in the transverse-traceless (radiation) gauge, expressed in terms of $h_{\mu\nu}$ by

$$\partial_\mu \bar{h}^{\mu\nu} = 0 \quad \text{and} \quad h^\mu{}_\mu = 0 = h_{0\mu} \implies h^{TT} = 0 = \bar{h}^{TT} \quad \text{and} \quad h_{\mu\nu}^{TT} = \bar{h}_{\mu\nu}^{TT} \quad (1.11)$$

If we choose a coordinate system where the wave travels along the z axis then the plane GW can be expressed as

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} e^{ik_\alpha x^\alpha}, \quad (1.12)$$

where the two polarizations are

$$\epsilon_{(+)}^{\mu\nu} = h_+ \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \epsilon_{(\times)}^{\mu\nu} = h_\times \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \quad (1.13)$$

Consider now the two polarization modes separately and discuss the effect of a gravitational plane wave on two test masses, separated by an infinitesimal coordinate displacement. The corresponding proper separation length between the two masses in the x-direction (i.e. $dx^\mu = (0, d\zeta, 0, 0)$), as a GW propagates in the z-direction with +-mode, is

$$ds = \sqrt{g_{\mu\nu} dx^\mu dx^\nu} = \sqrt{g_{11}} dx^1 = \sqrt{\eta_{11} + h_{11}} d\zeta \approx \left(1 + \frac{h_{11}}{2}\right) d\zeta \quad (1.14)$$

Similarly, for two test masses separated only in the y-direction (i.e. $dx^\mu = (0, 0, d\zeta, 0)$) we would have

$$ds = \sqrt{g_{\mu\nu} dx^\mu dx^\nu} = \sqrt{g_{22}} dx^2 = \sqrt{\eta_{22} + h_{22}} d\zeta \approx \left(1 + \frac{h_{22}}{2}\right) d\zeta \quad (1.15)$$

with $h_{11} = -h_{22} = h_+ e^{ik_\alpha x^\alpha} = h_+ \cos(\omega t - kz)$. For the GW with \times -mode, we should instead consider $dx^\mu = (0, d\zeta/\sqrt{2}, \pm d\zeta/\sqrt{2}, 0)$:

$$ds \approx \left(1 \pm \frac{h_{12}}{2}\right) d\zeta, \quad \text{with} \quad h_{12} = h_{21} = h_\times e^{ik_\alpha x^\alpha} = h_\times \cos(\omega t - kz) \quad (1.16)$$

It is worth noting that GWs, regardless of the propagation mode, will cause the proper distance to increase for events separated in the x-direction, while it decreases for events separated in the y-direction.

In figure 1.1 I show the changing distance between a ring of test particles due to the effect of a +-polarized (top) and a \times -polarized (bottom) GW propagating out of the page.

To discuss a GW generated by a source, we should consider the equation 1.9 and its general solution, which is formally expressed in terms of the retarded Green's function

$$\bar{h}_{\mu\nu}(\vec{x}) = \frac{16\pi G_N}{c^4} \int d^3y G_{ret}(\vec{x} - \vec{y}) T_{\mu\nu}^{(1)}(\vec{y}) = \frac{4G_N}{c^4} \int d^3y \frac{T_{\mu\nu}^{(1)}(\vec{y}, t_{ret})}{|\vec{x} - \vec{y}|}, \quad (1.17)$$




Figure 1.1: Tidal effect produced by the passage of a gravitational wave propagating along the z -axis on a ring of particles at rest on the xy -plane. The top (bottom) line shows effect of a $+$ -polarized (\times -polarized) GW. The oscillation axis of the $+$ -mode is rotated by $\pi/4$ with respect to that of the \times -mode. This image is reproduced from [4].

with the retarded time

$$t_{ret} = t - \frac{|\vec{x} - \vec{y}|}{c}. \quad (1.18)$$

Now imagine we are very far from the source (i.e. $|\vec{x} - \vec{y}| = r \gg R$, with R the size of the source, let's say, a binary star system) and the radiation is at low frequency ($\lambda \gg \delta R$), then

$$h_{ij}^{TT} \sim \frac{2G_N}{rc^4} \ddot{Q}_{ij}^{TT}(t - r/c), \quad Q_{ij} := I_{ij} - \frac{1}{3} \delta_{ij} I_k^k \quad (1.19)$$

where $I^{ij} \equiv \int \rho x^i x^j d^3x$ is the well-known mass quadrupole moment, from whose time-variation a GW is generated. The pre-factor of $G_N/c^4 \sim 10^{-44} N^{-1}$ means that we need very energetic events with the ability to generate GWs that could be detectable by ground-based detectors, which are now sensitive to strains around 10^{-23} in their most sensitive frequency band. In the following section, I will describe the main known sources of GWs while in section 1.2 I will address the problem of detecting these faint signals in detail, by introducing the principles of the interferometric detection.

1.1.2 Astrophysical and cosmological sources of GWs

At this point we can say that only astrophysical sources can generate strong enough gravitational signals to be detectable on earth. Usually, we classify these sources depending on the time evolution of the signal they generate: periodic sources like

spinning neutron stars; quasi-periodic sources such as coalescing binary systems of neutron stars or black holes; impulsive sources like Supernovae (SN); stochastic background of cosmological or astrophysical origin. They all have the potential of being detected by ground-based interferometers.

Compact binary coalescences

Compact Binary Coalescences (CBC) are likely to be the “loudest” and most frequent sources for ground-based GW detectors. They consist of pairs of Neutron Stars (NS) or Black Holes (BH) or NS-BH couples and are governed by a mechanism of extraction of rotational energy and angular momentum from the system, which are converted into gravitational radiation (this is called radiation reaction). When the two objects are widely separated, they orbit relatively slow and emit weak GWs. As the process continues, the two bodies inspiral toward each other at increasingly high orbital frequencies and eventually merge into a single object. The overall dynamics is summarized in the following three phases:

- The inspiral phase: the system spends hundreds of millions of years in this phase. During it, the GW radiation has pour luminosity and we can solve the dynamics using approximation methods, the most popular being the post-Newtonian (PN) approximation. Inspiralling compact binaries are very relativistic systems, with orbital velocities touching $0.5c$ in the last rotations. Measurements have demonstrated that the PN precision needed to implement successfully the optimal filtering technique for the LIGO/Virgo detectors corresponds grossly, in the case of Binary Neutron Stars (BNS), to the 3PN approximation [5–7]. The orbit of most inspiralling compact binaries can be considered to be circular because the radiation reaction mechanism tend to decrease its eccentricity. In first approximation, a system of two structure-less point particles, characterized by two mass parameters and possibly two spins is sufficient to model the inspiral phase. The main justification for a model of point particles is that the effects due to the finite size of the compact bodies are roughly comparable to a 5PN correction and they can in general be neglected in comparison with purely GW damping effects. The signal has a characteristic chirp-like shape, meaning that both the amplitude and the frequency slowly increase in time. With the current technology development, we are able to capture only the very last part of this phase, but in the near future a BNS system would stay in the Einstein Telescope’s (ET) sensitivity band for nearly a week starting from 1 Hz (see section 1.5 for more details). Interestingly, the inspiral phase is essential to achieve some of the scientific goals of GW astronomy. When the binaries are spinning, the signal is modulated due to spin-orbit and spin-spin couplings. The modulations of the signal due to spin-orbit and spin-spin couplings encode some parameters of the source, including the masses, inclination of the orbit and spins [8,9]. When the inspiral phase moves towards the merger phase, the PN expansion lose accuracy and it is taken over by numerical relativity computations [10–12].
- The merger phase: the two objects are now getting close to each other, moving at around $0.3c$ and experiencing extreme gravitational fields. In a time-frequency plot, the merging phase is identified by a sudden chirp signal lasting

only for a few milliseconds in the case of stellar mass BHs. The PN approximation is no more accurate and a full non-linear structure of Einstein's equations is required to predict the dynamics of the bodies, characterized by tidal deformations and disruptions. Numerical simulations of Binary Black Hole (BBH) coalescences have been highly successful [13–15], allowing full GR waveforms to be useful in the search for BBHs of ever lower masses (see figure 3 in [16]). However, the understanding of the merger phase of BNSs is still poor, due to some additional non-linear effects, such as the equation of state and the magnetic fields.

- The ring-down phase: the two bodies have finally merged to form either a new stable NS or BH. Immediately after the merging phase, what remains is a BH with large distortions in its shape, which are quickly radiated away, leaving behind an undistorted and rotating BH. Here the emitted GW radiation is computed using perturbation theory. The GW signal looks like an oscillating function with an exponentially damped amplitude and, just as in the merger phase, it lasts for a very short duration ($\approx ms - s$ depending on the mass of the final object). Once again, the unique signature depends on the mass and spin angular momentum in the case of BHs, but also on the equation of state in the case of NSs [17, 18].

The three phases are summarized in figure 1.2, where we can see a plot of the GW signal vs time of the BBH event GW150914 [19]. During the inspiral, this looks like an oscillating function, but the amplitude is increasing with time, along with the frequency of the oscillation.

Spinning neutron stars

Besides CBC signals, another class of GW signals potentially observable by ground-based advanced detectors are the periodic signals emitted by rotating, distorted neutron stars. These so-called Continuous Waves (CW) are particularly appealing, but the price to pay is that we expect them to be extremely weak. So far, we have only set stringent limits on their attributes.

NSs are remnants of SN explosions, which leave behind an astrophysical compact object with a typical mass $M \sim 1,4M_{\odot}$, radius of order $R \sim 10km$ and average density $\rho \sim 10^{14}g/cm^3$. Models of the internal structure of a NS suggest a sequence of spherical shells with different density and composition: the outer/inner crust and the outer/inner core. Finding the equation of state (relation between pressure and energy density, $p = p(\epsilon)$) of cold, ultra-dense matter is an open problem in nuclear physics [20].

The peculiarity of these isolated spinning NSs is that they are asymmetric with respect to the rotation axis and emit CWs at a frequency of two times the rotational frequency. If the star's rotation speed remains constant, so does the emitted GW. Hence the GW maintains continuously the same frequency and amplitude. The GW emission occurs when the NS mass distribution displays a time-varying quadrupolar component. What could cause the deviation from spherical shape in a compact object like a NS is still a subject of study. It could be due to:

- Asymmetric distortions (i.e. mountains), supported either by elastic stresses in the crust or by magnetic fields [21–24].




Figure 1.2: Top: numerical relativity estimation of the GW strain amplitude from the BBH event GW150914 and the reconstructed template from LIGO Hanford. Bottom: the Keplerian effective black hole separation in units of Schwarzschild radii and the effective relative velocity given by the post-Newtonian expansion. Figure reproduced from [19].

- Unstable oscillation modes in the fluid part of the star [25, 26]. The most promising candidate might be the family of r-modes, which are toroidal fluid oscillations driven by the Coriolis force [27]. The r-mode instability has been proposed as a source of GWs from newborn NSs [28] and from rapidly accreting NSs [24].
- Free precession, or the simply called ‘wobble’ of a NS with a misaligned rotation axis with respect to its symmetry axis [29, 30].

Besides these, there are also finite-size tidal effects arising on extended bodies when they are immersed in external gravitational fields, like the case of a BNS system. The overall result is that the structure of the body is deformed from the spherical shape and stretched in the direction of the gravitational source (see figure 1.3 for the simple case of a point-like gravitational source at distance r).

The finite-size tidal corrections are encoded in the so-called rotational tidal Love numbers, a set of coupling constants which characterize the multipolar deformation of a spinning compact object [32, 33]. It was shown that these quantities are precisely zero in the case of a slowly rotating Kerr BH, while those of a slowly spinning NS depend on the equation of state [34, 35].




Figure 1.3: Schematic representation of the quadrupolar tidal deformation of an extended body, together with the vector field lines representing the tidal field, induced by a pointlike gravitational source at distance r . Figure reproduced from [31].

Core-collapse supernovae

The rapid collapse of the core of a massive star (at least $8M_{\odot}$), the so-called type II SN, is one of the most energetic events in the Universe, releasing $\sim 10^{53}$ erg of gravitational energy in the compression of a massive star's iron core to a NS. Neutrinos radiate most of this energy ($\sim 99\%$). In the past, the type II SN event was considered a very promising source of GW emission. If the collapse is non-spherical, as a result of strong rotation or magnetic field mechanisms, then GWs could carry away some of the binding energy and angular momentum. More recent estimates of the energy emitted in GWs have reduced the prospects of detection by first-generation interferometric detectors to galactic events, at best. Galactic events are known to be rare: the estimates are less than one event per two decades in the Milky Way and in the local group of galaxies. A recent review [36] recalls, however, that there may be 1 SN occurring about every year within 3-5 Mpc from Earth. Within a few years of observation, the advanced detectors will possibly detect a few such events in our local universe up to a few Mpc [37]. Recent simulations suggest that the GW emission processes in a core collapse event, 1 Mpc far away from us, would give rise to strains $h \sim 10^{-24} - 10^{-22}$ [38], with an emission concentrated mainly at frequencies $f \sim 200 - 1000$ Hz. The shape of the signals emitted during SN events is not as well modeled as for coalescences: the examples shown in figure 1.4 are based on complex simulations of the core collapse [39].

Stochastic gravitational wave background

Stochastic Gravitational Wave Background (SGWB) is the superposition of many incoherent sources. It is mostly composed by unresolved GW sources such as distant mergers of NSs or BHs, but there could be also primordial GWs emitted in the early stages of the universe evolution [41] (by inflation [42], electroweak phase transition [43] or cosmic strings [44]). A SGWB from CBCs [19, 45], magnetars [46] and core collapse SN [47] is sometimes regarded as background of astrophysical origin and it is potentially detectable by the second-generation detectors at their design sensitivity [48, 49].




Figure 1.4: Examples of burst GW signals emitted by a simulated core collapse SN. Model A (upper panel) is a strongly rotating core that collapses rapidly. On the other hand, model B (lower panel) is initially an almost rigid rotator, which bounces later than model A. Relativistic models are denoted by solid lines, Newtonian ones by dotted lines. Image taken from [40].

Since the SGWB radiation is a superposition of signals from all possible propagation directions, the strain amplitude seen by a detector at time t and position \vec{x} can be written as:

$$h_{ij}(t, \vec{x}) = \sum_p \int d\Omega F_{ij}^p(\hat{\Omega}) \int_{-\infty}^{+\infty} df e^{i2\pi f(t - \frac{\vec{x} \cdot \hat{\Omega}}{c})} h_p(f, \hat{\Omega}) \quad (1.20)$$

where $\hat{\Omega}$ is a unit vector pointing to a source and $F^p(\hat{\Omega})$ is the detector's antenna pattern for polarization p , which specifies how a GW signal of given polarization translates into the output of the detector; therefore $h_p(f, \hat{\Omega})$ represents the strain amplitude that characterizes the p component of the stochastic signal from direction $\hat{\Omega}$ at frequency f . The stochastic background is assumed to be stationary, isotropic and Gaussian, hence the statistical properties of the strain are specified entirely by the second-order moment

$$\langle h_p(f, \hat{\Omega}) h_q^*(f', \hat{\Omega}') \rangle = \frac{1}{4\pi} \delta(f - f') \delta^2(\hat{\Omega} - \hat{\Omega}') \delta_{pq} \frac{1}{2} S_{GW}(f) \quad (1.21)$$

where $S_{GW}(f)$ is the spectral density. It can be seen that

$$\langle h_{ij}h^{ij} \rangle \approx \int_0^\infty df S_{GW}(f), \quad (1.22)$$

meaning that $S_{GW}(f)$ directly contributes to the detector's spectral noise density. It is customary to characterize the magnitude of the SGWB in terms of its normalized logarithmic energy density parameter:

$$\Omega_{GW}(f) = \frac{1}{\rho_c} \frac{d\rho_{GW}}{d \log f} = \frac{f}{\rho_c} \frac{d\rho_{GW}}{df}, \quad (1.23)$$

with respect to the critical energy density ρ_c required to close the Universe [50]:

$$\rho_c = \frac{3H_0^2}{8\pi G} \quad (1.24)$$

The ρ_{GW} in equation 1.23 is the energy density of a GW field:

$$\rho_{gw} = T_{00} = \frac{1}{32\pi G} \langle \dot{h}_{ij}\dot{h}^{ij} \rangle. \quad (1.25)$$

By substituting the two quantities in 1.25 and 1.24 in equation 1.23 and by taking the two temporal derivatives of h (from equation 1.20, we get a factor $(2\pi f)^2$), it can be directly shown that

$$\Omega_{GW}(f) = \frac{4\pi^2}{3H_0^2} f^3 S_{GW}(f). \quad (1.26)$$

In general, the models of the energy density of the cosmological SGWB are well approximated by a power law, at least in the frequency band accessible to ground-based interferometric detectors:

$$\Omega_{GW}(f) = \Omega_n \left(\frac{f}{f_0} \right)^n. \quad (1.27)$$

The standard inflationary theory foresees a flat spectrum $n = 0$, while other models, including string models, foresee $n > 1$ and therefore spectra increasing with the frequency. Taking into account current models (see figure 1.6 and 1.7), one should aim at $\Omega_{GW}(f) \leq 10^{-6}$. In terms of the detector amplitude spectral density, this means that one should detect at 100 Hz: $\sqrt{S_{GW}(f)} \simeq 10^{-24} Hz^{-1/2}$, about a factor 10 below the best sensitivity attained so far by advanced Virgo.

A SGWB appears in a single detector as an additional noise, and we learned that this would be below the instruments noise floor. However, the SGWB signal is coherent across different instruments, so the issue is faced by cross-correlating their outputs. The correlation between two detectors l and m can be written as

$$\langle h_l(t)h_m(t') \rangle = \int df e^{i2\pi f(t-t')} \gamma(r_{lm}, f) S_{GW}(f) \quad (1.28)$$

where γ is the overlap reduction function in terms of the distance r_{lm} between the two detectors: for distant detectors, it is similar to a *sinc* function of the argument, displaying an oscillating pattern, and tends to 1 for parallel, co-located detectors. Examples of overlap functions representing the LIGO Hanford - LIGO Livingston (H1-L1) and LIGO Hanford - Virgo (H1-V1) network's sensitivity to

stochastic background for incoming GWs of 3 possible kinds of polarization are shown in figure 1.5.

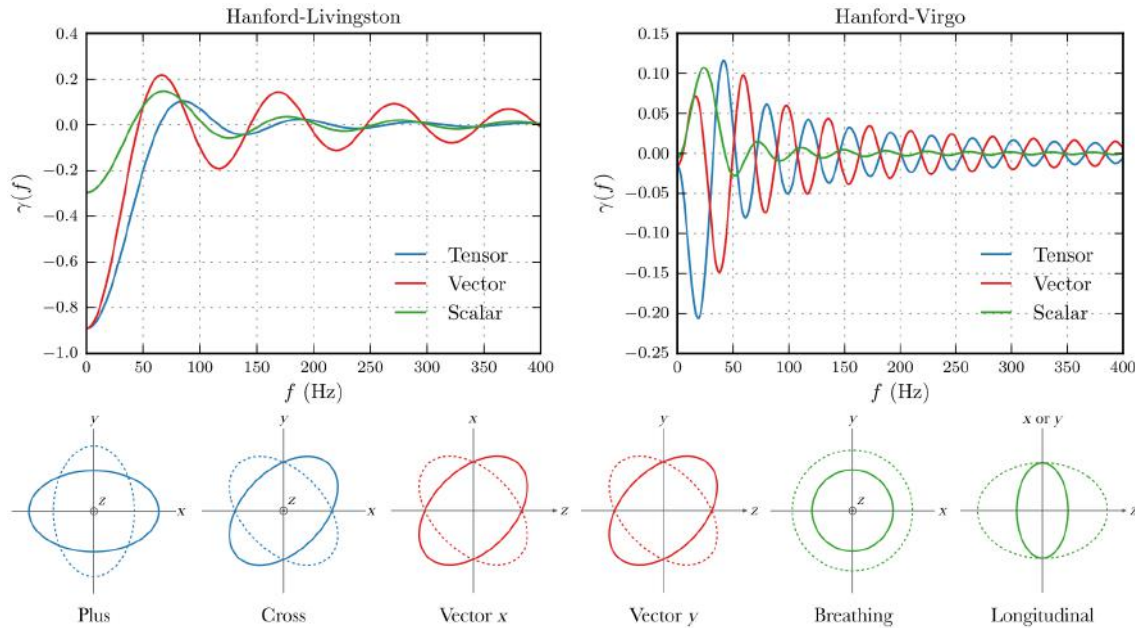


Figure 1.5: Overlap reduction functions representing the LIGO Hanford - LIGO Livingston and LIGO Hanford - Virgo network's sensitivity to stochastic background of tensor (blue), vector (red), and scalar-polarized (green) GWs. The lower panel reviews the 6 GW polarizations allowed in general metric theories of gravity, in terms of the deformation of a ring of freely falling test particles: 2 tensor (plus and cross), 2 vector (x and y) and 2 scalar (breathing and longitudinal) polarization modes. Figure reproduced from [51].

Interestingly, the H1-L1 standard tensor overlap reduction function decays rapidly to zero above $f \approx 64$ Hz and this results in a negligible sensitivity for tensor background searches with only the LIGO network. On the other hand, if we consider the H1-V1 network, we see an oscillating pattern of the tensor overlap function, which remains larger in amplitude than H1-L1's at frequencies $f \geq 200$ Hz, implying heightened relative sensitivity.

In figure 1.6 I also present the energy density constraints on the background, with bounds from ground-based interferometers compared with constraints from other measurements [52]: Cosmic Microwave Background (CMB) at low multi-pole moments, big-bang nucleosynthesis, pulsar timing and the ringing of Earth's normal modes.

It is also shown a projected limit from the space-based detector LISA. In addition there are some representative models of BNS and BBH stochastic background, built by using a $f^{2/3}$ power-law: a zoom of the black box region is shown in figure 1.7, together with the expected sensitivity of the network of Advanced Virgo (AdV) and Advanced LIGO (AdL) detectors. The presented implications for the SGWB from CBCs are the result of parameter estimations from the BBH and BNS merger events we have detected so far [19, 53].

A detection of a SGWB would be a significant result for GW astronomy, having ambitious implications for cosmology and astrophysics. By performing a year-long




Figure 1.6: Energy density constraints on the stochastic background, as well as some representative models. We compare the limits from ground-based interferometers, also with the projected design sensitivity of the network assuming two years of coincident data, with constraints from other measurements: CMB at low multipole moments, big bang nucleosynthesis, pulsar timing and the ringing of Earth’s normal modes. We also show projected limits from LISA. We also show models of the BNS and BBH distributions, using an $f^{2/3}$ power-law.

integration of data, it is possible to increase sensitivity to signals that would normally be hidden under the strain noise. A key assumption is that the noise in each detector is uncorrelated and hence reduced with continued integration, whereas correlated noise is not, producing a systematic bias.

However, the SGWB search is also sensitive to noise spectral lines and large transients, forcing us to use data quality vetos to reject the most noisy time periods for the analysis [48, 54].

Moreover, global electromagnetic (EM) noise from Schumann resonances is a source of correlated noise for such distant detectors [55]. It may be reduced by minimizing magnetic coupling to the strain channel or by using magnetometers to subtract correlated noise [56]. Effective subtraction requires precise measurement of the Schumann EM field near detectors of interest [57, 58]. This demands low local EM contamination around the measurement station. Measurements and the subtraction of Schumann resonances are part of my work, which will be treated in chapter 2.1.

1.1.3 Data Analysis techniques

The interferometric detectors have performed several GW searches over the last decade, sometimes in coincidence. Many noise events and non-stationarities spoils the fundamental Gaussian component of the detector output, with the consequence that searches for GW events require signals to be observed in multiple detectors, in order to reduce the number of false-alarm events due to instrumental and/or environmental disturbances. Moreover, a multi-detector network offers a better sky coverage and the possibility of localizing the position of the source. The strain amplitude signal may also contain rare GW transients, such as weak continuous signals or a stochastic background, buried in the instrumental noise. In the previous section, I reviewed the main GW astrophysical sources that we expect to emit short

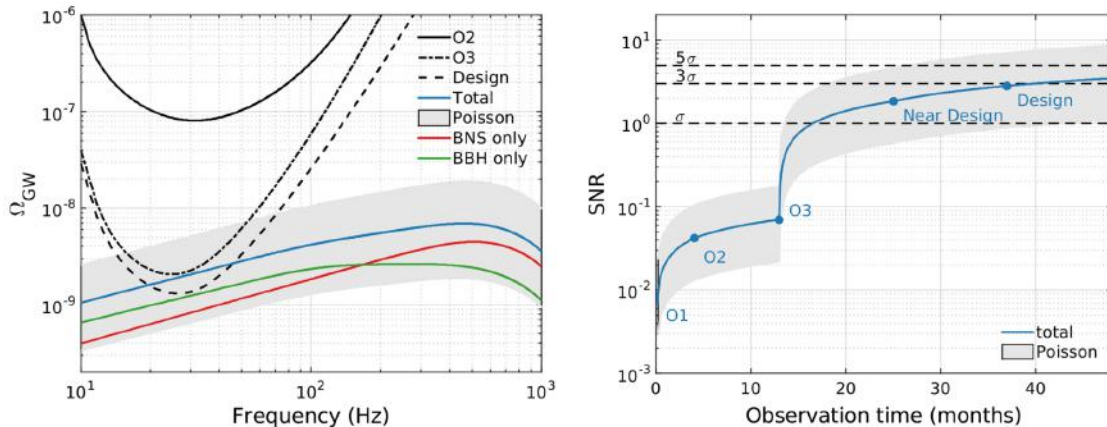


Figure 1.7: Left: predicted median background for the BNS (red) and BBH (green) models, the total combined background (blue), and the Poisson error bars (gray shaded region) for the total background. We also show expected power-law integrated curves for observing runs O2, O3, and design sensitivity. Virgo is included in O3 and beyond. The black curves for O3 and beyond cross the Poisson error region, indicating the possibility of detecting this background or placing interesting upper limits. Right: SNR median total background (blue) and associated uncertainty (shaded region). The median of the predicted total background can be detected with $\text{SNR}=3$ after 40 months of observation time, with LIGO-Virgo reaching design sensitivity (2022-2024). The markers indicate the transition between observing runs. The plots are reproduced from [53].

duration signals. When the latter are well modeled, as in the case of the CBCs, we can use template-based matched filtering techniques to search for GW signals [48, 59]. In the other cases, we rely on robust methods to detect a “burst” of excess energy in the detector network [60, 61], with the drawback that we become particularly susceptible to the presence of transient noise events (or glitches). Some of them are easily identifiable and the source can be localized and mitigated. However, for many transients, the noise source cannot be eliminated or it requires to wait until the end of the data acquisition period to make a move. In the worst case scenario, we have to exclude short periods of time surrounding these noise events (the so-called “vetoing”). Now I review the basic knowledge concerning data analysis of various types of GW signals.

- **Compact binary coalescences.** The very weak signals from compact binary inspiral require matched filtering of a set of accurate theoretical template waveforms against the detector output. The information about the properties of these sources, such as the masses and spins, can be extracted via Bayesian inference from the inspiral, merger and ring-down phases of BBHs [18, 62] and BNSs [63, 64]. Theoretical models are built by combining PN calculations [65, 66], the effective-one-body formalism [67, 68] and numerical relativity [13, 69]. Analyses in [70] have shown that the inclusion of PN effects up to the 3 PN level could suffice to get sufficiently accurate theoretical templates. In practice, the PN templates for the inspiral phase must match numerical relativity results for the subsequent merger and ring-down phases [71]. The search results are obtained by two different, largely independent matched filter

pipelines: PyCBC, GstLAL, and the burst search cWB.

- Continuous waves. It is expected that rapidly rotating asymmetric NSs emit CWs, as opposed to transient-like GW signals. We use matched filtering techniques, restricted to a narrow frequency band Δf centered on f_0 ($\Delta f \sim 10^{-4} f_0$), to target known pulsars [72–74]. On the contrary, an all-sky broadband (20 Hz - 2 kHz) search is performed for CW signals with unknown parameters, with the drawback of reducing the sensitivity and increasing the detector noise level [75–77]. Recently, a study for a possible deviation of CW signals from GR, by including non-tensorial modes, has been proposed [78].
- Stochastic gravitational wave background. SGWB searches are based on the assumption that these signals from unresolved binary mergers correlate two detector’s strains over a wide frequency range [79,80]. This searches are particularly sensitive to both noise spectral lines and large transients, which distort the estimation of the strain spectrum. That is why the published SGWB searches make use of the data-quality works to reject the most noisy time periods for the analysis [48].

Since many noise sources cannot be clearly understood nor mitigated, they must be identified and tagged in the strain data. These events will be vetoed during the application of the search pipelines with data quality flags. These flags are usually defined by using auxiliary channels witnessing some malfunctioning of a certain part of the interferometer or the presence of external disturbances. Any GW event occurring during flagged times is vetoed [81,82]. The AdV detector records thousands of these auxiliary channels that monitor instrument behavior and environmental conditions. Examples include sensors measuring angular drift of optics, light transmitted or reflected by a test mass, and actuation signals used to control the position of optical components. A crucial class for my experimental activity is that comprising the so-called Physical Environment Monitor (PEM) sensors, which are divided into slow and fast response sensors. The former include temperature and humidity probes, while the latter are seismic, magnetic, acoustic, radio-frequency, voltage and current sensors, all deployed inside (sometimes also outside) the interferometer buildings. During my stay at Virgo site, I had the opportunity to build an almost complete map of the fast sensor locations, which is shown in appendix C.

1.2 The Advanced Virgo detector

In section 1.1.1, I showed that the passage of a GW produces a tidal effect on a ring of particles at rest on the xy-plane, perpendicular to the wave propagation direction. Moreover, this effect is proportional to the mutual distance between the test masses and to the GW amplitude. Hence, a smart technique to detect directly the gravitational radiation is definitely large-scale interferometry. First generation detectors, like Virgo [83] and LIGO [84], have demonstrated the feasibility of a GW detection with interferometry and yet, in 2015, the first BBH coalescence was detected [85] by the advanced instruments. In this section, I discuss the layout and configuration of the Advanced Virgo detector, starting from the basic principles on which it is built. A separate subsection will be devoted to the suspension system,

which is a crucial component to understand the functioning mechanisms of both noises that I addressed during my PhD work.

1.2.1 Interferometer layout

An interferometric detector uses the interference of two coherent light beams along an optical path. Even a small change in their phase difference is directly associated with the relative displacement of the reflective elements of the detector, which are the mirrors. The simplest interferometric optical layout is the Michelson configuration, which is used as the basic scheme of current GW detectors.

Advanced Virgo is a ground-based GW detector built in Cascina, near Pisa. It is a dual recycled Michelson interferometer with 3 km long Fabry-Perot (FP) arm cavities [86]. A depiction of the main optical layout is given in figure 1.8.

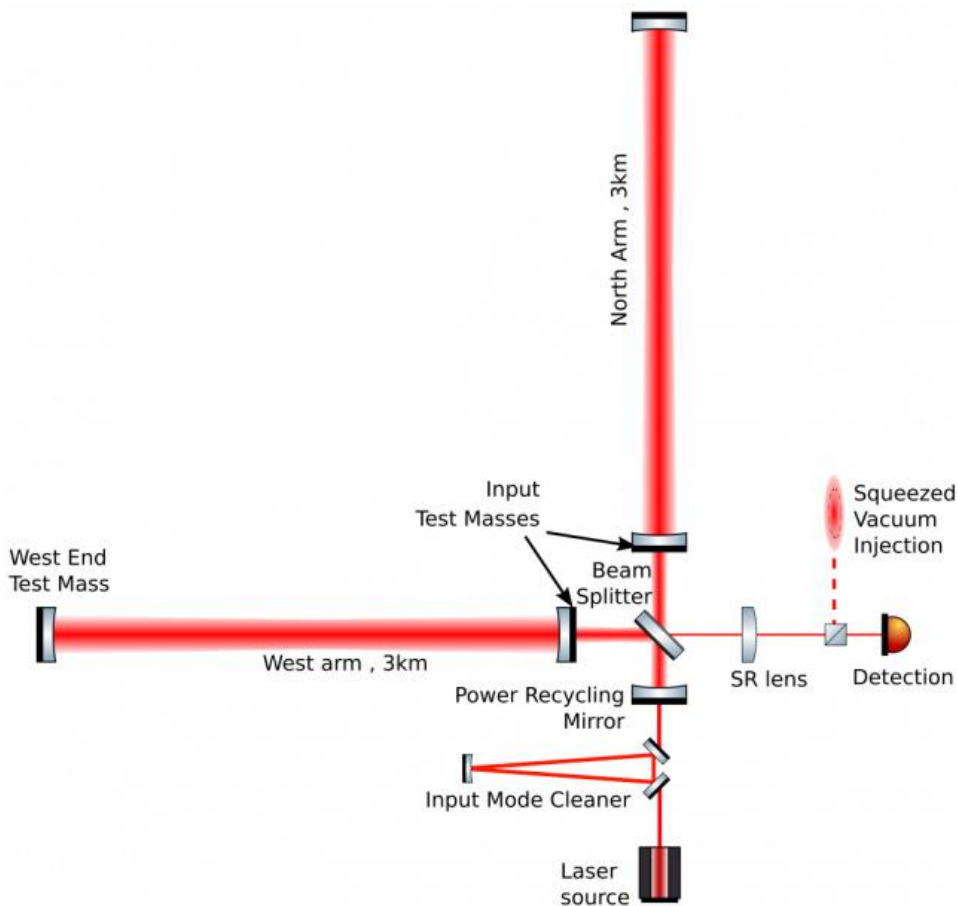


Figure 1.8: Simplified optical layout of the AdV interferometer during the third observing run, O3 (2019-2020). Each 3 km long arm cavity is formed by an Input Mirror (IM) and an End Mirror (EM). The recycling cavities are instead 12 m long and are formed by the Power Recycling mirror, the Signal Recycling lens (in anticipation of the actual mirror) and the two IM. Credits: The Virgo Collaboration.

The FP cavities are designed such that laser light, divided into 2 equal beams by a Beam Splitter (BS) mirror (55 cm in diameter and 6.5 cm thick), bounces many times within the Input and End Mirrors (IM and EM) of the cavity, before exiting. This design increases the power stored in the arms and their effective length by a

factor proportional to the cavity finesse (a quantity which depends on the reflectivity of the cavity mirrors), also increasing the effect due to a passing GW.

In order to be able to detect a GW, all the mirrors of the interferometer must be in a free-falling condition, making them effective gravitational test masses. Therefore, they are suspended to multistage passive attenuators designed to filter the ground micro-seismic vibrations. The 4 mirrors of the FP cavity are made of fused silica, with a mass of nearly 40 kg and a cylindrical shape with diameter of 35 cm and a thickness of 20 cm. The atomic level mirror polishing and the multi-layer coating allow to reach reflectivities of the order of 99.999%.

The instrument uses two main input laser sources: the high-power Nd:TVO4 slave laser, stabilized by the high-stability solid-state Nd:YAG master laser. The resulting beam is characterized by a wavelength $\lambda = 1064 \text{ nm}$ and an output power up to 200 W in the final design. Before entering the main parts of the detector, the laser beam passes through an Input Mode Cleaner (IMC), a 144 m long triangular cavity to suppress the unwanted transverse modes different from the fundamental Gaussian TEM00 mode.

To reduce the quantum fundamental noise of photon counting, which limits the sensitivity at high frequencies, AdV employs the squeezing technique [87]: a vacuum state of the optical field, the so-called squeezed vacuum, is injected into the interferometer.

A part of the optical power stored in the FP cavities is absorbed by the mirrors, producing a temperature gradient inside them. This absorption leads to thermo-optic distortions (due to the temperature dependence of the refraction index) and thermo-elastic effects (due to thermal expansion coefficient) on the mirrors. To contrast these effects, a Thermal Compensation System (TCS) is installed [88], consisting of additional components: a shielded ring heater placed around each mirror, to control the change in its radius of curvature due to thermo-elastic deformations; an independent double-axicon pattern CO_2 laser to heat a transmissive compensation plate placed in front of the mirror, and contrast the thermo-optic distortions. The TCS sensing is performed by phase cameras and Hartmann sensors [89].

To a first approximation, the working point of the interferometer is at its dark fringe mode. It means that no light reaches the output photodiode, thanks to a destructive interference. Conversely, in the presence of a GW, a very little light signal forms. As a consequence, the majority of the light leaves the interferometer propagating back towards the laser. Hence, a Power Recycling (PR) mirror reflects back into the FP cavities the light that would otherwise be lost. In a similar fashion, Advanced Virgo has been prepared to accommodate also a Signal Recycling (SR) mirror, to recycle light right before the Output Mode Cleaner (OMC). The latter is a 2.5 cm long resonant cavity used to further clean the laser signal from transverse electromagnetic modes which might be caused by interferometer's misalignments, and from radio-frequency sidebands that are introduced in the arm cavities for the Pound-Drever-Hall locking procedure. Tuning the length of the SR cavity can also broaden the frequency response of the interferometer [90].

The correct resonance conditions inside the cavities are maintained by controlling the distances between the main optical parts with active feedback systems, with an accuracy of about 10^{-16} m RMS. The longitudinal control deals with four main degrees of freedom, the most relevant of which is the differential change in length of the two FP cavities (DARM), namely the main degree of freedom sensitive to the

effect of a GW signal [91]. More in detail, the control (locking) of the interferometer is achieved acting on the mirrors with electromagnetic actuators at various stages of the suspension systems. The control system relies on the response of a variety of sensors, such as optical levers, accelerometers and displacement probes.

1.2.2 Suspension system

Seismic noise gives the most important contribution at low frequency where GW emitted by pulsars and coalescing binaries are expected.

For these reasons, since their first generations, ground based and broadband interferometers have been equipped with an appropriate suspension system of the test masses. The design of the suspension system for test masses is essentially based on the harmonic oscillator working principle. For frequencies higher than the oscillator fundamental mode ω_0 ($\omega \gg \omega_0$) the system behaves as a second order low pass mechanical filter. With this approach, the attenuation response of the test mass can be enhanced by cascading a certain number of this base oscillators. In a n -stage pendulum, the horizontal displacement of the suspension point is transmitted to the lowest mass with an attenuation factor of ω^{-2n} . A similar logic path is followed for the noise transmission along the vertical direction. Adopting a multistage pendulum this goal can be achieved by replacing each suspension wire with a spring to form a cascade of vertical oscillators. This mechanical attenuator is based on a particular geometry of cantilever blades called monolithic geometric anti springs (MGAS) [93]. The spring support heavy load and, at the same time, it is soft enough to exhibit a low resonant frequency. It is possible to confine the vertical modes of the chain in a very low frequency region obtaining a strong attenuation starting from a few Hz.




Figure 1.9: The AdV Superattenuator. We can distinguish, from top to bottom, the three legs of the Inverted Pendulum (IP), the Filter 0, the top ring, the passive filters and the mirror suspension. The figure is reproduced from [92].

The Superattenuator

The Superattenuator (SA) is the mechanical structure developed by the INFN to isolate the test mass of the Virgo interferometer from seismic noise and local disturbances [94–96]. The system is a hybrid one formed by: a chain of 6 passive mechanical filters linked to each other by a thin metallic wire; a three 6 m-long leg structure based on the working principle of an Inverted Pendulum and connected to ground by means of a flexural joint [97]; a payload as last stage formed by a marionette, a recoil mass and the mirror. The overall structure is presented in figure 1.9.

The Advanced Virgo Payloads




Figure 1.10: a) CAD (Computer-Aided Design) drawings of the AdV Input Mirror Payload (IMP) integrated in the approximately 10 metres Superattenuator suspension tower; b) CAD drawing of the IMP, including the main metallic assemblies surrounding the TM; c) photo of the integrated IMP.

The AdV payload (PAY) consists of two suspension stages: the *Marionette*, which is a structure holding the test mass (TM) and the *Actuator Cage* [98]. The PAY is suspended to the last stage of the so-called Virgo *Superattenuator*, which is a series of six vertical and six horizontal passive mechanical filters (Figure 1.10a). The overall system is designed to suppress the seismic vibrations by many orders of magnitude, starting from a few Hz [94]. The typical configuration of the PAY is shown in Figure 1.10b. The main structures close to the magnets are the actuator Cage, the Marionette and a set of ring-shaped components surrounding the TM (*Baffles*, *Ring Heater*, *Compensation Plates - CP*, etc...). The Cage is directly connected to the last stage of the Superattenuator through the *Coil Disk*. It also supports the set of driving coils that act on a total of 8 permanent magnets (Sm_2Co_{17} magnets of 1 T, 8 mm in diameter and 4 mm thick) glued on the Marionette [86]. The magnetic mount of the 8 actuators have horizontal and vertical orientation, in anti-parallel configuration (Figure 1.11a). This coil-magnet system steers the PAY in three degrees of freedom: the translation along the beam (roll) and the rotations (pitch and yaw) around the other two orthogonal axes. Starting from this general structure, each suspension chain is optimized in different ways so that we have 4 different types of PAY (e.g. the Input Mirror Payload - IMP - in Figure 1.10c). Other four coil-magnet pairs act directly on the TM, with the magnets (properties

in Table 1.1) glued in a cross anti-parallel configuration (Figure 1.11b). This solution is expected to be the most effective against electromagnetic disturbances.




Figure 1.11: a) Drawings of the final stage of the TM suspension system, which includes the Marionette with eight actuation magnets, four fused silica suspension fibres and the TM itself; b) schematic view of the anti-parallel cross configuration of the four magnets glued on the TM.

| Properties | Values |
|-------------------------------------|----------------|
| Material | Sm_2Co_{17} |
| Diameter [mm] | 1.5 ± 0.1 |
| Thickness [mm] | 1.5 ± 0.1 |
| Residual Induction [T] | 1.0 ± 0.15 |
| Magnetic moment [$mA \times m^2$] | 2.0 ± 0.3 |

Table 1.1: Properties of the magnets glued on the surface of each TM.

Indeed these magnets are not only sensitive to the magnetic field produced by the driving coils but also to any (noisy) external magnetic gradient. While the anti-parallel configuration of all the magnets located on both Marionette and TM should be insensitive to any contribution that is spatially uniform, asymmetries due to a distinctive field-structure can produce a net total force, causing a displacement noise. In addition the magnetic forces act on the TM directly, by-passing all the seismic isolation provided by the suspension system.

The Monolithic Suspensions

In the initial Virgo configuration the test masses were suspended by means of two C85 steel wires. In order to reduce the thermal noise, which is one of the main noise sources limiting the sensitivity of GW interferometers [99], we now use fused silica (SiO_2) wires. They are attached to the mirror by welding or using silica bonding, which can reproduce the connection between materials at the molecular level [100]. We refer to this design as monolithic suspension, since the suspension wires and their clamps can be considered as a single monolithic block. In particular, the ends of each fiber are welded into two T-shaped fused silica blocks (the anchors), which

are then connected to the marionette on one side and to the mirror on the other side. On the mirror side, the anchor is glued with a silica bonding technique to a section of fused silica protruding from the mirror (the ear). The other end of the fiber, welded to another anchor, is connected to the marionette with a stainless steel interface (the upper clamp assembly). The main steps followed to realize a monolithic suspension are depicted in figure 1.12. The monolithic fibers presently suspend all the 4 FP cavity mirrors, after solving some issues about the sudden failure of the fibers consequent to the tower venting procedure.




Figure 1.12: Sketch of the steps followed to realize a monolithic suspension. a) two silica ears are silicate bonded on the side of the mirror, with grooves for the insertion of suspension elements; b) a silica wire is produced with the lower tip welded to a silica piece named anchor, while the upper end is welded to a silica cone; c) four wires are placed on the mirror by inserting the anchor under the ear, within the groove, and silicate bonding them; d) the upper cones of the wires are clamped to the payload making use of tight steel clamps. Figure reproduced from [101].

1.2.3 Sensitivity and horizon distance

The sensitivity curves, such as those shown in figure 1.13, are computed by considering the different noise sources in the interferometer, depending on the frequency. In other words, they indicate the minimum detectable intensity of a GW, usually in units of equivalent spectral density h , computed using dedicated codes such as GWINC (GW Interferometer Noise Calculator) [102]. In figure 1.13, there is reported an example of AdV strain sensitivity curve at the beginning of O3 (blue line). Also the O2 reference sensitivity and the AdV observing scenario are reported.

In practice, The AdV spectral sensitivity to GW signals is obtained with a precise calibration of the detector response to mirror motion and the measure of the effect of a differential arm elongation on the dark fringe signal (transfer function, in the range 10 Hz to 10 kHz). The displacement sensitivity in m/\sqrt{Hz} is obtained by multiplying the dark fringe noise linear power spectrum in W/\sqrt{Hz} by the Virgo transfer function in m/W . The strain is obtained by dividing the result by the arm length.

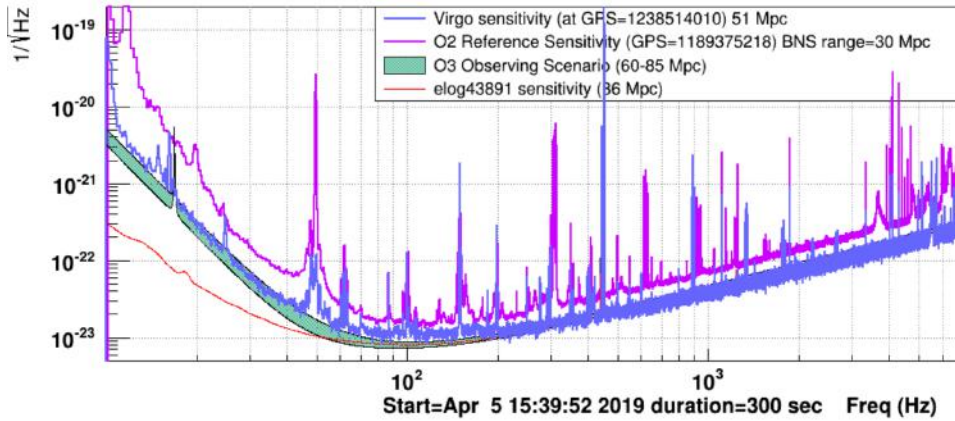


Figure 1.13: AdV strain sensitivity at the beginning of O3 (blue curve, 51 Mpc). The broadband improvement over the previous O2 sensitivity (purple curve, 30 Mpc) is evident, but it not sufficient to comply with the predefined O3 observing scenario (green band, 60-85 Mpc).

But the sensitivity of a GW detector is also a function of the properties of the source of interest. Therefore, it is worth introducing another quantity indicative of the detector sensitivity: the horizon distance d_h . It is defined as the maximum luminosity distance out to which a typical BNS inspiral event (NSs with $M = 1.4M_\odot$, optimally oriented and located in the sky) can be detected by the interferometer with a SNR equal to 8 [103, 104].




Figure 1.14: Binary neutron star averaged horizon (range) from 01 April to 01 October 2019, which is the first half of the third observing run (O3a). The variations are mainly due to weather conditions and control noises. The average BNS range is equal to 45.5 Mpc, which corresponds to 537 Mpc for $30M_\odot$ BBHs.

Figure 1.14 shows the BNS averaged horizon (range) during the first part of the third observing run (April to October 2019). Both the sensitivity and the observing horizon are regularly computed in near real-time, making it possible to monitor the noise status and act quickly in case of problems. The instantaneous output of the AdV detector, called h reconstructed, is then computed from the dark fringe signal and the calibration.

1.2.4 Noise sources

An interferometric detector is influenced by a number of different causes that impact on its sensitivity, as explained in the previous section. These undesirable signals take the name of noises. In this section, I will describe the main sources of noise encountered in a GW interferometer, referring to [105] for a more comprehensive coverage. Several designations are possible to group these noises. With regard to their origin, we can distinguish between environmental, technical or fundamental noises. Another useful classification takes into account the level of stationarity. Then, we could have low-frequency or high-frequency noises, depending on the frequency band in which they are dominant. Finally we could sort them on the basis of the effect they produce, so that we have displacement noises (effect in the mirror motion) or sensing noises (effect in the detection channels).

For instance, the low-frequency band (1-100 Hz) is limited by a lot of disturbances, for the most part from displacement noises. This noise class is faced both by improving existing facilities (e.g. minor upgrades, off-line subtractions) or with a complete redesign of the infrastructure.

Before listing all these noise sources and characterizing them, according to the above mentioned distinctions, first I want to better explain the heuristic definitions of fundamental and technical noises. Fundamental noises are those that demand a heavy infrastructure redesign or new technological achievements, in order to suppress them. These are thermal, shot and seismic noises. Technical noises require instead minor investments and they are mainly related to scattered light or electronic issues [92]. To understand more fully, consider an ideal interferometer, in vacuum space, zero kelvin temperature, with fluctuating test masses and a perfect monochromatic and stable laser source. This instrument would only be limited by the interaction between light and matter (quantum noise). In the real world, the detector is built on the ground, which is subject to seismic vibrations and density fluctuations (seismic and newtonian noises), at ambient temperature (thermal noise), not in vacuum space (gas excess noise). Furthermore we have to hold the instrument in a particularly stable regime (control noise), we have to inject and detect light (electronic and scattered light noises), not to mention more technical details of how and where the instrument is built and works, which bring with them a variety of extra noises (environmental noises, oscillator phase noise etc...). All together, they compose the detector noise budget, evaluated from the theoretical knowledge, that must overlap with the actual sensitivity. Otherwise, the models have to be refined or some effects are simply hard to account for. It is crucial to assist the theoretical models with careful experimental measurements, in order to validate them. Seen in this light, my PhD work aim to find a scalable solution to some of the most elusive noise sources that occur in the low-frequency range.

In figure 1.15, I show the projected contributions of the main noise sources affecting the design sensitivity of AdV. As we can see, noise amplitudes are usually expressed by their equivalent spectral densities, and therefore calculated in terms of the strain \tilde{h} : it is defined as the amplitude that an incoming GW must have in order to produce the same output as an internal noise.

environmental: seismic and Newtonian instrumental: scattered light fundamental: quantum low-frequency: seismic, Newtonian and thermal higher frequencies: quantum




Figure 1.15: AdV noise budget and the projected design sensitivity curve for a input laser power of 125 W (black plain curve). Credits: The Virgo Collaboration.

Shot noise

Shot noise directly originates from the discrete nature of the laser light, hence it is a fundamental noise [106]. It is also a sensing noise, since it affects the detector output channel. The equivalent spectral density of the shot noise \tilde{h}_{shot} is frequency independent (white noise):

$$\tilde{h}_{shot} = \frac{1}{L} \sqrt{\frac{\lambda \hbar c}{4\pi P_{in}}}, \quad (1.29)$$

where P_{in} is the power incident on the beam-splitter. However, a more accurate computation which considers the effect of the Fabry-Perot cavities gives

$$\tilde{h}_{shot}^{FP}(f) = \frac{1}{8FL} \left\{ \frac{4\pi\lambda\hbar c}{\eta G_{pr} P_{in}} \left[1 + \left(\frac{f}{f_c} \right)^2 \right] \right\}^{1/2} \quad (1.30)$$

where η is the photodetector efficiency, F is the finesse, f_c is the cut-off frequency of the cavities and G_{pr} is the PR gain factor. Now it is evident that the shot noise is reduced by increasing F , L and the circulating power. Moreover, the shot noise is a limiting factor at high frequencies for GW interferometers, since at lower frequencies it is outclassed by the radiation pressure noise.

Radiation pressure noise

This noise arises from the pressure exerted by the incident EM radiation on the Virgo test masses. It originates from the quantum nature of light like the shot noise

and so it is a fundamental noise too, as well as a displacement noise. Its dependence on frequency and laser power is inverse to that of shot noise [106]:

$$\tilde{h}_{rad}^{FP}(f) = \frac{4F}{ML} \left[\frac{\hbar G_{pr} P_{in}}{\pi^5 \lambda c} \cdot \frac{1}{f^4} \cdot \frac{1}{1 + \left(\frac{f}{f_c}\right)^2} \right]^{1/2} \quad (1.31)$$

where M is the mass of the mirrors. Radiation pressure noise contribution is crucial at low-frequencies, as encrypted on the frequency dependence $\propto 1/f^2$, and at high laser input power. Therefore, its reduction is closely interconnected with that of shot noise, forcing to a trade-off. The optimal condition is obtained with a perfect balance between the two spectral density contributions, leading to the so-called Standard Quantum Limit (SQL) [107] (see figure 1.16). The current mitigation actions against quantum noise are achieved by means of squeezed coherent state of light [87, 108] and signal recycling [90].

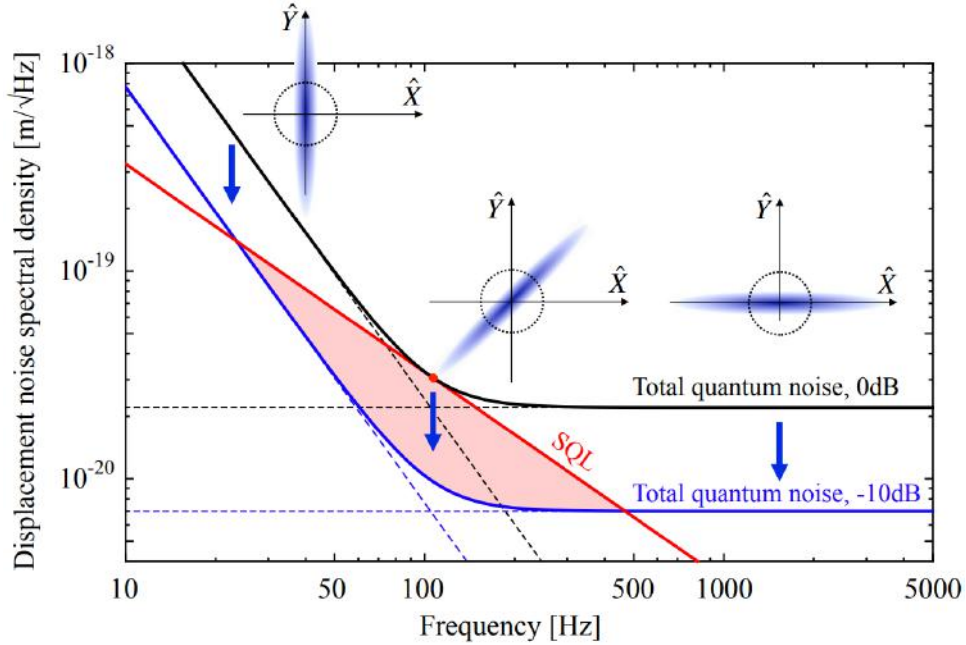


Figure 1.16: Squeezed light injection to surpass the SQL. Here, the squeeze angle is optimized for all frequencies resulting in a broadband quantum noise reduction. Dashed horizontal lines represent the shot noise, while dashed straight lines with negative slope represent the radiation pressure noise. The calculations use 10 dB of squeezing, a conventional Michelson interferometer with neither arm resonators nor signal recycling, a light power at the beam splitter of 1 MW at a wavelength of 1550 nm, and mirror masses of 1 kg [108].

Seismic noise

The combination of ground motion, tidal effects (tidal waves, wind etc.) and nearby anthropogenic activity cause a displacement seismic noise. Its impact on ground-based GW detectors is relevant only at low-frequencies (below 10 Hz), with a frequency dependence of $\propto 1/f^2$. Different seismic noise sources can be recognized, producing effects at various frequencies:

- Short-period seismic noise ($1 < f < 50$ Hz) is generated by meteorological phenomena, such as the wind friction over rough terrain, trees or vegetation and vibrating man-made objects. The dominant part comes from anthropogenic activities and machines (e.g. rotating machinery, road traffic, railways etc.). This activity, with clear day and night time variations, produces mainly surface (Rayleigh) waves that propagate along the surface (1-50 Hz). Therefore, it decreases with population density and depth, and it is related to the local soil geology.
- Micro-seismic noise ($0.1 < f < 0.8$ Hz) is related to the activity of oceans and seas.
- Long-period seismic noise ($f < 0.1$ Hz) is produced by barometric pressure fluctuations, which generates a background seismic noise due to the elastic response of the ground to these fluctuations [109]. This effect is unavoidable, but still it decreases with depth. A second source of long-period seismic noise is due to the Earth body tide, which is particularly evident at a frequency of $\approx 20 \mu\text{Hz}$.

Almost all the seismic noise effects can be reduced by means of two complementary approaches:

- Underground facilities, such as KAGRA [110] and ET [38, 111]. The Sardinia candidate site for the ET has already been the subject of seismic underground studies [112].
- Suitable seismic attenuation systems, such as the Virgo multi-stage inverted pendulums (see section 1.2.2 for more details) or the LIGO active actuators [84, 113].

All the mirrors and the main optics of the advanced interferometers are currently suspended to a certain kind of seismic attenuators and placed in ultra-high vacuum, but yet all the injection and detection components are situated on benches in air, standing on the ground. In order to make a distinction between the mirror seismic motion, which is filtered by the SA, and the one caused by the benches and the other detector components, we usually refer to the latter as acoustic and scattered light noises.

- Acoustic noise. It is important mainly from some Hz to some thousands Hz and one of its main origin is real acoustic signal close to the apparatus. A way to study the coupling of acoustic noise is to inject loud signals (e.g. with loudspeakers or shakers) and look at the effect on the dark fringe. However, the only way to identify directly the noise source between all the components in the interferometer is to carry out tapping tests on single parts. Projections of the environmental noise in the Virgo laboratories have been computed performing white noise injections of acoustic and seismic noise (I will use the same procedure with magnetic noise in chapter 2.2.4). Results of this analysis show that the estimated contribution from the acoustic noise is well below the actual sensitivity. Nevertheless, several interventions were required in the past: acoustic isolation enclosures were placed around many optical benches; several optical mounts were replaced with more rigid one.

- Scattered light. The light scattered by the mirrors is mainly due to small-scale defects on its surface (e.g. roughness), resulting in diffraction at large angles, and large-scale imperfections (e.g. due to polishing), causing scattering at low angles [114]. Also residual gas molecules can prompt some photons of the laser beam to scatter, bounce on the vibrating walls of the vacuum chamber and eventually recombine with the main beam with a different phase. This phase noise mostly affects the low frequencies, since it is linked to the surrounding seismic activity [115]. Noise from diffused and scattered light is very difficult to identify with coherent analysis, due to the highly non-stationary and non-linear coupling, which spoils the coherence between dark fringe and sensors, even if the noise is dominant. Scattered light noise can in principle be reduced by installing baffles made of absorbing material inside the vacuum tubes, close to the mirrors and the optical mounts. Still, some of the noise lines visible in the sensitivity between 100 and 1000 Hz are likely due to scattered light.

Newtonian noise

Current and future GW antennas are designed to be sensitive down to quite low frequencies (i.e. a few Hz). In this framework, random gravitational forces represent a noise source which is going to limit the useful bandpass of the detectors. Their consequence is a net TM displacement which bypasses the entire seismic attenuation.

The effect of these random gravitational forces has been studied since when the laser interferometry as a tool for GW detection was only a design draft, laid out by Rainer Weiss in 1972 [116]. At that time, gravity perturbations from seismic fields were recognized as a potential noise source, which, from now on, I will call Seismic Newtonian Noise (SNN). Semi-accurate analytical models of the effect produced by the SNN to the test mass of an interferometric detector appeared one decade later with Peter Saulson calculations [117]. After another decade, the two groups from LIGO and Virgo communities independently managed to square the circle in 1998 [118, 119]. In their calculations, they also took into account the wave nature of the seismic field, producing the accurate predictions of Newtonian noise that this day researchers still continue to rely on.

Gravity fluctuations are predominantly from two sources: density perturbations in the atmosphere or from seismic fields [117–122]. Seismic surface fields are predicted to dominate the newtonian noise contribution [123], although a recent work showed that the atmosphere can be relevant as well [124].

In [123], it was pointed out that the average stationary component of newtonian noise is likely to lie below the instrumental noise of advanced detectors, but, in the case of a particularly high level of non-stationary environmental noise, it can be dominant. An example of non-stationary noise in the newtonian noise band coming from strong seismic transients could be due to local sources such as wheeled vehicles. Hence, since model errors might be significant, a cancellation system will be required. The paper [121] suggest that a proper system able to mitigate newtonian noise should monitor the environmental field with sensor arrays.

In particular, SNN can be limited by going underground (i.e. for KAGRA and ET) or with a smart design of the infrastructure surrounding the test masses. The SNN is the higher impending newtonian noise component to limit the detector future sensitivity and a significant part of my work focused on it (detailed description in chapter 3).

Concerning the atmospheric newtonian noise, models are either associated with infrasound waves, or with quasi-stationary temperature fields advected by wind. For surface detectors, atmospheric newtonian noise starts to be significant below 10 Hz according to these models, but this is only true for sound spectra representative of remote places. Sound levels in laboratory buildings can exceed sound levels at remote locations by more than an order of magnitude even without heavy machinery running, e.g., due to ventilation systems or pumps. In this case, sound newtonian noise can become significant above 10 Hz already, and become a potentially limiting noise contribution in advanced detectors. Comparing with SNN (see figure 1.17), we see that atmospheric spectra are steeper and therefore potentially the dominating gravity perturbation in low-frequency detectors. On the contrary, SNN contribution is already above the safety level for technical noises in AdV and therefore it demands a prompt intervention.

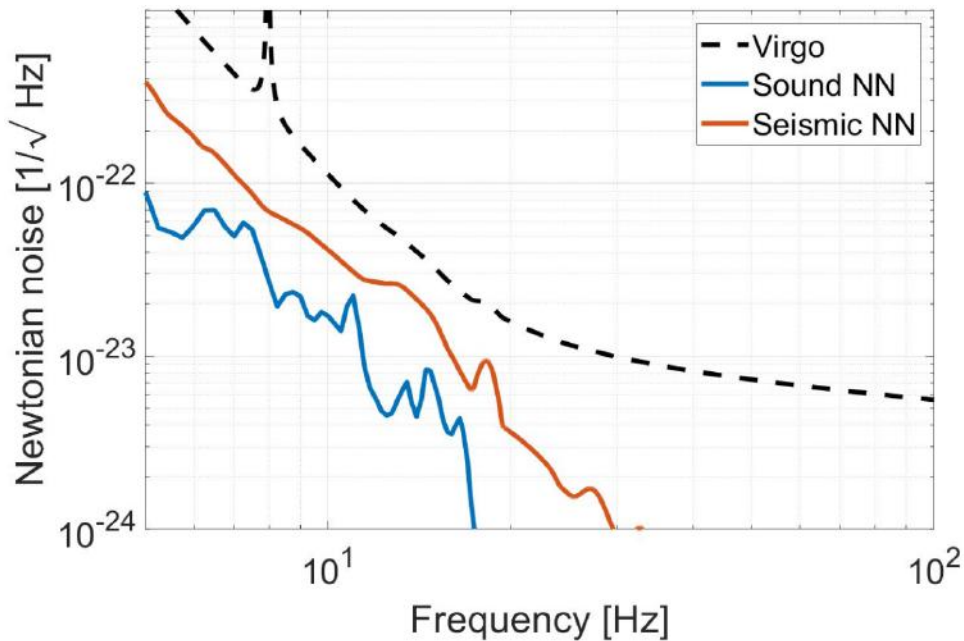


Figure 1.17: Seismic (red line) and atmospheric (blue line) newtonian strain noise contributions to the Advanced Virgo design sensitivity (dashed black line), as a function of the frequency. The figure is reproduced from [125].

Finally, gravity perturbations can also be produced by objects of finite size (in uniform motion or oscillating objects or rotating objects such as exhaust fans or motors) or objects that experience deformations (vibrations of vacuum chambers [126]). A first analytical study of gravity perturbations from objects was performed by Thorne and Winstein who investigated disturbances of anthropogenic origin [127]). So far, none of the potential sources turned out to be relevant. [120] presents a rule of thumb that make it possible to estimate the relevance of perturbations from an object “by eye” before carrying out any calculation. We learn that a solid object with 1 m diameter at a distance of 5 m, oscillating with amplitude equal to seismic amplitudes and equal density to the ground, would produce newtonian noise, which is about a factor 100 weaker than SNN.

Thermal noise

The thermal noise is a fundamental limit to the interferometric detector sensitivity and causes a displacement noise in the middle frequency range ($\approx 30 - 500$ Hz). As the definition suggests, it affects the mechanical parts of the TM suspension and the mirror itself, which all have a finite temperature. This fact is generally known as Brownian motion, a process driven by thermally-induced molecular random collisions, as Einstein exquisitely described in 1905 [128]. We can distinguish 5 different types of noises, related to thermal effects:

- Suspension thermal noise ($\approx 10 - 50$ Hz). It acts on the position of each TM and is caused by the thermal fluctuations in the mechanical suspension. Each stage of the suspension system contributes to the overall thermal noise. Its frequency dependence is $\propto 1/f^{5/2}$ [129].
- Mirror Brownian thermal noise ($\approx 50 - 100$ Hz). It is defined by internal vibrational modes of the TM bulk, with a spectral strain contribution $\propto 1/f^{1/2}$. Other temperature-dependent noise sources in the mirror are the shape deformation due to the so-called thermal lensing, and the thermodynamical fluctuations in the bulk [130].
- Coating Brownian thermal noise. It is due to mechanical losses of the multi-layer dielectric coatings. The dissipation of the commonly used silica/tantala (SiO_2/Ta_2O_5) coatings is found to be several orders of magnitude higher than the dissipation of bulk fused silica [131].
- Coating thermo-optic noise. Equilibrium heat fluctuations in the TM surface cause variations in the coating physical parameters: a change of the thermal expansion coefficient α leads to thermo-elastic noise; a change of the thermo-refractive coefficient β leads to thermo-refractive noise. Thermo-optic noise is the coherent sum of thermo-elastic and thermo-refractive contributions [132].
- Residual gas. Residual molecules in the vacuum tubes are also a second order source of thermal noise, produced by the momentum transfer between the mirror suspensions and the residual gas molecules in vacuum regime, causing a viscous damping.

Possible mitigation techniques involve the use of low-loss materials for the mirrors and the coatings and high-Q materials for the suspensions, to shrink the frequency range of the violin modes, at the cost of increasing the noise in that same range. In third generation interferometric detectors, thermally induced noises will be addressed with cryogenic techniques, as KAGRA already accomplished [110].

Mirror charging noise

Electrostatic charge on the mirror surface can form due to high speed charged particles flowing out of the pumping system port into the vacuum chambers, especially when large ion pumps without a proper shield are used [133]. This issue was first noticed at LIGO and solved by adding proper shields to prevent a direct line of sight between the pump and the chamber walls, nearby optical elements. Recently, the risk of electrostatic charging of mirrors has become a concern also for Virgo [134,135],

with a net effect on the sensitivity $\propto 1/f^{5/2}$, up to 100 Hz. Mirror charging noise is dealt with non-invasive upgrades of the Virgo infrastructure. Since the charge is likely non-uniformly distributed (patches), a solution would be to generate a bath of ions (negative and positive) in a controlled environment (i.e. diffusion regime). Also, different mechanical shielding solutions have been evaluated and discussed, in order to protect the critical optics. In addition, ion pumps with/without shields are presently under test, in collaboration with experts from major manufacturing company.

Magnetic noise

The main coupling channel between environmental magnetic activity and the AdV strain channel has been identified in coil-magnet pairs acting on the TM and on the upper stages of the suspension system for accurate positioning (see sec 1.2.2 for details about the overall assembly around the TM).

Voice-coil actuators are very sensitive to the presence of the environmental magnetic field, whose fluctuations engage directly with the magnets and induce a TM displacement noise. Although magnets are glued on the mirror surface in an anti-parallel cross configuration (to compensate for the effect of a uniform external magnetic field), in practice, the external field interacts with all metal structures around the mirrors. These, in turn, cause eddy currents and distort the magnetic field with uneven gradients.

This type of coupling was already observed during the first Virgo Scientific Run (2008) and by then they had been able to reduce the magnetic noise by replacing the magnets by five times smaller ones [136].

Magnetic coupling may also account for the correlated noise caused by global Schumann resonances. A careful study of this correlated signal between distant detectors is crucial in order to allow the observation of a Stochastic Gravitational-Wave Background (SGWB) [55, 58].

In addition, magnetic field glitches could also be taken for transient GW signals, as reported in [137].

A tentative estimation of this magnetic coupling dates back to 2014 but it was suffering from large uncertainty [138]. Therefore, the magnetic noise issue is still open.

In chapter 2 I will present the study on the global (sec. 2.1) and local (sec. 2.2) magnetic disturbances, together with a magnetic coupling modeling (sec. 2.3), the impact of MN on the detector sensitivity (sec. 2.4) and a proposal for its mitigation [139, 140] (sec. 2.5).

Technical noises

Technical noise is a general term to refer to a number of noises ascribed to the interferometer subsystems [141]. These are: control noises due to electronics and feedback loops (photodetector electronic noise, length control noise, angular control noise, mirror and suspension systems control noises); environmental disturbances caused by EM field contamination by power supplies and amplifiers (related to the operation of monitoring probes like temperature, pressure and humidity sensors, seismometers, magnetometers etc.); noises linked to the injection system (laser intensity and beam jitter noises) and to the detection system (OMC length noise);

phase noise in photodiode digital demodulation; noises connected to the detector calibration; noises related to the vacuum system or to the infrastructure (HVAC, IMC heater and chiller noise), to the TCS or the squeezing apparatus.

Noise spectral lines

Here I talk about narrow-band frequency disturbances, the so-called spectral lines, whose presence in the strain spectrum can reduce the sensitivity of CW searches. Their origin is often well-known: some are associated with non-removable resonances of detector components (mirror “drum” modes, suspension “violin” modes), others with constant signals for calibration and control purposes. However, a series of lines have an unknown origin or simply they are not stationary, since they fluctuate in amplitude and/or frequency. Figure 1.18 shows a collection of non-stationary and frequency drifting spectral lines, some of them associated to service infrastructure needed for the detector operation. Examples are machines for air cleaning and conditioning of the experimental facilities, vacuum pumps, cooling fluid pumps, cooling fans for electronic devices, digital clocks and the mains power supply.

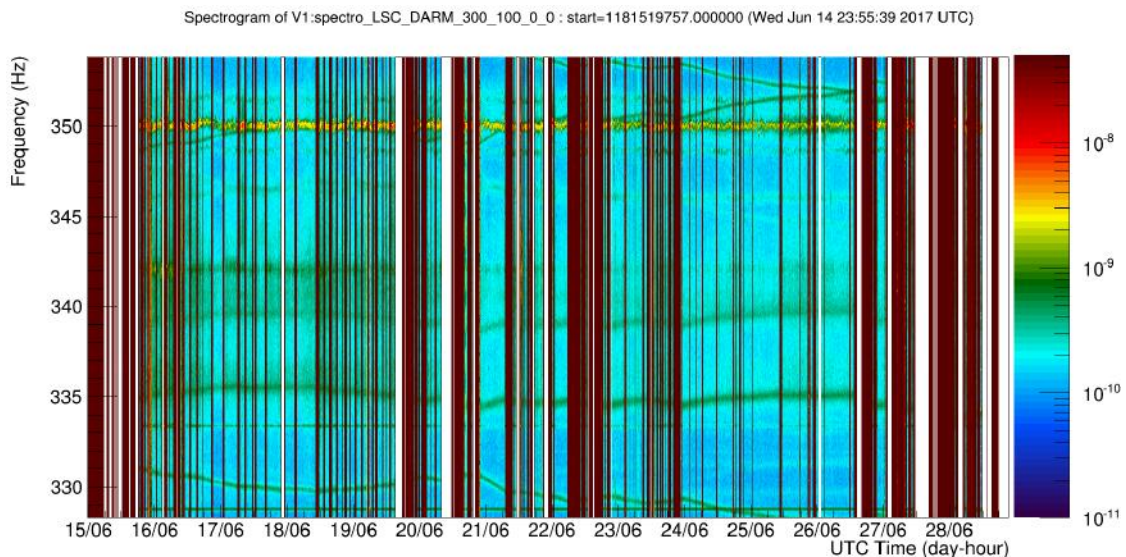


Figure 1.18: Time-frequency map of Virgo strain channel at a time when a great variety of noise spectral lines were present, in the frequency range of 328-354 Hz. For instance, the line at 342 Hz shows a non-stationary behavior, while some other lines move in frequency by up to 5 Hz during the entire 2-weeks interval.

Non-stationary lines become particularly troublesome as soon as they trigger transient events. With these conditions, it is crucial to track the spectral lines and monitor their behavior. At Virgo, a variety of line hunting and mitigation tools are used, such as the NoEMi (Noise Frequency Event Miner) algorithm [142], dedicated to the quasi-real time analysis of Virgo data. In particular, NoEMi looks for coincidences between noise lines in auxiliary channels and the dark fringe. Other tools are then dedicated to spectral noise [143], such as BruCo (brute-force coherence) for coherence analysis and a series of scripts for the time-frequency visualization of the lines (spectrograms).

Transient noise

A transient noise, also called “glitch”, is often generated by a sudden environmental disturbance that couples with the detector strain signal. The coupling path of a transient noise can be established by monitoring the auxiliary channels and their mutual correlations (glitch-to-glitch identification). “Omicron” is the main algorithm designed to detect and characterize this kind of noise in all AdV channels [144]. In figure 1.19, I present six families of typical glitches we can find in the strain channel. A glitch family is identifiable by its unique morphology, as seen in a spectrogram.




Figure 1.19: “Omicron” time-frequency spectrograms of six different glitches seen in the Virgo strain channel. The insets show the glitches in the auxiliary channel, when identified. From top-left to bottom-right: a 50 Hz power-line glitch also detected by the magnetometers; a series of glitches caused by scattered light induced by seismic activity; a glitch caused by a TCS instability; an airplane event with a clear Doppler effect; a glitch in the laser stabilization loop; a seismic event up-converted to higher frequencies. Figure reproduced from [54].

Correlated noise

With the term correlated noise are intended many disturbances that may affect more than a detector, almost simultaneously. This is a nasty issue, as the effect of such a noise could potentially imitate a GW event and would not be captured by time shift based searches. The most influencing correlated noise class is the one attributable to EM global events like lightnings, solar events, solar-wind driven noise (solar radio flares and currents of charged particles) and radio-frequency (RF) communications. In particular, lightning strikes, which occur tens of times per second globally, can excite magnetic Schumann resonances [145, 146] (see chapter 2.1). Although their amplitudes are of the order of pT, hence not enough to trigger a discernible signal in $h(t)$, they represent a correlated noise between LIGO and Virgo, which can rise out of the background during a coherent search. Indeed, the Schumann resonances could be detected by low-noisy magnetometers and therefore they can be subsequently subtracted from the strain data, using proper filters. Moreover, this correlated

noise threatens to contaminate the observation of a SGWB. In the work [57], the first efforts to eliminate global correlated magnetic noise using Wiener filtering are reported, showing approximately a factor of two reduction in the coherence between magnetometers on different continents. In chapter 2.1.3, I will present my work on dedicated magnetometer measurements at the Virgo site, which leads to the first results for correlated noise subtraction using data from distant GW detector sites. Concerning EM signals outside the GW detection band, LIGO and Virgo uses RF modulation for the optical cavity control, which could be influenced by RF background noise. However, this coupling appears to be of no concern, even at its maximum for the 9 and 45 MHz modulation frequencies. Finally, the EM radiation from cosmic ray showers is also monitored, but a coincident event between two distant detectors is highly unlikely, given that the flux drops to zero within roughly 10 km from the original cosmic ray decade [147].

Noise hunting

It is now clear that transient effects can reduce the sensitivity of GW searches. The detector signal can be disturbed by a large variety of noise sources originating from within the detector or from its environment. The path that connects the noise source to the detector output affects the characteristics of the noise. The process called “noise hunting” consists of tracking down each noise source and understanding the conversion mechanisms which occur between the source and the interferometer. To achieve this task, the GW detectors are equipped with hundreds of sensors, including microphones, seismometers, magnetometers, photo-diodes, current and voltage monitors, thermometers and cameras. The signals from these auxiliary channels are used to monitor external disturbances to help determine whether a candidate event was produced by a GW or by an instrumental artifact. The AdV noise hunting process can be summarized in the following way [54]:

- Event identification: transient or spectral noises, or a family of events with similar properties.
- Correlation with some unusual activity (human intrusions, earthquakes, thunderstorms, etc.).
- Check for the time matching between the event and external scheduled events, such as the stop/start of infrastructure machinery.
- Study of time coincidence with an event in one or several auxiliary channels (statistical algorithms are used to quantify correlations). Auxiliary signals are monitored in quasi-real time.
- Sometimes, the previous studies cannot differentiate whether the noise has been identified at its source or somewhere along its propagation. Experiments are performed to understand how the noise couples with the detector. For example, one can artificially inject noise in a hardware component and study the response of the detector [148]. Another possibility is to switch off a potential noise source to see if the noise disappears.
- If a noise source is identified, we try to eliminate or reduce it. If it is not possible, we try to reduce the coupling to the detector.

The Physical Environment Monitor (PEM) sensors are a very important set of auxiliary channels, which monitor the local surroundings for potential disturbances that may affect the GW strain data [149]. Artificial noise injections have been crucial during my PhD work. They investigate the relationship between environmental noise as witnessed by a PEM sensor array and the GW strain signal $h(t)$. An example is the one in figure 1.20, which illustrates a magnetic field injection test at the LIGO-Hanford detector. It measured the magnetic field coupling to $h(t)$ as well as the response of the local magnetic sensor to the injected field. The frequency-dependent coupling between noise and $h(t)$ can be calculated from these measurements and used to accurately predict the response of $h(t)$ to the presence of the noise.

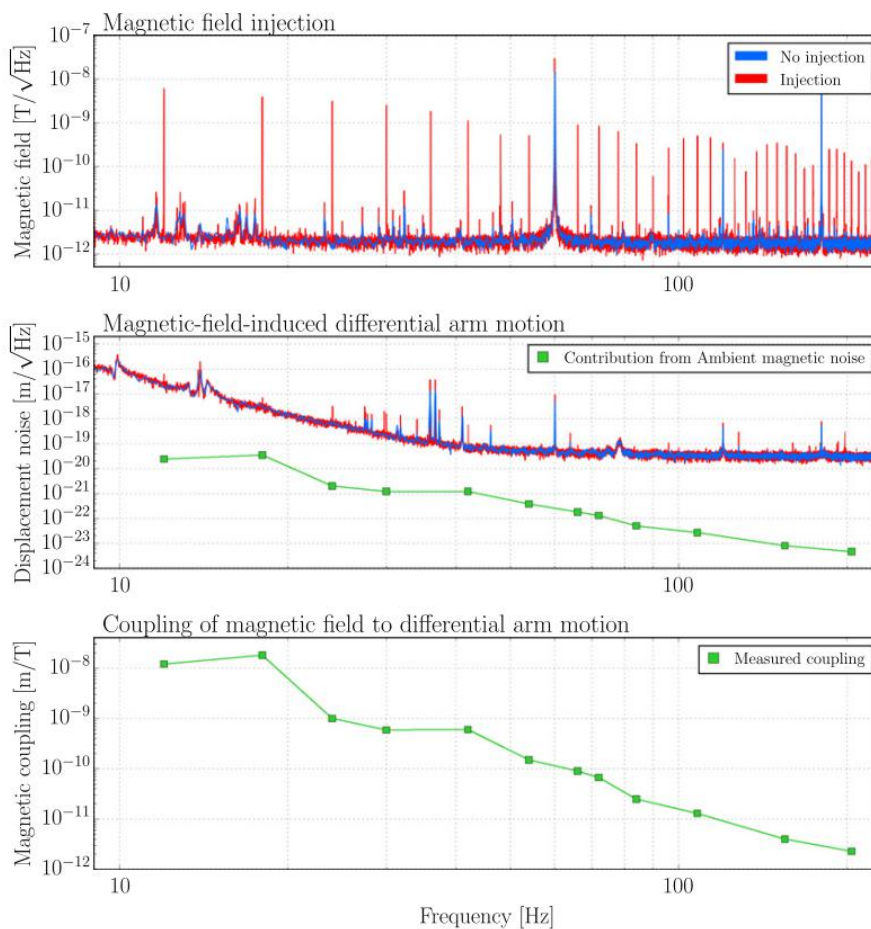


Figure 1.20: Magnetic noise coupling example at LIGO-Hanford. The top panel shows the output of a magnetic sensor during the injection of a series of single frequency oscillating magnetic fields at 6 Hz intervals (in red) and at a nominally quiet time (in blue). The middle panel shows the noise contribution to the strain. The bottom panel shows the magnetic noise transfer function. This figure only shows data for one location, but similar injections are repeated at all locations where magnetic coupling might be of concern. The figure is reproduced from [150].

1.3 Observing Runs and Detections

The AdL detectors' first observing run (O1) lasted about four months, from September 2015 to January 2016, and collected 51.5 days of data. In this dataset, they found two clear detections and one tentative signal (85% confidence), all three coming from BBH mergers [151].




Figure 1.21: Left: BNS range for each instrument during O2; Virgo joined O2 in week 34. Right: Amplitude spectral density of the total strain noise of the Virgo, LHO and LLO detectors. The curves are representative of the best performance of each detector during O2.

The second observing run (O2) began on the 30th of November, 2016 and went on (with some downtime for maintenance) until the end of summer 2017 [151].

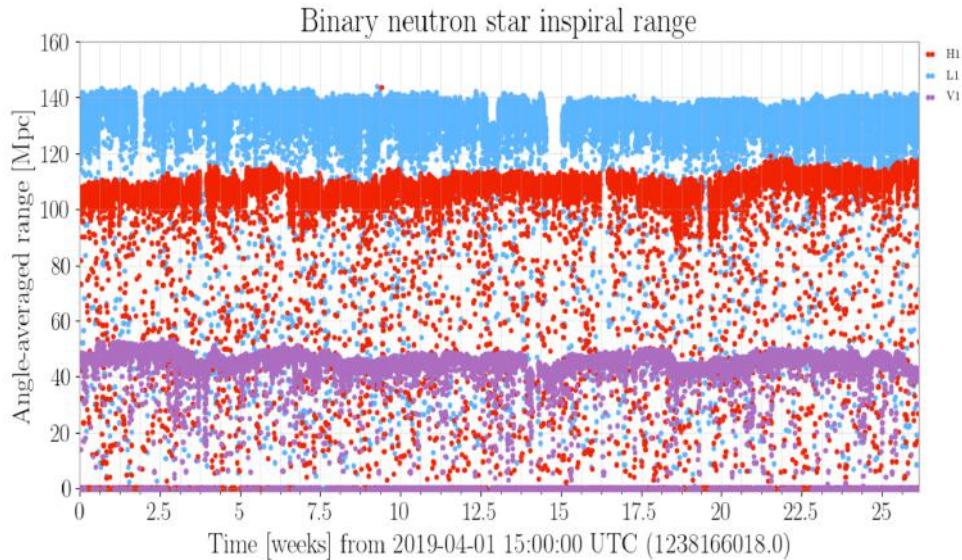


Figure 1.22: BNS range for the three detectors during O3a. A significant improvement ($> 90\%$) with respect to the average O2 sensitivity (26 Mpc) is obtained with AdV during O3a. Here, the laser power amounts to 18 W, compared to 10 W in O2.

This was the second observing run of AdL, and the first one of Advanced Virgo, which joined O2 on the 1st of August, 2017 and collected 20 days of clean data at a BNS sensitivity range of nearly 30 Mpc. Figure 1.21 shows the BNS ranges of

AdL and AdV instruments over the course of O2, and the representative amplitude spectral density plots of the total strain noise for each detector. During O2 the AdL network duty factor was of $\approx 45\%$, resulting in 118 days of data suitable for coincident analysis [66]. Of this data 15 days were collected in coincident operation with AdV, which after joining O2 operated with a duty factor of $\approx 80\%$. Observations in O2 include 7 BBH mergers, as well as the first BNS merger (GW170817) observed in gravitational waves [152]. As these neutron stars spiraled together, they emitted GWs that were detectable for about 100 seconds; when they collided, a flash of light in the form of gamma rays was emitted and seen on Earth about two seconds after the GWs. In the days and weeks following the merging, other forms of electromagnetic radiation (X-ray, ultraviolet, optical, infrared, and radio waves) were detected. Across the entirety of O1 and O2, the GW170817 remains the event with the highest network SNR.




Figure 1.23: The planned sensitivity evolution and observing runs of the AdL, AdV, KAGRA and LIGO-India detectors over the coming years. The colored bars show the observing runs, with the expected sensitivities for future runs, and the achieved sensitivities in O1, O2 and O3 [153].

The AdV joining also enabled the first 3-detector observations of GWs. As the sky location is primarily determined by the differences in the times of arrival of the GW signal at the different detector sites, LIGO-Virgo coincident events have a vastly improved sky localization, which is crucial for EM follow-up campaigns [154, 155].

Finally the third scientific run (O3) started on April, 2019 and saw LIGO and Virgo improving further their sensitivity and KAGRA trying to join the network at the very end. O3 has been divided in two parts, separated by one month of commissioning in October 2019. The plot on figure 1.22 is referred exactly to the BNS range achieved by the 3-detector network in the first part of O3 (O3a). A comprehensive view of the three scientific data takings is shown in figure 1.23, together with the projected future sensitivities and timelines.

1.4 Multimessenger Astronomy

BNS mergers are expected to produce isotropic electromagnetic signals, which include (i) early optical and infrared emission, a so-called kilonova due to radioactive decay of rapid neutron-capture process (r-process) nuclei synthesized in dynamical and accretion-disk-wind ejecta during the merger; and (ii) delayed radio emission from the interaction of the merger ejecta with the ambient medium. The follow-up program established by the LIGO-Virgo collaboration led to the first joint detection of gravitational and electromagnetic radiation from a single source [156]: on 2017 August 17 an approximately 100 s long GW signal (GW170817) was followed by a short GRB (GRB170817A) and an optical transient (see figure 1.24).




Figure 1.24: Multimessenger detection of GW170817 and GRB170817A. Top: lightcurve for GRB170817A between 50 and 300 keV (the background estimate is overlaid in red) from the FGST. Bottom: time-frequency plot of GW170817, obtained by coherently combining LIGO-Hanford and LIGO-Livingston data. Figure reproduced from [157].

The source was detected across the electromagnetic spectrum in the X-ray, ultraviolet, optical, infrared, and radio bands over hours, days, and weeks. What follows is a brief chronicle of the detection. On 2017 August 17 12:41:06 UTC the Fermi Gamma-ray Space Telescope (FGST) localized a GRB. A notice was issued at 12:41:20 UTC announcing the detection. Approximately 6 minutes later, a GW candidate (later designated GW170817) was registered in low latency based on a single-detector analysis of the LIGO-Hanford data. The signal was consistent with a Binary NS coalescence with merger time 12:41:04 UTC, less than 2s before the GRB. A notice was issued at 13:08:16 UTC. Given the temporal coincidence with the GRB, a circular was issued at 13:21:42 UTC reporting that a highly significant GW candidate event consistent with a binary NS coalescence was associated with the time of the GRB. An extensive observing campaign was launched across the electromagnetic spectrum in response to the Fermi and LIGO-Virgo detections, and

especially the subsequent well-constrained, 3D LIGO-Virgo localization (see figure 1.25). A bright optical transient was discovered in the Galaxy NGC 4993 (at ~ 40 Mpc) by the 1M2H team (August 18 01:05 UTC) less than 11 hours after the merger.

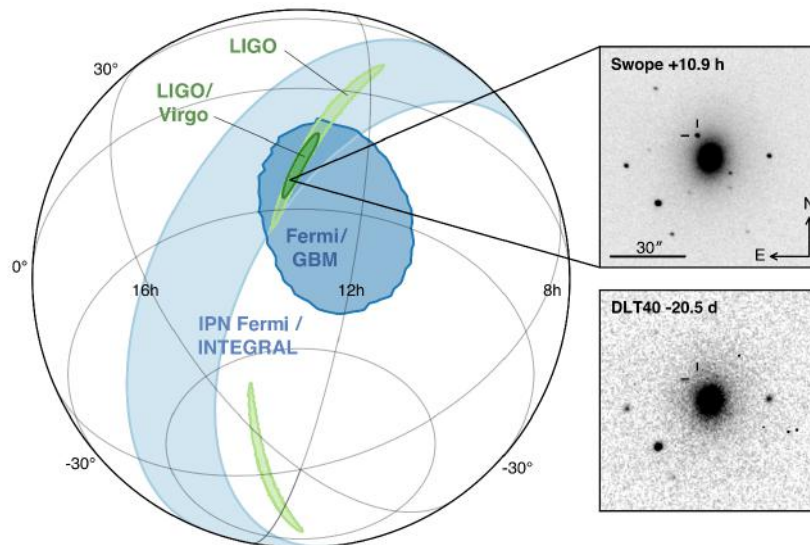


Figure 1.25: Localization of the GW, gamma-ray, and optical signals. The left panel shows an orthographic projection of the 90% credible regions from LIGO alone (190 deg^2 ; light green), the initial LIGO-Virgo localization (31 deg^2 ; dark green), IPN Fermi/INTEGRAL (light blue), and Fermi/GBM (dark blue). The inset shows the location of the apparent host galaxy NGC 4993 in the Swope optical discovery image at 10.9 hours after the merger (top right) and the DLT40 pre-discovery image from 20.5 days prior to merger (bottom right). The pointer marks the position of the transient in both images. Figure reproduced from [157].

1.5 Future prospects

The number of detected GW events is proportional to the observation time and to the cube of the observing horizon. This means that 1 day of data with a BNS range of 60 Mpc (projected O3) is equivalent to 125 days at 12 Mpc (initial-Virgo best). Hence improving the sensitivity is better than observing for a long time period. Figure 1.26 shows the observing scenario for the LIGO-Virgo-KAGRA network, with LIGO-India planning to contribute starting from 2024.

The sky localization capabilities of the second generation (2G) network at mid 2020s report that more than 60% of the sources will be localized within 10 deg^2 (see figure 1.27).

Moreover a set of upgrades capable of enhancing the performances of the current detectors, are currently being commissioned (AdV+ and AdL+), with a timeline of 4 year from now. The AdV+ project involves a two (three) phases upgrade strategy:

Phase 0 (completion of AdV): BNS range up to 120 Mpc

- High power laser (in-fiber 200 W laser system so to have 125 W after the Input Mode Cleaner)

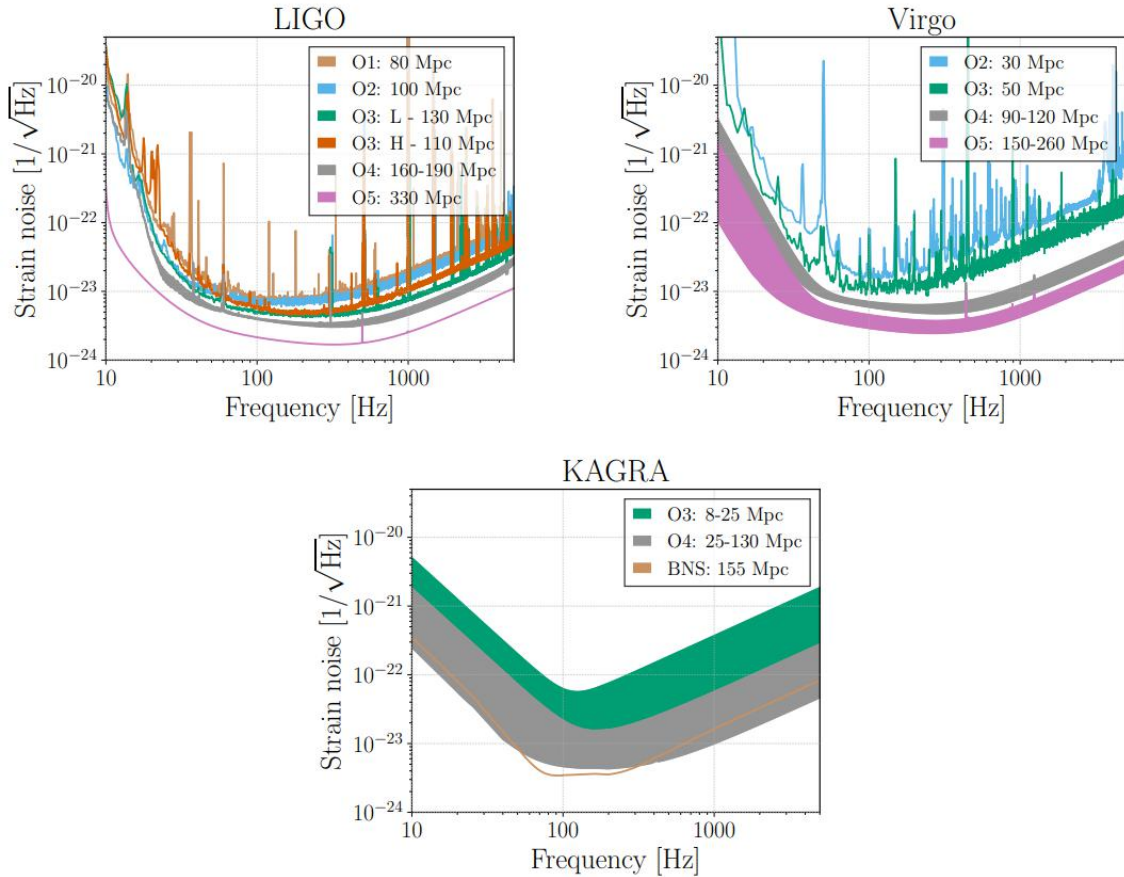


Figure 1.26: AdL (top left), AdV (top right) and KAGRA (bottom) target strain sensitivities as a function of frequency. The quoted range is for a $1.4M_{\odot}$ BNS merger. The BNS range for past and future observing runs is shown [153].

– Signal Recycling

Phase 1: BNS range up to 160 Mpc

- Frequency dependent squeezing (Squeezed light source and filter cavity)
- Newtonian noise cancellation (Seismic sensor networks, SNN modeling, HVAC modification)

Phase 2: BNS range up to 260 (300) Mpc

- New, larger mirrors (diameter: 550 mm; thickness: 200 mm; mass: 105 kg)
- New coatings (Flatness < 0.5 nm rms over central 160 mm of mirrors; $Ti : Ta_2O_5$ and SiO_2 stacks with optical absorption about 0.3 ppm) for a factor 3 of coating thermal noise reduction
- General upgrade of vacuum system, infrastructure, payloads and suspensions

So this is what it concerns the 2G detectors. Next we are thinking of new infrastructures of third generation (3G) to reach the early universe, with a one




Figure 1.27: 5-detector network localization accuracy for face-on BNS systems (detectable signals with a network SNR exceeding 12) with all advanced detectors at final design sensitivity. The ellipses show 90% confidence localization areas [158].

order of magnitude gain in sensitivity, compared to the previous generation. This facilities would be the European Einstein Telescope (ET) [111, 159], which would implement an assembly of 6 co-located interferometers (each 10 km long), and the US Cosmic Explorer project (CE) for a 40 km interferometer [160]. From a scientific point of view, the 3G detectors would answer to Fundamental physics and Gravity questions, together to astrophysics and multimessenger astronomy and cosmology ones. From a technological point of view, the xylophone concept of the ET (figure 1.28) would overcome the main contradicting requirement in a general technical detector design: to reduce shot noise we have to increase the light power, which in turn will reduce the sensitivity at low frequencies due to higher radiation pressure noise.




Figure 1.28: The ET xylophone design, in which one GW detector is composed of two individual interferometers: A low-power cryogenic low-frequency interferometer and a high-power room temperature high-frequency one. The full observatory will consist of three detectors with a 60° opening angle, arranged in a triangular shape.

For that we would need cryogenic mirrors to damp thermal effects, but due to residual absorption it is hard to combine cryogenics with high power lasers. The

solution chosen for ET is the conservative one with a 2-band xylophone: a low-power, cryogenic low-frequency detector, limited by Newtonian and Quantum noises (figure 1.29 left) and a high-power, room-temperature high-frequency one, limited by thermal noises (figure 1.29 right). The final infrastructure would consist of three of these detectors couples, where the second xylophone to fully resolve polarization and the third one for redundancy and null-streams. In figure 1.29 I also show the combined ET projected sensitivity.

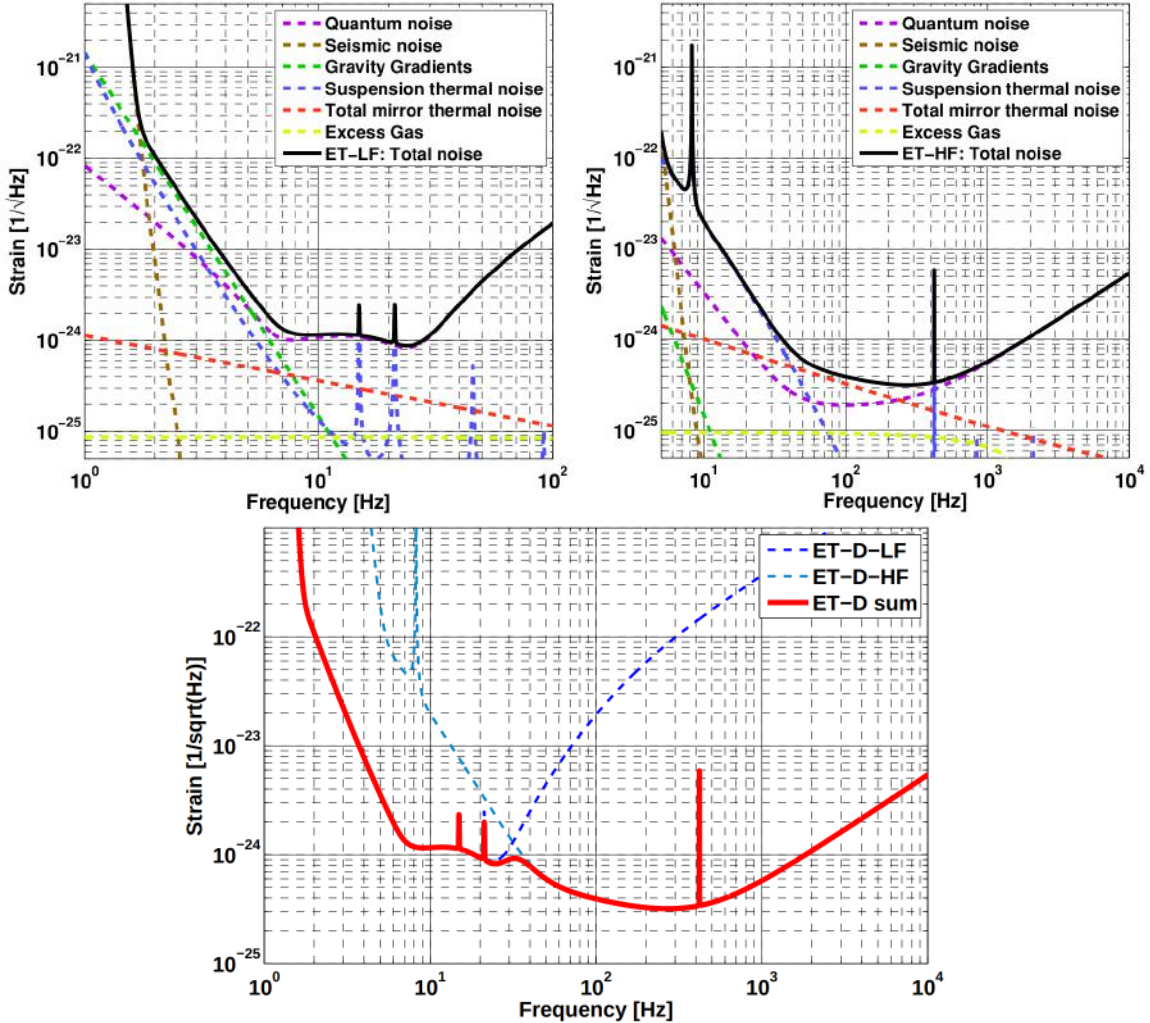


Figure 1.29: Top: noise budgets for the ET low- and high-frequency interferometers. The former is able to suppress the 2 thermal noise components and hence is limited by Newtonian and quantum noises; the latter manage to decrease quantum effects but it is limited at low-mid frequencies by the thermal ones. Bottom: combined sensitivity of the 2 quasi-complementary ET infrastructures.

On the contrary, The CE will focus on vastly increasing the length of the 2 main optical cavities and on a cooling strategy based on radiation instead of conduction through the suspension fibers. On paper CE might push further the limits of the sensitivity at mid-high frequencies, as quantum and thermal noises scales with the arm length (figure 1.30).

In summary the ground-based 3G GW astronomy might be dominated by these 2 large facilities starting from the third decade of this century. Until then, the

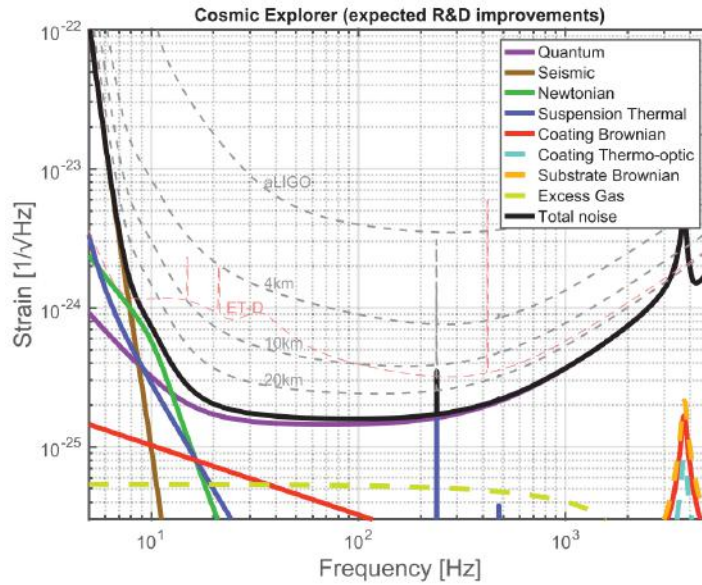


Figure 1.30: CE target sensitivity. The solid curves are for a 40 km long detector, while the dashed gray curves show the sensitivity of shorter, but technologically similar detectors; lengths are 4, 10 and 20 km. The AdL and ET design sensitivities are also shown for reference.

2G interferometers will continue to upgrade and hopefully make ground-breaking discoveries, made possible by a continuous improvement of the detection and infrastructure technology and data analysis techniques. There are also efforts to realize sub-Hertz GW detectors based on atom interferometry [161], superconduction [162], and torsion bars [163, 164]. Last but not least, the Laser Interferometer Space Antenna (LISA), a constellation of three spacecrafts, arranged in an equilateral triangle with sides

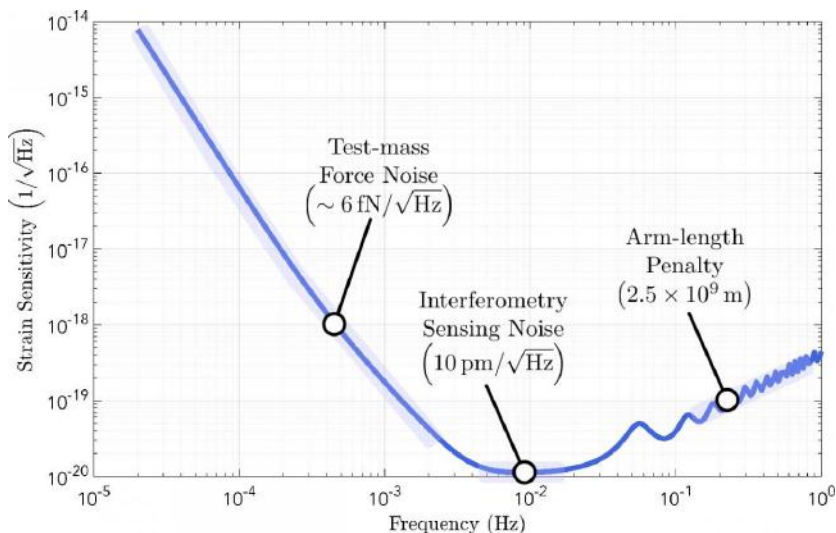


Figure 1.31: Time, sky and polarization averaged LISA sensitivity as defined in the mission proposal in January 2017. Credits: LISA consortium.

2.5 million km long, would eventually observe, among others, signals from Galactic Compact Binaries, Supermassive Black Hole Mergers and Extreme Mass Ratio

Inspirals. This would be possible by extending the detectable horizon down to low-frequencies (\approx mHz, see figure 1.31).

Gravitational Astronomy is starting to evolve faster and faster and the goal is to go beyond political and economical differences between countries and start to address future questions as a global scientific community.

Chapter 2

Magnetic noise

In section 1.2.4 I introduced the main noise sources affecting the AdV sensitivity, including the magnetic displacement noise. Now I will go into details of my activity regarding Magnetic Noise (MN) reduction, carried out both with experimental and statistical data analysis tools.

Section 2.1 is about the correlated MN between distant GW interferometers due to the global Schumann resonances and our attempt to subtract this mutual noise component by employing the Wiener filter. The broadly peaked Schumann resonances, which arise in the Earth-ionosphere wave guide from the tiny attenuation of Extremely Low-Frequency (ELF) EM waves, are potentially problematic for the Stochastic Gravitational Wave Background (SGWB) searches. A detailed search for a low-noise location to measure the Schumann peaks at Virgo site is included (section 2.1.2).

In section 2.2, I will present the extended magnetic field measurement campaign performed inside the main AdV buildings, in order to isolate the most powerful localized MN sources. This activity served as a preparation for the actual experimental measurement of the transfer function between the ambient MN and the interferometer output. This measurement campaign was performed through a series of magnetic field evaluations and artificial noise injections [140].

Then in section 2.3, I will present the work on modeling the magnetic coupling to the AdV payloads, from whom another possible way to evaluate the relationship between the environmental noise and the gravitational wave strain signal $h(t)$ is derived [139].

In section 2.4 all the magnetic noise projections are compared and finally in section 2.5 I will propose a feasible MN mitigation strategy [140].

2.1 Global magnetic noise

2.1.1 Schumann resonances

The Schumann resonances (SR) are the kind of signal we usually refer to when we talk about global electromagnetic (EM) noise. They are global EM resonances excited within the Earth-ionosphere waveguide, primarily by lightning discharges, which radiate EM energy at frequencies below 100 Hz [165]. The mechanism involves the lower part of the atmosphere, which is like a thin dielectric layer located between two conductors, the plasma of the ionosphere and the surface of the Earth. Thus, a

cavity with the properties of a spherical waveguide is formed. When the wavelength of EM waves in the cavity becomes of the order of the Earth's circumference ($\approx 40,000 \text{ km}$), a global resonance may occur. Despite their weakness, as the distance from the discharge grows, they travel across the globe thanks to the fact that the waveguide behaves like a resonator at extremely low frequencies (ELF, 0.03-300 Hz). A sketch of the phenomenon is presented in figure 2.1.




Figure 2.1: A picture showing the spherical cavity at the origin of the Schumann resonances.

The original idea of natural global EM resonances was presented in 1893 by George F. Fitzgerald and then again by Nikola Tesla in 1905 [166]. However, it was Schumann who first tried to measure them, unsuccessfully [145, 146, 167, 168]. In order to extract the SR signal from the background, we had to wait until adequate analysis techniques were available [169]. Now we know that a relatively long integration time ($\approx 15 \text{ min}$) is needed to detect the SR clearly in the spectrum, as the natural signal is actually random “noise”.

Being a global phenomenon, Schumann resonances had numerous applications over the years. As an example, in the last century, they had been used to monitor the enemy's nuclear explosions in remote parts of the globe [170] and as a long range communication system with submarines [171]. Besides the military uses, the primary interest was in lightning research, such as global lightning activity tracking [172–175] and climate changes monitoring [176, 177] (tropical land surface temperature, tropical upper tropospheric water vapor etc.). An additional application of SR is extraterrestrial lightning research on the planets and moons of the solar system [178, 179].

The simple theory behind the SR's propagation consider the terrestrial waveguide as an ideal one, made of two perfectly conducting concentric spheres separated by height h , which is much smaller than the Earth's radius a . The resonant frequencies f_n are thus determined by the Schumann formula [180]:

$$f_n = \frac{c}{2\pi a} \sqrt{n(n+1)} \quad (2.1)$$

However, the observed peak frequencies (7.8, 14.3, 20.8, 27.3, 33.8, ... Hz) are lower than those predicted by equation 2.1, because the Earth-ionosphere waveguide is not a perfect EM resonant cavity. Moreover, the observed peaks are wider than expected, with a spectral width of $\approx 20\%$ which vary seasonally and with proximity to lightning storms.

In the time domain, the SR appear as a constant background signal, which is a superposition of individual pulses arriving from about 50-100 lightning strikes per second around the world [181]. The amplitude of these intense transients (tens of picotesla) is often ten times higher than that of the background noise (Figure 2.2a). For further details, see [182].




Figure 2.2: a) Typical time series of 5 min. for one magnetic field coil. The signal consists of background noise caused by global lightning activity, with amplitude of 3-5 pT, together with intense transients from individual powerful lightning discharges, with amplitudes of tens of pT. b) Spectrogram showing the dominant SR mode at 8 Hz, with some higher modes at 14 Hz and 20 Hz. The intense transients show up clearly as enhancements of the SR spectra.

A Spectrogram of the same time series is shown in figure 2.2b, where both Schumann resonances (8, 14, and 20 Hz) and transients' signatures are easily recognizable.

The potential for the contamination of searches for a SGWB from this global SR signal is strong [55, 56] and hence its proper detection at each interferometer site becomes crucial for an optimal correlated noise suppression.

2.1.2 Search for a low-noise location

Typically, the sensors used to measure Schumann resonances consist of two horizontal induction coils to detect the horizontal magnetic field in the north-south (B_{NS}) and the east-west (B_{EW}) directions and additionally a vertical antenna for observing the vertical electric field (E_z). This arises from the assumption that in the far field, the horizontal electric field (and the vertical magnetic field) tends to zero. For

our purpose, we only employed the two horizontal magnetic probes, essentially for the difficulties in the realization of an electric antenna setup and the higher local interferences in the electric field, with respect to the magnetic field.

The magnetic sensors are the single axis magnetometers MFS-06 by Metronix, with a very low intrinsic noise: $\approx 0.01 \text{ pT}/\sqrt{\text{Hz}}$ at 10 Hz. They consist of several thousand copper turns around a high permeable ferrite core, all enclosed in a one meter long, shock resistant, waterproof, cylindrical plastic tube, that acts as a protection against mechanical stress and ultraviolet radiation. The measured fields are very small (pT) compared with the average Earth's magnetic field of $50 \mu\text{T}$.

Measurements in the ELF band are usually affected by man-made noise, which forces us to locate the SR measuring stations in isolated rural areas, away from high power supply lines, traffic, pedestrians and industrial activity [183]. Moreover, the sensors are sensible to external static EM fields, so that even the slightest vibration of an antenna will result in high signals induced at the input of the receiver. For this reason, the horizontal magnetometers should be buried in the ground to dampen the interference induced by ground vibrations or wind.

At Virgo, we tested many different sites, both in the proximity of the main buildings and in the surrounding countryside, in order to find a proper location to settle a permanent SR measuring station.

First tests in the countryside near Virgo

Preliminary tests involved the construction of a temporary station at Villa Cristina, which is a magnetically quiet site 12.72 km southwest from Virgo. First of all, we needed to assess if the presence of some magnetometer misalignments or non-linearities could affect a SR detection. With this in mind, we deployed two co-located and co-aligned sensors, at varying mutual distances. This sanity check represents the ideal situation, as it seeks to coherently remove both Schumann and local noise contributions, leaving only some possible spurious unwanted ones. In addition, we should detect any potential magnetometer mutual couplings, which are possible as they use induction coils to probe magnetic fields. Eventually, it is important that we are able to distinguish potential coherence between nearby sensors from true coherence due to the SR magnetic field.

The experimental setup consist of two single axis magnetometers MFS-06, two sensor cables with ruggedized military standard connectors, the data acquisition system (Centaur digital recorder by Nanometrics [184]), a 12V battery, two switching devices and a GPS antenna receiver (to synchronize the acquisition with GPS time). For the so called parallel sensor test, these two magnetometers are positioned horizontally and parallel with a distance ranging from 0.5 to 10 m from each other. Then we recorded time series. The time-frequency plots on the left of figure 2.3 include five different measurements, each lasting ≈ 15 min and each at increasing magnetometer mutual distance: 0.5, 1, 2, 5 and 10 m. The four transient features between each measurement are simply due to the act of moving one of the two probes away. In addition, we recognize some other low-frequency transients explained by seismically induced vibrations of the sensors coming from relatively close car passing ($\approx 50\text{m}$). The first four Schumann peaks are always visible in both magnetometers. On the right of figure 2.3, we plot the coherence ($1 - \text{coherence}$) of the magnetometers.

Except for the case of the 0.5 m measurement (blue line), for which we suspect some mutual coupling between the two coils, coherence is sufficiently high to exclude




Figure 2.3: left: spectrograms of the 2 co-aligned magnetometers at a mutual distance of 1m. Right: Coherence (1-coherence) between the 2 magnetometers at increasing mutual distances (up to 10m).

a bad sensor manufacturing. Coherence is maximized at 1 m of mutual distance (red line), while the 2 m measurement (yellow line) is probably affected by many external noise transients. Finally, we argue that 10 m are likely too much (green line), as the magnetometers start to sense different local disturbances and the coherence decrease, especially under 15 Hz. Hence the magnetometers should always be placed at least one meter and no more than five meters from one another.



Figure 2.4: A photo of the Villa Cristina scout house interiors, where we deployed the magnetometers for a 2 days measurement of the Schumann resonances.

Villa Cristina proved to be an optimal location for measuring ELF signals, which was enough for us to plan a longer measurement and understand how well we can detect the SR. Therefore, we put the two magnetometers inside the scout house, in order to protect the equipment from the external agents, taking care to turn off the

power of the building. Figure 2.4 is a picture of the house interiors, together with the aligned sensors.

After 2 days, I collected the data and found the first six Schumann peaks with a great resolution on both magnetometers. Moreover, I also remark the presence of the 7th resonance and some other resonances higher than 50 Hz in the spectrum of the sensor oriented North-South (see figure 2.5). In addition, there are spectral features which could be related to the data acquisition system or to nearby power lines.




Figure 2.5: Magnetic spectra of two magnetometers, during a 2 days measurement inside the Villa Cristina scout house.

I also report time-frequency spectrograms obtained with the two magnetic probes and a third seismic probe (Guralp CMG-3TD tri-axial seismometer, with a flat velocity response over the 0.0083–100 Hz frequency band [185]) which we put at the same location to monitor the seismic activity during the measurement (see figure 2.6). As expected, the seismic activity is lower during the night, even in a place so far away from a populated area. Interestingly, the background magnetic noise follow a similar trend (albeit with some time delay), as proof of the influence of the seismic vibrations on the cleaning of the probed magnetic signal.

Measures at Virgo site

Similarly to what has been done at Villa Cristina, I also performed an extended magnetic field survey at Virgo, looking for quiet locations along the detector arms and around the main buildings. This with the aim of building a permanent measuring station of the SR within Virgo boundaries, which would be crucial for the detection of the Stochastic GW background. During the years, we explored many different locations but one of the quietest was North-East of the Central Building, where we buried two orthogonal magnetometers at 1.5 m mutual distance and at half a meter of depth. The whole setup is protected from rain and direct sun exposure with an enclosure. Data were acquired with a Centaur Digital Recorder at 200 Hz sampling frequency. This measurement was done in August 2017, in coincidence

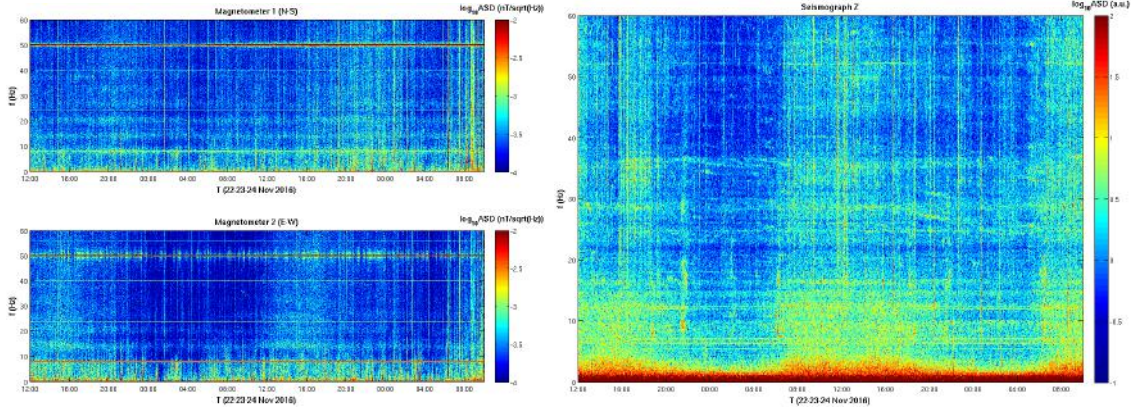


Figure 2.6: On the left, the time-frequency plots of the two magnetometers, aligned along the North-South and East-West directions; on the right a seismic spectrogram of the same time period from a seismometer deployed very close to the magnetic sensors.

with the first official combined Advanced LIGO - Advanced Virgo observing run and with a corresponding measurement of the SR at LIGO site. Schumann resonances are seen clearly in the data, as shown in the blue trace of figure 2.7, where the first six Schumann peaks are easily recognizable.




Figure 2.7: The two best Schumann resonances experimental evaluations compared: the measurement outside the Central Virgo building (blue trace) and the measurement inside the Villa Cristina scout house (13 km far from Virgo - green trace). They are both acquired with a magnetometer oriented North-South.

The figure also shows the magnetic spectrum taken at Villa Cristina by the co-aligned NS magnetometer (green line). The difference is evident below 10 Hz, mainly due to the peculiar natural/anthropogenic seismic activity of the sites.

I also show an example of the type of undesirable transient features that can be found in the magnetometers. While the most common ones are caused by nearby

storms or seismic activity in the surrounding area, others may result from probe failures or bad tuning. An example of the latter is the characteristic V-shape we note in the spectrogram of figure 2.8. After many tests, we recognized the source in the “chopper mode” feature of the magnetometers: it uses a free-running local oscillator whose frequency is temperature dependent, as proved by the fact that these features are evident 2-4 times per day, only during daylight hours. In addition, the noise intensity is highly variable. These disturbances are easily mitigated by switching off the chopper mode.




Figure 2.8: Example of the V-shaped transient in magnetometers data, attributable to the chopper mode of the probes.

From then on, the low-noise location North-East of the CEB will be employed as permanent station to measure the Schumann magnetic field at Virgo site, in order to pursue the coherent subtraction of the global EM noise.

2.1.3 Coherence measurements

In the work [57], the authors carried out a demonstration of Wiener filtering with a goal of reducing the coherence between widely separated magnetometers. My work carry on and finalize this project. A major point to mention is that we use the magnetometer output as a surrogate for a GW interferometer strain channel, as currently the latter has not sufficient coherence with the MN. It is also important to underline that the Wiener filter assumes a high level of coherence between the input and the target signals. With this in mind, it was necessary to measure the coherence between quiet magnetometers stationed near GW detectors and initiate a close collaboration with additional measuring stations from all over the world, forming a network of magnetometers with optimal sensitivity. We used a variety of permanent and temporary extremely low frequency (ELF) magnetometer installations. The medium sensitive band of these magnetometers is 3-300 Hz, with a sensitivity at 14 Hz better than $0.01 pT/\sqrt{Hz}$. Here is the complete list:

- Virgo (V1): permanent; located in the very close area around the interferometer (see the previous section for further details); magnetometers are MFS-06 by Metronix [186]; Virgo 1 is the North-South (NS) sensor and Virgo 2 is the East-West (EW) sensor.
- Villa Cristina (VC): temporary; located deep in the countryside near Virgo; NS and EW facing magnetometers MFS-06 by Metronix; measurements were made between July 20-22, 2016, and additional measurements between November 22-24, 2016.
- LIGO Hanford (H1): permanent; LEMI-120 sensors [187], oriented along the X-arm and Y-arm of the interferometer.
- LIGO Livingston (L1): permanent; LEMI-120 sensors; LIGO 1 is the X-arm direction magnetometer, LIGO 2 is the Y-arm direction magnetometer.
- KAGRA: temporary; magnetometers MFS-06 by Metronix; KAGRA 1 (K1) is located outside of the cave of KAGRA site at Kamioka (NS, EW and vertical sensors), KAGRA 2 is in-cave (NS, EW sensors); measurements were made between July 20-22, 2016.
- Hylaty station (POL): permanent; located in the Bieszczady Mountains, in the southeast of Poland [188]; AAS1130 sensors, oriented NS and EW.
- Hugo station (COL): permanent; located in the Hugo Wildlife Area in Colorado, USA; AAS1130 sensors, oriented NS and EW.
- Patagonia station (PAT): permanent; located in Rio Gallegos in Patagonia, Argentina; AAS1130 sensors, oriented NS and EW.

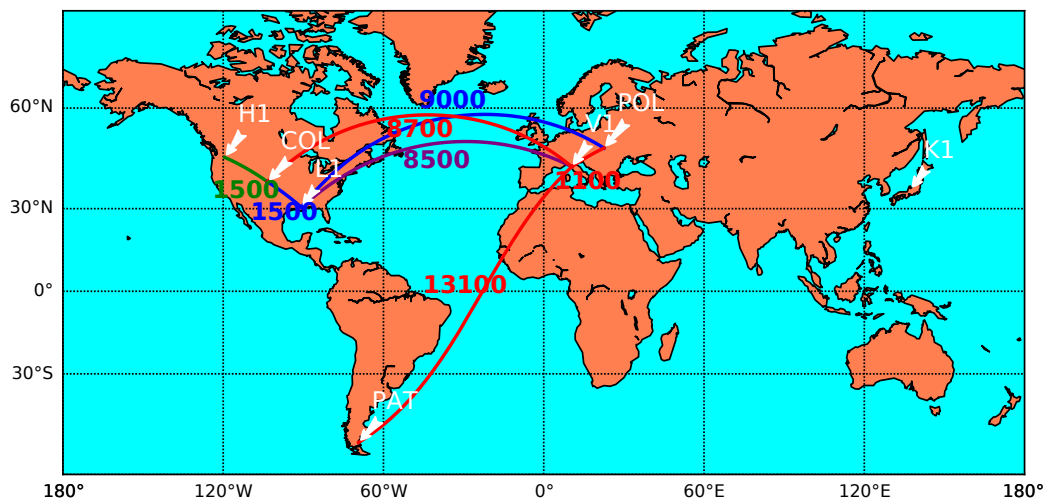


Figure 2.9: Location of the 7 permanent/temporary magnetometer stations (here V1 and Villa Cristina are co-located). The colored lines indicate the networks used in the analysis, with the distance in kilometers shown.

The latter 3 permanent installations are part of the Polish WERA project (World ELF Radiolocation Array) [189], an intercontinental radio-location system, working

within the ELF range and serving to obtain the measurement of discharges in the Earth's atmosphere.




Figure 2.10: Locations of the magnetometers at LVK detector sites. Their coordinates in the interferometer system (x, y) , relative to the interferometer vertex, are: $V1 = (80, -72)$ m; $L1 = (120, 3000)$ m; $H1 = (1030, 195)$ m; $K1 = (400, -600)$ m.

The LIGO-Virgo-KAGRA (LVK) sensors are placed far enough from the interferometers so as to not be sensitive to local magnetic noise but close enough to measure approximately the same Schumann field as the detectors do. Figure 2.9 shows the location of the magnetometer stations in the list. The colors indicate the networks we used in the coherence analysis and the numbers refer to the distances in kilometers between the pairs. Figure 2.10 shows the sensors location on the individual LVK detector sites. The power spectral densities showing the Schumann resonances detected in the sites of interest in the current analysis are shown in figure 2.11.

The metric we use to evaluate the correlation between magnetic sensors is the coherence $C(f)$:

$$C(f) = \frac{|S_{12}(f)|^2}{S_{11}(f)S_{22}(f)} \quad (2.2)$$

where $S_{12}(f)$ is the cross spectral density between the channels, while $S_{11}(f)$ and $S_{22}(f)$ are the power spectral density of channels 1 and 2.

Figure 2.12 shows the coherence of three magnetometers at KAGRA, LIGO Livingston and Virgo with other magnetometers belonging to the network. Clear peaks with coherence values of about 0.2-0.6 are visible in all pairs, showing statistically significant magnetic correlation over the entire Schumann band. The co-located and co-aligned KAGRA magnetometers (KAGRA1-KAGRA2) show broadband coherence at the level of 0.5-0.9 and several schumann peaks, providing a test for the best subtraction possible. The co-located LIGO and Virgo curves are instead more influenced by local magnetic noise. However, it is important to include these channels in the subtraction network as it will help to remove the local noise itself in the target channel.




Figure 2.11: Median power spectral density of co-oriented magnetometers, computed using 128 s segments at the all 8 stations of the network. In addition to some sharp instrumental line features, the Schumann resonances are visible in all of the magnetometers. The 20 Hz line at LIGO Livingston is likely due to power lines which cross the Y-arm.

2.1.4 Noise subtraction with Wiener Filter

In this particular case, we use Wiener filtering to reduce Schumann correlated noise from a magnetometer (the “target” sensor), assuming that the Schumann noise is shared between the time series of other magnetometers (the “witness” sensors). In the future, the target sensor would be no more a magnetic probe but a GW strain channel, for example.

Figure 2.13 shows the ratio between the residual spectrum, after the coherent subtraction, and the original spectral density of the target sensor, using a few different target-witness sensor combinations. The green curve employs one KAGRA magnetometer as the target sensor and a second KAGRA magnetometer as witness. The remaining lines use a Villa Cristina magnetometer as target and all available sensors (dashed black line) as witness or alternatively the sensor with the highest coherence in each frequency bin (blue line). The witness sensors in this test include POL and COL local orthogonal sensors. The VC plots shows that using the entire network improves the subtraction by about 10% compared to when using the best channel only, adding evidence to the fact that magnetic correlations are significant over the entire considered frequency band.

Finally we present the results in the most appropriate metric for SGWB searches, which is the correlation between magnetometers, before and after Wiener filter subtraction, to measure the effect that the WF has had on the correlations. In this sense, we can see a magnetic sensor pair as a proxy for a 2-interferometer network. In figure 2.14, we show the coherence between the 3 different magnetometer cou-

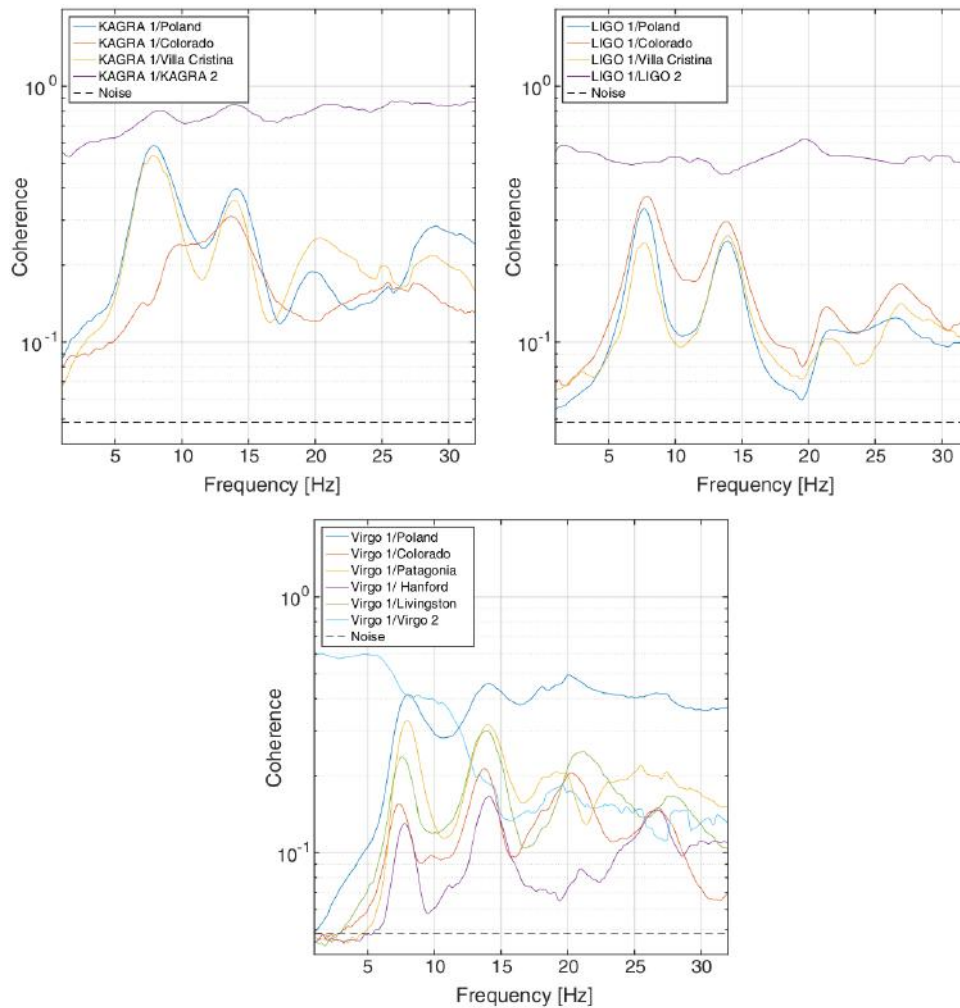


Figure 2.12: On the left is the coherence between KAGRA and the NS Poland, NS Colorado, and Villa Cristina magnetometers over 2 days of coincident data. The dashed line corresponds to the expected correlation given Gaussian noise. On the right is the same for LIGO 1. On the bottom is the coherence between Virgo 1 and the NS Poland, NS Colorado, NS Patagonia, LIGO 1, and LIGO Hanford X-arm direction magnetometers over a week of coincident data.

ples. On the left is the coherence between the KAGRA and VC magnetometers. There is no measurable remaining correlated noise contribution from the Schumann resonances. On the right is the coherence between LIGO Livingston and VC and we find that the coherence is reduced to the expected noise floor level. On the bottom of figure 2.14, we also show the coherence between the LIGO Hanford and Virgo magnetometers, in order to test the permanent installations: we managed to achieve a reduction in coherence of about 60%, while some residual coherence visibly remains, mostly due to the distances between target and witness sensors in the Virgo network.

This is the first time we show how dedicated EM measurements at the interferometers can reduce correlated noise, making a potential detection of a SGWB possible.




Figure 2.13: Ratio between the residual spectrum after the WF subtraction and the target sensors PSD before the WF subtraction, for 3 different target-witness sensor combinations: KAGRA - KAGRA (green line); Villa Cristina - sensor with the highest coherence (blue line); Villa Cristina - all available sensors (black line).

2.2 Local environmental magnetic noise

Besides the global magnetic noise, a significant amount of MN comes from the local Virgo environment. In this sense we need to characterize the ambient MN in the close area around the TMs, inside the three main AdV buildings: the Central Building (CEB), hosting the two input TMs, the West-End Building (WEB) and the North-End Building (NEB), hosting the two end TMs. In figure 2.15a-b I show the blueprint of NEB, with all the main interferometer infrastructure, including the vacuum chamber (North-End Tower, NET) hosting the NE mirror and its seismic attenuation system, and all other ancillary devices. I will explicitly make measurements either in a terminal or the central building, depending on the availability of the area and the type of measurement I want to perform, knowing that the CEB has a much more different design than the others.

2.2.1 Typical magnetic sources

Many electric and electronic devices, which we can encounter throughout the experimental halls, could represent a source of ambient MN. They include electronic boards, pumps, motors, lights, electrical power circuits and, in principle, any wire where current flows. Furthermore, we must deal with eddy currents induced by the conductive nature of materials. Because of eddy currents, external magnetic fields are spatially distorted and their magnitude is affected by the presence of conductive

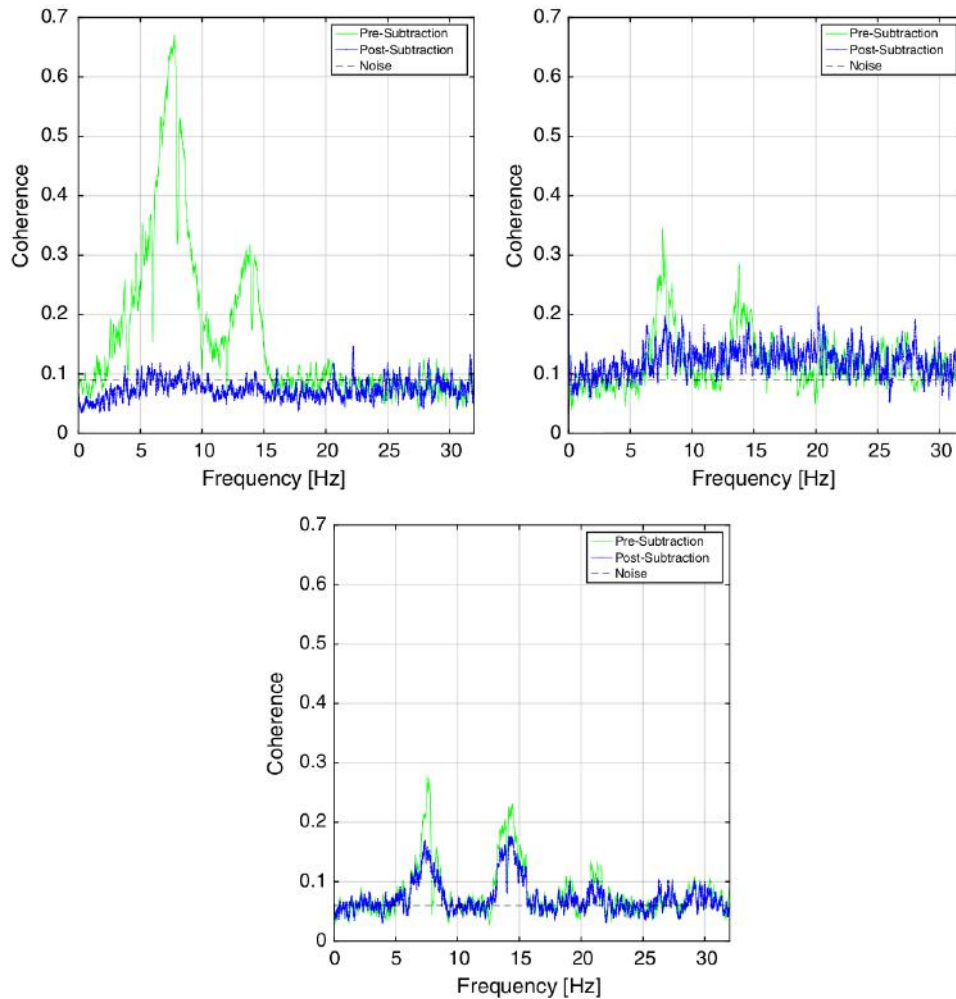


Figure 2.14: On the left is the coherence between the KAGRA and the Villa Cristina magnetometer before and after Wiener filter subtraction. On the right is the same between LIGO Livingston and the Villa Cristina magnetometers. On the bottom is the same between LIGO Hanford and Virgo magnetometers.

parts. Examples of magnetic coupling mechanisms between a MN source and the interferometer are a direct coupling to the magnets used for the TM control, but also couplings through cables carrying signals to the voice-coil actuators to which the magnets belong. A series of countermeasures have been put in operation during the years, in order to reduce the magnetic coupling, including anti-aligned magnets and Shielded Twisted Pair cabling. However, a residual signature is unavoidable, which would become a severe noise source at the time when Virgo will reach its design sensitivity. This kind of noise is very difficult to deal with, since the EM field can have a strong spatial dependency and each source must be singularly tracked down. Not to mention that it is usually hard to find any coherence between the strain signal and magnetic probes, except when strong transients occur. When a coherence is found, the best practice is to act on the noise source and try to cure it. Typical sources are mains transformers and cables carrying high current, motors, impulsive electronics driving heavy loads (e.g. Pulse-Width Modulation regulators). Indeed, any sufficiently close low power supply module and electronic board can also trigger a detectable disturbance. Other than that, many recurrent features are present in

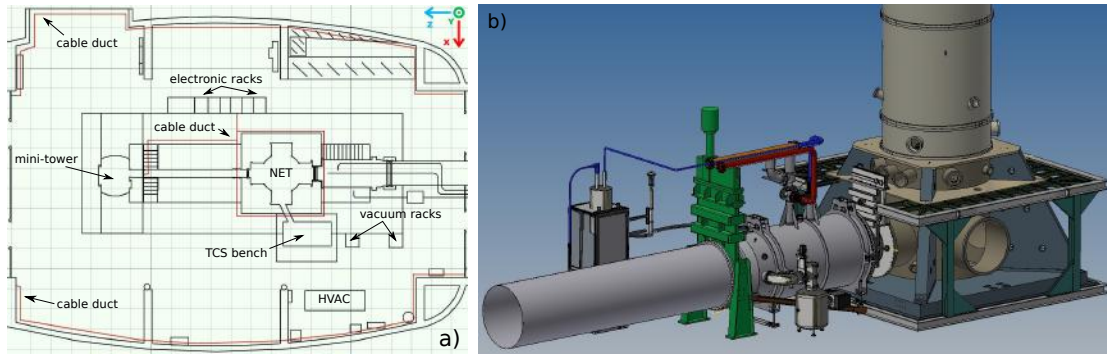


Figure 2.15: a) Top-view drawing of the NEB experimental area. TCS, HVAC and NET stand respectively for “Thermal Compensation System”, “Heating, Ventilation and Air Conditioning” and “North-End Tower”. b) 3D CAD picture of the vacuum-tight enclosure of the interferometer components.

the AdV magnetic spectra (figure 2.16), ranging from several lines in the 4-10 Hz interval caused by the air conditioning system machinery to the 50 Hz power line and its harmonics, associated to the frequency of the electrical system. Lines due to the main power supply could be broadened by the presence of side-bands: they are barely discernible in the plot of the vertical magnetometer (V axis) in figure 2.16 (e.g. around the 50Hz line, at 36 Hz and 64 Hz; around the 100 Hz harmonic, at 86 Hz and 114 Hz). Sometimes, the side-bands are rather intense and can be visible as separate bumps adjacent to the main one, causing undue coupling to the interferometer output and impairing the spectral sensitivity.




Figure 2.16: Typical magnetic field spectra detected inside the central building by the 3 permanent single-axis magnetometers used to monitor the magnetic activity. Periodic features and broader structures are easily identifiable as known sources. N, W and V stand respectively for North-South, West-East and Vertical directions.

2.2.2 Evaluation of the vacuum chamber shielding properties

Inside the WEB, on February 2018, I carried out a new evaluation of the magnetic filtering effect produced by the steel tank surrounding the TM. The tank is a grounded austenitic stainless steel 304L shell with a diameter of 2 m and a thickness of 15 mm, which isolates the vacuum-sealed environment from the outside. During a previous unpublished measurement (2006), a transfer function between the inside and outside of the shell with frequency slope of $f^{-1.3} - f^{-2}$ was found, but there were high uncertainties on the shape of the filter response, its order (first or second order) and its pole (from about 5 to 30 Hz). The new measurement required the positioning of an home-made injection coil, which will act as magnetic field source, and two magnetic probes, one inside the Tower (fixed) and the other outside. This was possible because at that time the test mass was not suspended and hence particular care has been given to the positioning of the probe, right where the mirror has its geometrical center. Figure 2.17a shows a map of the different configurations: the 3 “injection coil - external magnetic probe” position pairs are numbered consecutively.

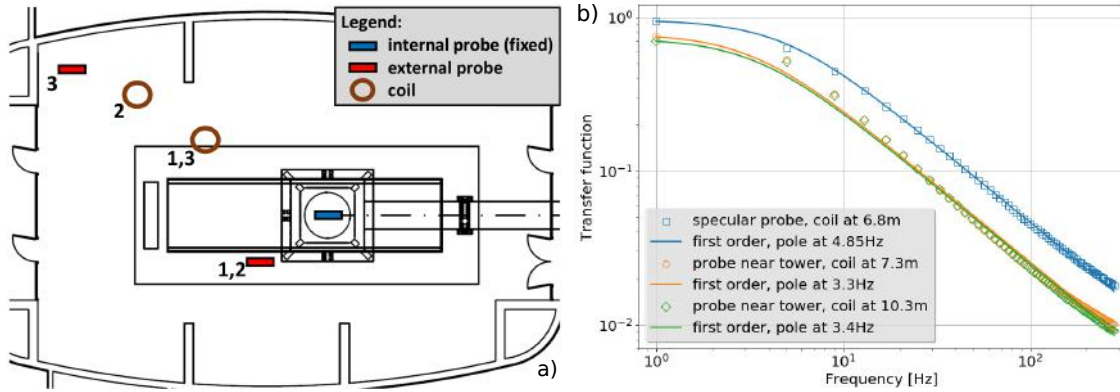


Figure 2.17: Frequency evolution of the West-End Tower (WET) shielding power: a) The 3 “injection coil - external magnetic probe” configurations for their positioning inside the experimental area: coil at 7.3 m, probe near the WET (config. 1); coil at 10.3 m, probe near the WET (config. 2); coil at 6.8 m, specular probe (config. 3). b) Data fit to a first order lower pass filter model of Butterworth type.

I quantify the shielding power of the steel enclosure by computing the filter frequency response

$$H(f) = \frac{B_{mod}^{in}}{B_{mod}^{out}} = \frac{k}{\sqrt{1 + (f/f_{pole})^2}}, \quad (2.3)$$

where B_{mod} is the magnetic field modulus, k is a constant and f_{pole} is the frequency of the filter pole or in other words the cut-off frequency. The trend of the experimental data fits the first order lower pass filter curve of Butterworth type with pole ranging from 3.3 to 4.85 Hz, depending on the setup configuration (figure 2.17b). Above 100 Hz we note a mild inflection and a consequent change in the slope. This points to unknown effects, summing to that of a pure low pass filter (whose frequency response would indefinitely continue to decreasing). The difference between the measurements with the probe near the tank and the measurement

with the specular probe suggests that the magnetic field is more concentrated in the proximity of the test mass. Note that the frequency evolution is nearly the same for all the investigated configurations.

2.2.3 Identification of point magnetic sources

It is common knowledge that ferromagnetic materials are easily magnetized during various machining phases. Depending on the processing, this residual magnetism may have different causes: forging, bending, welding, cutting, hardness checks, transport, handling and even mechanical vibrations. In section 1.2.4, I highlighted the main coupling route between a magnetic noise and the interferometer, which results in a direct effect on the EM actuators of the TM and a potential disturbance on the magnetic probes. As the NET can be mechanically excited by sound waves and seismic vibrations, the presence of magnetic patches on it could be a dangerous MN source. For this reason, we measured the DC magnetic field on the NET external walls to be sure there were no significant magnetization patches on it. In figure 2.18a is the experimental setup we used for most of the measures inside the Virgo buildings, which includes the portable triaxial magnetic field sensor FL3-100 by Stefan Mayer Instruments with a measurement range of $\pm 100 \mu\text{T}$ and an intrinsic noise of few $\text{pT}/\sqrt{\text{Hz}}$ at 10 Hz, mounted on a perch 1.3m high, a 12V battery and the signal cables.

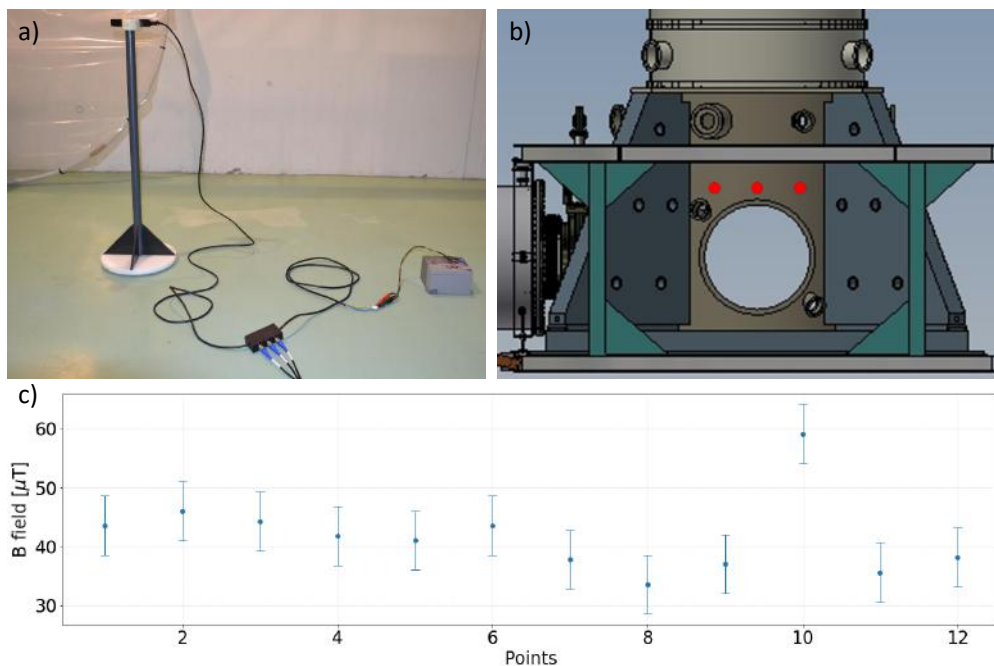


Figure 2.18: Details about the DC magnetic field measures on the NET external walls. a) the experimental setup; b) example of 3 points at which the B field is evaluated; c) plot of the B field modulus at 12 locations around the tank.

The red dots in figure 2.18b mark the height at which the measures around the NET were conducted. What we found was an average field strength of about 35-50 μT in modulus on 24 locations around the tank, within a 1 m and a half height from the platform ground. Note that figure 2.18c displays the B field modulus evaluated

only on 12 points on the surface of the tank. These figures correspond quite well to the average magnetic field we encountered in other internal and external building locations and to the typical magnetic field of the Earth. Therefore we do not see any evidence of alarming magnetic patches on the surface and the core of the steel tank at the height level of the TM. Nevertheless, there may be other high MN sources in the proximity of the tank. In particular, we found a rather intense and periodic EM activity around the electronic racks, which are situated almost 2m away from the tank. However, its intensity drops to the standard quiet magnetic activity ($0.01 - 0.1 \text{ nT}/\sqrt{\text{Hz}}$) before reaching the tank.

A supplementary example of noise coming from fields irradiated by electronic devices was tracked down to power supply boxes, with nearly a 3 orders of magnitude higher magnetic noise than the standard quiet magnetic activity. Already during the DC measures around the NET, a particularly high noise source was identified (see the B field magnitude found in point 10 of the plot in figure 2.18c). These boxes were situated close to each tank, on the first floor platform. Their final tagging as the noise we were looking for comes from a dedicated search of many days, passing through some incorrect identifications. We performed some on-off tests on the noisy power supply boxes and then we removed them, since they were residue from past interferometer configurations and no longer used. In figure 2.19 we compare the noise level near the localized MN source, before (orange) and after its removal (green), with the environmental low-noise activity in the area (blue).




Figure 2.19: Comparison between three different magnetic noise levels, found near a localized noisy source, before (orange) and after its removal (green) and in correspondence of a low-noise location (blue), inside the North-End Building.

This preliminary characterization and mitigation work was essential for the evaluation of the ambient noise contribution to the AdV sensitivity.

2.2.4 Far-field magnetic noise injections

Virgo uses stimulated noise injection on a large scale to assess if some noise would compromise the sensitivity of the instrument. This approach is called “active” be-

cause we stimulate the interferometer (or only small sections of it) with an active probe to drive a perceptible signal at the detector output.

To study MN, we had to build a drive coil from scratch. The coil consists of a copper wire with a cross section of 1 mm, wrapped 50 times to create a winding with a diameter of 1 m around the full PVC frame as a supporting structure.

We extract the magnitude of the environmental or the injected magnetic field using two different types of probes. First, a single-axis, one-meter long, tube-like magnetometer MFS-06 by Metronix, with a very low intrinsic noise (better than $0.01 \text{ pT}/\sqrt{\text{Hz}}$ at 10 Hz), which makes it suited to measure the field strength variations of the Earth. Secondly, the portable triaxial magnetic field sensor FL3-100 by Stefan Mayer Instruments with a measuring range of $\pm 100 \text{ } \mu\text{T}$ and an intrinsic noise of few $\text{pT}/\sqrt{\text{Hz}}$ at 10 Hz.

A Very important aspect to consider is the position of the coil and magnetic probes inside the experimental areas. More specifically, we want the interferometer and the magnetometers to detect approximately the same injected field for a faithful projection. In addition, we want to be in the regime of the far-field injections, where the source is ideally located at an infinite distance from the detection device (coil dimensions \ll distance). If we consider the space constraints of the area surrounding the vacuum chambers, we have limited choices for the positioning of the instrumentation. A quick calculation shows that a current of about 10A flowing in the coil create a magnetic field of about a few tens of μT at 20 m from the coil.

The first measurement campaign was carried out in the second half of 2017 in the three main AdV buildings: NEB, WEB and CEB. The same procedure is repeated every time, adapting the configuration to the building environment. The three actual coil-probe mutual positions are depicted in figure 2.20.




Figure 2.20: From left to right the WEB, NEB and CEB floor plans. The thunderbolt and the blue cylinder are respectively the coil and the probe, inside the three facilities; the red spots identifies the four TMs and the red lines follow the main laser path.

We injected a series of monochromatic magnetic signals (magnetic lines) lasting 600 seconds each and field intensity of few mT at the magnetometer locations. This applies to each line and to the entire explored frequency range (from 14 to 140 Hz).

For comparison, this is nearly the same field strength we encountered in correspondence of the noisy source of figure 2.19, below 100 Hz. There are several different conditions that must be met during the procedures: interferometer in science mode¹; quiet environmental magnetic activity²; sensors readout far from saturation; the amplifier tuned in such a way that the output measured on the spectral density had the same amplitude for all injected lines, regardless of the input frequency (from several Hz to several hundred Hz); good coherence between injected lines and strain signal $h(t)$ for almost every tested frequency. In figure 2.21 we show the injection of 114.8 Hz line at CEB, which is visible by the magnetic sensors (2.21a), as well as by the detector (2.21b). The overall coherence is above 0.8 for almost all other frequencies, except for the first two ones due to poorer detector sensitivity (see figure 2.22).

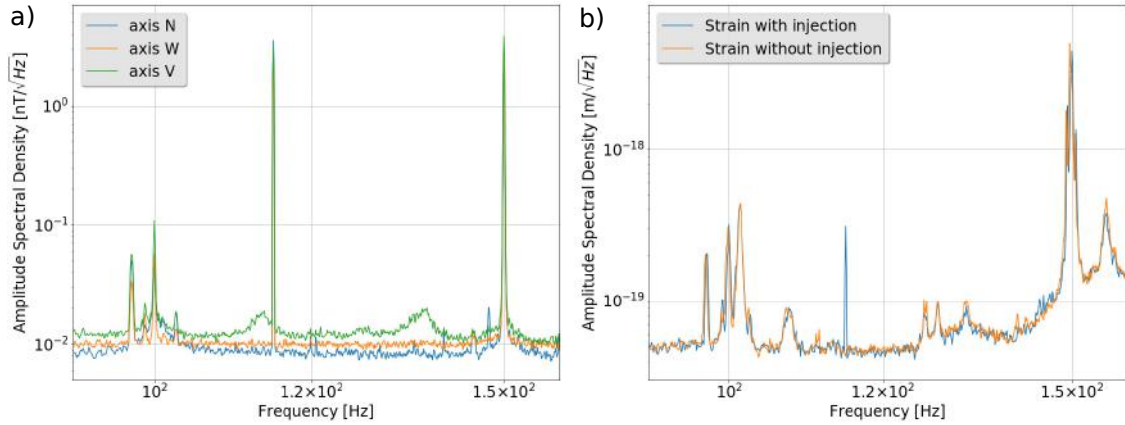


Figure 2.21: 114.8 Hz line injection at CEB. It is readily visible by the 3 single-axis magnetic probes (a) and the interferometer (b), where the blue peak stands out against the reference orange curve.

The injection scan have been repeated several times to make sure we have always stable interferometer conditions. After that, the responses of the magnetometers and the instrument are analyzed to quantify the coupling between the environmental MN (m) and the strain (h). At the time of the injections (inj), the transfer function corresponds to

$$TF_{mh}(f) = \frac{ASD_{hm}^{inj}(f)}{ASD_{mm}^{inj}(f)}, \quad (2.4)$$

where ASD stands for Amplitude Spectral Density (cross spectral density as a numerator and power spectral density as a denominator). The overall coupling for all measurements inside the 3 buildings is computed by taking the incoherent sum to obtain a single TF, which is displayed in figure 2.23. The coupling curve is fitted to a power law and the model parameters are estimated with a least squares fitting to the data. The results show that the points follow the model

$$TF_{mh}(f) = TF_{mh} \left(\frac{f - f_0}{10 \text{ Hz}} \right)^\alpha + \beta, \quad (2.5)$$

¹Fully locked interferometer in the foreseen configuration for observation, with no human or automated intervention on the machine.

²Absence of known noisy magnetic field sources close to the TM, for instance due to accessory/monitoring electronics and/or during a thunderstorm.




Figure 2.22: Coherence between injected lines (at CEB) and strain signal $h(t)$ for every tested frequency. The values are such that they lead to a faithful projection of magnetic noise.

with amplitude TF_{mh} and pole f_0 . Notably, the exponent $\alpha = -3.3 \pm 0.4$ is consistent with the expectations regarding the decreasing trend of the coupling function. Indeed, it is expected a factor -1 from the tank shielding, plus a factor -2 from the TM pendulum mechanical contribution. Above 100 Hz, spurious phenomena such as the electronic coupling through cabling are probably dominant and add up to the descending pattern given by the low-pass filtering of the tank. The data at 14 Hz is considered to be just an upper limit, since the magnetic excitation of the interferometer at these low frequencies is rather challenging.

For this reason, the coherence between the injected lines and $h(t)$ is considered to be reliable over the frequency range [20, 140] Hz, while at lower and higher frequencies the lines are not always perfectly visible by the interferometer and this could lead to an overestimate of the projection. However, higher frequency coupling was also investigated and the results can be found in section 2.4, in terms of comprehensive magnetic noise projections.

The last step involves the computation of the magnetic noise projection, which is done by multiplying the coupling function of equation 2.5 by the “quiet” environmental magnetic spectrum (no injections):

$$P(f) = TF_{mh}(f) \cdot ASD_{mm}^{quiet}(f). \quad (2.6)$$

The values of $P(f)$ from equation 2.6 work under the assumption that the magnetic sensors and the tank containing the TM experience the same field intensity, which is true in the far-field regime. However, in our case the distance between the driving coil and the tank is hardly greater than 20 m and hence the positioning of the magnetometers becomes a crucial aspect. At CEB, it has been verified that the transfer function slope is conserved even when the location of the magnetic sensors is changed. For this test, the coil location was fixed and three different set of magnetic sensors were placed inside the building, including one nearer to the coil and another one farther away from the coil with respect to the TM. Results show that




Figure 2.23: Magnetic coupling results. The orange diamond indicates that the corresponding injected line is detected in the strain channel with a good SNR, while the green star means that only an upper limit can be set, due to bad coherence between the injected field and the strain output. The fit to the coupling function is also included.

the transfer function slope is conserved, while the amplitude undergoes a shift so that the nearest sensor underestimates the intensity, the farthest one overestimates it.

As a I mentioned before, the MN effect in different buildings is also assessed. It turns out that the coupling level at CEB is comparable to the one measured at the WEB and NEB. With a view to keep monitoring the Virgo MN level, it is believed that MN injections at CEB, in a time when the sensitivity is optimal, could be confidently take as a proxy of the measures at the other 2 main locations (NEB and WEB). A strategy to ensure a periodic and cost-effective MN injection procedure is currently on the agenda.

2.3 Magnetic coupling modeling

In the current section I will describe and carry out the extensive work on the coupling between the environmental magnetic field and the metallic frame that encloses and suspend the Virgo TM (i.e. the payload), started by the PhD student of my research group, Martina Neri. It is well known that any time-varying magnetic field interacts with metallic structures such as the payload to produce local gradients, which exert forces on the four magnets glued on the TM (see section 1.2.2 for details). These gradients can induce a worsening of the intrinsic displacement noise of the detector and, in the presence of high magnetic transients, create a glitch, which can be misinterpreted as a GW signal.

That provided the idea of studying the magnetic response of the payload to a given environmental field. Since analytical solutions of this problem exist only for

very simple geometries and a direct measurement of the induced field is impractical, Finite Element (FE) simulations are performed. They intimately rely on the language of mathematics to describe the laws of physics, reducing all to partial differential equations. Further details of the theoretical background of the technique can be found in the original works [190,191]). More recent reviews are set out in the books by Zienkiewicz [192] and Strang and Fix [193]. The paper [139] summarizes the work I did on this topic, while here I present it more in detail.

Several steps are necessary on the system in order to reduce the solution time and simplify the FE model. After that, an optimization procedure based on Design of Experiment (DoE) techniques is developed to find the optimal electric configuration of the payload. The magnetic field gradient is then calculated by taking into account magnetic moment and magnet position dis-homogeneities. Monte Carlo simulations are used to compute the total force on the TM as a function of the magnetic moment tolerance and of the magnet position error, with respect to the payload structure. The strain noise contribution was estimated for both translational and rotational forces and it was compared with the AdV sensitivity curve (section 2.4.1).

2.3.1 Magnetic coupling to the payloads

In the previous sections I showed that magnetic disturbances may arise from the superposition of both known point-like sources (local) and unknown ones (environmental and/or global). The search and mitigation of local magnetic sources nearby the TM (section 2.2.3) allow to consider only the environmental noise component for the estimation of the magnetic noise. Under this hypothesis, the contribution of MN to the AdV noise budget is relevant only in the low frequency range because at high frequencies (roughly above 150 Hz) the interferometer sensitivity is limited by other noises and the environmental magnetic field (at the payload location) is filtered out by the metallic enclosure which surrounds the payload. The latter will be explicitly taken into account in the noise budget estimation.

Starting with the basics, a magnetic field acts on a permanent magnet causing the following two effects: a translational force associated to the magnetic gradient: $\vec{F} = \nabla(\vec{\mu} \cdot \vec{B})$, where $\vec{\mu}$ is the magnet's magnetic moment; a torque produced by the magnetic force: $\tau = \vec{F} \times \vec{r}$. More specifically, one should refer to two distinct components of torque: one related to the TM and the other one related to each single magnet. However, in point-like magnet approximation, the second contribution is neglected, especially because the magnets are glued on the TM and the whole system is treated as a rigid body.

In the AdV configuration, all the magnets used for the TM actuation have approximately the same magnetic moment, directed along the beam propagation direction (i.e. \hat{z} in the AdV reference system). The force on each magnet can be written as $F_z = \mu(\partial B/\partial z)$ and hence the total force on the TM is simply the sum of the forces on the four magnets. In the ideal situation, the four magnets have exactly the same magnetic moment and the magnetic field gradient is the same on each magnet. With these assumptions, the anti-parallel cross configuration should guarantee a null force on the TM. In the real-world case, $\vec{\mu}$ has a tolerance of about 15% around its nominal value and the magnetic field gradient spatial symmetry is no longer guaranteed.

In the first place it is argued that the main cause for a non-uniform gradient is the

interaction between the external field and the metallic structure of the payload. This structure couples with any time-varying magnetic field generating eddy currents, which in turn warp the field and produce the unwanted gradient.

In what follows, I present the numerical FE study on the Input Mirror Payload (IMP), since it is the most sensitive to magnetic coupling, together with the End Mirror Payload (EMP). The procedure consists of 2 steps: the construction of the FE model and its validation.

2.3.2 System modeling

The goal here is to determine the magnetic field around the payloads, which are complex, composite objects consisting of several parts (10 main parts; volume of about $1.2 m^3$) of different materials that are assembled by bolts, welds and screws. For this reason, the detailed electrical configuration of the payload and the eddy currents flow in it have to be estimated. This is a situation where we lack a reference standard and the analytical solution is not trivial.

The COMSOL Multiphysics FE simulation software [194] is employed, alongside the built-in AC-DC module for the computation of the time-varying magnetic field study. Now I will get into the details of some delicate simulation steps such as the geometry simplification and the model meshing.

Building the geometry from sketches can be very challenging, therefore it was imported through the COMSOL CAD Module. Because CAD drawings are made for machining tools, most features and details are a nuisance for the FE analysis (figure 2.24a).




Figure 2.24: a) CAD drawing of the Input Mirror Payload (IMP), which is imported in the FE software COMSOL; b) - c) finite element models of the IMP: two levels of simplification are applied; d) meshed IMP: the small parts are meshed finer than the big ones.

The first step consists in the de-featuring, which means repairing and pruning all the entities (lines, surfaces and volume domains) that are not correctly recognized by the Import Module, or that are irrelevant for the simulation. Then, for the sake of mesh simplification and in order to save simulation time, two layers of abstraction to the model are applied. The first layer deals with critical sections on which the meshing algorithm often does not converge (figure 2.24b). These entities include screw holes, bolts, very small features, threadings, etc. The second layer targets domains that can be replaced with a sketched version of itself (figure 2.24c),

essentially to simplify the mesh at a higher level and trying to reduce the solving time. In the end, the simplified model is validated against the original drawing.

The model meshing step is a delicate one, during which all the elemental volumes, up to a million small elements that together form the shape of the macroscopic structure, are generated, in order to make the simulation [195]. Calculations are performed for every single element and combining the individual results gives the final result. The entire process is an approximation, consisting in consecutive interpolations over the element(s), which can be linear, quadratic, cubic etc. Generally speaking, a balanced and fine mesh often leads to good and stable solutions. A free mesh made of tetrahedral elements is chosen for the overall payload structure and the air volume surrounding it, up to the infinite element boundaries, due to the highly irregular geometries (figure 2.24d). On the other hand, volumes around the mirror magnets are meshed with a more refined swept quadrangular elements distribution. It turns out that the total number of elements in the mesh is about $7 \cdot 10^5$.

A subtle but significant problem lies in simulating two adjacent metallic domains that are assembled together (i.e. in close physical contact). The software implicitly assumes electrical conductivity between the entities. In practice though, when two metallic parts are assembled without welding, the electrical contact can be impaired by a thin insulating film (e.g. oxides, contaminants, reaction products), which introduces an energy barrier that can limit the current flow [196,197]. Electrons may not have enough energy to tunnel across the barrier, depending on its thickness. Moreover, eddy currents are very feeble, making them especially susceptible to energy barriers. The conductivity across two surfaces in electric contact is determined by several factors: the type of material, the surface finishing and the pressure applied between the two surfaces. I designed and validated a simple model (open coil) which verifies that the eddy currents flow critically depends on the force applied to keep two aluminum surfaces in contact. A sharp transition in the intensity of the eddy currents with respect to the force applied was experimentally measured in that case. This simple experiment leads to consider each connection among the metallic parts of a composite object as a two-state system (open-closed) in order to take into account the (unknown) state of the surfaces and the applied pressure. The splitting of the payload to independent parts can be parametrized by the conductivity at the level of each mechanical connection p as 1 or 0, leading to a number of 2^p configurations. Therefore, I employed a hierarchical analysis to capture only the most relevant contact points in the payload assembly. Eventually, only $p = 7$ connections with a relevant impact on the gradient nearby the magnets are found. The corresponding parametrized volumes are inserted in correspondence of the real mechanical connections (screws and bolts) and they are modeled in COMSOL using either air (insulation) or aluminum (same material as the bulk). These 7 parts are highlighted in figure 2.25.

In a FE mathematical model, The constitutive laws involve physical properties of the materials. In practice, this implies estimating the correct validity intervals for each material properties [198]. The payload assembly is made of different materials: aluminum, steel and titanium. In the FE simulations, only the aluminum components are modeled, as they are the most relevant ones in terms of number of parts and they have the highest conductivity, $3.03 \cdot 10^7 S/m$, compared to $1.4 \cdot 10^6 S/m$ for steel and $0.6 \cdot 10^6 S/m$ for titanium. The model was then embedded in a




Figure 2.25: The 7 mechanical connections used to simulate two-state electrical systems: a) Cage Legs - Baffle Holder, Cage Legs - Arch, Arch - Baffle Holder and Arch - open Baffle Holder; b) Cage Legs - Compensation Plate and Cage Legs - Arch; c) Compensation Plate - Baffle Holder.

uniform, sinusoidal magnetic field to mimic the real environmental one.

2.3.3 Model validation

The FE model delivers the EM field at any point but we still have to determine the electrical connection states among the relevant parts of the assembly, which is called the model validation step. This activity implies to find out which of the $2^7 = 128$ possible models optimally represents the real payload and it involves comparing each FE simulation with an experimental measurement. In addition, the latter can be exploited to tune the geometrical representation and the uncertainties on the material properties.

The validation procedure requires the development of an accurate setup, where the driving magnetic field and the geometrical properties are kept under control, so that discrepancies between the simulation and the measurements can be exclusively assigned to the model assumptions.

Experimental set-up

A team was set up at Virgo site in order to build a driving system with two coils (“Big” and “Small”, with an external radius of respectively 1.040 ± 0.001 m and 0.540 ± 0.001 m) and a structural PVC cubic frame that aligns the coils with the payload (figure 2.26).




Figure 2.26: The PVC frame built for the FE model validation: a) CAD drawing; b) the actual object inside the Virgo clean room. The structure embeds two integrated coils (“Big” and “Small”) on one side and two gridded panels on the opposite sides for accurate positioning of the magnetic probe.

The coils are made of 100 turns of copper wire with a diameter of $0.95 \pm 0.02 \text{ mm}$. Whenever needed, the payload is housed between the coils and two gridded panels with several reading slots, where a magnetic probe can be inserted: this is a triaxial magnetic field sensor FL3-100 (Stefan Mayer Instruments), with intrinsic noise $< 20 \text{ pT}/\sqrt{\text{Hz}}$ at 1 Hz and a measurement range of $\pm 100 \text{ } \mu\text{T}$. The slots are geometrically tailored to the probe in order to minimize positional errors and vibrations during the measures. They are designed to get a fairly dense sampling of the magnetic field. Figure 2.27 on the left shows a focus on the placement of the magnetic probe inside a frame slot. Some of them are bigger so that a probe-holder rod can be inserted, in order to make spot measurements along a line inside the volume of the payload. A photo of the actual setup is provided in figure 2.27 on the right.



Figure 2.27: On the left is a focus on the placement of the magnetic probe inside the PVC frame. In addition (on the right), a probe-holder rod can make spot measurements along a line inside the volume of the payload.

The two coils generate spatially different magnetic field configurations and therefore they provide a more accurate validation. They can be independently driven by

an AC current generator (CoCo80, by Crystal Instruments) coupled to a linear amplifier (BAA 120, by TIRA). The cubic frame is placed in the Virgo Central Building (CEB), in a class 100 clean room, right under the input towers. It is ensured that there are no significant metallic objects or cabling in a radius of $\approx 2\text{ m}$ around the apparatus (figure 2.28).



Figure 2.28: View of the class 100 clean room and the surrounding site inside which all the experimental apparatus has been set up.

Reference measurements

The first validation step consist in the “zero” measure, that is a set of reference measurements to study the contribution of the surrounding environment and the cubic frame alignment to the evaluated magnetic coupling.

For this purpose, the measurements are provisionally performed without the use of the payload inside the cubic frame. It is worth mentioning that the amplifier exhibits a small drift in the amplitude of the AC current signal, whose nominal value was of 1 A, which is taken into account in the post processing analysis by performing a current optimization procedure. The three spatial components of the magnetic field are measured on a set of 65 reading slots for each gridded panel.

At the same time, the structural frame and the coils are simulated with COMSOL. The PVC frame is transparent to magnetic fields, which allows to consider only two materials for the simulation: copper for the coil winding and “air” for the remaining parts. The magnetic field is calculated with the FE analysis in correspondence of the experimental measurement points, for all the explored configurations. Finally, the experimental and simulated data are compared by minimizing the relative difference function m_k between the measured and the simulated magnetic field values over the input current

$$m_{(k)} = \left| \frac{B_{meas}^{(k)} - cB_{sim}^{(k)}}{B_{meas}^{(k)}} \right| \quad (2.7)$$

where c is the optimization parameter for the current and $k = [1, \text{number of measuring points}]$. Measurements and simulation agree within $\approx 5\%$, after the

optimization of the current. Looking at the box-plot in figure 2.29, it can be assessed that the distribution of the relative differences is compatible with zero for both the “Big” and “Small” coil configurations and at both 33 Hz and 333 Hz.




Figure 2.29: Distributions of the relative difference function m_k between the measured and the simulated magnetic field values, without the payload inside the test cubic frame (Eq. 2.7, reference measure). The statistic in the box-plot is computed for all the measured positions, for both coils and two magnetic field frequencies. The red horizontal line is the median value of the distribution, while the bottom and top edges of the blue boxes indicate the 25th and 75th percentiles, respectively. The whiskers extend up to 1.5 of the box range and the outliers are plotted individually with single red dots.

Input Mirror Payload measurement

The entire procedure is replicated with the Input Mirror Payload (IMP) inside the PVC cubic frame. The IMP is kept in place by a support structure with an aluminum base and four steel legs. Figure 2.30a shows a photo of the apparatus. The IMP is laser-aligned with the coil axis (estimated accuracy of $\simeq 1$ cm) and measurements are taken both in correspondence of the two gridded panels and with the help of the extension rod (see figure 2.27 on the right), to get as close as possible to the payload assembly.

The validation goal is to identify which electrical configuration epitomizes the real payload. Clearly it is impractical to FE simulate all the possible 128 configurations for each optimization step. A more feasible approach would be to simulate only the 7 conditions (the diagonal matrix of the connections), where just one connection is closed and all the others are open. Then, the best approximation to the measurements can be found by linear combination of the diagonal solutions. Unfortunately, the superposition principle does not apply to solutions with eddy currents. In order to solve this issue, it is resorted to a Design of Experiment (DoE) technique [199, 200], combined with a heuristic algorithm of minimum search.

The DoE approach considers the system under scrutiny as a black box, where only categorical inputs and continuous outputs are measurable. It provides a general




Figure 2.30: The Input Mirror Payload (IMP) inside the PVC cubic frame. a) the photo of the experimental apparatus; b) the FE model of the parts relevant for the simulation.

rule for dealing with inputs and outputs in order to estimate the system behavior. A Plackett-Burman implementation is followed [201], with the rows of the Hadamard matrix, a square matrix whose entries are only $+1$ and -1 (or 1 and 0) and whose rows are mutually orthogonal, as inputs. Such a design allows to sample the parameter space, which has the effect of reducing the number of tests to perform, but taking into account as many interactions as possible. In addition, it approaches the problem without the necessity to define the kind of interaction in the system. This is exactly what it is needed, since the different electrical configurations of the payload are not independent due to the fact that the current flow changes a lot with the connections and the magnetic effect cannot be calculated as a superposition.

Furthermore, a heuristic optimization procedure is adapted to a discrete parameter space, as each connection can assume only the values 1 or 0 . The cost function to be minimized is the one of equation 2.7, with the optimization parameter c equal to 1 , in order to compare the experimental and simulated magnetic field values. The optimization algorithm is described hereunder:

1. The Hadamard matrix spanning p parameters is used to find N combinations to test.
2. A FE simulation is built and solved for each N combination.
3. The magnetic field obtained from each N simulation is compared with the measured one.
4. The tested N combinations are sorted on the basis of the metric value (Eq. 2.7).
5. The n parameters ($n < p$) that have the same values in the first two (sorted) combinations are fixed.

6. As long as $n \neq 0$, the Hadamard matrix at $p - n$ parameters is used to find new combinations to test.

This procedure is iterated until all the parameters are fixed, defining the temporary best configuration. A further refinement step allows to check a parameter space subset surrounding the best configuration: all the configurations that differ from the best one at least by two parameters are tested. This optimization process was first tested with Monte Carlo simulations on complex polynomial functions. These showed that the optimization algorithm finds the global minimum in $\geq 99\%$ of cases. Still, it cannot be demonstrated that optimization over the payload electrical configurations finds the global minimum without an exhaustive exploration of the parameter space.

Eventually, 56 configurations were tested out of the possible 128 with the optimization procedure. For each configuration, the cost function in equation 2.7 consists of a set of values calculated on the 120 positions where the magnetic probe has been placed. All the configurations are sorted according to both the 75th percentile and the median of their distribution. The use of the 75th percentile is a conservative choice to select the configurations whose distribution has a low relative error for the majority of positions. In order to avoid degeneracy in the configurations, the optimization is carried out using two magnetic field sources (“Big” and “Small” coils). Thus it is less likely that two or more configurations can yield to the same measurement results. It is worth noting that, while the heuristic minimization does not guarantee to reach the global minimum, the optimization looks for a solution set whose relative error is below 10%.

Finally, 21 statistically equivalent configurations that can represent the real payload are identified: the equivalence was defined by using the non-parametric Kolmogorov-Smirnov test, applied to the sorted distribution versus the best one. Since all these configurations are statistically equivalent, all of them have to be considered in the evaluation of the magnetic noise contribution to the AdV sensitivity. The agreement between measurements and simulations is presented in Table 2.1, where also the reference case at 33 Hz is reported once more for comparison. These values are the overall statistics of the 21 equivalent configurations.

| | $(B_{meas} - B_{sim})/B_{meas}$ |
|--------------------------|---------------------------------|
| payload - “Small” Coil | -0.031 ± 0.021 |
| payload - “Big” Coil | -0.0057 ± 0.0220 |
| reference - “Small” Coil | 0.0052 ± 0.0090 |
| reference - “Big” Coil | -0.00093 ± 0.00740 |

Table 2.1: Mean and standard deviation of the relative difference function m between the measured and the simulated magnetic field values (Eq. 2.7), with and without the payload structure. Values combine measurements at both frequencies (73 Hz and 33 Hz) and they are given for the two source configurations (“Big” and “Small” coils).

2.3.4 Magnetic strain noise calculation

Within a frequency domain representation, a force F acting on the interferometer Test Mass (TM) displaces it, along the laser beam propagation direction (namely

\hat{z}), by $\Delta L = F/(M\omega^2)$, where M is the TM mass (kg) and $\omega = 2\pi f$ is the frequency in *rad/s*. The associated strain noise in the detector is

$$h_{mag} = \frac{\Delta L}{L_0}, \quad (2.8)$$

where $L_0 = 3000$ m is the interferometer arm length. The total h due to the payload magnetic coupling is the sum of a translational term plus a rotational one. Assuming linear superposition and that magnetic forces acting on the TM are uncorrelated over long range (so that the sum over the 4 TM is treated as incoherent), the term associated to the translational force is

$$h_{transl} = \frac{1}{L_0 M (2\pi f)^2} \sqrt{\sum_{j=1}^4 \left(\sum_{i=1}^4 F_i \right)_j^2}, \quad (2.9)$$

with $M = 42$ kg the mass of the AdV TM [86] and f the frequency. For each j -th TM, the total magnetic force is the sum of the force on the i -th magnet. Considering that all the magnetic moments are directed along the \hat{z} direction and that they are constant, the translational magnetic contribution is reduced to

$$h_{transl} = \frac{1}{L_0 M (2\pi f)^2} \sqrt{\sum_{j=1}^4 \left[\sum_{i=1}^4 \mu_i \left(\frac{\partial B_z}{\partial z} \right)_i \right]_j^2}, \quad (2.10)$$

where B_z is the magnetic field component along \hat{z} . As the force is directed along \hat{z} , only two kind of torques exist: τ_x and τ_y . They are respectively equal to

$$\tau_x = [(F_z^1 + F_z^2) - (F_z^3 + F_z^4)]y \quad (2.11)$$

$$\tau_y = [(F_z^1 + F_z^4) - (F_z^2 + F_z^3)]x, \quad (2.12)$$

where F_z^j are the magnetic forces on the four magnets ($j = [1, 4]$), as defined in figure 1.11b), while x and y are the components of the magnet position vector or, in other words, the force's application point relative to the TM center mass. Hence, the rotational contribution to the strain is

$$h_{rot} = \frac{D}{L_0 (2\pi f)^2} \sqrt{\sum_{j=1}^4 \left(\frac{\tau_{xj}}{I_{xx}} + \frac{\tau_{yj}}{I_{yy}} \right)^2}, \quad (2.13)$$

where I_{xx} and I_{yy} are the moments of inertia of the TM and D is the laser misalignment from the center of mass of the TM (assumed the same for each test mass). Finally, the total magnetic strain noise is computed, considering the same contribution on each TM:

$$h_{mag} = h_{transl} + h_{rot} = \frac{2F}{ML_0\omega^2} + \frac{2\sqrt{2}D}{IL_0\omega^2}(\tau_x + \tau_y), \quad (2.14)$$

with F , τ_x and τ_y respectively the total force and torques calculated on each TM. We also assumed identical inertia moments ($I_{xx} = I_{yy} = I$) and a conservative beam off-centering of $D = 1$ mm.

Force calculation

The FE model of the payload is placed in a uniform magnetic field of modulus 1T, directed along the \hat{x} , \hat{y} and \hat{z} axis alternately and with frequency range $f = [10, 2000]$ Hz.

Because of the magnets layout described in figure 1.11b, the total net force on the mirror is null. In practice, neither the total force nor the total torque are null due to non-ideal conditions: (i) dis-homogeneity of the magnetic field across the mirror area; (ii) dis-homogeneity of the magnetic moments; (iii) magnets position uncertainty. The last two effects are computed with Monte Carlo computational algorithms, with respect to the payload structure. The first Monte Carlo is used to extract a population of magnetic moments from a normal distribution centered on the nominal value $\mu = 2.20 \text{ mA} \times \text{m}^2$, with σ_μ being the accuracy provided by the manufacturer (15% of the nominal value), that is $\sigma_\mu = 0.33 \text{ mA} \times \text{m}^2$. The contribution to the net translational force is:

$$F_z = \nabla B \sum_{i=1}^4 \mu_i \sim 2(\nabla B)\delta\mu \quad (2.15)$$

The second Monte Carlo provides a population of magnet positions, uniformly sampled within a cubic volume of 2 mm per side, centered around the magnet nominal position. This takes into account the gradient spatial distribution. It can be similarly written

$$F_z = \mu \sum_{i=1}^4 \left(\frac{\partial B_z}{\partial z}\right)_i (-1)^i = 2\mu\delta(\nabla B), \quad (2.16)$$

this time at fixed μ .

From the set of equations 2.9 - 2.14, it can be seen that the magnetic strain noise calculation only requires knowledge of the force on the TM magnets, which in turn depends only on the magnetic gradient and the magnetic moment.

For what it concerns the magnetic gradient, it is observed that the main contribution to it come from B_z . The FE calculation proceeds by creating small volumes ($2 \times 2 \times 2 \text{ mm}^3$) around the nominal position of each magnets, which are meshed much finer than the remaining geometry, in order to improve the computation accuracy. The magnetic field gradient is computed in each point of the meshed volume using a finite difference method, for each of the four magnets on the mirror, for each of the 21 equivalent configurations found in the validation step and as a function of frequency. The magnets position uncertainty previously evaluated are also taken into account.

The net total force is computed from the force on each magnet (see the beginning of section 2.3.1) and it is the mean value of a Monte Carlo procedure to take into account deviation from the nominal parameters both in the magnet moments and in the magnet positioning. Three different case studies are considered in the Monte Carlo simulations: one takes into account both the contributions of the magnetic moment tolerance and the position error, while the other two study the effect of each single uncertainty. The two individual contributions turns out to be approximately equivalent and hence it is necessary to optimize both of them.

Figure 2.31 shows the results of the force simulation, based on both the contributions coming from magnetic moment and magnet position uncertainties, as mean




Figure 2.31: Monte Carlo simulations of the magnetic force on the AdV test masses, taking into account both magnetic moment and magnet position errors. The diamonds represent the mean values, while the grey band is the associated confidence level.

values and confidence levels as a function of the frequency. It can be noticed that, within the frequency range of interest, from 10 to 100 Hz, the magnetic force level remains almost constant around 2 mN.

Up until now, the force calculation has been carried out with a spatially uniform driving field. The actual force values, instead, depend on the real environmental magnetic field, which will be considered in the magnetic noise estimation.

Environmental magnetic field

The estimation of the real environmental magnetic field is based on a set of measurements carried out in August 2017. The magnetic field has been sampled at several locations inside the three experimental Virgo buildings. For each of them, the magnitude of the magnetic field vector is computed (figure 2.32).

The filtering effect produced by the steel tank that surrounds the suspension system is also taken into account. From section 2.2.2 it is known that the steel enclosure effectively acts as a first order low pass filter of Butterworth type, with a cut-off frequency of $f_{pole} = 4.85$ Hz (the most reliable evaluation is considered here). Thus, the estimation of the equivalent B_{env} is

$$B_{env} = \frac{\sqrt{B_{CEB}^2 + B_{CEB}^2 + B_{NEB}^2 + B_{WEB}^2}}{\sqrt{1 + (f/f_{pole})^2}} \quad (2.17)$$

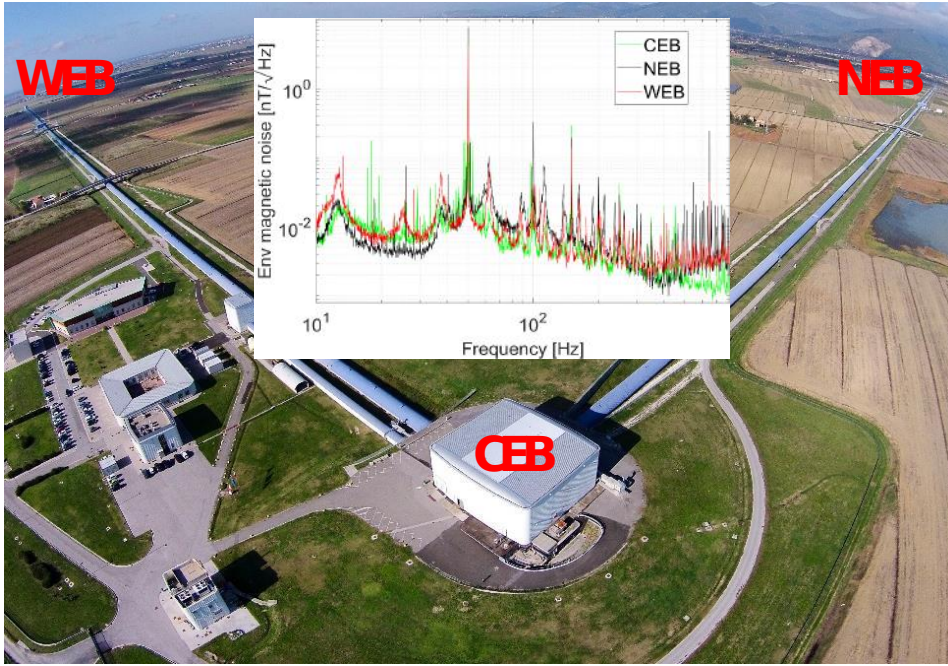


Figure 2.32: Aerial photo of the Virgo site. The environmental magnetic field activity spectrum, acquired in the three main buildings (CEB, WEB and NEB), is shown in the inset.

2.4 Magnetic noise projections

2.4.1 Projection from FE analysis

In figure 2.33a, the estimation of the magnetic noise contribution to the Advanced Virgo strain noise is shown and it is calculated by multiplying h_{mag} in equation 2.14 by the real magnetic field spectrum B_{env} of equation 2.17. The force values used to compute h_{mag} are an average on the 21 equivalent electrical configurations (see figure 2.31).

The translational and rotational contributions to h_{mag} are plotted separately for comparison. Clearly, the translational force contribution (blue line) is dominant on both rotational ones (red and orange lines). The comparison with the projected AdV sensitivity curves shows that magnetic noise (translational plus rotational) could be considered a nuisance only for the latest interferometer configurations (gray “Design” curve). In general though, the magnetic contribution stays well below the safety requirements, which are one order of magnitude below the best sensitivity limit.

Also, when the uncertainty due to the electrical configurations is taken into account, it can be noticed that the spread in the calculated force values is well within safe limits. Figure 2.33b shows the envelope region due to slightly different force values computed on the 21 electrical configurations. The contributions from both translational and torque forces are again considered in a combined way for the plot in figure 2.33b. Naturally, the upper expected magnetic noise gets closer to the AdV sensitivity curves, although it is still of moderate concern for the latest design specification only.

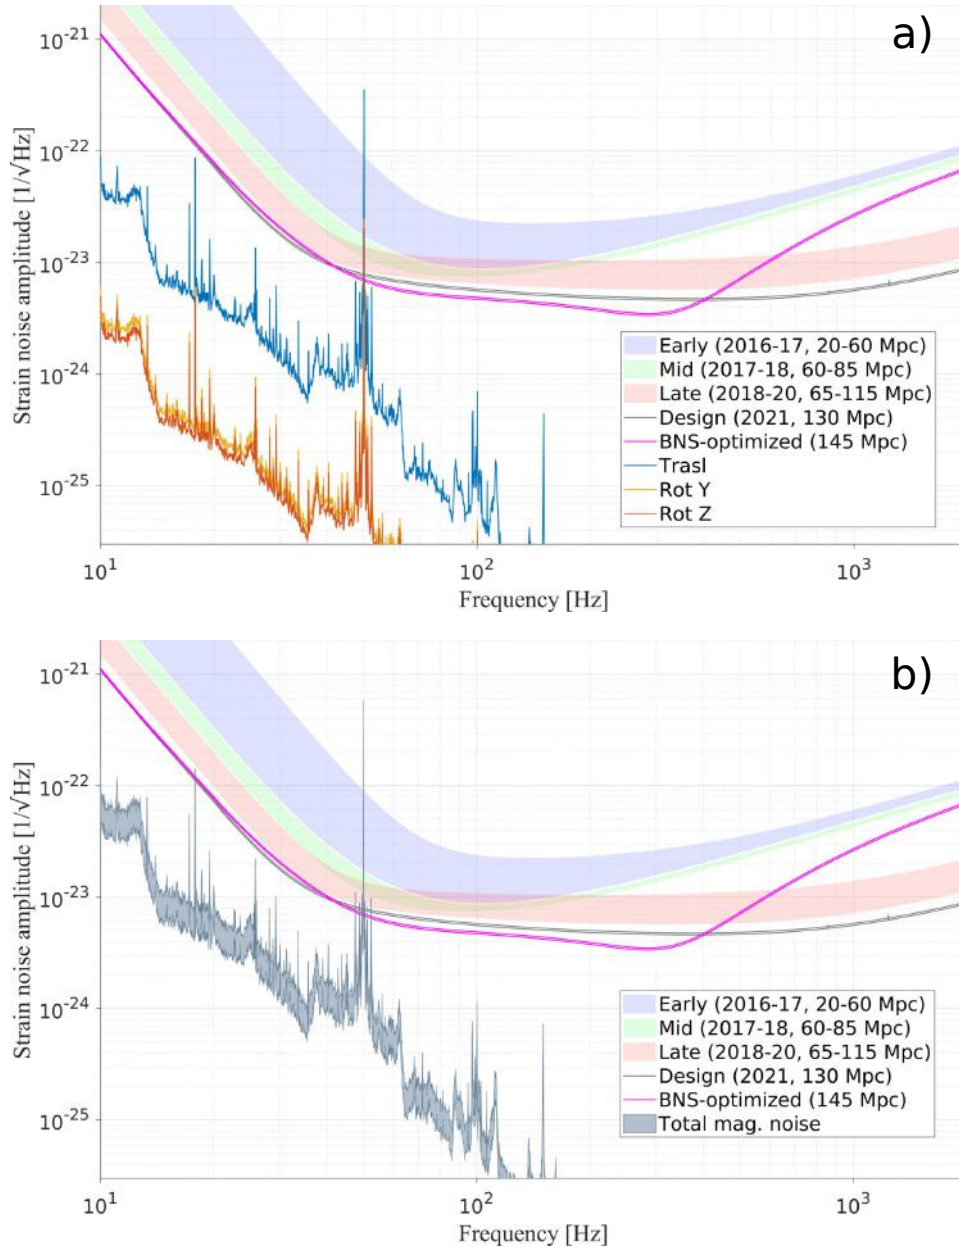


Figure 2.33: Contribution of the magnetic noise to the AdV sensitivity, in the frequency range of interest for astrophysical binary merger detection. The colored patches display the outcomes of the AdV sensitivity over the years, from the “Early” to the “Design” interferometer configurations. a) Spectra of both translational (blue line) and rotational (red and orange lines) magnetic effect on the TMs; b) Total spectra envelope (translational plus rotational) on all 21 electrical configurations.

2.4.2 Projection from far-field injections

On the wave of data analysis following the experimental magnetic noise injections (section 2.2.4), the MN budget can be evaluated once more. Figure 2.34 shows a comprehensive view of the AdV and AdV+ observing scenarios and the MN estimates. I specifically compare two completely independent evaluations: the FE analysis done in 2016 [139], described in section 2.3 and the experimental noise injections (both the 2017 and 2019 ones). The FE simulation study takes into account

a large number of possible electrical configurations of the payloads and eventually the sum of translational and rotational effects are taken. This simulations consider the impact on the 4 mirrors of the 2 Fabry-Pèrot cavities. Hence, the comparison with the magnetic injections is made possible by incoherently summing in quadrature the injections performed in the three main AdV buildings.

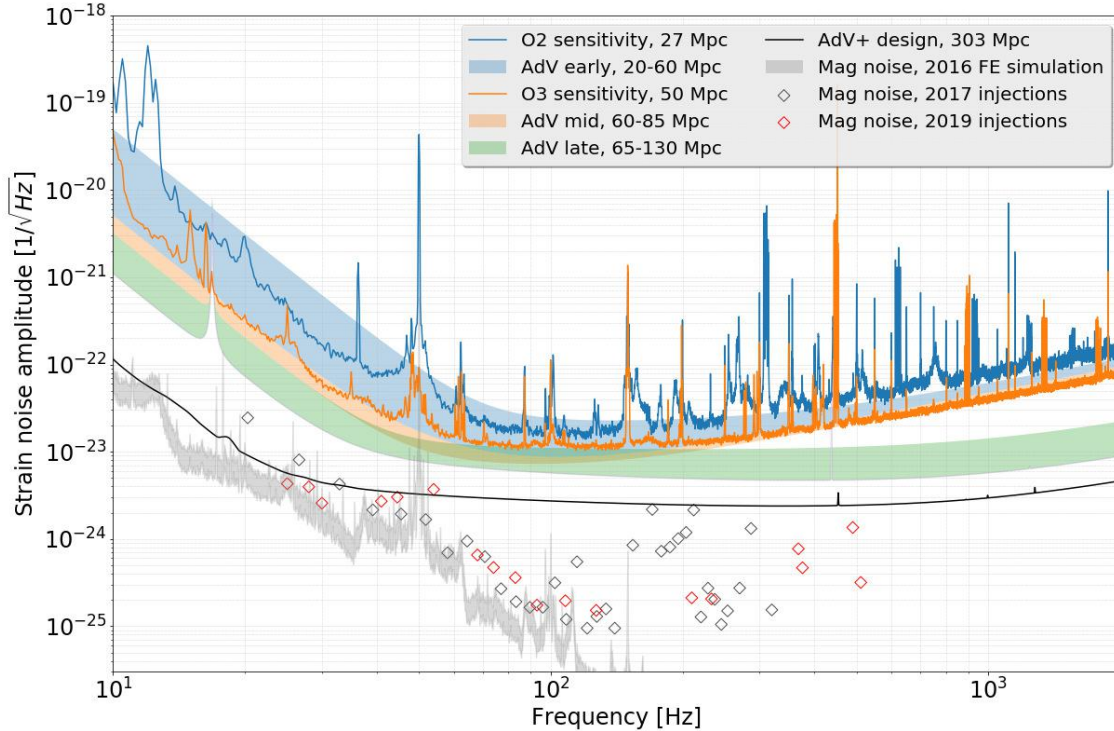


Figure 2.34: Magnetic noise budget. The relevant AdV and AdV+ present and future sensitivity outcomes are shown as reference. In particular, the blue and orange curves correspond to the best sensitivity obtained so far by AdV, during the O2 and O3 scientific runs. The gray band represents the FE simulation uncertainty among all the investigated electrical configurations of the payload. Diamonds represent the far-field magnetic noise injections, performed 2 years apart (gray for 2017, red for 2019).

The FE analysis provides a slightly more optimistic estimate than the one obtained with the current MN injections but, on the other hand, the slope of the strain noise spectrum (below 100 Hz) appears to be the same around f^{-3} . I also point out the substantial consistency of the MN level during a 2-year period, from 2017 to 2019.

At frequencies higher than 100 Hz, simulation and injections don't agree anymore, with the former continuously decreasing and the latter spreadly flattening. The rationale for this behavior is currently under investigation but it seems plausible to attribute it to the magnetic coupling with nearby electrical cables, an effect not taken into account in the FE simulations. However the influence of the MN on the instrument sensitivity is not concerning for the O3 observing run yet, but it might be for the next interferometer design phases. A similar study of the magnetic coupling was carried out at the Advanced LIGO (AdL) [113], with comparable results [150] (see section 2.6 for further details).

2.4.3 Projection with coherence

The far-field projections and the simulation-based approach are not the only ways to estimate the noise contribution to the sensitivity: a third possibility is to use the coherence between the magnetic and the GW channels as if the background magnetic activity were the (external) injected field. The main limitation is that a long time interval is necessary in order to get a meaningful value of the coherence, as the MN does not contribute significantly to the GW signal. With this method, it should be obtained an upper limit of the MN level, which can be compared with the projected noise contribution already got from the injections.

The background coherence analysis lasts for 10 hours, as data are acquired during a stable functioning interval of the interferometer (i.e. with the optical cavities in a locked state). The starting time point was on Monday 08 Apr 2019, at 12:00:00 UTC, and during the data gathering period the Binary Neutron Star observing range was around 50 Mpc. These figures are sufficient to calculate the coherence under stationary conditions by averaging over 360 data chunks, each one 100 seconds long (50% overlapping).

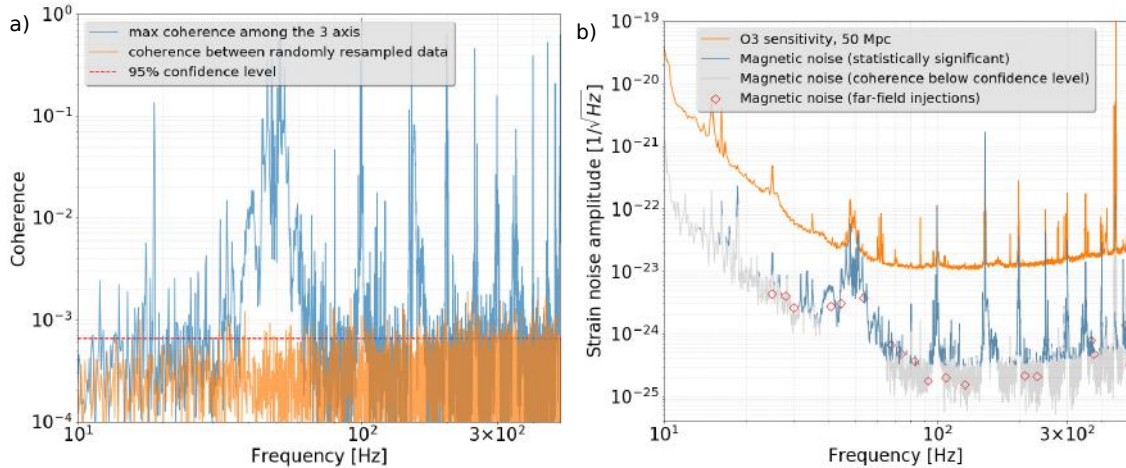


Figure 2.35: Magnetic noise projection with coherence. a) Maximum coherence at each frequency between the strain channel and the 3 magnetic probes in the CEB (blue line); the dashed red line indicates the statistical significance level of the coherence, calculated according to the 95% confidence level of the coherence distribution between randomly re-sampled data (orange line). b) Related magnetic noise projection (blue for statistically reliable bins, gray for coherence below significance level), compared to the most recent one obtained by the far-field injections (red diamonds) and the AdV sensitivity (orange line).

Figure 2.35a shows the maximum coherence value among the GW output channel $h(t)$ and the 3 magnetometers in the CEB (blue line), that is C_{mh} . The statistical significance of the computed coherence (dashed red line) is the 95% confidence level of the coherence distribution between randomly re-sampled data (orange line). Figure 2.35b shows a noise projection calculated according to:

$$P_{cohe} = \sqrt{C_{mh}} \cdot ASD_{hh}, \quad (2.18)$$

where ASD_{hh} is the amplitude spectral density of the magnetic activity spectrum. Here the statistically reliable coherence bins are drawn in blue, while the

remaining ones in gray. Results from the background coherence analysis are in good agreement with the ones obtained by the far-field injections acquired in the CEB during the latest measurement campaign (red diamonds). Incidentally, these results come from the analysis of the same time period. The advantage of the background coherence method is the possibility to obtaining a noise projection over all frequencies, as opposed to the discrete sampling of the frequency interval in the case of the artificial injections. The drawbacks are the need for long time data series in order to get an acceptable level of coherence and the influence of glitches that might occur during this particularly long time period. For this reason, the data sections are vetoed by using a glitch monitoring channel. Indeed, it could be interesting in the future to further investigate the broadening of the coherence around 50 Hz. For now, I can only speculate on the origin of these structures: possible causes could stem from significant phase-noise fluctuations of the mains (I remind the acquisition time of 10 hours) and/or non-linear behavior of the various power supplies and electronics.

2.5 Proposed mitigation strategy

As a last step, I present the study on a mitigation strategy to reduce the MN, on the chance it would be a limiting noise source in the future. In this regard, I have investigated various passive shielding options consisting of Helmholtz coils and / or spherical metal grids to be placed in the proximity of each TM vacuum chamber.

The idea is to passively reduce the intensity of the EM field near the TMs to reduce the gradients in the vicinity of the magnet positions. This works with noisy sources both on the external and internal side of the screen represented by the steel vacuum chamber. Ideally, a passive magnetic field cancellation device could decrease the field as an active cancellation system do, without requiring any power sources and feedback systems.

I used the COMSOL Multiphysics FE analysis software to simulate the system geometry and physics. The vacuum camera that surrounds the TMs is recreated with the same identical dimensions as the real one. Here I show the results on the two best solutions coming from different tested shielding layouts: the Helmholtz coils and a spherical mesh shield (figure 2.36a). The first ones have an external radius of 2 m / 5 m, an internal radius of 1.2 m, a thickness of 0.1 m and are positioned at 2 m distance from the TM.

On the contrary, the metal sphere is centered on the TM position, it has an external radius of 2 m and a thickness of 0.1 m. The FE model is then exposed to a spatially homogeneous and harmonic magnetic field of intensity 1T along \hat{z} . It has been discovered that the Helmholtz coils of radius 5 m are able to lower the external magnetic field up to a factor 10, with respect to the intrinsic shielding of the steel Tower, and homogenize it within 1 m from the center of the mirror along the axis y and z (figure 2.36b-c). Also, the screening effect (rated in the virtual TM position) increases with frequency (figure 2.36d) and the windings of the coil, as expected.

On the other hand, the configuration with the spherical grid has reached a better reduction level up to a factor of 20 but in a more restricted region. This point is supported by both longitudinal and orthogonal profiles. In this sense, the integration of the Helmholtz coils should be considered as a good possible future mitigation strategy, essentially in terms of enhanced homogenization power, simple implemen-

tation and limited overall space required. As a further validation, it is obtained a good agreement between the shielding experimental measures of the Tower (with the filter pole at 4.85 Hz, see section 2.2.2) and the numerical simulations (without further shielding, see figure 2.36d).

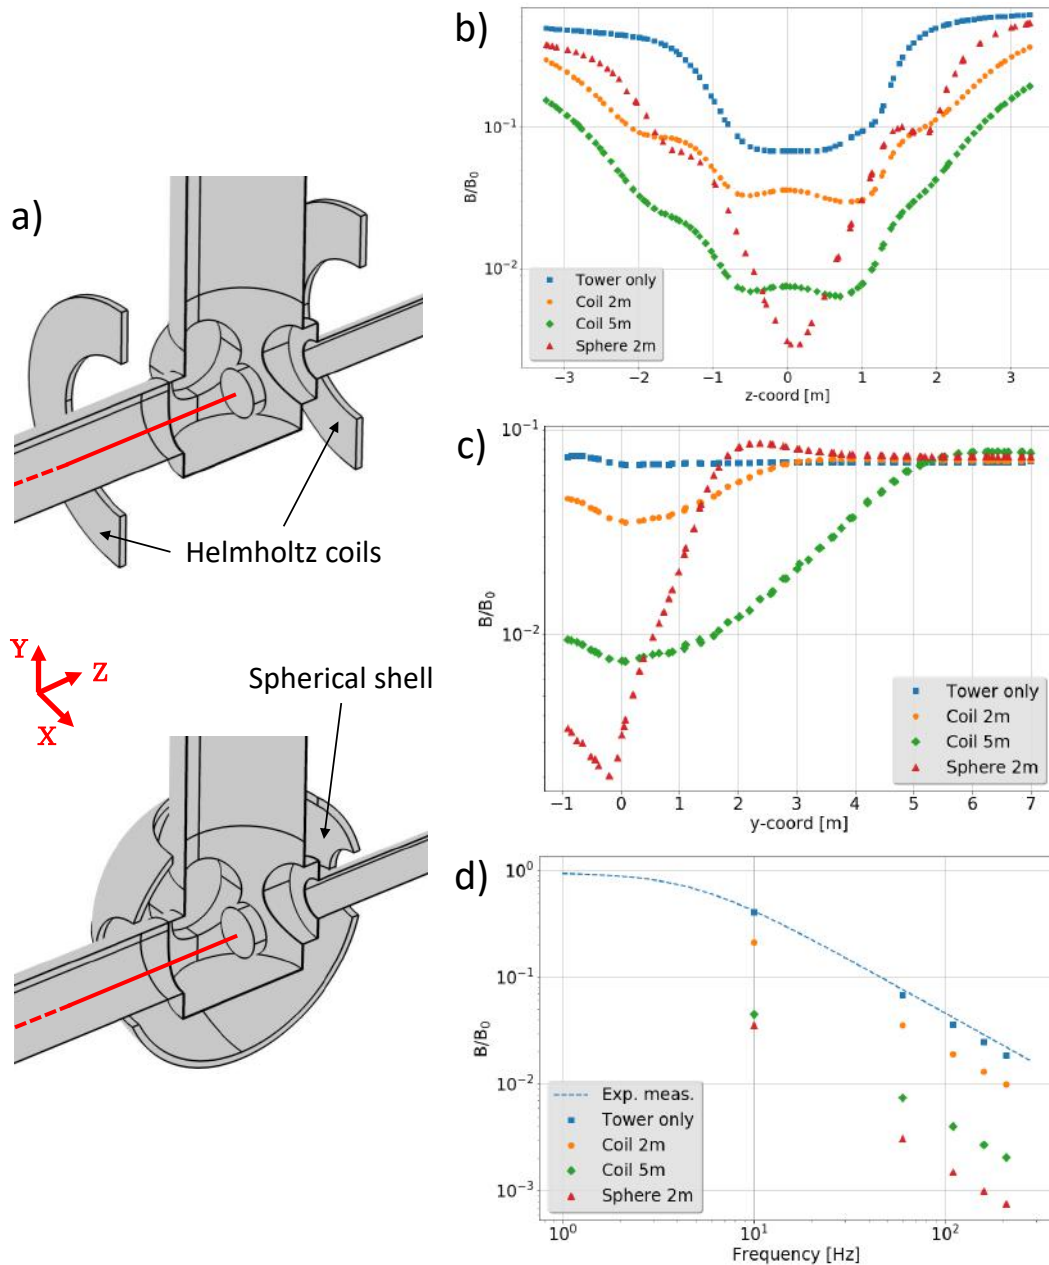


Figure 2.36: Finite Element analysis of the two investigated solutions for shielding the magnetic field: a) the Helmholtz coils and a spherical shell (different shapes and sizes); the red line follow the main laser path and the axes follow the Virgo convention. b) Longitudinal studies (along z) on the normalized reduction of the magnetic field. c) Orthogonal studies (along y) within a height range of $[-1.7]$ m, along the Tower axis. d) Frequency evolution of the normalized field at the TM location.

This is an exploratory mitigation analysis that can continue further in order to comply with the sensitivity objectives of the detector projects. Of course, separate

strategies could also be considered, such as a new TM actuation method that does not require parts with magnetic susceptibility. MN mitigation will also be important for the imminent construction of the Einstein Telescope GW facility, since its spectral sensitivity would be at least one order of magnitude higher than that of AdV+.

2.6 Discussion

Global magnetic noise

In this chapter, I started by describing ELF magnetic field measurements at various GW detector sites. The issue is the correlated MN from Schumann resonances, which is relevant to the observation of a Stochastic Gravitational Wave Background (SGWB) in GW detectors [55, 58].

I argued that the cavity formed by the surface of the Earth and the ionosphere is excited by the lightning activity around the world. A magnetic field of $0.5 - 1.0 \text{ pT}/\sqrt{\text{Hz}}$ is produced on the Earth's surface, which is roughly one order of magnitude lower than the environmental MN at CEB (see figure 2.16). Correlated noise, however, cannot be reduced through integration nor can it be mitigated through instrumental (re)design and/or background subtraction. For this reason a LIGO-Virgo joint experimental evaluation of the global MN contribution is worth investigating. In anticipation of that, it has been decided to keep monitoring the global MN activity at each GW site with dedicated low-noise magnetometers [58].

In summary, it has been shown that the magnetic fields associated with Schumann resonances are a possible source of correlated noise between advanced GW detectors. For example, without subtraction, the Schumann resonances induce correlated noise in advanced LIGO such that $\Omega_{mag} = 1 \times 10^{-9}$ [57], which is a potential limit. It is argued that, in order to do coherence measurements, the magnetometers should be located close to the interferometer sites but far enough from the local varying magnetic field activity. A careful treatment of the measurement setup would prevent the occurrence of unwanted transient noise and stationary spectral lines.

Optimal filters to perform subtraction of the Schumann correlated noise between magnetometers have been computed. In the recent paper by Coughlin et al [57], the authors show that magnetometer pairs thousands of kilometers apart (in that case, the NS Colorado, Villa Cristina and Poland magnetic antennas) are capable of reducing magnetic correlations by about a factor of 2 at the fundamental Schumann peak. This gave hope that magnetometers near to the interferometers can effectively subtract magnetic noise with Wiener filtering. From that study on, the level of reduction expected from an uncorrelated time series has been recently achieved and described in the current work. Going forward, it will be important to compute magnetometer correlations with GW detector data in order to measure the effect from the Schumann resonances. From there, subtraction using magnetometers can be performed. Bayesian techniques that aim to separate magnetic contamination from GW signals in cross-correlation search statistics are also being developed in parallel to those presented in this work. It is important to approach the issue of magnetic contamination with many different methods as it promises to be a significant problem for cross-correlation-based SGWB searches in the future.

As it shall be discussed, a supplementary approach to the subtraction of the

global MN would come from the application of machine learning based techniques, such as Artificial Neural Networks and Genetic Algorithms, which I already applied to the similar issue of newtonian noise cancellation.

Finally, it will be important to perform similar magnetometer measurements where the probes are aligned and separated by nearly 3 km distance, with the aim of mimicking more closely the scenario when magnetometers are used as witness sensors for a GW detector.

The great work of magnetic characterization performed at all of the test masses at the AdV detector has the additional goal of determining the level of the correlations of the Schumann resonances at multiple test masses.

Magnetic coupling modeling

The second piece towards the understanding and mitigation of the magnetic noise is represented by the magnetic coupling FE modeling.

In the text, I showed how the validation step has shed light on the definition of the electric configuration of the payload. The agreement between the simulated and the measured field is found to be within 5%. This value is the current limit of the FE model in representing the real payload and it is due to all the uncertainties and simplifications embodied both in simulations and in the experimental step.

The hierarchical and simplification steps are a reasonable heuristic approach to a problem with this complexity. Yet, these steps could have introduced further uncertainties, hence the need for further refinement and study of different decompositions. Moreover, the steel and titanium parts of the payload are neglected in the simulation, given that their contribution is suppressed by a factor $\approx 20 - 30$ due to the electrical conductivity.

Forces and displacements are calculated supposing a uniform environmental magnetic field in the area around the payload. This assumption is almost surely not verified in the real conditions. The only practical way to test the simulation forecasts was to measure the transfer function during AdV working conditions (i.e. during the commissioning phase).

The displacement produced by the force on the TMs summarizes all the forces acting on the four TMs of the two Fabry-Perot cavities. Identical payload structures and uncorrelated contributions are assumed, so that the total displacement is calculated in quadrature. These assumptions stem from the close similarity among the Input Mirror Payload (IMP) and the End Mirror Payload (the main difference is in the compensation plate, which is present only in the IMP) and the significant distance between them. That said, the TMs in the central building are rather close so that correlation in the magnetic response cannot be excluded. This issue is not considered here and needs further testing. Moreover, the Schumann resonances common-mode contribution is going to be addressed in a separate study, which lays its foundation on the paper [58].

Results show that the magnetic noise budget should not impact the AdV initial observations (2017-2020). Nevertheless, in a couple of small frequency ranges, the estimated noise is higher than the desired level (technical noises are required to be ≤ 0.1 of the incoherent sum of all fundamental noises). As I showed, these simulations are confirmed by experimental data on MN injections on the interferometer and this prompted the drafting of mitigation strategies. For instance, the environmental field was already addressed by past mitigation actions carried out on

initial Virgo (2009). These included the size reduction of TM magnets and power cable routing optimization. As a direct consequence of this work, the design of the payload structure was optimized to reduce gradients and shorten eddy currents paths. The agreement between simulation and experimental injections underlines the importance of an extensive validation phase that would imply a more realistic modeling of the payload environment.

Local environmental magnetic noise

In parallel, I have studied the overall MN contribution to the AdV sensitivity. The approach is experimental, and it is compared to simulation studies of various EM phenomena and their impact on the interferometer configuration. For this study, I had to deepen the understanding of the various local EM sources and their effects on the measuring process. The evaluation of the EM shielding properties of the steel tanks embedding the TMs was the starting point. The tank mostly behaves as a low-pass filter as expected, albeit with a mild slope inflection above 100 Hz. Then I identified and mitigated whenever possible the most annoying EM sources nearby the TM Towers, including the verification of the presence of any residual magnetism on the metallic enclosure walls. Indeed, a MN source like that of figure 2.19 would affect the AdV sensitivity below 100 Hz, if not removed.

With this preliminary work, the environment has been prepared for the MN injection campaign, in order to directly assess the ambient MN effect on the AdV output. Results show that there is substantial agreement between simulations and noise injections, and that MN is currently not an issue for the AdV sensitivity. There are, however, some discrepancies between the simulation and the noise injections at frequencies > 100 Hz, which suggest the presence of additional effects. The slope flattening in the transfer function is real (see section 2.2.2) and it can't be ascribed to the shielding properties of the metallic vacuum chamber. A possible explanation is that the simulation is limited to the payload structure and does not consider all the extra components, passive or active, in the area. Interestingly in Advanced LIGO (AdL) they notice a very similar behavior, therefore the general thought is to ascribe it to the magnetic coupling to all the electronic cables in the proximity of the vacuum chambers. While AdL uses a very different actuation scheme, with mixed electrostatic and EM actuators [113], AdV uses exclusively EM actuators for all the positioning corrections. Accordingly, AdV is expected to be more plagued by any magnetic disturbance. As a result, AdL is probably more dominated by coupling to cables and electronics, cable connectors and so on, which could be influenced by cabling issues and could enter the noise path through the marionette actuation stage. In this case, the magnetic noise impact should drastically change along with the general maintenance and commissioning of the interferometer on short time scales. For further reference, see the AdL technical note [202].

The small discrepancy found below 100 Hz between the payload magnetic simulations and the magnetic injections could also be explained by the updated payload design, whose details are not included in the simulation study and cannot be easily accounted for due to their small details and intricacies (for instance the installation of Fiber Guards to shield fused silica fiber from high speed particles [203]).

Similarly to what has been proposed to monitor the global Schumann activity, also routine MN injections on a weekly/monthly basis, during the normal scientific run, are foreseen to keep track of the magnetic transfer function. The injections

would be made with a series of permanent coils at each building, lasting less than an hour per injection cycle.

In the next few months, AdL and AdV will complete the third observing run with increasingly large sensitivity and number of GW detections. The development of noise reduction techniques to improve the opportunities of detection, especially for low SNR events, is therefore a crucial effort. The MN mitigation strategy, consisting in the integration of a passive Helmholtz coil system around the most sensible TM Towers, may thus prove to be a valuable one. The present study has shown that, if the GW detectors reach their design sensitivity, some MN mitigation strategy might become relevant. That is also the case for the future third generation ground-based GW detectors as the European Einstein Telescope (ET) [159], whose low frequency spectral sensitivity would be certainly limited by the above documented MN level.

Chapter 3

Seismic Newtonian noise

In this chapter I will discuss the main features of the newtonian noise, and the strategies necessary to suppress it, especially in the first phase of the AdV+ project.

I will start from the basic principles regarding exclusively the seismic component of the newtonian noise (see section 3.1), including a number of analytical and numerical calculations, in order to make the reader aware of the progress that has been made on the physical interpretation of newtonian noise.

The conventional approach towards the seismic newtonian noise (SNN) cancellation is to use Wiener filters on sensor arrays data, similar to the scheme already used for global Schumann resonances subtraction (see section 2.1) and other detector noises [204]. The foundations on which the SNN cancellation system is based will be uncovered in section 3.3, focusing on seismic sensor types and array configuration. A general overview of the SNN cancellation framework is also presented in section 3.2.

During my work, I started to investigate on a novel noise subtraction approach based on Machine Learning (ML) techniques. I will extensively return on this topic throughout the current chapter, specifically in section 3.4.

Further and current efforts are dedicated to the crucial issue of the optimization of the placement of the seismic sensors [205], which will be the subject of discussion in section 3.5.

3.1 Basic principles

3.1.1 Seismic waves

Among the many types of seismic waves, one can generally divide them into body waves, which travel through the Earth, and surface waves, which travel at the Earth's surface. A comprehensive review is given in [206]. Here, I summarize the gist of the theoretical framework. Two types of particle motion result in two types of body waves: the primary compressional waves and the secondary shear waves.

In the case of compressional waves, the displacement is parallel to the direction of wave propagation. They propagate through a material by alternately compressing and expanding the medium and they are the first to arrive after an earthquake (i.e., they are the fastest waves). For this reason, compressional waves are sometimes called “P-waves”, where P stands for primary. P-waves are characterized by a frequency $\omega = k^P \alpha$ with a wave vector \vec{k}^P and a wave speed $\alpha = \sqrt{(\lambda + 2\mu)/\rho}$ in

terms of the Lamé constants λ , μ and the density of the medium ρ . The displacement field of a plane P-wave can be written as

$$\vec{\xi}^P(\vec{r}, t) = \hat{e}_k \xi_0^P(\vec{k}^P, \omega) e^{i(\vec{k}^P \cdot \vec{r} - \omega t)}. \quad (3.1)$$

On the contrary, shear waves produce transversal displacement and are denoted as “S-waves”, where S stands for secondary as they always follow P-waves during earthquakes. The S-wave speed is equal to $\beta = \sqrt{\mu/\rho}$ and the displacement

$$\vec{\xi}^S(\vec{r}, t) = \hat{e}_p \xi_0^S(\vec{k}^S, \omega) e^{i(\vec{k}^S \cdot \vec{r} - \omega t)} \quad (3.2)$$

where \hat{e}_p is the polarization vector and $\hat{e}_p \cdot \vec{k}^S = 0$. Together, P-waves and S-waves are denoted as body waves, since they propagate in all directions through media. Typical wave velocities values for clay-like soils range from 500-1000 m/s for α and 250-600 m/s for β . The overall seismic displacement produced by body waves can be represented in terms of the Lamé potentials as

$$\vec{\xi}(\vec{r}, t) = \nabla \phi_s(\vec{r}, t) + \nabla \times \psi_s(\vec{r}, t) \quad (3.3)$$

where $\phi_s(\vec{r}, t)$ and $\psi_s(\vec{r}, t)$ are respectively called P-wave and S-wave potentials. Besides bulk waves, Rayleigh waves arise in correspondence of the surface of media. They are produced by the interaction of P-waves and S-waves, with a 90° phase difference between these two components that gives rise to elliptical particle motion. Surface waves are characterized by lower frequencies and larger amplitudes than those of body waves (see an example of an earthquake seismogram in figure 3.1). In particular, Love-waves are surface waves with a zero P-wave component.

A useful quantity to link the speeds of body waves (α, β) and surface Rayleigh waves (c_R) is the Poisson’s ratio ν of a medium, defined as the ratio of transverse contraction strain to longitudinal extension strain in the direction of the acting force:

$$\beta = \alpha \cdot \sqrt{\frac{1 - 2\nu}{2 - 2\nu}}, \quad (3.4)$$

$$c_R \rightarrow R((c_R/\beta)^2) = 0, \quad R(x) = x^3 - 8x^2 + 8x \frac{2 - \nu}{1 - \nu} - \frac{8}{1 - \nu}. \quad (3.5)$$

The horizontal Rayleigh wave speed c_R is slightly lower than the S-wave speed. In particular, seismic body P-waves and surface Rayleigh waves, by producing alternating compressions and dilatations, result in local density fluctuations of the crustal rocks.

In 1995, Miller and colleagues studied the percentage of the three kinds of elastic wave of the total seismic energy input, originated from single-element sources [207]. They found that Rayleigh waves accounted for 67.3%, S-wave accounted for 25.8% and P-wave accounted for 6.9%. Also from seismic field measurements at LIGO Hanford, it is known that the major seismic wave component among the frequencies relevant for SNN cancellation is Rayleigh waves. This identification comes from the calculation of the seismic wave speed at the TM location [208]. As one can see in the plot of figure 3.2, the histograms are centered at average velocities of about 300m/s at 10 and 15 Hz and 380m/s at 20 Hz. Outliers might be from body waves or aliasing effects, which can originate from the overlap of several waves at the same frequency simultaneously propagate through the array. However, these measurements are consistent with Rayleigh waves, which simplifies the SNN modeling.

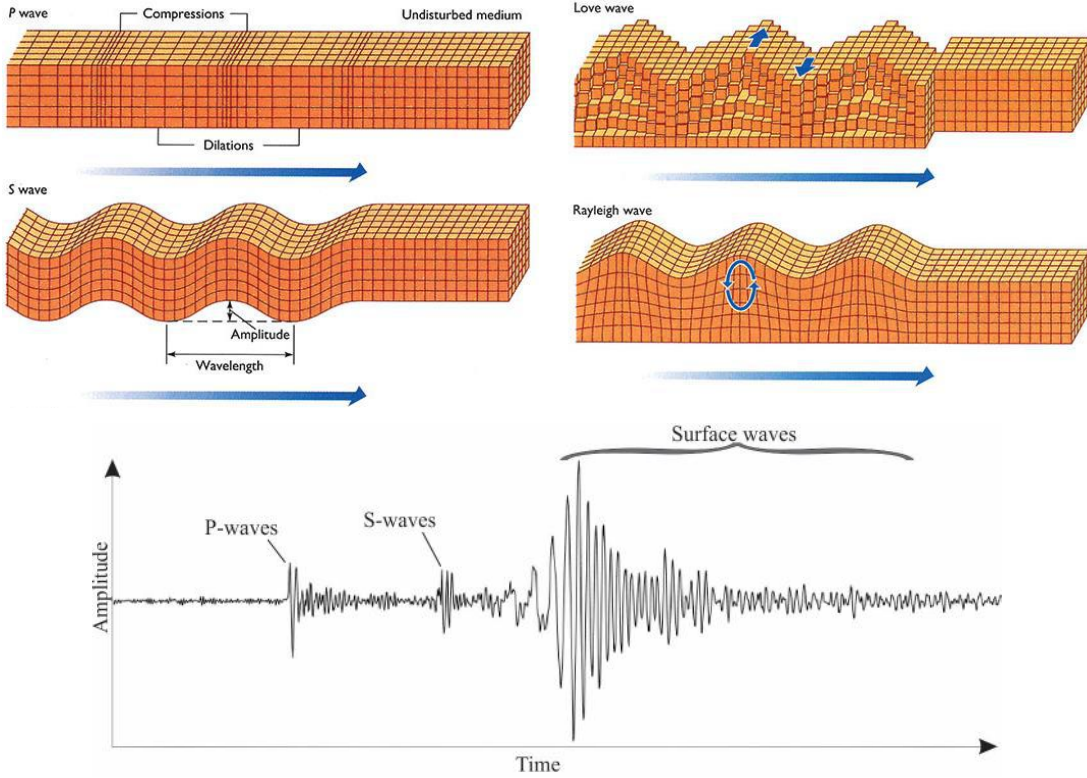


Figure 3.1: Illustration of the 4 main kinds of seismic waves in a solid: body waves (P and S) on the left and surface waves (Love and Rayleigh) on the right. The particular waveform of an earthquake detected at a point on the earth surface is shown in the lower time-amplitude plot: this is called Seismogram.

3.1.2 Seismic gravity perturbations

A local density modification leads to a perturbation of the gravity field. As these variations are tiny compared to the typical wavelength of the seismic field, newtonian gravity is enough to describe them. For a density variation $\delta\rho(\vec{r}, t)$, the gravitational potential at point r_0 behaves as

$$\delta\Phi(\vec{r}_0, t) = -G \int dV \frac{\delta\rho(\vec{r}, t)}{|\vec{r} - \vec{r}_0|}. \quad (3.6)$$

On Earth, different sources of fluctuations of the gravity field, associated to different mechanisms of density variation, are observed (see section 1.2.4 in the for a complete list). In the case of ground-based GW interferometers, these gravity field perturbations induce motions of the test masses generating a displacement noise. Among these mechanisms of density variation, the one of interest here is the perturbation of the gravity field induced by the variation of density of Earth's crust caused by a seismic wave passing through. I will refer to it as Seismic Newtonian Noise, or with the acronym SNN.

The continuity equation gives an expression for the density perturbation $\delta\rho(\vec{r}, t)$ caused by a seismic displacement

$$\delta\rho(\vec{r}, t) = -\nabla \cdot (\rho(\vec{r})\vec{\xi}(\vec{r}, t)) \quad (3.7)$$




Figure 3.2: Histograms of seismic speed measurements using LIGO Hanford array data at 10, 15, and 20 Hz. Figure reproduced from [208].

where the assumption that the density perturbation is much smaller than the unperturbed density is made (e.g. $\delta\rho(\vec{r}, t) \ll \rho(\vec{r})$). Equation 3.6 now reads

$$\begin{aligned} \delta\Phi(\vec{r}_0, t) &= -G \int dV \rho(\vec{r}) \vec{\xi}(\vec{r}, t) \cdot \nabla \frac{1}{|\vec{r} - \vec{r}_0|} \\ &= G \int dV \rho(\vec{r}) \vec{\xi}(\vec{r}, t) \cdot \frac{\vec{r} - \vec{r}_0}{|\vec{r} - \vec{r}_0|^3}. \end{aligned} \quad (3.8)$$

Then the perturbation of the gravity acceleration is obtained by taking the gradient of equation 3.8 with respect to \vec{r}_0 :

$$\begin{aligned} \delta\vec{a}(\vec{r}_0, t) &= -G \int dV \rho(\vec{r}) (\vec{\xi}(\vec{r}, t) \cdot \nabla_0) \cdot \frac{\vec{r} - \vec{r}_0}{|\vec{r} - \vec{r}_0|^3} \\ &= G \int dV \rho(\vec{r}) \frac{1}{|\vec{r} - \vec{r}_0|^3} (\vec{\xi}(\vec{r}, t) - 3(\vec{e}_{rr_0} \cdot \vec{\xi}(\vec{r}, t)) \vec{e}_{rr_0}), \end{aligned} \quad (3.9)$$

where $\vec{e}_{rr_0} \equiv (\vec{r} - \vec{r}_0)/|\vec{r} - \vec{r}_0|$. The equations set in this form provide the basis for the FE models of the seismic gravity perturbations (see section 3.5.2 and section 3.5), where each meshed element has a mass $\rho(\vec{r})\delta V$. Note that equation 3.9 describes seismic gravity perturbations inside infinite bulk media (suited for underground-based detectors, e.g. ET and KAGRA) as well as media with surfaces (suited for surface-based detectors, e.g. AdV and AdL).

3.1.3 The Saulson model

In this section, the forces due to density fluctuations produced by plane P-waves in a homogeneous half space are calculated. I refer to the case of figure 3.3, which is a simple scheme of Newtonian noise generation in an interferometric detector with L-long arms: the coherently fluctuating mass region $M(t)$ induces a varying gravitational acceleration on the test masses. Imagine that in this region there is a

fluctuation $\Delta M(t) = M(t) - \langle M(t) \rangle$, the force induced on a test mass (m) is equal to

$$\frac{\vec{F}(t)}{m} = \frac{G\Delta M(t)}{r^2} \hat{r}. \quad (3.10)$$




Figure 3.3: Basic scheme of Newtonian noise generation in an interferometric detector with L-long arms, used by Saulson to present his analytical model. The figure is reproduced from [117].

Transforming into the frequency domain and taking the projection along the x direction

$$\frac{F_x}{m} = G\Delta M(\omega) \frac{\cos \theta}{r^2} \quad (3.11)$$

If the test mass is considered to be suspended from a pendulum, the system can be described as an oscillator with resonant frequency ω_0 and damping time τ . Hence, equation 3.11 can be rewritten in the form of $a = \omega^2 x$, expressing the magnitude

$$\left[(\omega_0^2 - \omega^2)^2 + \frac{\omega^2}{\tau^2} \right] |x(\omega)|^2 = G^2 |\Delta M(\omega)|^2 \frac{\cos^2 \theta}{r^4}. \quad (3.12)$$

This is true for only a single coherently fluctuating region, but the formula can be extended to the case of many equivalent regions of fluctuating mass, all within a coherence radius of order $\lambda_R/2$, with $\lambda_R = c_R/f$ the wavelength of Rayleigh waves. A second assumption to make is that the gravitational forces generated by fluctuating masses in different regions are uncorrelated and independent, so that they can be added in quadrature (verified if $\lambda_R \ll L$). The sum over all the regions is reduced to an integral over θ and r of the last term in equation 3.12:

$$\sum \frac{\cos^2 \theta}{r^4} \approx \frac{1}{\lambda_R} \int_{r_{min}}^{\infty} \int_0^{2\pi} \frac{\cos^2 \theta}{r^4} d\theta dr = \frac{64\pi}{3\lambda_R^4} = \frac{64\pi}{3c_R^4} \left[\frac{\omega}{2\pi} \right]^4. \quad (3.13)$$

The integral converges if an inner cutoff radius $r_{min} = \lambda_R/4$ is introduced. This fact indicates that the random gravitational force is strongly dominated by the

nearest few coherently fluctuating regions. Substituting equation 3.13 into equation 3.12

$$\left[(\omega_0^2 - \omega^2)^2 + \frac{\omega^2}{\tau^2} \right] |x(\omega)|^2 = \frac{4G^2}{3\pi^3 c_R^4} \omega^4 |\Delta M(\omega)|^2 \quad (3.14)$$

The next step is to cast the formula in terms of an observable: $|\Delta M(\omega)|$ can be expressed as a function of the displacement of a point on the surface from its equilibrium position, produced by the passing seismic wave:

$$|\Delta M(\omega)|^2 = \frac{\pi \rho^2 \lambda_R^4}{16} |\Delta X(\omega)|^2 = \frac{\pi^5 c_R^4 \rho^2}{\omega^4} |\Delta X(\omega)|^2 \quad (3.15)$$

where ρ is the local density of the medium. If we consider $|x(\omega)|$ as the differential displacement between the test masses $|\Delta x(\omega)|$ and we take into account that the contribution from each of the 4 test masses is added in quadrature, we obtain the following relationship

$$\left[(\omega_0^2 - \omega^2)^2 + \frac{\omega^2}{\tau^2} \right] |\Delta x(\omega)|^2 = \frac{16\pi^2 G^2 \rho^2}{3} |\Delta X(\omega)|^2. \quad (3.16)$$

In other words, the previous equation gives the ASD of the oscillation transmitted through the SNN to the test masses $|\Delta x(\omega)| \equiv \tilde{x}_{NN}(\omega)$ in terms of the ASD of the ground oscillation induced by the seismic noise $|\Delta X(\omega)| \equiv \tilde{x}_{seis}(\omega)$:

$$\tilde{x}_{NN}(\omega) = \frac{4\pi G \rho}{\sqrt{3(\omega^2 - \omega_0^2)^2 + \frac{\omega^2}{\tau^2}}} \tilde{x}_{seis}(\omega) \approx \frac{G \rho}{\sqrt{3\pi} f^2} \tilde{x}_{seis}(\omega) \quad (3.17)$$

where the approximation is based on the assumption that the resonance and the damping of the TM down to 1 Hz can be ignored, for simplicity. For an average density of the soil near the AdV site of $\rho \approx 2 \times 10^3 \text{ kg/m}^3$, it is found that

$$\tilde{x}_{NN}(f) \approx \frac{2.4 \times 10^{-8}}{f^2 / \text{Hz}^2} \tilde{x}_{seis}(f) \quad (3.18)$$

Since the main contribution to the SNN comes from compressional P-waves, which have a larger amplitude on surface (e.g. Rayleigh waves), it is argued that SNN is attenuated in underground sites, where also the anthropogenic seismic noise is reduced.

The equation 3.18 is however strongly model dependent and in the absence of SNN observations, studies of noise cancellation schemes rely on precise modeling. Models of seismic NN have been gradually refined over the past decades [117–120].

A predicted SNN is shown in figure 3.4 on the right, together with a reference sensitivity of AdV. While the average SNN spectrum lies below other instrumental noise and therefore it is not a limiting noise source for AdV, it is believed that, due to occasional stronger transients in the seismic field, SNN can exceed other low frequency noises [123]. This is why it is always preferable to plot the histograms instead of an average spectrum, since seismic noise (figure 3.4 on the left) is in general non-stationary in the long term (especially below 10 Hz) and therefore relatively large variations in spectra are often witnessed. Moreover, it has been predicted that SNN will be relevant at low frequencies in the future upgrades of advanced interferometers [120].




Figure 3.4: Histograms of seismic spectra acquired at CEB (left) and modeled Rayleigh wave SNN (right). A sensitivity target of AdV is plotted for comparison. Seismic data are taken from the time period between June 4, 2011, UTC 00:00 and September 3, 2011, UTC 00:00. The figure is reproduced from [120].

3.2 Newtonian noise cancellation

Newtonian noise is a name to refer to gravity gradients caused by density fluctuations of the surrounding media (rocks, air, water etc...). Those gradients eventually cause a TM displacement, completely similar to that resulting from a GW arrival. Seismic waves are the most impending sources of newtonian noise for the next Virgo configurations and hence a prompt cancellation strategy is required. Subtraction techniques based on active/passive shielding of the sensible components of the interferometer are ineffective on seismic newtonian noise (SNN). Currently, the most accredited procedure involves the deployment of a seismic sensor array around each TM location, in order to evaluate the surface seismic displacement and put in operation some filtering method to predict the newtonian forces on the TM. This procedure require the calculation of the subtraction residuals by analytically evaluating the correlation between sensors and the TM as a function of frequency and sensor locations:

$$R(f) = 1 - \frac{\vec{C}_{SN}^\dagger(f) \cdot (C_{SS}(f))^{-1} \cdot \vec{C}_{SN}(f)}{C_{NN}(f)}. \quad (3.19)$$

Here, frequency-domain correlations are expressed as cross power spectral densities (CSD). The subscripts S and N stands respectively for sensor and newtonian noise. The performance of a subtraction filter is indeed quantified by $R(f)$, which represents the relative residual SNN spectrum that a subtraction filter leaves in the GW data. Equation 3.19 implies that, in order to achieve a good subtraction efficiency, the sensors should correlate as much as possible with SNN (large C_{SN}) but as little as possible with different sensors (small C_{SS}). Moreover, the intrinsic noise of the sensors should be also small. The best possible cancellation using N equal sensors with a certain SNR value is then

$$R_{min}(f) = 1 - \frac{1}{1 + 1/(N \cdot SNR(f)^2)} \approx \frac{1}{N \cdot SNR(f)^2}. \quad (3.20)$$

Accordingly, the residuals from optimal arrays will always decrease at least as $1/N$, but they can only be achieved if geometrical limitations are ideally overcome,

which is not the case even for an infinite number of sensors. Interestingly, equation 3.19 does not depend on any details about the way subtraction is implemented and hence it describes the performance of all subtraction methods that are based on linear filtering. Indeed, a substantial part of my PhD was devoted to the development of a subtraction technique that can be alternative and/or complementary to the standard Wiener filter. Nevertheless, the type of filter and the algorithm used to calculate it are supplementary constraints on the noise cancellation quality.

The optimization of the sensor array is not trivial to find. Moreover, it is already known that a step-wise optimization by placing one sensor at a time leads to a sub-optimal array. Therefore, only a global optimization of all seismometer locations simultaneously can lead to optimized configurations. An example of optimal array configuration for 10 noiseless sensors can be found in figure 3.5 (reproduced from [123]), where a specific plane-wave composition of the seismic field and the minimization of equation 3.19 are run. Although the optimal array should have some kind of symmetry, the authors argue that the subtraction performance they obtained should be still sufficient for advanced LIGO detectors. Some rules of thumb are uncovered: very small arrays does not perform well at low frequencies (1-20 Hz) due to high degeneration; a small number of sensors poorly performs at high frequencies ($f > 10$ Hz).

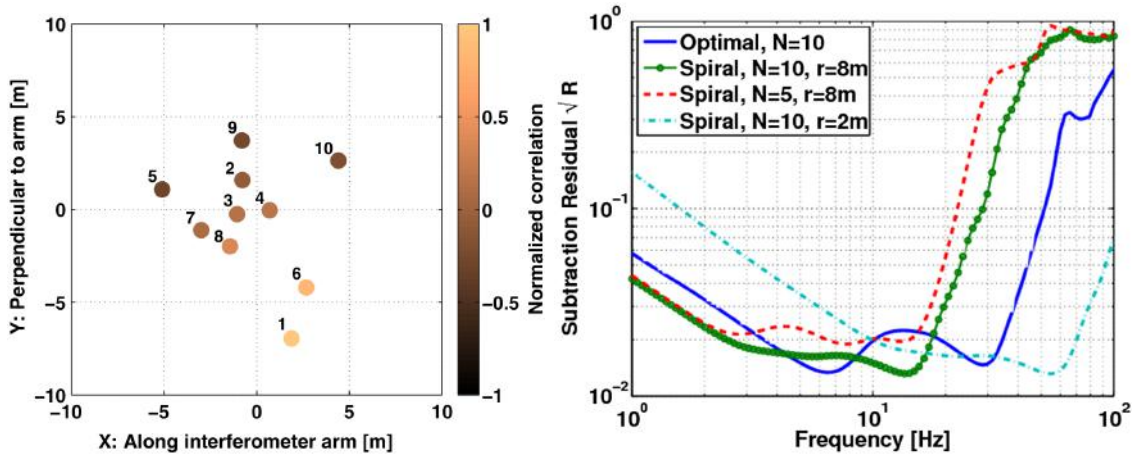


Figure 3.5: Left: locations of 10 sensors resulting from numerically minimizing the subtraction residual. The TM is located at (0,0). The colors indicate the normalized seismic correlation between sensor 1 and all other sensor. Right: subtraction residual as defined in equation 3.19 for the array shown on the left and three other spiral configurations. It is assumed that the SNR of the sensors is constant at all frequencies and equal to 100.

A limit of these models is that the residual is a function of seismic wave correlation length and hence frequency. There, for example, the residuals were minimized at 10 Hz for a Rayleigh wave speed of 200 m/s. For this reason a broadband subtraction needs to be investigated. Secondly, the seismic activity at Virgo is further diversified and the sub-surface level infrastructures distort the seismic field in unpredictable ways. With these preconditions, the plane-wave models have only the merit for analytical understanding, since the complexity of the real seismic field (seismic scattering plus local sources) makes them impractical. Consequently, the sensor array optimization is still an unresolved problem. An approach based on finite ele-

ment (FE) modeling of the system can be the answer, mainly for the following two reasons.

- A FE model is able to simulate the effect of exceptionally complex geometries. In this case, it should be possible to take into account the building structures and foundations and different soil layers.
- As long as the SNN is not the most limiting noise at low frequency for the Virgo GW detection, the C_{SN} vector in equation 3.19 can not be analytically calculated, unless a SNN model is used. Indeed, the alternative is to consider an additional seismic sensor as a proxy for the TM. In this respect, a sensor measuring the seismic tilt is maximally correlated with surface seismic waves and hence it could prove to be an optimal choice, provided that it is deployed very close to the TM. On the contrary, the FE analysis allows to evaluate all the terms in the residual's equation, including the direct newtonian forces on the TM.

One main ingredient for this kind of FE simulation is the driving seismic field, which can be represented by point-like sources of sinusoidal signals. However, the simulated field has little to do with the real field. So the idea is to build a high number of temporal FE models, each time by changing the excitation seismic field configuration, and train a machine learning system to recognize the one that fits better with the real seismic measurements.

The general scheme of the present study is summarized in figure 3.6. The two followed lines are the target seismic sensor prediction and the optimization of the sensor array. The former is carried out by the application of standard Wiener filtering (section 3.3.4) and machine learning with optimized neural network architectures (section 3.4), both in time and frequency domain. The latter is instead tackled with FE analysis and the standard analytical approach (section 3.5).

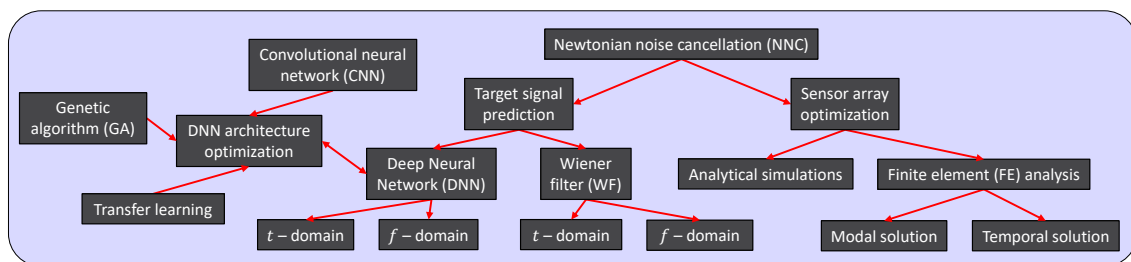


Figure 3.6: Seismic newtonian noise cancellation framework.

3.3 Cancellation system infrastructure

Seismic field measurements at the GW detector sites suggest that SNN from surface Rayleigh waves is the major concern in the band between 10 and 20 Hz (see the end of section 3.1.1). Hence, a mitigation system can be designed by deploying a surface array of sensors around each test mass (inside and outside the infrastructure main buildings) monitoring vertical ground displacements [123]. Such a cancellation scheme needs to address:

1. The optimization of the filter used to estimate the SNN from seismic data
2. The optimal placement of the seismic sensors

The optimization of the array placement is important for maximizing the efficacy of the cancellation and it will be discussed in section 3.3.2. Furthermore, the traditional cancellation scheme is based on the linear MISO (multiple-input single-output) Wiener filter (see section 3.3.4), typically calculated from observed correlations between sensors and target channels. However, the cancellation system infrastructure described here is suited for a number of different noise subtraction techniques, according to the type of the analyzed data. Thereafter, I will present a deep learning driven approach to the subtraction of SNN, based on artificial neural networks.

3.3.1 Seismic sensors

The Data Acquisition (DAQ) system is provided with a simple Synchronization and Powering Unit (SPU) for each sensor array, whose data processing is done by real-time at a central computer room. The role of the DAQ system is to manage, acquire and buffer data from the seismic sensors installed at Virgo's facilities. The system assure synchronization of sensors being located in different areas around each facility. The mutual sensor distance and the distance from the central unit of the system never exceed 15 meters. A single SNN DAQ system is able to host up to 64 sensors connected in a mixed star-chain topology. Seismic data gathered from the sensors are transferred on to Virgo's data storage system via Ethernet. The working principle of the SNN DAQ system is presented as a system block diagram in figure 3.7.

Technically speaking, a seismometer measures the offset or motion of the ground, while an accelerometer measures its proper acceleration. A simple seismometer is a mass hanging from a spring, both suspended from a frame that moves along with any motion detected. The relative motion between the mass and the frame provides a measurement of the vertical ground motion. Seismograph output is generally given as displacement vs time, whereas accelerometer output is given as acceleration vs time. For this reason, the reference mass of a seismometer operates above its resonant frequency and stays relatively still, while measuring acceleration usually requires working at frequencies well below the mass/spring resonance, as the mass tries to follow the instrument's case. The mechanics involved in both devices are generally similar (conceptually measuring the response of a mass on a spring), but optimized for the application of interest. The accelerometer targets a stronger motion with respect to a common seismometer because it is much less sensitive, but have a much greater range. Hence, the data from accelerometers can be essential to study earthquakes, since the instrument remains responsive during the strongest seismic shaking. Seismometers also generally do not move very far or change orientations rapidly, as they are fixed to a specific piece of rock or ground. On the contrary, accelerometers are easier to handle and they can be placed onto vertical or uneven surfaces. Modern accelerometers are often small micro electro-mechanical systems (MEMS), consisting of little more than a cantilever beam with a proof mass [209]. The following should be close to a complete list of the seismic sensors available at Virgo site:

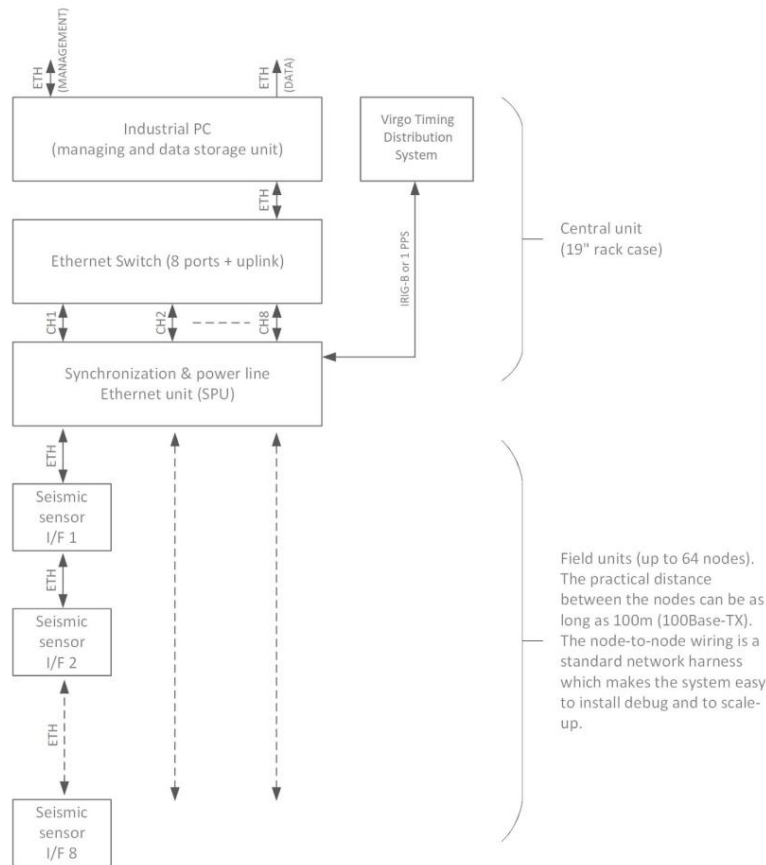


Figure 3.7: Block diagram of the seismic newtonian noise DAQ system, already in operation at Virgo’s facilities.

- Güralp CMG 40-T ultra-lightweight triaxial seismometer (figure 3.8a); force feedback velocity technology; rugged, waterproof stainless steel design for ease of installation; sensitivity band of 0.017-50 Hz. It is currently used to monitor the seismic activity inside the main Virgo Buildings.
- Trillium 240 triaxial broadband seismometer by Nanometrics (figure 3.8b); force feedback design; self-noise below the NLNM (New Low Noise Model) from 100 seconds to 10 Hz. It has been used mainly for outdoor and underground measurements to seismically characterize GW detectors sites.
- EpiSensor FBA ES-T (figure 3.8c); force balance triaxial surface accelerometer by Kinematics; full scale recording ranges of ± 0.25 to $\pm 4g$ (user selectable); enhanced bandwidth of DC to 200 Hz; one of these is currently installed inside the Input Mode Cleaner Building.
- Portable ICP accelerometer 393B12 (figure 3.8d) by PCB Piezotronics (MTS Systems Corporation); thanks to its compactness and maneuverability, it is frequently used to probe seismic vibrations of the mechanical parts of the interferometer and to make spot measurements.

The majority of the seismometers in the market, as a consequence of their physical parameters and costs, are unsuitable for a large scale array installation to sample the seismic wave-field with sufficient density. In order to achieve affordable



Figure 3.8: Seismic sensors used in Virgo: a) Güralp CMG 40-T seismometer; b) Nanometrics Trillium 240 seismometer; c) Kinometrics EpiSensor FBA ES-T accelerometer; d) ICP accelerometer 393B12 by PCB Piezotronics (MTS Systems Corporation).

dense sampling, an autonomous, integrated seismic recording system called Tremor-net [210] was developed for flexible deployment around the GW detector’s key components, aiming to the desired SNN cancellation. This novel sensor technology was provided by the National Institute for Subatomic Physics (Nikhef) group through its spin-off company InnoSeis. The accelerometer is depicted in figure 3.9, together with the main specifications and a photo of the actual installed instrument.




Figure 3.9: Accelerometer designed by InnoSeis to be part of large arrays for dense seismic field sampling. On the left, some important specifications are listed. On the right, an example of indoor placement.

As a matter of fact, the accelerometers can be alternatively put in the ground, taking advantage of a special tip, or upon a flat steel frame for indoor installations. The required sensitivity is achieved by using a high-sensitivity, 5 Hz geophone in combination with a low noise recording channel, resulting in a self-noise level of less than $1 \text{ ng}/\sqrt{\text{Hz}}$ at 1 Hz and as low as $0.1 \text{ ng}/\sqrt{\text{Hz}}$ at the sensor resonance frequency of 5 Hz (see the plot in figure 3.10 in terms of velocity PSD). Individual accelerometers of the array show to have the long battery life required to be practical in their intended surveys, showing only 50% battery consumption after a 21-day survey at high-gain setting. In addition, they are GPS synchronized to about $2 \mu\text{s}$.




Figure 3.10: Instrument noise of the InnoSeis sensors. At high frequencies, noise was measured on an isolated platform (MultiSAS), while, at low frequencies, tests were performed with a high-quality broadband sensor (HGN).

Tiltmeters

Besides the accelerometers array, in [211] the authors argue that a single seismic tiltmeter per TM can be used to reduce SNN, provided that it is placed in a location where the seismic field at the TM position is accurately reproduced. They also investigate the case in which the tilt measurement is obtained with a pair of seismometers used to read differential ground motion between two locations. The main difference is that the two seismometers are also sensitive to a common-mode signal, while tiltmeters are designed to respond to differential ground motion only.

The underlying idea is in the fact that surface seismic displacement along the horizontal direction and gravity acceleration are maximally correlated. Unfortunately Love waves also contribute to horizontal displacement. The solution is to use vertical displacement, which only has contributions from Rayleigh waves and it is already measured with a seismometer array at GW detector sites. The alternative solution is seismic tilt, which is maximally correlated with Rayleigh NN, and at the same time, it is not responsive to Love waves.

Tiltmeters would be ideally deployed directly under each of the 4 TM and they are currently being built in Naples laboratories (figure 3.11) as an optional measure to reduce the number of seismic sensors. In figure 3.11 on the right, the first installed version of the tiltmeter is showed and it is placed in the NEB, currently acquiring data.

An optimal subtraction scheme would use the array of seismometers as input to the filter and a tiltmeter as target channel, which is therefore the best available proxy for testing SNN cancellation schemes. [212]

On August 2019, the Precision Laser Inclinometer (PLI), an innovative instru-




Figure 3.11: The tiltmeter prototype at Naples laboratories (on the left) and the final version (on the right) installed at Virgo NEB.

ment for monitoring angular microseismic activity, was installed and put into operation at Virgo. Nevertheless, this specific instrument is not useful for SNN since it is not sufficiently sensitive at SNN frequencies. Apart from that, an inclinometer is just a tiltmeter that targets lower frequencies.

3.3.2 Sensor array configuration and first tests

The InnoSeis sensor array configuration is determined in terms of required number and proper location. The number of sensors required to achieve the desired SNN reduction factor is estimated using data from the seismic characterization of the Virgo buildings, and predicting the SNN cancellation performance of optimized sensor arrays. Preliminary measurements at WEB were performed on January 2018, when an array of 38 seismometers was deployed inside the building and a second one of 9 seismometers was installed outdoor [213]. The data acquisition covered a period of 14 days. In figure 3.12, the actual positions of the indoor accelerometers used for the test are superimposed on the WEB blueprint. All sensors were fixed on the floor thanks to an heavy mount plate with double-sided adhesive tape (see the photo on the right of figure 3.9). Outdoor sensors were partially dug into the ground to protect them from the wind. A central data-acquisition unit was used for the readout, synchronization and powering of the entire array. Seismic data were sampled at 500 Hz for indoor sensors and 250 Hz for outdoor sensors.

From figure 3.12, it can be noticed that 14 of the indoor sensors are located along the internal perimeter, other 9 in the proximity of the tower platform (the light blue box), 13 other sensors directly on the platform and the remaining 2 at the sub-surface level. This particular choice is to address the question about the influence that a different soil composition has on the propagation of the seismic waves. Indeed, the ASD of four seismometers positioned on North side of the building is reported in figure 3.12 on the right. They are chosen as pairs, so that sensors 5 and 7 (blue and red curves) are located on the building platform while sensors 25 and 26 (yellow and violet curves) on the tower platform, with a distance of about 0.5 m between them. For frequency values above ~ 10 Hz, a significant difference in ASD starts to rise between tower platform sensors and building platform ones. The reason for this

distinction stems from the fact that dominant seismic sources are located outside the tower platform, and the seismic waves they produce are reflected when coming across the soil discontinuity, which is also 1 cm maximum wide. Note that this gap between platforms does not go deep enough to efficiently reflect seismic waves below 10 Hz. Interestingly, a step of nearly 9 cm has formed over time between the building and tower platforms, as shown in figure 3.9 on the right. This is due to subsidence of the building platform with respect to the tower platform, whose deep foundation pillars are longer and reach a more stable gravel soil layer. These aspects will be taken into account in the following section for the FE numerical simulations.

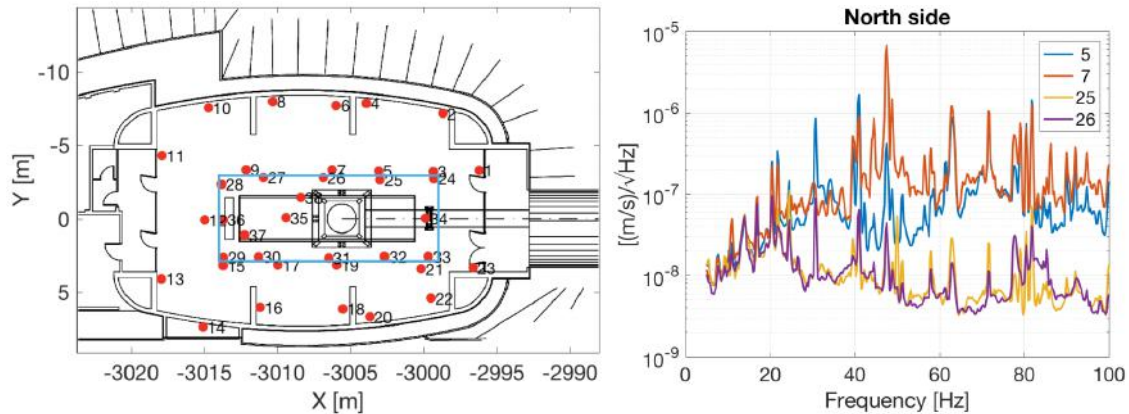


Figure 3.12: On the left, a map of the seismometer positions at Virgo WEB. Light blue rectangular perimeter denotes the extent of the tower platform. On the right, the power spectral densities of two couples of seismometers positioned at North side of WEB. Sensors 5 and 7 (blue and red curves) lies on the building platform while sensors 25 and 26 (yellow and violet curves) on the tower platform. Figures are reproduced from [213].

A very similar seismic survey was performed at the Virgo NEB, whose structure is specular to the WEB. Thirty eight sensors probed the environmental seismic field for more than 4 months, starting from 20 February 2019. Results are compatible with the WEB ones.

These test measurements and SNN models have proven that no more than 38 seismometers are needed for the future noise cancellation. Likely, a number near to 30 seismometers per End Building would be sufficient, since the level of degeneracy of the tested array turned out to be very high [213]. With that in mind, a conservative estimate of 60 seismometers to be deployed at CEB is obtained. A characterization of the seismic field at CEB is scheduled to obtain a better estimate of the required sensors. Currently, the baseline design of the array is to use only indoor sensors. If necessary, outdoor ones will be considered in later configuration improvements, same for borehole seismometers.

In the near future, the array data will be part of Virgo data frames. Sensor electronics is currently being modified by the Polgraw group, in order to be compatible with central data acquisition and distributed timing signal. One data acquisition unit will be positioned at each building implementing TOLM interface (Timing and Optical Links Mezzanine) for data transfer and synchronization. Each array data unit is designed to host up to 64 sensors, into chains of 8 (further details in section 3.3.1).

3.3.3 Online and offline subtraction schemes

In this framework, offline subtraction denotes the cancellation of noise in recorded data. The noise cancellation filter can therefore be causal or non-causal. The scheme developed for offline cancellation of beam-jitter noise at the LIGO Hanford detector [214] is frequently followed. However, additional features include fully automatic and low-latency data processing, including whitening and band-pass filters focusing on the 8 to 30 Hz SNN band.

For what it concerns online noise cancellation schemes, it is possible to implement SNN subtraction by exerting a force directly on the test masses or, alternatively, by acting directly on interferometer data. The main difference between online and offline subtraction is that the former can only be done with causal filters, whose coefficients can only change slowly in time. The online subtraction scheme presented in [123] is based on a multi-input, single-output (MISO) finite impulse response (FIR) Wiener filter [215] that is continuously applied to the interferometer output to filter out the SNN, as was already demonstrated successfully for seismic noise cancellation schemes [204]. The inputs consist of the accelerometer channels, and the single filter output is the SNN estimate. This is exactly the same baseline scheme I will use when implementing a causal machine learning based subtraction. However, in the absence of the real SNN, the cancellation can only be verified on a proxy of it, which could be another accelerometer, not previously chosen as input, or even better a tiltmeter (see section 3.3.1).

The only filter parameter that is predefined is the order of the FIR filter, i.e., the number of filter coefficients. The filter order essentially determines the time span of the filter. For instance, if the time series are 100 s long, a filter order of $N=50$ corresponds to a time span of 0.5 s. It is argued that the order is too high when the seismic array cannot provide sufficient information to disentangle SNN contributions from individual seismic waves. On the contrary, it is too low when an insufficient amount of data is used to accurately estimate the SNN from individual, resolved waves [123].

Among the total predicted number of 120 sensors, a few will likely fail every year and a single broken sensor would completely spoil the SNN cancellation. Hence, failure detection needs to be done continuously. The plan is to temporarily recalculate the filter excluding broken sensors using stored sensor correlations. Eventually, the broken sensors can be replaced with spare units. In addition, it is also necessary to do performance checks to make sure that no excess noise is produced, to identify problems with sensors and to restart the SNN cancellation in case. Moreover, the application of auxiliary filters (e.g. band-pass filters), data decimation and re-sampling may be required. Excess noise detection is an important feature since the cancellation system creates an additional shortcut between environmental noise and Virgo data. For example, if there is evidence of an unexpected strong local seismic source, it is possible that the associated SNN will not be suppressed, causing on the contrary significant excess noise in the residual GW data. In such cases the preferred action would be to temporarily deactivate the cancellation pipeline.

3.3.4 Cancellation with Wiener filter

Coherent noise cancellation with WF leverages the correlations between reference seismic data flows and a target data flow to provide a consistent estimate of SNN

contribution to the target sensor [216]. The theory behind the Wiener filter implementation is introduced in Appendix A. Wiener filters are the optimal linear noise cancellation filters, provided that the data are described by wide-sense stationary and random processes [217]. Assuming this to be the case, the mutual spectral densities between input sensors are denoted with C_{SS} , which is a $n \times n$ matrix, with n the number of input sensors, containing the cross power spectral densities of the vertical surface seismic displacements. Considering one TM for simplicity, I denote with C_{NN} the SNN spectral density acting on this TM. The spectral density between SNN and the input sensors is denoted with C_{SN} , which is a n -component vector. The coherent SNN estimate is now obtained by multiplying the Wiener filter \vec{w} with the amplitudes of seismic displacement observed by the n seismic sensors, as all defined in the Appendix A. Subtracting the coherent estimate from the actual target channel, a residual $R(f)$ is achieved, which is determined by equation 3.19.

The WF cancellation performance can be limited by sensor instrumental noise, or by the amount of information extracted from the seismic field. For example, the mutual distance between sensors of the array has to be comparable with the wavelength of the dominant component of the seismic field. It is argued that the input sensors must provide information about the seismic field with sufficiently high signal-to-noise ratio. Moreover, there may be limitations for the maximum number of input sensors, for the possible array configurations, and for the amount of data that can be used to calculate the optimal filter. In general, so, many constraints and a little understanding of how they limit the performance exist.

Finally, the calculation of the Wiener filter does not address the problem whether the sensor locations are optimized to extract information most efficiently. The optimization of array configurations for noise cancellation directly affects the matrix C_{SS} and the vector C_{SN} , and it will be discussed in section 3.5.

3.4 Cancellation with Machine Learning

It is very likely that different SNN cancellation techniques will be combined in the near future. In this perspective, I have developed a noise subtraction method based on ML. The theory behind this is presented in Appendix B. The peculiar ML subtraction scheme I use, avails itself of Deep Neural Networking (DNN). The usual application of an artificial neural network is pattern recognition and/or classification, whether it is a picture, a sound, or a video the subject in question. For the purpose of predicting the forces on the test masses generated by SNN, or simply the output of a seismic sensor from an observational dataset, I use the DNN as a numerical variable predictor or, in another words, for regression analysis. Regression algorithms are ML techniques for predicting continuous numerical values and they require labelled training examples.

During this section, I will go through all the necessary steps of the analysis, trying to address the following issues. First, the possibility to build a robust and simple architecture of neural network with adequate predictive performances; then, the ability to estimate the quality of the neural network as a function of the number of input sensors and their particular choice, searching for an optimal layout; third, the potential to get comparable results with the gold standard for stationary noise reduction, the MISO Wiener filter (see the Appendix A for the reference theory).

Finally, I will present the results from the comparison between DNN and WF

by doing a Monte Carlo (MC) simulation over a high number of choices of selected input sensors. The target sensor is rather left unchanged. I also vary the number of selected input sensors from a minimum of 3 to a maximum of 35, in order to uncover the trend of the filter residuals. The scheme in figure 3.13 presents an example of the followed framework with 4 randomly chosen input sensors.

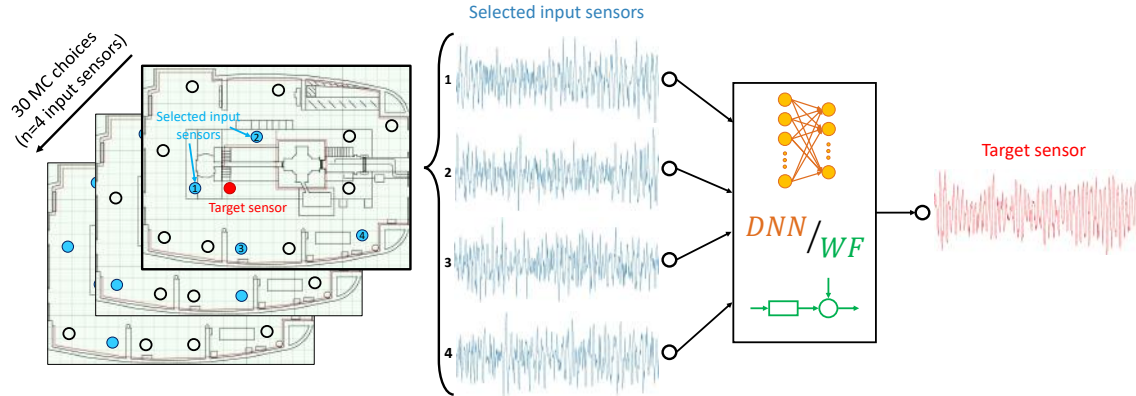


Figure 3.13: Monte Carlo simulation framework for the prediction of a target seismic sensor. From left to right, the MC random choices (i.e. a number of 30) for a fixed number (i.e. $n = 4$) of selected input sensors and then the data analysis flowchart (DNN or WF). The maps on the left are blueprints of the WEB and the circles are seismic sensors, taken as input (blue) and output (red) of the subtraction filters.

Note that there is no real hope to subtract the SNN without having a detector sensitive to this noise. This also applies to Wiener filter, since only the auto-correlation between sensors can be measured easily, which is not the case for the covariance between sensors and newtonian forces.

3.4.1 Data processing

In the context of SNN cancellation, the goal is to estimate its contribution to the AdV strain noise using data from a seismic array, and to subtract it from the GW data. Since the SNN lies still below the GW noise floor, I will test the subtraction potential by reconstructing the signal from a target seismic sensor, using the remaining sensors as inputs to the filter.

A picture to keep in mind from beginning to end is the one shown in figure 3.14, which represents the input/output sensor selection criteria in a very straightforward way (it is just a magnification of the blueprints in figure 3.13). The colored circles denote the seismic sensors inside the WEB: the red circle is the target sensor, the blue circles are the sensors selected to be the inputs of the filter and the empty black circles are the remaining unchosen ones.

The data acquisition relevant for the analysis covered 14 days, from January 24 to February 6, 2018. The seismic sensor array is the one deployed at WEB, comprised of 38 accelerometer units. The dataset is provided as a series of CSV files, one for each hour of data. I load the input data into a Pandas data frame, which is handy for importing big complex arrays. A conversion factor of $5/(2^{23} \cdot 77.3)$ from V to

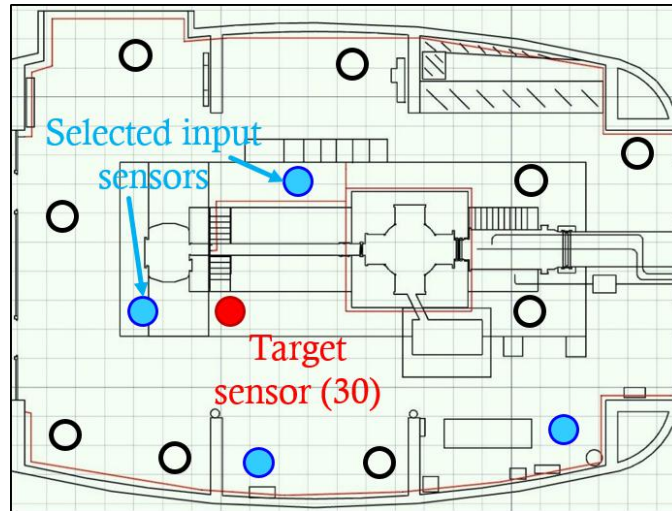


Figure 3.14: Blueprint of the West End Building. Some colored circles are drawn on top of it in random positions to give the idea of what is the input and output content (blue and red circles) of the neural network (or Wiener filter).

m/s for indoor sensors has to be applied, although the subsequent scaling (z-score normalization), which is necessary when data are provided to the neural network.

Before any data manipulation, it is good practice to always check for stationarity of the time series. This is crucial if the goal is to make future forecasting from past time instances. I used the following two tests for univariate time series: Augmented Dickey Fuller (ADF) Test and Kwiatkowski-Phillips-Schmidt-Shin (KPSS) Test. They are almost interchangeable, with the exception of Trend Stationarity and Difference Stationarity, which are two kinds of pseudo-stationarity that could cause the two tests to result in opposite outcomes. Nevertheless, they both confirm the stationarity of the seismic data, already apparently clear from a “visual” approach, for every sensor.

Seismic spectra are obtained using Welch’s method with section length of $L = 4s$, 50% overlap and a Hamming window. The employed time interval is 1 hour long, sampled at 500 Hz, and subsequently decimated at 250 Hz. This means that the data vector has 9×10^5 time instances and it is finally divided into 900 time segments.

Below, I present a characterization of the seismic field in the WEB. In the three plots of figure 3.15, I report seismic spectra respective of the target sensor on 3 February during every hour of the day (a), the target sensor at 1:00 pm during every day of the data taking period (b) and finally a few selected sensors on 3 February at 1:00 pm (c). From the first plot, it is clear that the output of a single sensor remains almost constant over an entire day, while it visibly changes during a two-weeks period, as witnessed by the second plot. This variability is even more evident if different sensor outputs are investigated, as in the third plot. There, it can be noticed that the data from sensor 35 (yellow dashed curve), which is located on a metal sheet, clearly stands out above the level of all other units, which however spans a pretty high amplitude range, approximately from 10^{-8} to $10^{-7} m/s/\sqrt{Hz}$. As already pointed out, there is evidence for a clear distinction between sensors on the building platform (solid lines) and sensors that are on top of the tower platform (dashed lines), especially at higher frequencies. Moreover, there are several lines in the spectra that affect the entire building. Specifically, lines at 17 Hz, 21 Hz, 22

Hz, 24 Hz, 31 Hz and 39 Hz can be easily identified. The majority of them are persistent, and can be ascribed to disturbances from vacuum pumps and/or cooling fans.

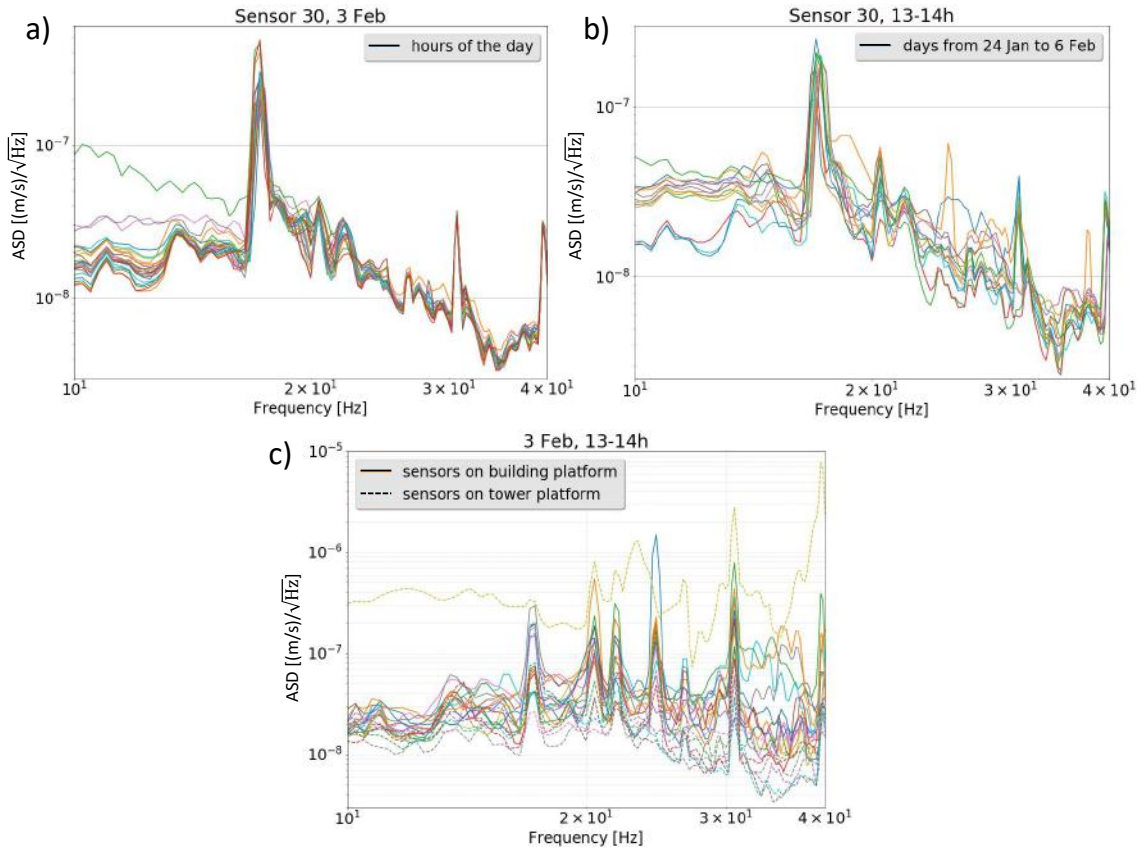


Figure 3.15: Amplitude spectral densities of WEB sensor array. a) Spectra of the target sensor on 3 February at every hour of the day; b) spectra of the target sensor computed on one hour (1:00-2:00 pm) from 24 January to 6 February; c) spectra of a selection of sensors, both near the building walls, near the tower platform and on top of it. The outlier (yellow dashed line) is a sensor (35) located on a metal sheet and hence it is dominated by high vibrational noise.

At this point, the process consist of the random selection of 2 hours of data from the complete WEB dataset of 2-weeks: the first hour is used to train and test the neural network, while the second hour is needed to evaluate its performances on brand new data. Two things must be pointed out: first, the 2 hours have to be temporally consecutive, in order to preserve roughly the same seismic field conditions, not counting unpredictable transients; second, the same data sections are chosen for both neural network and Wiener filter, for a correct comparison. After that, n input reference sensors out of the total of $N = 36$ (the target sensor is not included among the input and neither is the sensor 35) are selected. A smaller time interval, usually 1 minute long, is randomly sampled from the total first hour of data. The main reason for this stems from the great memory demand for storing such heavy matrices.

It is generally a good practice to scale the data at the standard score (z-score), calculated by subtracting from each value the population mean and then dividing

the difference by the population standard deviation. The result is a dimensionless quantity, but it is necessary for the neural network to be able to speed up the learning process. More specifically, normalization helps because it ensures firstly that there are both positive and negative values used as inputs for the next layer which makes learning more flexible and secondly that the network's learning regards all input features to a similar extent. The last step before feeding the data to the network is the splitting into train and test datasets, which is generally performed by randomly dividing the first hour of data into 70% for training and 30% for testing.

Programming infrastructure

The programming language I use to implement the SNN cancellation is Python and the code is developed by setting up a Conda environment and the TensorFlow Python backend. TensorFlow is an open-source software library, first developed by the Google Brain team for internal Google use and publicly released on 2015. It is also widely used for machine learning applications such as neural networks. The actual code writing is done on the integrated development environment (IDE) Spyder. I define my neural network type of models using Keras, which is a high-level neural networks API, written in Python and capable of running on top of TensorFlow. For other types of coding needs, I regularly use the Python numerical and scientific libraries NumPy and SciPy, the plotting library Matplotlib, the machine learning library Scikit-learn and the data manipulation library Pandas. The computation is done either on a desktop computer provided of 128 GB of RAM and 20 available CPUs and on a computer cluster, supplied by the INFN Genoa section, consisting of 5 individual machines and 120 cores in total, all accessible for parallel computing. The latter is managed with the Python software module SCOOP (Scalable Concurrent Operations in Python) for distributing concurrent tasks on heterogeneous grids of workstations, all accessible through SSH protocol. Finally, the Genetic algorithm is designed with the evolutionary computation framework DEAP [218] (Distributed Evolutionary Algorithms in Python).

3.4.2 Neural Network building

Next I define the neural network. This is a simple sequential model that takes a vector of $n = [3, 36]$ (number of sensors) features as input and outputs a single prediction (target sensor), all iterated for thousands of time instances. This is still a basic architecture, made with only fully connected layers (FC).

The Keras Python library makes creating deep learning models fast and easy, since it allows to generate sequential models layer-by-layer, where they can be added piece-wise in a linear stack. The model needs to have knowledge about what input shape it should expect, so this information is supplied once and for all to the first layer (the next layers can automatically infer the shape). Each layer has an activation function that defines its output given the input (see the appendix B). I use the Leaky ReLU activation function for all the hidden layers and the Linear function just for the output FC layer. This stems from the regression nature of this particular machine learning problem: I am interested in predicting numerical values directly without any transform or bounding.

Before training a model, the learning process has to be configured, which is done via the compile method. It receives two important arguments: the loss function,

which is the objective that the model will try to minimize, and the optimizer, which helps to minimize the loss function in order to produce slightly better and faster results by updating the weights and biases.

Here I choose the Mean Squared Error (MSE) loss function

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i^{pred} - y_i^{true})^2, \quad (3.21)$$

which evaluates the mean squared difference between the true and predicted output values y .

I also use the adaptive ‘‘Adam’’ optimization algorithm [219], which is a more robust extension of ‘‘Adagrad’’ and ‘‘Adadelata’’ [220]. It was designed specifically for training deep neural networks. Adam is an adaptive learning rate method, which means, it computes individual learning rates for different parameters. Its name is derived from adaptive moment estimation, since it uses estimations of first and second moments of gradient to adapt the learning rate for each weight of the neural network. I used the parameters provided in the original paper, which are a learning rate of $\lambda = 0.001$ and the 2 moments $\beta_1 = 0.9$ and $\beta_2 = 0.999$.

Before starting to train the network, weights and biases have to be initialized. I initialize weight matrices for all the dense layers (fully connected layers in Keras) with a specific initialization scheme called glorot uniform. It follows a distribution which is zero-centered and within ± 0.2 . The biases are instead initialized at zero. These parameters are meant to be the standard good ones for this type of input data.

Models are trained on NumPy arrays of input data and target output. Keras offers the fit function to train the network for a fixed number of epochs (iterations on a dataset). However, I introduced the possibility to stop early the training process, once peculiar predefined conditions on the new train loss and its gradient are met. Early stopping also prevent the so-called over-fitting (see the appendix B).

Models are trained for a number of epochs not smaller than 100 and rarely bigger than 1000. The latter case likely means that something went bad during the training process and hence either the network architecture or the input data have an issue.

Another crucial training parameter is the batch size, which refers to the number of training samples utilized in one iteration. It can be equal to the total dataset (batch mode), it can be set to 1 (stochastic mode), or it can be a fraction of the dataset size (mini-batch mode). It is known that mini-batch gradient descent is the recommended variant for deep learning. However, its choice is a tradeoff between training time and convergence quality. It is often tuned to the power of two that fits the memory requirements of the CPU hardware like 32, 64, 128, 256, and so on: small values give a learning process that converges quickly at the cost of noise in the training process; large values give a learning process that converges slowly with accurate estimates of the error gradient. At first, I used a batch size of 128 or 256, since it requires less memory. Later, I tried to use bigger values like 512 and 1024, aided by the distributed computing.

At this point, the Keras function ‘‘Evaluate’’ gives the performance of the model on new data and allow to compute the test loss by comparing true and predicted outputs. Finally, the function ‘‘Predict’’ returns the actual predicted values.

Preliminary tests

I briefly present some preliminary tests of supervised learning on a sample of 100 networks with different architectures, on the plots of figure 3.16. The only features that distinguish these neural networks are the number of neurons per layer and the number of FC layers. An average prediction performance on test data of 10% is found (see figure 3.16b), but it is not sufficient to discriminate between good and bad network architectures. The remaining plots show examples of the difference between performances on train and test data: figure 3.16a is about the difference distribution between true and predicted values, while plots 3.16c-d are the normalized network outputs.

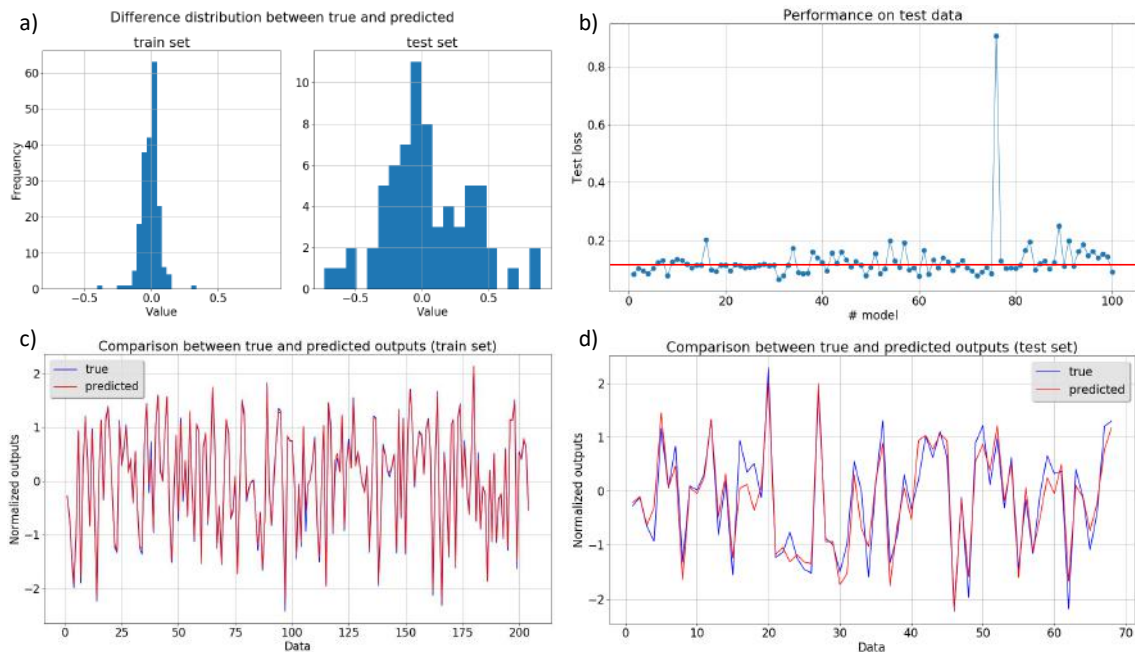


Figure 3.16: Preliminary tests on finding a proper network architecture. a) Example of difference distributions between true and predicted outputs for train and test dataset. b) Performance on test dataset (mean squared error) for each of the 100 different networks; the horizontal red line is at 10% of accuracy. The remaining plots show a portion of normalized outputs and the corresponding predicted ones, for train (c) and test (d) dataset.

It is clear that these basic architectures achieve good results only on training dataset, but they do not succeed in generalize on data never seen before.

Convolutional Neural Networks

Although traditionally developed for two-dimensional image data as feature extractors, Convolutional Neural Networks (CNN) can be used to model univariate time series forecasting problems. My case is still a little bit different, because I do not want to predict future time instances $t + \tau$ for now, but exactly at that time t . This way, I organize the input data to build consecutive input matrices made of a series of past observations $(\dots, t - 3\tau, t - 2\tau, t - \tau, t)$ for each seismic sensor. In this sense,

it is not a multivariate time series problem, even if there are multiple input vectors, since different sensors probe the same physical quantity. The output of each input matrix is a single time instance (t) of the target sensor. In summary, the data has been pre-processed in such a way that each data record contains an history h of time instances between 3 and 100 (sometimes selected with a step of 2 or 3). Within each time window, the n accelerometer values are stored. This results in an $h \times n$ matrix.

In order to build a CNN, I have to incorporate into the model a series of additional layers, first of all the one-dimensional Convolutional layer (Conv1D). This layer creates a convolution kernel (a filter with a certain kernel size s) that is convolved with the input layer over a single temporal dimension to produce a tensor of outputs. The total number k of filters have to be defined, resulting a $k \times (h - s + 1)$ output matrix. Each column of the output matrix holds the $h - s + 1$ weights of one single filter.

A Max Pooling layer (MaxPool) with a size of three is used after the CNN layer in order to reduce the complexity of the output and prevent overfitting of the data. This means that the size of the output matrix of this layer is only a third of the input matrix. There is the possibility of casting multiple Conv1D-MaxPool blocks in order to learn more complex connections in the data.

In the same way I normalize the NN input data to speed up learning, it is possible to do the same thing also for the values in the hidden layers, that are changing all the time. The Batch normalization layer (BatchNorm) do the job, allowing each layer of a network to learn by itself a little bit more independently of other layers. I add the BatchNorm transform immediately before the non-linearity of the activation layer, as to be sure to contain the activations within a stable value distribution [221]. Moreover, it reduces overfitting because it has a slight regularization effect, without losing any information, as opposed to Dropout layers.

Finally, a one-dimensional Global Average Pooling (GAP1D) is necessary in cases where the temporal dimension has to be reduced to 1, before the final FC layer.

At this stage, there are an enormous number of possible different NN architectures that can be valuable in terms of performance on the available seismic data. A sketch of this scenario is presented in figure 3.17 as a general NN architecture layout, taking into account all the pre-defined layers.

Let me consider a random example of NN architecture of that type and test its performances on real data. These are usually evaluated by looking at the train-test loss curves, as in figure 3.18. Both trends behave as expected from an optimal NN, to varying of the number of epochs: the train loss curve (blue line) strongly decreases just after a few epochs and keep diminishing more slowly until a plateau is reached. The curve is also quite stable and free from strong spikes, which denotes a smooth learning process. On the other hand, the test loss curve (orange line) is slightly noisier but very close to the blue one. More importantly, it reaches the plateau quickly and without leaving it by starting to rise back. The latter phenomenon would have meant the presence of overfitting and consequently the need to put in operation additional countermeasures.

Lastly, I want to report some preliminary studies on the increase of the convolutional temporal window (history) and on the possibility of predicting a near future time instant $t + \tau$. Predicting the future state of the target seismic sensor, proxy of the interferometer test mass, is particularly appealing for SNN cancellation pur-




Figure 3.17: The employed Convolutional Neural Network architecture, which consist of a sequential model with 2 initial Conv1D layers. Then there are a number of intermediate layers and finally the output fully connected layer. The core of the model is interchangeable so that very different architectures can be explored by leaving the initial and final steps untouched.

poses, since it could open to the possibility of making an online noise cancellation. Unfortunately, the performances identified by the residuals from the prediction, deteriorate seriously, both in the case of increasing the convolutional window (figure 3.19 on the left), and in the case of predicting a future instance (figure 3.19 on the right). Interestingly, the worsening due to the increase of the history is much more evident at frequencies roughly higher than 15 Hz, but the reason for this behavior is currently under investigation. The increasing trend to varying of τ was instead foreseeable, but the plateaux may be achievable at lower residuals if a longer convolutional window is used (the one used for the plot is equal to 6 time instances). Unfortunately, the neural networks training time rises fast with growing history.

3.4.3 Neural Network Optimization

At this point I have so many different configurations of deep neural networks with different layers and ideally I want to find the best architectures. The DNN architectures pass through an evolution process called Genetic Algorithm (GA), so that a subset of the best specimens can be found. GAs belong to the class of all those artificial intelligence algorithms that inherit their most distinctive feature from the natural world. In particular, GAs are inspired by the concept of Darwin's theory of evolution and the process of natural selection. The foundations of the GA theory are well presented in the book by J. H. Holland [222]. For my purposes, the use of GAs consist of identifying a collection of DNN parameters and structures which can be varied in order to test hundreds of combinations, by leaving the initial and output steps untouched. These parameters are the following:

- The layer type, randomly chosen from the following stacked structures:




Figure 3.18: Train-test performances of a typical employed CNN, to varying of the number of epochs.




Figure 3.19: On the left, the DNN prediction residuals to varying of the convolutional temporal window (history). On the right, the performances to varying of the τ for a prediction in the future at constant convolutional window, equal to 6 time instances.

- FC-act: set of a Fully Connected layer and the Leaky ReLU activation layer.
 - FC-BatchNorm-act: same as FC-act, with an intermediate Batch Normalization layer.
 - Conv1D-act: same as FC-act, with a one-dimensional convolutional layer in place of the FC one.
- The total number of hidden layers, chosen from the pre-defined types.
 - the number of neurons per each hidden layer, randomly chosen from a range of $(1, n)$, with n the total number of seismic sensors. A downward trend for the neurons number is preserved until the single-output layer.
 - the number of filters in both convolutional initial layers, which are stacked into two Conv1D-MaxPool-act layer structures. Filter numbers are randomly

selected from a range of [30,60] for the first Conv1D layer and [60,90] for the second one.

All these figures are implemented by creating a python class called *DNA*, which is instantiated every time a DNN is created. However, the actual GA is designed with the evolutionary computation utility DEAP. It allows to establish a custom GA framework by providing a comprehensive set of python functions. A detailed overview of a GA features can be found in [223], but here I want to describe how they can be implemented.




Figure 3.20: Schematic view of the genetic algorithm I used to optimize the architecture of the convolutional neural networks.

I start with a collection of initial guesses of different DNN architectures, in number of dozens or hundreds, which compose the population of individuals. The initial population is evaluated and the performance of each individual is computed. Outcomes are not only based on the test loss, but on a multiple number of metrics, including the DNN complexity (total amount of trainable parameters) and the mean and standard deviation of the DNN residuals. I specifically look for DNNs with small test losses, low complexity possible and residual distributions with almost zero-centered means and small standard deviations. These are the parameters the GA seeks to optimize as soon as the individuals enter the evolution process, which is a cycle that lasts for a relatively high number of generations. The first phase is a proper selection of the individuals that better fit the efficiency requirements. I apply the NSGA-II selection operator without repetition, with the size of selected individuals equals to the population size. In this way, the selection has only the effect of sorting the population according to their Pareto front rank. This specific choice is dictated by the need to preserve the numerosity of the population over the generations. After that, a series of crossover and mutation tools are applied at the neurons and layers level, in order to navigate through the parameter space with small adjustments. In particular, I created custom mate and mutate functions that fit with the complexity of the established *DNA* class-based framework. The

idea of the mate function is that two selected individuals (parents) generate two brand new individuals (offspring) that inherit a portion of DNN structure from the first parent and the remaining part from the other. The crossover cut is randomly applied between two hidden layers of the parents. Therefore, the number of layers of the DNN offspring can be lower or higher than the number of layers of the parents, according to the crossover cut application point. The dimensionality of the population after the mate algorithm is preserved here too, since the process is applied again by reversing the two parents. In addition, after the crossover step, I included the possibility of having a small mutation with a certain probability (set to 0.2) for the number of neurons in each hidden layer. To keep track of the best individuals that ever lived in the evolutionary process, an Hall of Fame (HoF) is provided as the cycle goes along. The importance of the HoF is crucial, because otherwise any possibility to retrieve the best DNNs would be lost. Indeed, the particular chosen selection algorithm allows only to explore the parameters space in a wide way, but it does not result in any single champion at the end. That is why the informations have to be extracted during the evolutionary process.

That was the flowchart of a standard genetic algorithm, modified to meet special needs. A schematic view of the process is also presented in figure 3.20. Qualitatively speaking, it's nothing else but pure survival of the fittest or Darwinian selection of species: the best individuals survive, make children who inherit some characters and the dominant gene is carried out.

At the end, the GA is able to find the best network architectures, in terms of high performance and simplicity. Indeed the plot of figure 3.21 shows that after 20 generations, the final DNNs have a lower test loss on average. This is shown clearly by considering the normal probability density functions (dashed red curves) that fit the two distributions after 4 and 20 evolutionary generations.

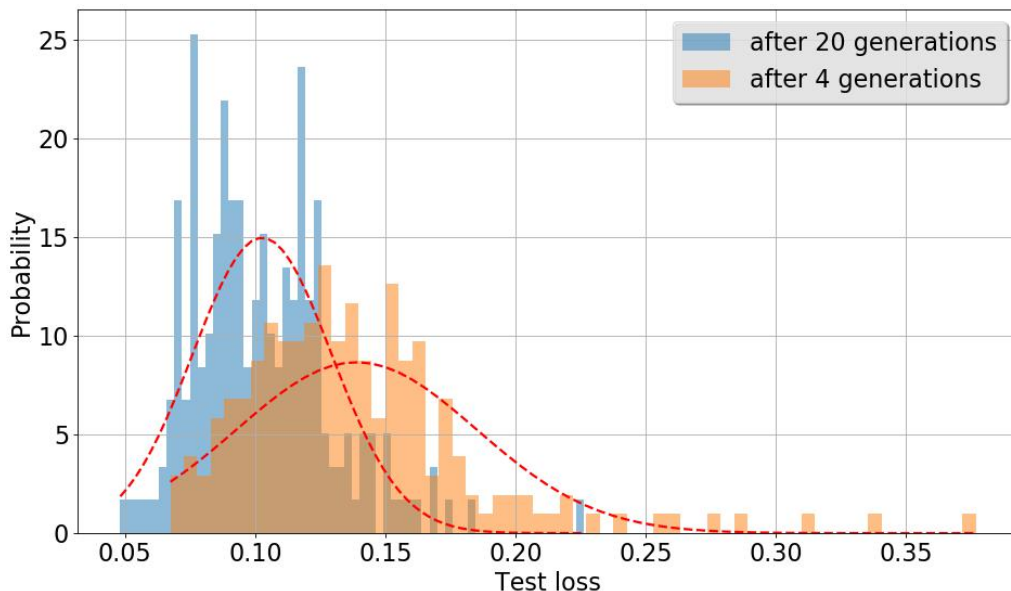


Figure 3.21: Test loss probability distributions belonging to the population of DNNs, after 4 (orange histogram) and 20 (blue histogram) generations of the evolutionary cycle. The normal probability density functions (dashed red curves) are superimposed for a clearer distinction between the two distributions.

A comprehensive view of the mean test loss trend to varying of the number of

generations is also presented in figure 3.22. The mean is computed on the temporary population resulting from each GA step.




Figure 3.22: Mean test loss trend to varying of the number of generations of the evolutionary cycle. The mean is computed on the temporary population resulting from each GA step.

From the shown trend, it is clear, in the first analysis, that just a few generations are enough to rule out a very inefficient class of DNN architectures. These are networks that usually have more than 10 hidden layers, no intermediate BatchNorm after a FC layer and a sequence of layers with a very irregular number of neurons. With the latter I mean that a network with the succession of hidden layers with a number of neurons like [30,20,10,2], would perform better than [30,7,5,5,4]. Interestingly, a BatchNorm layer makes the DNN worse if it is placed right before the output layer. On the other hand, the plot also shows that a substantial plateau is reached within the first ten generations whereupon the mean test loss constantly decreases slowly. This is an indication of the presence of a dense class of DNNs that could be used to evaluate a target seismic sensor knowing a fair amount of other seismic sensors in the vicinity. The GA-based optimization has the purpose of identifying that particular class, which can be used as initial population of further evolutionary cycles and deeper optimizations.

Transfer learning

Another subtle aspect of the learning process during an evolutionary cycle is the transfer learning (TL). In the deep learning context, TL refers to a process where a DNN model trained on one problem is used in some way on a second related problem. This means reusing the weights in one or more layers from a pre-trained DNN model in a new model and either keeping the weights fixed, fine tuning them, or adapting the weights entirely when training the model. TL has the benefit of decreasing the training time for a DNN model and resulting in lower generalization error.

However, I apply the TL in a different way, which consists of transferring the weights of the parents in the evolutionary process, already trained the first time, to the offspring, right before it is trained. In this way, when a child is born, it

has the majority of the weights already initialized in an optimal way, except for the interconnection between the two intermediate layers at the boundary of the crossover cut. Indeed these intermediate weights have to be initialized from zero. An instructive view of this particular TL is presented in figure 3.23, for the specific case of an equal crossover for the parents.

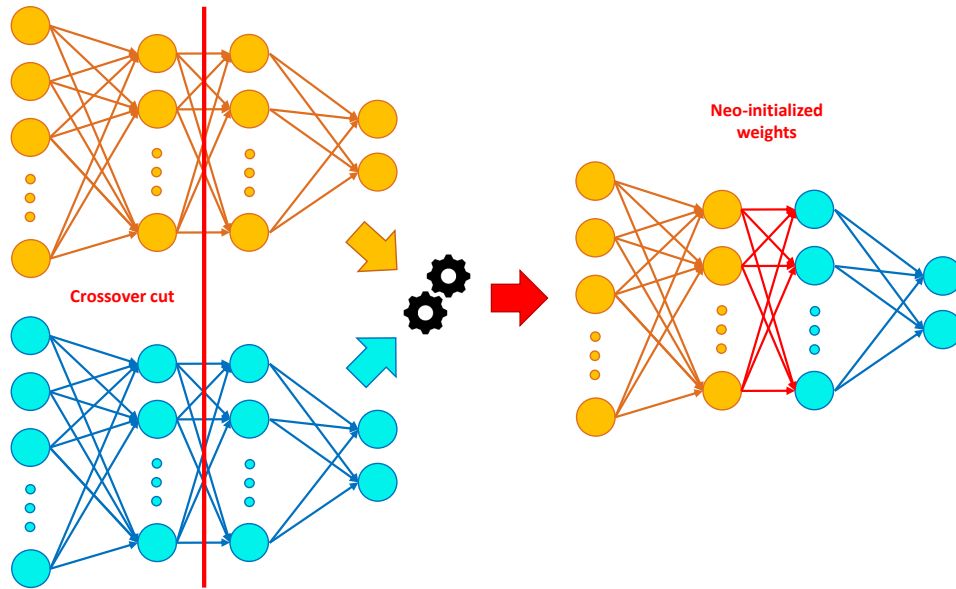


Figure 3.23: Transfer learning scheme at its simplest. On the left, the two neural network parents, on which a crossover cut is applied, during the evolutionary process, in order to generate a new network offspring. The improvement in the training procedure comes from the fact that the majority of the offspring weights are inherited from the parents ones, already initialized in an optimal way. The only neo-initialized training parameters are in correspondence of the interconnection between the two parts.

The obtained improvement is mostly related to the DNN learning speed and goodness, as is apparent from the blue (random weights initialization) and red (TL) plots in figure 3.24: below 200 epochs, the learning process is strongly suppressed for the network with random weights, while the TL allows to reach the training loss plateaux even within 100 epochs, without any discontinuity. The two curves starts to behave in the same way only after 500 epochs. This means that the networks can get the same learning skills in one fifth of the epochs.

I also investigated the possibility of performing a partial training (PT) of each brand new network that originates from the mating process. I tried two slightly different approaches, one consisting of training only the intermediate layer with the neo-initialized weights and the other that, in addition, provides the training also for the nearest neighbor layers. In this way, the total training time decreases remarkably. Figure 3.24 also shows the training loss curve for these two PT examples. The first one (purple line) inherit the TL quick start, but fails at reaching the learning plateaux within a reasonable number of epochs. On the contrary, the second approach (green line) manage to achieve a training loss trend almost as good as for the complete training with TL (red line), but in a smother way. In some sense, TL plus PT give comparable results to the batch normalization layer. This proves that

transfer learning and partial training together are ideal tools to further optimize the DNN training skills by decreasing the computation workload.

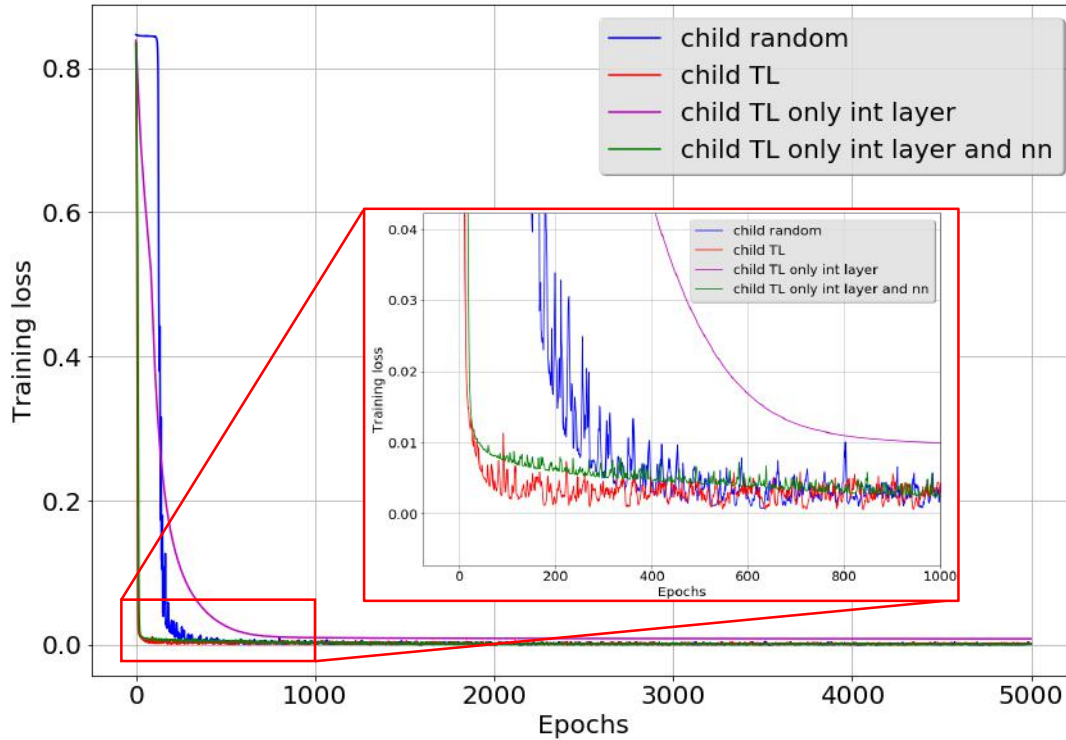


Figure 3.24: Training loss to varying of the number of epochs, for several transfer learning and partial training implementations: random weight initialization (blue line); transfer learning (red line); transfer learning plus partial training only for the intermediate layer (purple line); transfer learning plus partial training for the intermediate layer and its nearest neighbor layers (green line).

3.4.4 Results

The framework for a seismic newtonian noise cancellation based on deep neural networks has been defined. It remains to be seen whether the DNN is able to predict a target seismic sensor output as good as the Wiener filter do (see section 3.3.4).

Going back to figure 3.13, the two filtering techniques can be compared with a series of MC simulations, where one sensor is chosen as the target signal to predict and a subset of the other sensors are selected as inputs. The selection of these n input sensors is done 30 times on a random basis, at fixed n . The 30 MC simulations are repeated for DNN and for WF, for each n , from a minimum of $n = 3$ to a maximum of $n = 35$.

Two examples of seismic spectra of the target channel and the predicted outputs of the DNN and the WF, together with the residuals (entire array as input, so $n = 38$), are presented in figure 3.25. The discrepancy between the two plots is in the different time interval used for the analysis. Apart from that, the employed DNN is the same for them both and it has an optimized architecture. Note that the residuals are simply computed as differences between true and predicted time

series. These examples show that the specific choice of the dataset can influence the residuals a lot, with a larger variability for the WF, especially at frequencies higher than 15 Hz.




Figure 3.25: Two examples of seismic spectra of the target sensor, together with the predicted spectra with a DNN and a WF, and the residuals of the predictions. Both predictions used the entire array (38 sensors) as input. The horizontal lines stand for the mean of the residual spectra, over the plotted frequency range. The only relevant discrepancy between the two plots is in the different dataset used for the analysis.

For this reason, it was decided to do the comparison between the DNN and the WF by leaving the time interval of the input dataset unchanged.

All MC simulations are summarized in the box plot of figure 3.26, which shows the two mean residual distributions in function of the number of input sensors n . For each value of n , the plot shows two competing distributions of 30 random choices of input sensors. A single DNN point in the plot corresponds to the average performance among a set of best networks, whose architecture was optimized with a genetic algorithm. Moreover, the mean of the residual spectra (horizontal lines in the plots of figure 3.25), over the plotted frequency range, is taken. I also stress the point that those results are evaluated on a dataset that neither the DNN or the WF have seen before, with regard to the dataset on which the filters are trained. At first glance, the approach with DNN seem to produce statistically better predictions than the one with WF, and this leads to lower residuals.

There are also some very bad outliers for the Wiener filter, which are the 4 markers at the top of the figure with values of the residuals [3.34, 0.66, 0.61, 3.16], going from left to right.

What stands out is that even with few input sensors, a dozen for example, there are some particular combinations of them that give residuals comparable with the ones obtained from the full-array case. Indeed, figure 3.27 shows that 10 properly selected sensors, out of 38, can predict the target channel with an error less than 10% at some frequencies. These 10 sensors were chosen directly from the best combination among the 30 trials of the MC simulation with 10 input sensors.

One of the reasons for this behavior is that the higher the input-target correlation, the better the residuals. This proves that not all the deployed sensors have a good correlation with the target one and therefore they are, in a sense, redundant. Of course, a sensor which is poorly correlated with the target sensor can instead be highly correlated with the newtonian forces on the TM and hence it could still be

useful for the purpose of SNN cancellation.




Figure 3.26: Box plot displaying the two mean residual distributions for both cancellation techniques, the DNN (blue) and the WF (orange), to varying of the number of randomly selected input sensors, out of 35. The solid line connects the median values of the distributions and the “plus” markers are the outliers.

However, we need a more general optimization of the sensor array that is independent of the specific cancellation technique and that is based on a real SNN prediction, rather than that of another seismic sensor. This topic will be widely discussed in the last following section.

Another way to see the results of the MC simulations is shown in figure 3.28, from which it is clear that there is a high correlation between the DNN and the WF residuals. The correlation level is evaluated by computing the Pearson’s correlation coefficient, that measures the linear relationship between two datasets. It varies between -1 and +1 with 0 implying no correlation and ± 1 an exact linear relationship. Pearson’s correlation coefficient ρ_{xy} is defined as the covariance of the two variables x and y , divided by the product of their standard deviations σ_x and σ_y :

$$\rho_{xy} = \frac{\text{cov}(x, y)}{\sigma_x \sigma_y}. \quad (3.22)$$

A value of $\rho = 0.736$ is found, symptom of a particularly high correlation. It means that, when one of the two have bad performances, the other have bad performances too. There are, however, some cases in which the WF gives so much worse results. Nevertheless, almost all the points in the plot stand above the bisector red line, which seems to indicate a systematical success of the DNN approach against the standard WF. The reason is still under investigation, but it is true that the sum of the algorithms can provide a better noise cancellation than the WF alone would do, even in a sub-optimal sensor array configuration.




Figure 3.27: On the left, the spectra of the original target sensor and the prediction residuals, obtained from the DNN and WF. 10 optimal selected sensors are used as input. On the right, the same residual spectra, normalized to the original spectrum of the target sensor.

A possible sign of WF’s bad behavior could be found in correspondence of seismic transients, as in the case of figure 3.29, right after 400 seconds from the initial time. There, it can be seen that the presence of the transient deteriorates the WF output only, while the DNN residuals seem to be untouched.

Finally, I would like to draw the attention to two possible uses of the WF: one that operates on the same time series used to calculate the filter coefficients, and a second one where the same filter is applied on another time series, which is often temporally consecutive. I applied the two of them and it turned out that the subtraction performances are quite different for the two cases. Therefore, it can be assessed that subtraction performance also depends on the specific wave content of the seismic field and not only on the average correlations between sensors and the SNN. Note that the application of the WF on the same data on which it has been already trained makes it a non-causal filter and therefore it can not be applied online.

Frequency-domain filtering

So far, I have used the DNN and the WF in the time-domain, but now the question is if better performances can be found with the singular two algorithms in the frequency-domain. The Wiener filter has a very straightforward implementation in the frequency-domain, as described in the Appendix A. There, I say that it is numerically much less demanding and more accurate, since the dimensionality of the system of equations to resolve is reduced with respect to the time-domain filter. Therefore, I compared the DNN and the WF once more, but this time by feeding the filters with the discrete Fourier transform of the seismic time sequences. The major difference is precisely in the dataset: in the time-domain, the input convolutional windows have the dimensions of *number of sensors* \times *time history* ($\approx ms$) \times *number of samples* and the output is one single sensor at the time to predict; in the frequency-domain, instead, the time history is replaced with the fft (Fast Fourier Transform) of a larger time window (up to $1ks$) and the output is the entire fft of the target sensor. The scheme in figure 3.30 displays these characteristics. Another difference with the time-domain results is that here I fix the number of input sensors to 10, choosing them with the criterion of higher correlations with the




Figure 3.28: Correlation plot between the residuals coming from DNN and WF. The data points are the same as the ones in figure 3.26. The red dashed line is the bisector of the quadrant and stands for the perfect matching between the two filtering techniques.

target sensor.

Results are presented in figure 3.31 and figure 3.32, as comparisons between time-domain and frequency-domain filterings. First of all, the frequency-domain filters have slightly lower residuals (dashed lines) than the time-domain ones (solid lines). This is true for the example shown in figure 3.31, characterized by a relatively short time window of 72 seconds for the fft calculation, but also in general for every investigated case.

Interestingly, when I test the performances of the frequency-domain filters by increasing the time window within the range of $[40, 1680]$, the DNN and the WF behave in very different ways. Figure 3.32 report the results obtained from a series of MC simulations, similar to those already performed in the time-domain case. The main difference is that it takes a lot of time to train the neural networks and hence a proper genetic evolution of the DNN architectures is not feasible. Therefore, I simply train a small set of networks with random architectures and consider the best 5 ones for each value of the time window. It turns out that the frequency-domain WF strongly depends on the time window length, while the DNN residuals




Figure 3.29: A chunk of the target sensor time series (blue trace), where a particular kind of seismic transient, right after 400 seconds from start, deteriorates the WF output only (orange trace), leaving the DNN residuals unchanged (green trace).

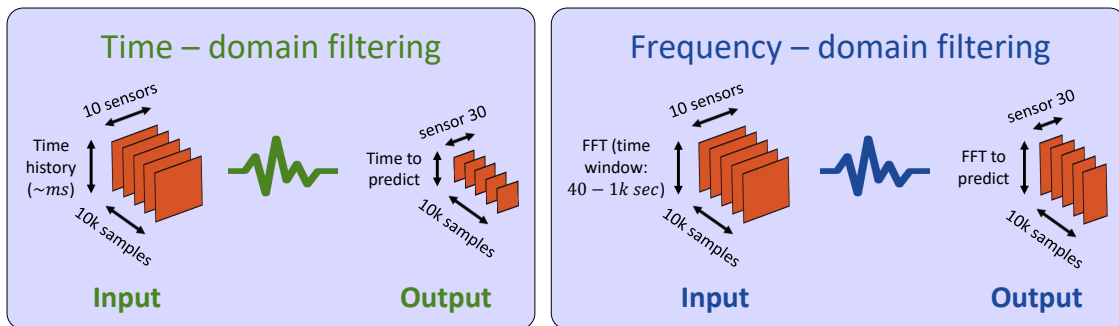


Figure 3.30: Differences in the input/output dataset, for both the cases of time-domain and frequency-domain filterings.

seem to follow an approximately constant trend. In particular, the residuals of the WF present a descending trend to increasing of the time window, resulting in better performances, compared to the DNN. Unfortunately, in order to explore higher time windows, the frequency-domain DNN would require more computational resources.

3.5 Sensor array optimization

The characterization of the seismic field performed at the Virgo end buildings did not take into account any sensor array optimization aiming to an ideal SNN subtraction, but it merely followed common sense rules. As shown in [123], sufficient SNN subtraction down to the level of other noise contributions can be achieved with a proper array configuration. The key design factors are the type and number of sensors and their displacement. These points are mainly discussed in [120, 122, 123, 205],




Figure 3.31: Prediction residual spectra for the following cases: DNN in the time-domain (solid green) and in the frequency-domain (dashed green); WF in the time-domain (solid red) and in the frequency-domain (dashed red). The horizontal lines correspond to the mean values of the residual spectra over the plotted frequency range.

while the first accurate study for underground detectors can be found in [224]. In the underground detectors case, the optimal array and residuals depend strongly on the seismic field composition in terms of P-wave and S-wave content [224]. A mixture of wave types would be very difficult for any SNN subtraction.

In the previous chapter, I showed that great filtering performances can be obtained from a dozen of sensors placed at optimized locations. Now, the goal is to find this optimum standard. I decided to tackle the issue by FE modeling the system. The reasons for this choice, instead of the analytical approach (as in section 3.2), are mainly two-fold: 1) the great number of complex structures around the TM demands to evaluate their contribution to the SNN; 2) FE analysis allows the calculation of the real SNN effect on the TM, starting from the seismic field of the area, which is not possible with the filtering techniques. An important drawback is that the simulated seismic field will never be truly representative of the real one. Hence, the idea of performing a high number of temporal FE simulations, changing the seismic field conditions each time. Before presenting the results, I want to introduce the well-known problem of calculating the soil response to a seismic excitation, the so-called Lamb's problem (section 3.5.1), together with the FE simulation framework to assess it (section 3.5.2).




Figure 3.32: MC simulations for the comparison between the DNN and the WF in the frequency-domain. For each value of the time window, 50 trials are shown for the WF, just 5 for the DNN.

3.5.1 The Lamb's problem

In order to calculate the soil response to a seismic excitation, it should be considered a well-known problem in seismology, which is the so-called Lamb's problem. It deals with the classical problem of horizontal and vertical point load applied suddenly onto the surface of an elastic, homogeneous infinite half-space. The equations are valid for the full range of Poisson's ratios ν from 0 to 0.5. The issue was posed in 1904 by Horace Lamb [225], but it took many decades to reach a full mathematical and physical understanding. The first truly complete solutions to Lamb's problem were given in 1855 by Pekeris [226] and in 1960 by Chao [227], who provided closed-form expressions for the components of motion but only for $\nu = 1/4$. This problem was taken up again by Mooney [228], who the solution to any arbitrary Poisson's ratio, but he did so only for the vertical component while ignoring the radial one. Then in 1979, Richards considered this problem once again and gave a complete set of exact formulae in a paper that has remained largely unknown within the community [229]. More recently, a technical note by Kausel [230] presented a reissue of Lamb's problem by means of a very compact set of formulae for all loading cases and Poisson's ratios. The complete solution to the 3D Lamb's problem is derived using the Cagniard-de Hoop method [231,232]. The solutions are quite amenable to numerical calculations. A complete collection of these solutions can be found with a uniform notation and in a form suitable for numerical calculations in [233], in terms of the Green's function for the elastic wave equation in a uniform half-space.

If the final goal is to FE model the gravity gradients response of a GW detector test mass to a seismic excitation, the related analytical solution to the Lamb's problem has to be considered. The task increase with complexity when ranging from an infinite domain to a layered half-space and when a full range of the Poisson's ratio is allowed. Here I mention the simple case of a point source in an infinite media [234]. The notation u_{ij} is used to represent the displacement in the i -direction due to a concentrated point force, specified as $f(\vec{x}, t; \vec{\xi}) = f_0(t)\delta(\vec{x} - \vec{\xi})\hat{e}_j$, with $f_0(t) = H(t)$

the Heaviside step function, applied in the j -direction at the point $\vec{\xi}$. Consequently, the displacements over a media of density ρ may be determined by solving the following equation

$$4\pi\rho u_{ij}(\vec{x}, t) = \frac{3\gamma_i\gamma_j - \delta_{ij}}{r^3} \int_{r/\alpha}^{r/\beta} \tau f_0(t - \tau) d\tau + \frac{\gamma_i\gamma_j}{\alpha^2 r} f_0\left(t - \frac{r}{\alpha}\right) - \frac{\gamma_i\gamma_j - \delta_{ij}}{\beta^2 r} f_0\left(t - \frac{r}{\beta}\right) \quad (3.23)$$

where $r = |\vec{x} - \vec{\xi}|$ is the distance from the source and $\gamma_i = (x_i - \xi_r)/r$ are the direction cosines of the vector. Now, a semi-infinite domain is considered to include surface waves, with the vertical stress acting on the free surface denoted as σ_{33} . From this loading, the vertical $w(t)$ and horizontal $u(t)$ displacements are calculated in the special case of $\nu = 0.25$:

$$\begin{Bmatrix} w(t) \\ u(t) \end{Bmatrix} = \frac{\sigma_{33}}{\pi^2 \mu r} \left(\frac{\alpha}{\beta}\right)^2 \frac{r}{\beta} \int_{-\infty}^r \frac{df}{dt} \Big|_{t=r\tau'/\beta} \begin{Bmatrix} G(\tau - \tau') \\ R(\tau - \tau') \end{Bmatrix} d\tau' \quad (3.24)$$

where $G(\tau)$ and $R(\tau)$ are functions of $\tau = (\beta/r)t$. The set of equations 3.23 and 3.24 are the basis for the FE simulations I will show in the next section.

3.5.2 Numerical simulations

Now I present a dynamical FE simulation framework to calculate the soil response to a seismic excitation. In this section, I will consider a state of the art model of the seismic waves in homogeneous media [235, 236]. Numerical simulations of the ground displacement are done with the COMSOL Multiphysics FE analysis software (see section 2.3.2 for further details on FE modeling). A homogeneous half-space of radius 2.2 km was simulated (a quarter of it is shown in figure 3.33a).

A single vertical excitation force was applied at the origin, with a time dependence of $f(t) = A \sin^2(\pi t/T)$ for $0 \leq t \leq T$. The Poisson's ratio is equal to 0.25 to match with the analytical solutions described in the previous section. It is important to note, for a later discussion, that no material damping was implemented in this model. Figure 3.33b shows the FE results of seismic displacements at a location on the surface. Note the phase difference of $\pi/2$ between the x and z displacements of the Rayleigh wave.

The FE gravity gradient modeling was investigated as well (see figure 3.33c). The nodal displacements were recorded as a function of time and the SNN was calculated again at the same distance as before. Only the SNN acceleration in the horizontal direction is shown since it has the largest effect on the performance of an interferometer. The expected arrival times of the different waves are indicated in the figures and show that the Rayleigh wave dominates the SNN contribution of a detector on the surface. The two plots on the right also show that SNN builds up before any seismic disturbance actually reaches the test mass and for short times is only dependent on surface contributions.

To confirm the ground motion response calculated by a FE simulation, a comparison was made with analytical solutions described in the previous section (Saulson and Hughes and Thorne [117, 118]). The two approaches are in good agreement (see figure 3.34).

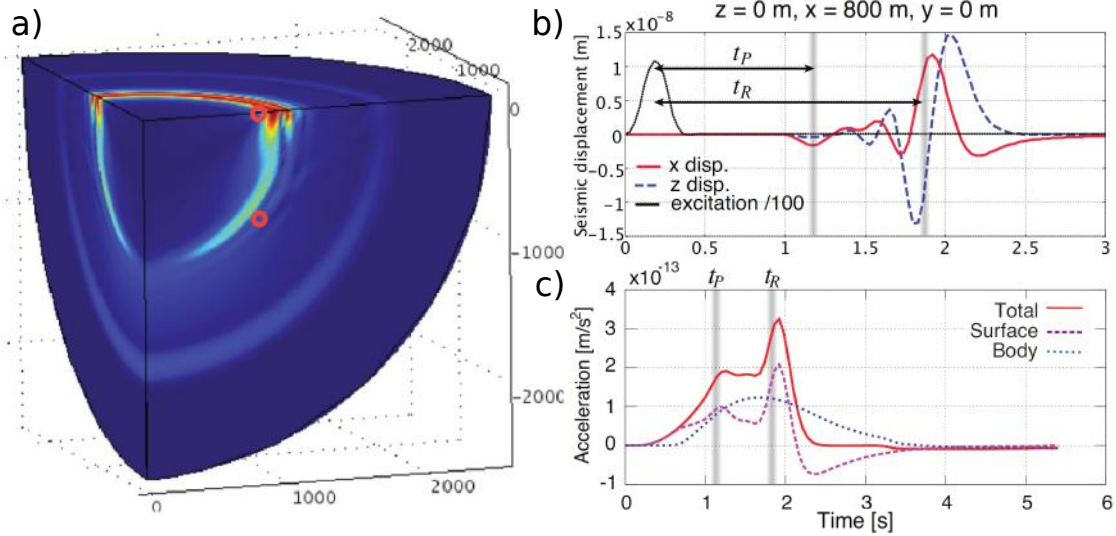


Figure 3.33: a) Total displacement (surface plus body wave contributions) for a time domain simulation at 2.34 seconds after a $1\mu\text{m}$ pulse excitation at the center of the half-sphere. The red circles point to the measurement locations. b) FE ground displacement (along the horizontal and vertical directions) at a surface location after the pulse excitation. Arrival times of Rayleigh and P-waves are also indicated. c) Time domain evolution of SNN acceleration at a surface test mass. Only the horizontal component of the SNN acceleration is shown. The plots are reproduced from [235].

However, analytical descriptions of natural seismic activity, such as the Saulson model and later ones, do not suffice when considering a non-homogeneous medium or complex geologies. Also for local sources, the analytical solutions of the wave fields become more challenging. In these cases, FE analysis could accurately predict seismic motion for a given excitation. However, realistic seismic FE simulations with a complex medium are far more difficult to implement.

One Limitation: Rayleigh damping

A current limitation to this kind of simulations is soil damping. With my group, I first studied the propagation of seismic waves in a simple homogeneous infinite/semi-infinite solid, by comparing the two FE softwares COMSOL and ANSYS, seeking for a cross-validation.

Rayleigh wave propagation in the soil is mainly influenced by geometry damping and material damping, so that the amplitude and energy of the wave will decline as the distance from focus increases. Soil material damping, specifically, is determined by the friction between the particles, the viscosity of water and air voids. Among all kinds of damping models, Rayleigh damping model is the most used one [237]. It sums up these effects in the Rayleigh damping matrix C , which is a linear combination of the mass matrix M and the stiffness matrix K [238]:

$$C = \alpha M + \beta K, \quad (3.25)$$

where α and β represent mass and stiffness proportional damping coefficients. An unappealing feature of Rayleigh damping is that the critical damping ratio ξ_n




Figure 3.34: FE calculation of the SNN for a surface detector. For comparison, the analytical results from Saulson, Hughes and Thorne, and the integral are shown. Figure reproduced from [235].

varies with the natural frequency ω_n of the system vibration modes n :

$$\xi_n = \frac{1}{2\omega_n}\alpha + \frac{\omega_n}{2}\beta. \quad (3.26)$$

In a simple system, the values of the two coefficients α and β are usually selected such that the critical damping ratio is given at two known frequencies. For complex geometries, however, the main exciting frequencies are unknown a priori and the optimal coefficients are chosen by a trial and error approach. Nevertheless, some golden rules exist:

- An high mass proportional term results in over-damping of the low frequency response.
- The mass proportional damping is connected to rigid body motions and therefore it is normally neglected for compliant structures undergoing large rigid body motions.
- For such systems, it is common practice to use the stiffness proportional term only, as the change of α makes little difference to Rayleigh wave's propagation characteristics.
- On the other hand, the change of β makes no difference to the propagation speed, but affects wave's amplitude, which consequently decreases. If only stiffness proportional damping is used then any high frequency response will be artificially over-damped.

Given that, a recommended procedure for applying a specified soil Rayleigh damping is the following:

- Identify a number of critical load cases (sinusoidal, Gaussian or Heaviside pulse).
- Run these cases without Rayleigh damping.
- Identify the minimum and maximum frequencies present in the system response.
- Apply mass and stiffness proportional damping with response period between f_{min} and f_{max} : the resultant damping ratio will be roughly constant in the frequency range of interest.

In this way, the chosen Rayleigh damping will be applied at frequencies f_{min} and f_{max} , while, for intermediate frequencies, the damping ratio will be less than the specified value, but at least it ensures that artificial over-damping is avoided. Note that, if the system response is concentrated at the wave frequency, rather than across a range of frequencies, then it is safe to apply only stiffness proportional damping. The graph in figure 3.35 illustrates the way in which the two damping terms contribute to the overall damping ratio. An example of soil damping numerical simulation can be found in [239], where a minor soil damping was introduced to avoid numerical waviness. There, Rayleigh damping coefficients $\alpha_R = 0.001$ and $\beta_R = 0.002$ was assigned to the soil mass.




Figure 3.35: Calculated frequency-dependent soil damping ratio, by combining stiffness proportional and mass proportional damping: this example case is for $\xi = 2.5\%$ at 50 Hz and 250 Hz, so that the value for intermediate frequencies do not oscillate too much. The figure is reproduced from [240].

Now we use the method of FE to study the influence regulation of mass ratio damping coefficient α and stiffness ratio damping coefficient β in Rayleigh damping on a very simple case. In figure 3.36a we first modeled an homogeneous semi-infinite solid, triggered by a sinusoidal impulse at 2 Hz, with ANSYS and COMSOL softwares. It turned out that, even with a rather extreme meshing and time resolutions

(which would be impractical for any realistic 3D model), we could not agree on the solution within acceptable errors, nor we could determine whether one software was at fault.

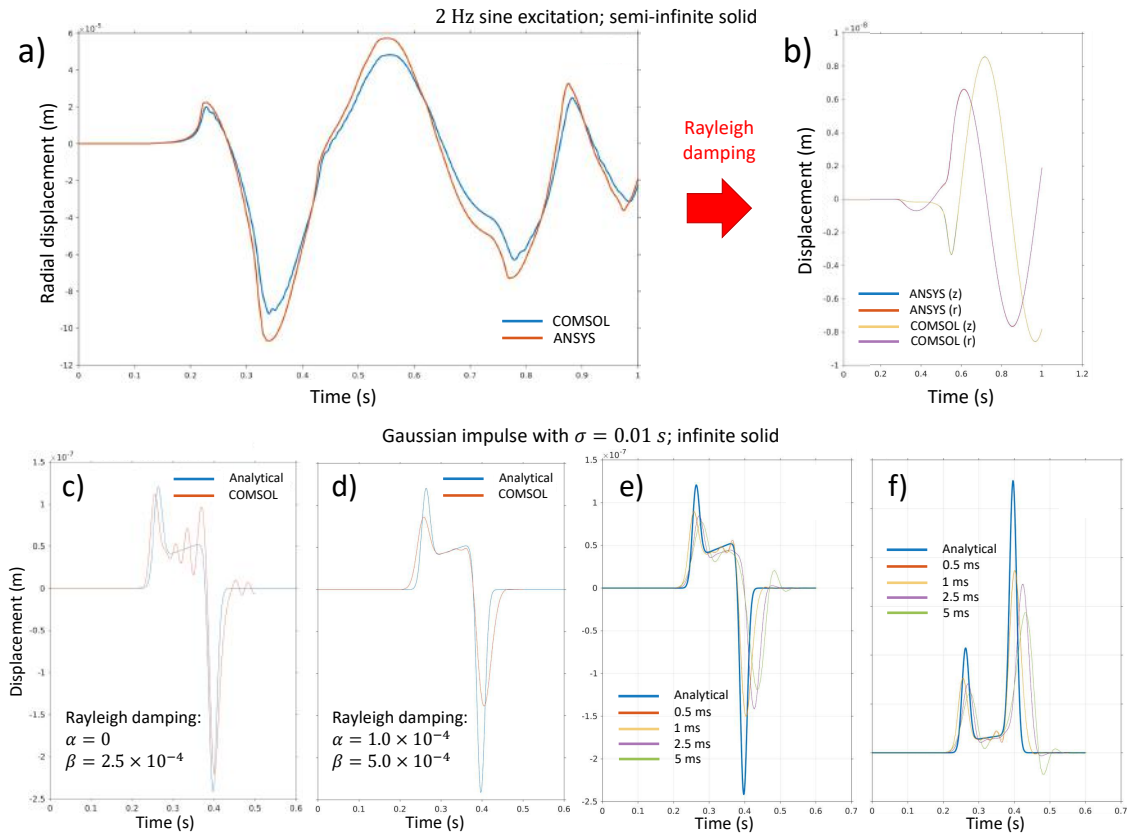


Figure 3.36: Tests of seismic waves FE modeling with added soil damping. a) The radial displacement of the soil (semi-infinite solid) at 5 m distance from the application point of a sinusoidal pulse with frequency 2 Hz is calculated against time. Here, ANSYS and COMSOL do not properly match. b) After introducing the soil damping parameters, the two softwares are in excellent agreement, which is $< 1\%$, for both radial and vertical displacement solution. c-f) further tests by varying the Rayleigh damping parameters and the solution time-step (from 0,5 ms to 5 ms), in the case of Gaussian impulse with $\sigma = 0,01s$ at 7 m distance on the surface of an infinite solid. Reducing the high-frequency ripples has the drawback of reducing the amplitude of the main oscillations (c and d), while increasing the time-step above 1 ms causes a time shift, a mild amplitude drop and a gradual onset of ripples with respect to the analytical solution, both for radial (e) and vertical displacement (f).

When we include Rayleigh damping, the two softwares return to be in excellent agreement, which is $< 1\%$ (see figure 3.36b). With this in mind, we searched for the optimal parameters α and β to match with the analytical calculations based on the Lamb's problem (see section 3.5.1): their choice is a trade-off between the over-damping (figure 3.36c) and the amplitude loss of the wave oscillations (figure 3.36d).

Finally, we also tried to increase the FE solution time-step from 0,5 ms to 5

ms and we noticed that, above 1 ms, the main displacement oscillations undergo a time shift, a mild amplitude drop and the gradual onset of ripples with respect to the analytical curve (both for radial and vertical displacement, see figure 3.36e-f). All these issues are unavoidable when you are dealing with FE analysis of seismic waves in complex systems. Researchers came up with different methods regarding to the value of Rayleigh damping, but a unified understanding has failed to be formed yet. Hence, the accredited way to proceed is optimizing all the parameters stage by stage.

In section 3.5.3, I will show the first efforts towards a semi-realistic FE simulation of an AdV end building and the non-homogeneous terrain around it. These efforts are directed toward the issue not yet solved of the sensor array optimization.

3.5.3 Temporal FE solution

It should be noticed that even the most reliable analytical simulation of the optimal array is not enough reproducible, hence its solutions can hardly trigger a real sensor disposition. Another remark is that they do not consider the ambient around the TM, where the usable space for a sensor array is quite limited. In this respect, it should be mentioned that the ground is partially hollow directly under the TM, which can be taken into account well with dynamical FE simulations. A SNN reduction thanks to the presence of this “Recess” has already been proven [241].

The displacement field induced by a seismic excitation is governed by equations of motion that connect neighboring finite elements, hence the only physical issue in a dynamical simulation is the accurate analytical modeling of the seismic sources. On the other hand, dynamical FE simulations are computationally very expensive, especially if one aim to accurately model the building infrastructure and foundations (end buildings NEB/WEB), different soil layers and heterogeneities in general. In addition, the non-negligible effect of the soil damping (see section 3.5.2) requires a precise and careful procedure, otherwise there is no real advantage over solely analytical calculations.

By the way, a dynamical FE simulation also requires the model and the system geometry to be validated against an experimental measurement, in complete analogy to the magnetic noise case (see section 2.3.3). For the experimental validation of a seismic simulation, one might think of stimulating the sensor array, already deployed inside some Virgo building, with a known seismic excitation at various locations, and see how they react to it. From there, in principle, it would be possible to optimize the FE geometry in order to make the simulations match with the experimental measurements. In practice though, solving the inverse problem is not trivial and it does not only depend on the experimental measure. Moreover, it has not yet been possible to perform such a seismic survey at the site and therefore I focused on the FE simulations.

What I propose to do is to FE model the system under scrutiny, in order to provide support for the sensor array placement. A variety of tools are used for the following simulations. The first one is the commercial software COMSOL Multiphysics. It conveniently interfaces with Matlab, which facilitates data analysis. Finer simulations are run with ANSYS. Results are analyzed both with Python and MATLAB.

Each ANSYS model contains four main elements: a) the terrain, b) the building

structures, c) the seismic sources and d) the virtual sensor grid. A coarse representation of the system is shown in figure 3.37.




Figure 3.37: Sketch of the semi-realistic FE model, built for the computation of the Newtonian noise resulting forces on the test mass, from the displacements calculated by k virtual seismic sensors on the surface (S_j). The simulation includes an homogeneous half-sphere terrain (or a rectangular box) and the relevant surrounding infrastructure. Underground, N point-like sources randomly distributed $P_i = f_0 \sin(\omega_i t + \phi_i)$ with $i = [1, N]$, random ϕ , $\omega = 2\pi f$ with $f = [5, 40]$ Hz, $f_0 = 1$ kN.

- a) The terrain is modeled as a box with rectangular base (dimensions: $-60 < x < 60$, $-60 < y < 60$, $-60 < z < 0$), centered at the TM horizontal location ($x = y = 0$) and vertically 31 m below it (surface at $z = 0$). The box size is such that the arrival time of reflected waves on the edges is greater than the simulation time, so that this effect can be neglected. It is logarithmically mapped (meshed) to reduce the number of nodes, with the idea that a higher density is needed at the center. Each mesh node can be provided with arbitrary mechanical properties, so that a sufficiently high number of configurations can be used, in principle, to adapt the results to the real settings.
- b) The infrastructure belonging to the Virgo buildings includes a significant number of reinforced concrete parts. A problem I had to face when dealing with the modeling of these structures is that, in principle, their mechanical properties are unknown a priori, due to the dual composition (concrete plus steel). Therefore, I divided the structure into macros (plinths, foundation beams, insoles etc...), each one sharing the same steel percentage of the total, which goes from a minimum of 0.5% to a maximum of 2.5%. Finally, the cross-section of each macro is homogenized and its dimensions enlarged, in order to associate to it the elastic modulus of solely concrete. From the NEB structural schemes and blueprints

(figure 3.38a-c), the simplification phase led to the sketched version in figure 3.38d.

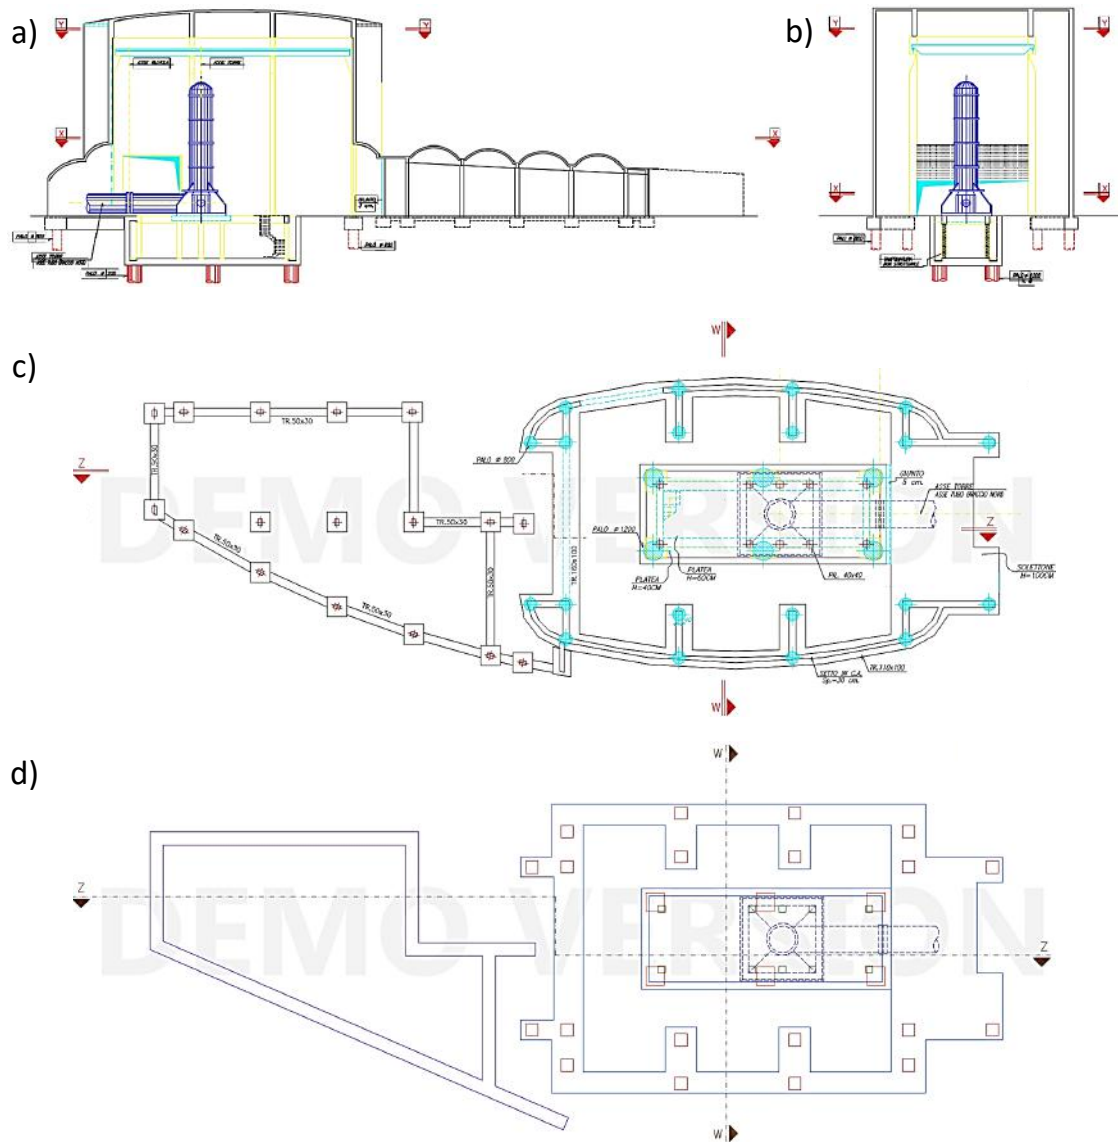


Figure 3.38: North End Building blueprints. a) Section z-z, scale 1:100; b) section w-w, scale 1:100; c) top view of the NEB building CAD drawing; d) geometry simplification of the building grid performed on the FE software ANSYS.

A portion of the building underground structures is depicted in the two insets of figure 3.39, in the form of meshed geometry inside the ANSYS software. Also the foundation pillars, 52 and 30 m long, have been modeled across these preliminary tests. As already mentioned, the “Recess” below the tower platform was modeled too by filling it with a custom material having very soft mechanical properties (Young modulus $E = 10^5$ Pa), in order to mimic an air volume.

- c) Driving seismic sources are randomly scattered in the underground terrain. They are all point-like sources that are simulated as sinusoidal generators with: Gaussian random phase for each simulation time step; intensity of 1 kN; random

direction, defined at the beginning once and for all. Their total number ranges from 1 to 50.

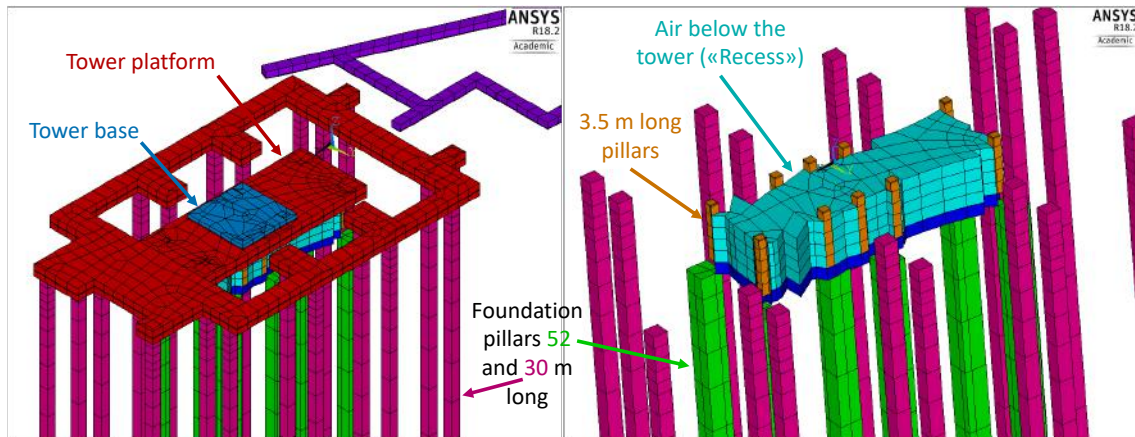


Figure 3.39: North End Building underground structures, consisting of the foundation pillars, the tower base, the tower platform and the “Recess” air mass. They are presented as meshed geometry inside the FE software ANSYS.

- d) Virtual sensors are placed on the surface mesh nodes and are gathered in a regular or irregular grid. These points are used to reproduce the experimental measurements of the seismic displacement. All the points on the surface can be used to evaluate the covariance between the vertical displacements and the newtonian forces on the TM caused by gravity gradients.

In order to prove the feasibility of the entire analysis, I built a toy-model on MATLAB by modeling all the necessary steps towards the extraction of the Newtonian noise forces and to address to the array optimization issue.

In a nutshell, a high number of full sensor covariance matrices (as if the sensors were installed in any grid node) are built on MATLAB from the same number of models (i.e. 200 models); from each model, several partial models with a variable number of sensors randomly picked from the original grid (covering of $0.1 < n/N < 0.99$) are generated (i.e. 10^4 variants computed from 200 models); a Principal Component Analysis algorithm (PCA) is used to extract important features and reduce the datasets; a Deep Neural Network (DNN) is used to predict the full covariance matrix from partial information ($n < N$ sensors); the array optimization is done simply by minimizing the residuals in equation 3.19. This basic demonstrator proved that, if the simulations are sufficiently varied, it can be assumed that an approximation of the real thing is actually simulated. Therefore, one could simply make a great number of ANSYS simulations by varying the grid of sensors and the stochastic distribution of the seismic sources and use the DNN to reconstruct the complete covariance matrix from partial information for array optimization. In other words, it is like building hundreds of partial covariance maps and finding the more plausible and complete one by taking the statistical average and variance over the sensor positions. All without experimentally validating the models. Unfortunately, the first FE simulation took roughly one month to run on a desktop computer station (Intel Xeon E5-2643, 3.3 GHz, 64 GB of RAM) and the

final temporal solution (30 s long and a time step of 10 ms - sampling of 100 Hz) was nearly 3.5 TB heavy, which is quite unfeasible to analyze. Hence, the solution time of a single simulation is clearly not compatible with the necessity of doing hundreds of repetitions.

3.5.4 Probability maps with normal modes

In order to overcome the necessity of doing hundreds of simulations to span the parameter space and match the model with the proposed experimental validation, I started thinking about a radically different approach. This involves pseudo-solutions from normal modes of frequency ω_m for the seismic displacement $u(x, y, z, t)$:

$$\vec{u}(x, y, z, t) = \sum_{m=1}^{\infty} A_m(t) \vec{v}_m(x, y, z) \quad (3.27)$$

where $A_m(t)$ are the unknown amplitudes of the m modes and $\vec{v}_m(x, y, z)$ are the seismic displacement eigenvectors. Here, I assume that all modes contribute to the total motion with the same energy.




Figure 3.40: Top: simulated system with the uni-dimensional massive rope and the probing point at coordinates (x_0, y_0) . The red arrow denote the newtonian force along the x direction. Both the amplitude of the rope oscillations and the intensity of the force are exaggerated in the plot. Bottom: normalized newtonian force to varying of the spatial wavelength of the rope normal mode (divided by the distance between probing point and the rope).

The feasibility of this procedure can be tested with a simple uni-dimensional case study, where I compute the newtonian force produced by a rope with mass M and finite-length L , at the location with coordinates (x_0, y_0) (see figure 3.40 in the upper

panel). From the plot on the bottom of figure 3.40, it turns out that the most relevant normal modes, in terms of normalized newtonian force $C_m = F_m^x / \sigma_{F_m^x}$ along the x direction, at point (x_0, y_0) , are those with spatial wavelength $\lambda \approx 5y_0$. Moving away from this magic value, the contribution decreases less and less, while some modes that are symmetric with respect to the probing point have a null contribution. The generalized effect can be found by doing a weighted average $w(x)$ of all the modes of the simulation (here 80 modes), where the weights are the newtonian forces C_m :

$$w(x) = \sum_m C_m y_m(x) \quad (3.28)$$

Figure 3.41 on the left allows to better visualize the outcome: the upper plot shows the $y_m(x)$, which are the vertical displacements of the rope at point x , for each normal mode m (the blank rows refer to the modes giving a null contribution); the bottom plot shows the weighted average of equation 3.28, highlighting that the most relevant parts of the rope are located very close to the vertical position of the probing point, in a symmetric way. These are the exact locations in which two seismic sensors would be deployed for an optimal SNN cancellation. In the real 3D system, however, all modes would give a certain contribution and the optimal sensor locations would result from the spatial statistic over the entire set of modes. Note that peaks and valleys both denote a maximum level of correlation with the newtonian force.

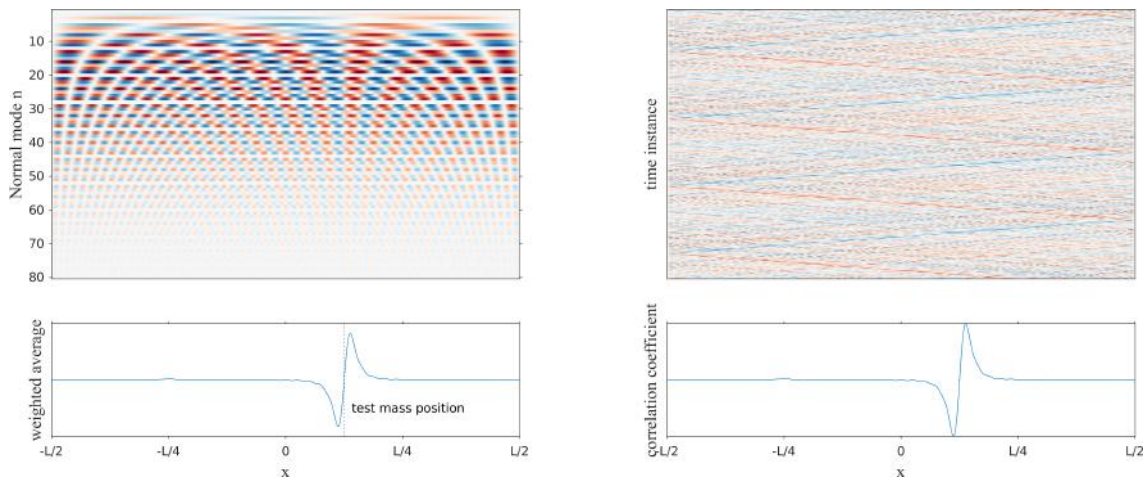


Figure 3.41: Left: vertical displacements $y_m(x)$ of the rope at point x , for each normal mode m (top), and the weighted average $w(x)$ of all the modes of the simulation (bottom). Right: $y_m(x)$, produced by all summed modes, at each time instance (top), and its correlation with the newtonian force (bottom).

Going more in-depth, we can now compute the time evolution of the rope oscillation, resulting from the superposition of all the m modes. The plot on the top right of figure 3.41 shows exactly the $y_m(x)$, produced by all summed modes, at each time instance (1000). In this way, I can calculate the correlation with the newtonian force and obtain the most relevant points, in complete analogy with the weighted average. Having found the same profile, suggests that the temporal simulation and the modal solution are equivalent, at least for the very simple uni-dimensional case.

This pushes us to move to the 3D simulated system. In figure 3.42, I show some examples of normal modes for the underground soil displacement.

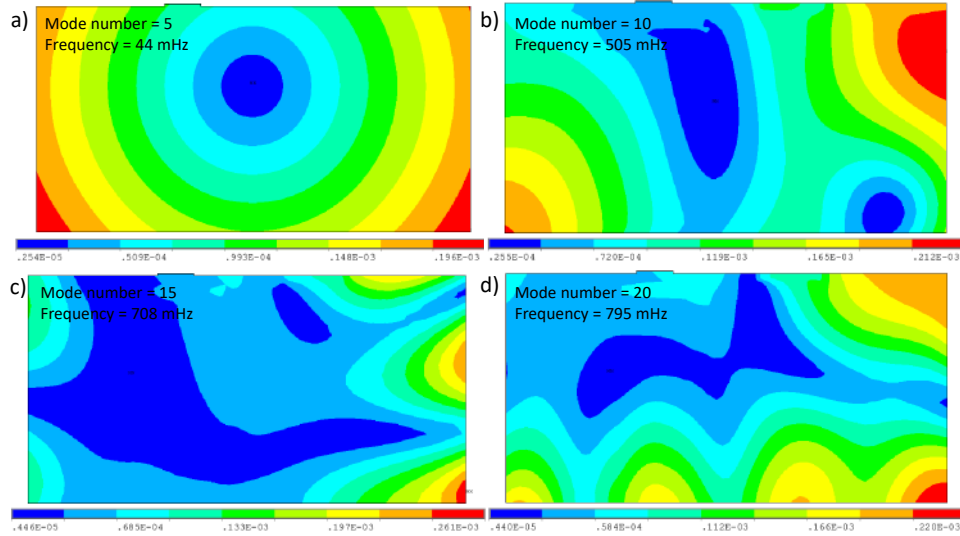


Figure 3.42: Transverse sections of the modeled terrain slab (at the top is the surface), one for each oscillation normal modes, with numbers in the range $[5,10,15,20]$, equivalent to frequencies of mHz. The color maps indicate the normalized energy content on the section plane.

The four color maps, one for each mode $m = 5, 10, 15, 20$, represent the modal energy contents on the transverse section $y = 0$ of the soil.

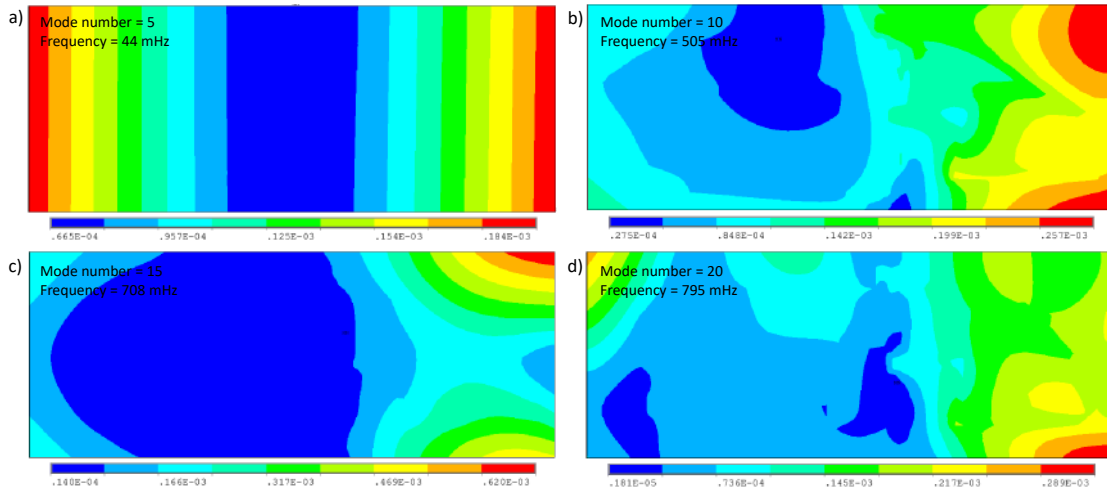


Figure 3.43: Color maps of the oscillation normal modes on the soil surface. Mode numbers are in the range $[5,10,15,20]$, with frequencies from 44 mHz to 795 mHz. The color scale refers to the normalized energy content on the surface plane.

There, the building foundations are not displayed for a clearer view. On the same note, figure 3.43 refers to the modal energy content on the soil surface. What we can tell from these figures is that the building foundation effect starts to contribute from a certain mode number onwards, while, at $m = 5$, the wave fronts are unperturbed. Be aware that fixed boundary conditions had to be introduced at the four underground corners of the rectangular FE box, in order to constrain the motion.




Figure 3.44: Monochromatic maps of the best surface locations at which the seismic sensors have to be deployed to correlate maximally with the newtonian forces on the test mass. The color maps are superimposed on the simulated geometry of the Virgo building foundations, which is not reported for clarity. Only a virtual test mass and incoming laser are displayed. The 12 maps are computed by averaging on the total number of normal modes, ranging from 10 to 9999 logarithmically equidistant modes. The contours of the building platform can be discerned from the first three maps.

That is why the mode frequency values displayed on the color maps are so small (mHz) compared to the typical frequency of the real oscillating system, as they depend strictly from the particular size of the modeled box. However, future tests must assess that the final results do not depend on the particular size of the box. At a final analysis, I plot the corresponding 2D surface maps of the weighted average of equation 3.28, to varying of the total number of normal modes used in the simulation. The results are shown in figure 3.44, in terms of $w^2(x)$. It is evident that by increasing the modal content, the optimal positions converge on narrower spatial regions, mostly situated on the external contour of the tower platform, right outside of it (see also figure 3.45 for a match between the maps and the actual infrastructure).

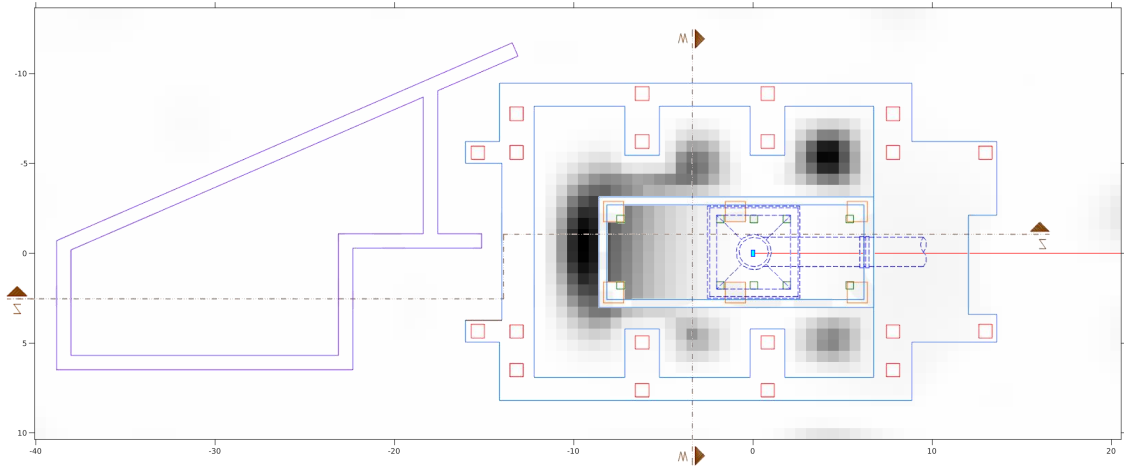


Figure 3.45: The map for the total number of normal modes (9999) is here superimposed on the actual Virgo NEB blueprint.

Moreover, the space occupied by the reinforced concrete parts is affected by a highly depressed modal content. Since the analogy between modal and temporal solution has been proved, we can conclude saying that also 3D simulations with normal modes may result in a conscientious seismic sensor deployment. In the previous section, I showed that the corresponding 3D temporal solution is unfeasible due to the heavy computation, but, if possible, a probability map of the sensor array positions inside the Virgo facilities could be built too, one for each time instance. Figure 3.46 is only a sketch of the final computation we need to perform. The contours in each map refer to the covariance values between the vertical ground displacement witnessed by the sensor array and the newtonian force on the TM, for each time instance. In other words, each map tells the locations at which a certain sensor is maximally correlated with the SNN-induced forces on the TM. The distributions of the covariance intensities for each time instance, over the covariance maps, is calculated once again as in figure 3.41. From this, an optimal location for the sensors can be found.

The main drawbacks of the normal modes approach are that they are taken individually and they are independent of each other. This is in contrast with the temporal solution philosophy of considering the overall seismic excitation of the entire system vibration modes, including the cross effects between them. Despite

the normal modes decomposition being a major simplification, it proved to be a good compromise between the calculation weight of a temporal simulation and a crude analytical solution.




Figure 3.46: Probability maps for the covariance C_{SN} between the vertical ground displacement witnessed by the sensor array and the newtonian force on the TM, for each time instance. The maps refer to the probability density of the sensor array positions inside the Virgo facilities. The red vertical line indicates the intention of averaging over the maps to find an optimal configuration.

However, it requires a validation against the analytical result and a meticulous experimental survey. The overall framework is ready and the calculations are still in progress, hence the complete results are not ready to be presented on this document.

3.6 Discussion

In this chapter, I showed the results obtained from a series of data analysis procedures aimed at proving the feasibility of a seismic newtonian noise cancellation. This kind of noise does not impact the AdV and AdL current GW observations, but it is considered to be one of the most impending noise source for future detector configurations. SNN is also one of the most difficult to mitigate, since newtonian forces are of the same type of signal that the detector aspires to reveal. In summary, it has been shown that deep neural networks perform better than the Wiener filter on average. This may be ascribed to transients or non-idealities in the data and/or a bad Wiener filter implementation. Nevertheless, the DNN-based technique is far from being ready to be used for coherent noise subtraction at Virgo. A non-trivial drawback of machine learning approaches for noise cancellation is linked to the availability of input and output data. The latter, specifically, have to be ideally the real newtonian forces on the Virgo TM, but this is not possible yet, until SNN is a limiting noise at low frequencies. As long as that is the case, the only possible class of target signals for prediction is composed of seismic sensors or, even better, gravity sensors. On the contrary, a Wiener filter can always rely on time integration, since it employs correlations between input seismic sensors and target channel, and hence it

can compensate for poor SNN contrast on the sensitivity curve. Another story is the possibility to make an online noise subtraction with machine learning, by increasing either the history (convolutional window) or the tau for a future prediction. These are crucial internal parameters of a neural network that are closely related to the nature of the data. However, the prediction performances decrease for these two actions, making an online SNN cancellation unfeasible for the time being with the available means.

The sensor array optimization do not directly influence the nature of the noise cancellation tool, but it seriously affects its efficiency. During the last section, I showed that a finite element optimization is preferred to a solely analytical one, on account of the real Virgo facilities complexity and the possibility of modeling the actual impact of the newtonian forces on the TM. However, a match between analytical and numerical simulations is far from being globally perfect, even for considerably simplified models, because of the soil damping. An option could be making a lot of FE simulations by varying the conditions on the underground excitation field, in a way that the real seismic displacement is contained in the statistic. It turns out that the idea of performing a large number of temporal solutions is not feasible in terms of the time taken by a single simulation to converge, at least on a desktop computer. Moreover, it was not possible to parallelize the computation on a computer cluster, in order to reduce the total simulation time. More specifically, a single temporal solution step underlie the computation of the previous temporal steps, and hence the total process have to run on the same computing node. In this sense, future simulations would require an optimization compatible with an experimental seismic validation. Besides that, at least a shortening of the foundation pillars would be necessary, because that would allow to decrease the infinite element boundary radius around the system, especially deep underground. A radically different approach could involve pseudo-solutions from normal modes for the seismic displacement. The issue that the team is currently facing is related to the interaction between the mesh step and the frequency of the normal modes, which, for now, is limited to 6 Hz.

Conclusion

Throughout this thesis I have covered a wide array of subjects related to noise in GW detectors. The strategy pursued the overarching goal of improving the low-frequency noise level of the ground-based interferometric detectors like Advanced Virgo.

In chapter 2, I laid out in detail how I come to understand and characterize different aspects of the instrumental magnetic noise (MN) in Virgo detector. This included a discussion of the tools and experimental methods I used in identifying and mitigating MN noise sources, both local and global. I also designed a method for computing the magnetic coupling to the AdV metallic structure surrounding the test masses and estimating its impact on the total AdV noise budget. As a result of this work, no further serious actions are required for local MN, besides routine monitoring of the electromagnetic activity. At the same time, I outlined a framework for global MN estimation and coherent subtraction to calculate the impact that improvement has on current and future searches for a stochastic gravitational-wave background (SGWB). Now, a permanent Schumann field measuring station operates at Virgo to look for global magnetic correlations with the LIGO and KAGRA counterparts.

In chapter 3, I described the second part of my activity focusing on the seismic newtonian noise (SNN) cancellation. SNN accounts for fluctuating density variations in the surrounding soil that result in tiny variations in the local gravitational field. Because of its nature, SNN is more elusive than MN and is impossible to shield mechanically. I developed some optimization strategies for SNN cancellation with a twofold investigation: first, I employed deep neural networks, whose architectures are tuned with evolutionary algorithms, to test whether machine learning can compete or complement the Wiener filter standard subtraction approach; second, I used numerical finite element simulations to investigate the optimal way in which the available sensor array must be displaced on the field to enhance the cancellation potential.

Seismic newtonian noise and magnetic noise become increasingly relevant towards lower frequencies and eventually they are predicted to limit the sensitivity of the advanced generation detectors at their design configuration (see figure 3.3 for a SNN estimate and figure 2.34 for a MN estimate). Moreover, they will be crucial issues for future generation detectors, which surely will be increasingly sensitive. The improvement of the measuring apparatus works up to a certain point, beyond which the most fundamental and less controllable noises will be the real protagonists to subtract. The planned European ET detector (see section 1.5 for further details) is going to target GW observations down to a few Hertz [111] and its construction will be particularly affected by the effect of MN and SNN. The idea is to build the detector underground to suppress the SNN foreground, which is typically stronger above ground by some orders of magnitude, and secondly to put in operation smart

passive TM shieldings to reduce the environmental MN. The Japanese GW detector KAGRA already adopted the former strategy, as it is built underground at the Mozumi mine [110].

Beside the infrastructure improvements, data analysis strategies are involved to mitigate these two noise sources, directly from the detector output. These include offline coherent noise cancellation based on Wiener filters [121] and/or machine learning. The underlying idea is to monitor the seismic and magnetic fields using auxiliary PEM sensors, and to use their data to generate a coherent prediction of the noise. From here, some challenging issues arise, such as determining the number, position and type of the sensors in order to optimize the cancellation performance [205]. These problems are set out and addressed herein. In addition, a second feasible subtraction algorithm aside from Wiener Filter, which must be robust with respect to terrain parameters, seismic sensor placement, number and type, has been tested for the SNN subtraction. Wiener filter is the best predictor for stationary linear stochastic signals, but here I show that deep neural networks can rival and complement it. Over time, deep neural networks could flank conventional filters more and more, in cases where the signal to be filtered suffer from poor stationarity. This can be the case for seismic and magnetic fields, both global and local kinds. Examples of non-stationary time periods for the seismic field are in correspondence of strong seismic shocks due to natural phenomena (wind, tides, waves, earthquakes ...) or transients related to human activities in the surroundings (wheeled vehicles, rail transport, aircraft, vibrations of machinery, etc...). Magnetic field transients are instead due to variations in the ionospheric electrical activity and all the electrical and electronic devices that are repeatedly turned off/on. the beneficial effects of all countermeasures put in place to reduce these disturbances to sensitivity will be felt very soon.

Appendices

Appendix A

Linear MISO Wiener filter

The Wiener filter is a linear filter that provides an evaluation of a random stationary signal, minimizing the deviation from the predicted target and the estimation [216]. The underlying idea is that the reference signals exhibit some sort of correlation to the target channel. Note that the random processes only need to be wide-sense stationary, which require that first moment (i.e. the mean) and correlations do not vary with respect to time. Wiener filters can be implemented both in the time-domain (real values) or in the frequency-domain (complex values).

Time-domain filter

I call the input random processes reference channels, and collect them as components of a vector \vec{x}_n , where n represents time $t_n = t_0 + n\Delta t$ and Δt is the sampling time. Since the latter is discrete, it is straightforward to introduce the convolution with a finite-impulse response filter (FIR) \vec{w} , characterized by the order N :

$$\hat{y}_n = \sum_{k=0}^N \vec{w}_k \cdot \vec{x}_{n-k} \equiv \vec{w} \circ \vec{x}_n \quad (\text{A.1})$$

where \vec{w} is a $(N + 1) \times M$ matrix with components w_{nm} , assuming M reference channels. This equation assumes that there is only one target channel y_n , which is the case of the multiple-input single-output (MISO) Wiener filter. The hat on the filter output only distinguish the real target channel y_n from its estimated prediction \hat{y}_n . By minimizing the mean-square deviation between target channel and predicted filter output, the coefficients of the Wiener filter can be calculated, directly leading to the Wiener-Hopf equations:

$$\vec{C}_{xx} \cdot \vec{w}(\cdot) = \vec{C}_{xy} \quad (\text{A.2})$$

Here, $\vec{w}(\cdot)$ is the NM -dimensional vector that is obtained by concatenating the M columns of the matrix \vec{w} . The $(N + 1) \times M \times (N + 1) \times M$ symmetric Toeplitz matrix \vec{C}_{xx} is the cross-correlation matrix between reference channels. The $(N + 1)M$ -dimensional vector \vec{C}_{xy} is a concatenation of correlations between each reference channel and the target channel. The correlations in the Wiener-Hopf equations are in general unknown and need to be estimated from measurements. The performance of the filter is assessed by computing the residual of the target channel after subtraction

of \hat{y}_n :

$$r_n = y_n - \vec{w} \circ \vec{x}_n \quad (\text{A.3})$$

This equation summarizes the concept of coherent noise cancellation, used in the text for the seismic newtonian noise subtraction.

Frequency-domain filter

The frequency-domain version of the Wiener filter is obtained by dividing the input data into segments and calculating their discrete Fourier transforms. Equation A.3 remains the same, except for n , which now specifies both the segment and the frequency bin. Interestingly, correlations between noise amplitudes at different frequencies are zero, in case of stationary random processes. This means that noise cancellation can be done on each N frequency bin separately, which is numerically much less demanding and more accurate: the dimensionality of equation A.2 is reduced from NM to M . Moreover, there are no issues associated with the particular choice of the filter order, as in the time-domain implementation.

Appendix B

Artificial Neural Networks

That of artificial neural networks (ANN) is the archetype of a rising technology: interconnected structures of basic computing units, the neurons. Each neuron can be a switch that goes from on to off or vice-versa, depending on the state of the other base units to which it is connected. More specifically, ANNs are computer simulations designed to learn from the available data. Possible uses are for data clustering and classification, pattern recognition, algorithm optimization, function approximation, and data prediction. Their structure is composed by the layers, with one input layer, a certain number of hidden layers and the output layer. Those that represent the real power of an ANN are the connections between all the layers, along with the weight coefficients [242]. The type of networks used in this work belongs to the class of the Deep Neural Networks (DNN). This means that the network has tens or hundreds of hidden layers and the number of nodes in them determines the ANN complexity. With thousands of neurons and millions of interconnections, there is no simple logical path in the system, which is intentionally designed in this way. A crucial aspect is that neural connections are not fixed in advance, but they adapt during a process of trial and error called training.

Several emerging methods give to artificial intelligence systems, such as ANNs, features once considered solely human (see figure B.1). I would mention, for example, the meta-learning that prepares a network to adapt rapidly, so it can learn new tasks without needing a huge amount of data; the so-called generative antagonist networks offer a kind of fantasy, allowing machines to reproduce the statistical characteristics of the datasets; finally, the disentanglement allows the ANNs to become sensitive to the underlying structure of the data and thus makes its internal mechanisms more understandable in human terms.

Supervised learning (when training samples have corresponding known labels) often distinguish between classification and regression techniques to develop predictive models [244]. Labels have a discrete nature in the classification case, and a continuous nature in the regression case (the one I have been interested in). While both classification and regression models have a dense layer at the end, the latter have only a single output and a linear activation. Thus the dense layer just returns the sum of the activations from the previous layer. In addition, the loss function used for regressions is typically the mean squared error (MSE). The analysis of time sequences of sensor data, such as those from gyroscopes, microphones or accelerometers, is suitable for the particular type of Convolutional Neural Networks (CNN). A 1D CNN is very effective when the goal is to extract interesting features

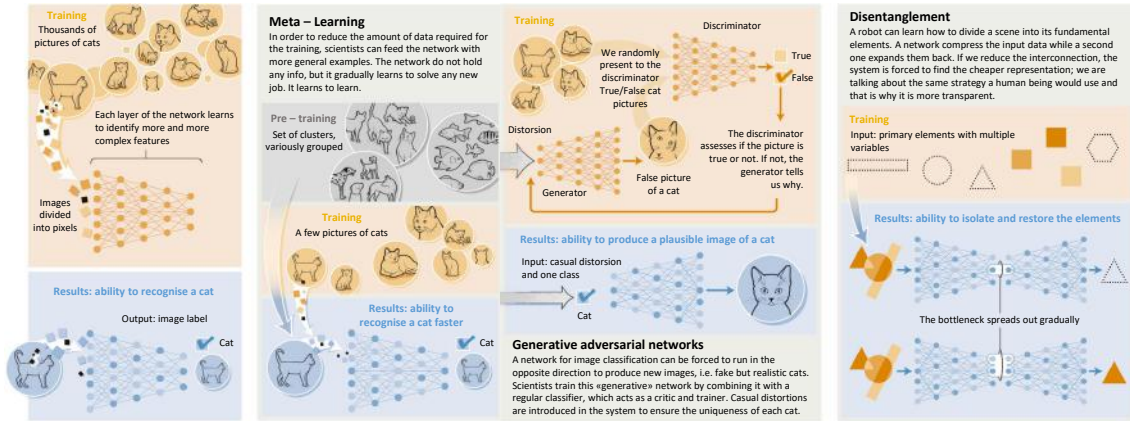


Figure B.1: Three emerging methods of artificial intelligence, applied to neural networks: meta-learning, generative adversarial networks and disentanglement. The figure is readapted from [243].

from shorter segments of the total data set and where the feature location is irrelevant [245]. Here I provide some additional basic concepts about ANNs, which could be useful to better appreciate the main work, without having in any way the intention to be exhaustive.

Activation function

The activation function is indispensable for the ANN to be able to learn non-linear relations within data, using only a small number of layers. Consider a neuron:

$$Y = \sum (weight * input) + bias. \quad (B.1)$$

Theoretically, Y can range from $-\infty$ to $+\infty$ and so the activation function decides whether the input information received by the neuron Y is relevant (activated neuron) or can be discarded. A basic example is a threshold-based activation function (step function), which only works for binary classifier. Intermediate values are obtained for instance with a linear function, whose derivative is however a constant, not allowing further improvements during the back propagation. The solution, at least for classification purposes, are the two almost equivalent “Sigmoid” and “Tanh” activation functions. The problem with them is that the output values at the function edges respond poorly to changes in the input (small gradients), giving rise to the problem of “vanishing gradients”. The Rectified Linear Unit (ReLU) function solves the problem of sparsity and is less computationally expensive, but still does not consider negative values. A further variation in ReLU simply makes negative values not zero by parametrizing the function: this is the Leaky ReLU, which is the one I extensively used in this work.

Gradient descend

An optimization algorithm such as gradient descent is often used for finding the internal parameters of a DNN. It works by allowing the model to make predictions on training samples and using the prediction error to update the model weights so

that this error can be reduced. Gradient descent can vary in terms of the number of training patterns used to calculate error. Depending on this, there are three main flavors of gradient descent: batch, stochastic, and mini-batch (the one used in this work) [246].

Overfitting

The training process normally uses some set of exemplary data for which the desired output is known (supervised learning). But the algorithm has to be able to generalize, meaning that it must perform well also on predicting new validation data that was not encountered during the training. When the performance on unseen data starts to get worse, while the training is still improving, the issue called overfitting arises. Usually, it occurs when learning is performed too long or where training examples are not sufficient, causing the network to learn very specific random features with no causal relation to the target function. This violates the Occam's razor principle, reducing the robustness of the algorithm [247].

Appendix C

Fast PEM sensor maps

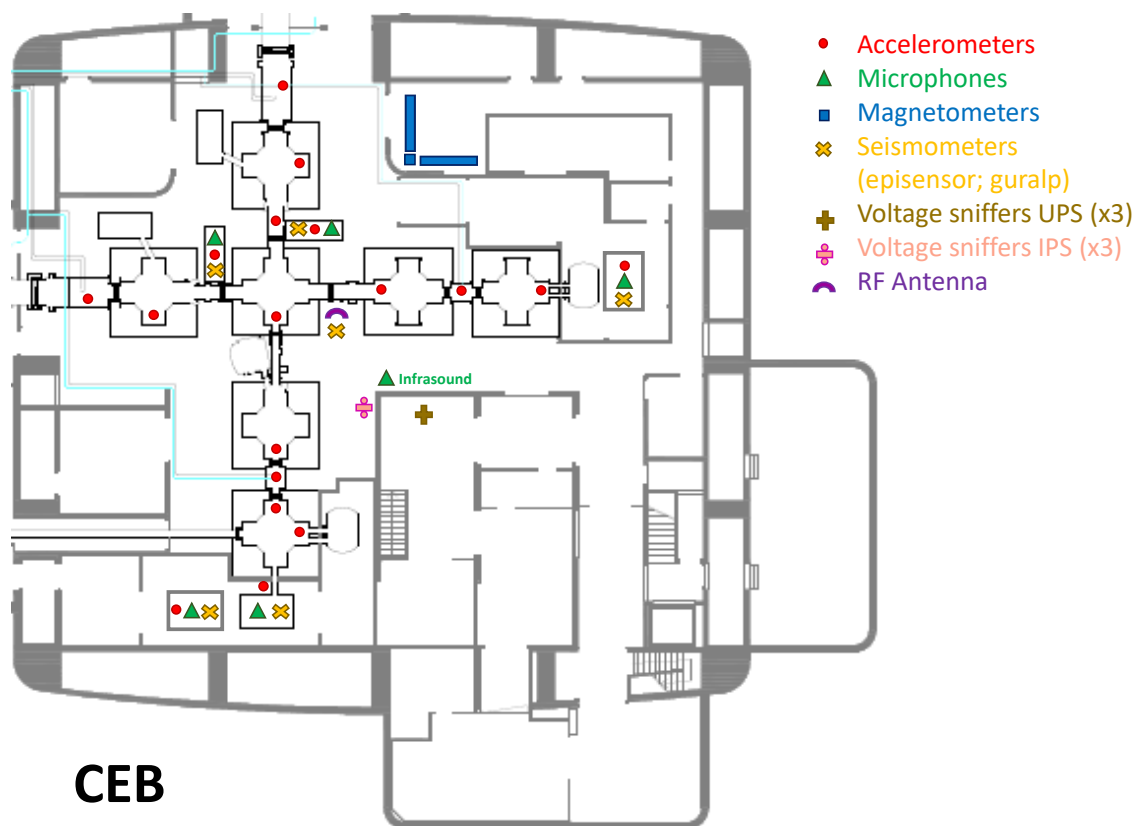


Figure C.1: Fast PEM sensor map at Virgo central building.

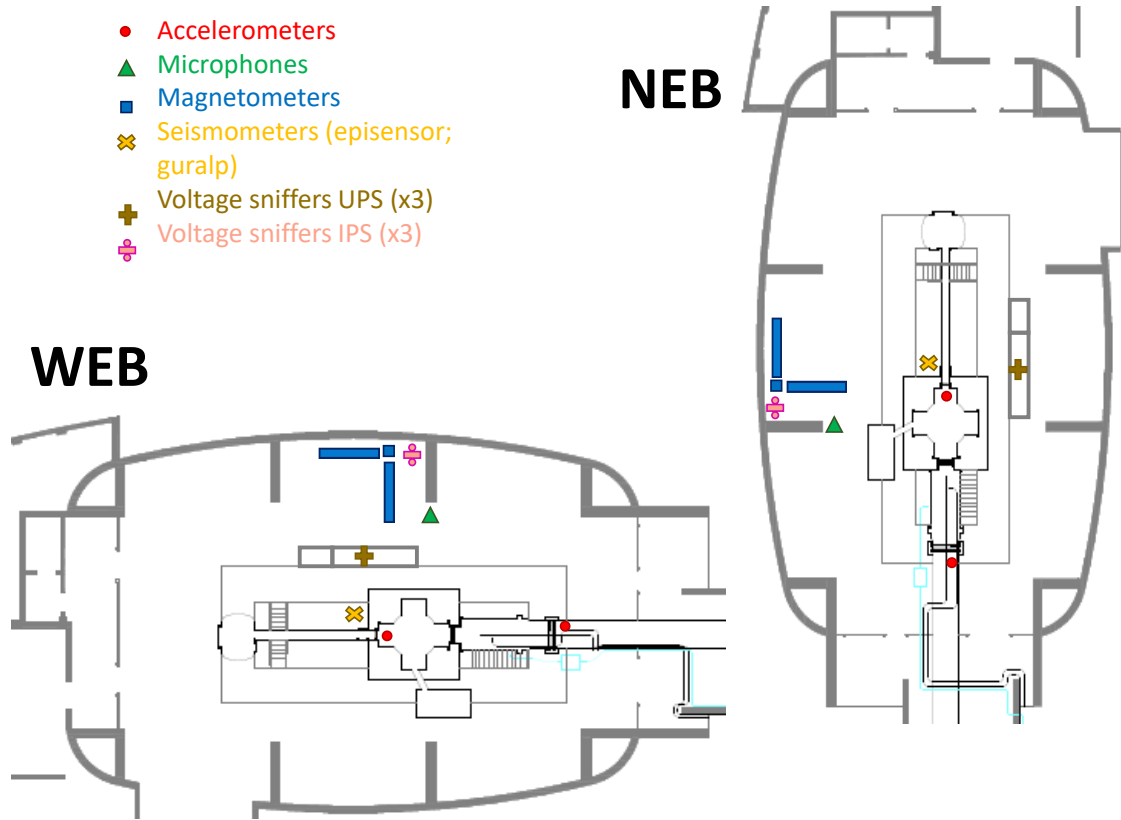


Figure C.2: Fast PEM sensor map at Virgo north-end and west-end buildings.

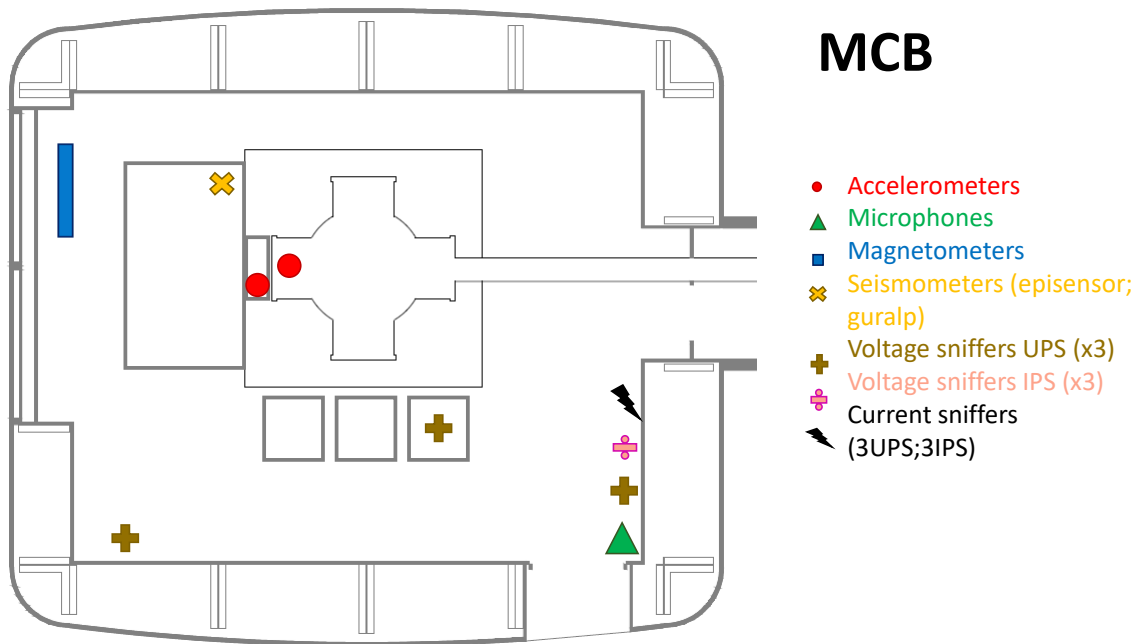


Figure C.3: Fast PEM sensor map at Virgo mode cleaner building.

Bibliography

- [1] A. Einstein. Naeherungsweise Integration der Feldgleichungen der Gravitation. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, pages 688–696, 1916.
- [2] A. Einstein. Über gravitationswellen. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, pages 154–167, 1918.
- [3] W. K. Clifford. On the space theory of matter. In *Proceedings of the Cambridge Philosophical Society*, volume 2, pages 157–158. 1876.
- [4] A. Le Tiec and J. Novak. Theory of gravitational waves. In *An Overview of Gravitational Waves: Theory, Sources and Detection*, pages 1–41. 2017.
- [5] T. Damour, B. R. Iyer, and B. S. Sathyaprakash. Improved filters for gravitational waves from inspiraling compact binaries. *Physical Review D*, 57(2):885–907, 1998.
- [6] C. Cutler and E. E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform? *Physical Review D*, 49(6):2658–2697, 1994.
- [7] E. Poisson. Gravitational radiation from a particle in circular orbit around a black hole. vi. accuracy of the post-newtonian expansion. *Physical Review D*, 52(10):5719–5723, 1995.
- [8] K. G. Arun, B. R. Iyer, B. S. Sathyaprakash, and P. A. Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-newtonian gravitational wave phasing: The nonspinning case. *Physical Review D*, 71(8):084008, 2005.
- [9] C. Van Den Broeck and A. S. Sengupta. Phenomenology of amplitude-corrected post-newtonian gravitational waveforms for compact binary inspiral: I. signal-to-noise ratios. *Classical and Quantum Gravity*, 24(1):155–176, 2006.
- [10] J. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter. Black-hole binaries, gravitational waves, and numerical relativity. *Reviews of Modern Physics*, 82(4):3069, 2010.
- [11] T. W. Baumgarte and S. L. Shapiro. Numerical relativity and compact binaries. *Physics Reports*, 376(2):41–131, 2003.
- [12] J. A. Faber and F. A. Rasio. Binary neutron star mergers. *Living Reviews in Relativity*, 15(1):8, 2012.

- [13] F. Pretorius. Evolution of binary black-hole spacetimes. *Physical review letters*, 95(12):121101, 2005.
- [14] M. Campanelli, C. O. Lousto, P. Marronetti, and Y. Zlochower. Accurate evolutions of orbiting black-hole binaries without excision. *Physical Review Letters*, 96(11):111101, 2006.
- [15] J. G. Baker, J. Centrella, D.-I. Choi, M. Koppitz, and J. van Meter. Gravitational-wave extraction from an inspiraling configuration of merging black holes. *Physical Review Letters*, 96(11):111102, 2006.
- [16] M. Hannam. Status of black-hole-binary simulations for gravitational-wave detection. *Classical and Quantum Gravity*, 26(11):114001, 2009.
- [17] E. Berti, J. Cardoso, V. Cardoso, and M. Cavaglia. Matched filtering and parameter estimation of ringdown waveforms. *Physical Review D*, 76(10):104044, 2007.
- [18] P. Ajith et al. Inspiral-merger-ringdown waveforms for black-hole binaries with nonprecessing spins. *Physical Review Letters*, 106(24):241101, 2011.
- [19] B. P. Abbott et al. Gw150914: Implications for the stochastic gravitational-wave background from binary black holes. *Physical review letters*, 116(13):131102, 2016.
- [20] J. M. Lattimer and M. Prakash. Neutron star structure and the equation of state. *The Astrophysical Journal*, 550(1):426, 2001.
- [21] P. Jaranowski, A. Krolak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Physical Review D*, 58(6):063001, 1998.
- [22] C. Cutler. Gravitational waves from neutron stars with large toroidal b fields. *Physical Review D*, 66(8):084025, 2002.
- [23] A. Mastrano, A. Melatos, A. Reisenegger, and T. Akgün. Gravitational wave emission from a magnetically deformed non-barotropic neutron star. *Monthly Notices of the Royal Astronomical Society*, 417(3):2288–2299, 2011.
- [24] L. Bildsten. Gravitational radiation and rotation of accreting neutron stars. *The Astrophysical Journal Letters*, 501(1):89, 1998.
- [25] N. Andersson. Gravitational waves from instabilities in relativistic stars. *Classical and Quantum Gravity*, 20(7):105, 2003.
- [26] N. Stergioulas. Rotating stars in relativity. *Living Reviews in Relativity*, 6(1):3, 2003.
- [27] W. Unno, Y. Osaki, H. Ando, and H. Shibahashi. Nonradial oscillations of stars. *Tokyo, University of Tokyo Press*, 1979.
- [28] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, and N. Andersson. Gravitational waves from hot young rapidly rotating neutron stars. *Physical Review D*, 58(8):084020, 1998.

- [29] M. Zimmermann. Revised estimate of gravitational radiation from crab and vela pulsars. *Nature*, 271(5645):524–525, 1978.
- [30] D. I. Jones and N. Andersson. Freely precessing neutron stars: model and observations. *Monthly Notices of the Royal Astronomical Society*, 324(4):811–824, 2001.
- [31] T. Abdelsalhin. Tidal deformations of compact objects and gravitational wave emission. *arXiv preprint arXiv:1905.00408*, 2019.
- [32] J. E. Vines and E. E. Flanagan. First-post-newtonian quadrupole tidal interactions in binary systems. *Physical Review D*, 88(2):024046, 2013.
- [33] T. Hinderer. Tidal love numbers of neutron stars. *The Astrophysical Journal*, 677(2):1216, 2008.
- [34] P. Pani, L. Gualtieri, A. Maselli, and V. Ferrari. Tidal deformations of a spinning compact object. *Physical Review D*, 92(2):024010, 2015.
- [35] J. Gagnon-Bischoff, S. R. Green, P. Landry, and N. Ortiz. Extended i-love relations for slowly rotating neutron stars. *Physical Review D*, 97(6):064042, 2018.
- [36] C. D. Ott. The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, 26(6):063001, 2009.
- [37] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Classical and Quantum Gravity*, 26(20):204015, 2009.
- [38] M. Abernathy et al. Einstein gravitational wave telescope conceptual design study. 2011.
- [39] H. Dimmelmeier, C. D. Ott, A. Marek, and H.-T. Janka. Gravitational wave burst signal from core collapse of rotating stars. *Physical Review D*, 78(6):064056, 2008.
- [40] H. Dimmelmeier, J. A. Font, and E. Müller. Gravitational waves from relativistic rotational core collapse. *The Astrophysical Journal Letters*, 560(2):L163, 2001.
- [41] B. P. Abbott et al. An upper limit on the stochastic gravitational-wave background of cosmological origin. *Nature*, 460(7258):990, 2009.
- [42] L. P. Grishchuk. Amplification of gravitational waves in an isotropic universe. *Zh. Eksp. Teor. Fiz*, 67:825, 1974.
- [43] A. Kosowsky, M. S. Turner, and R. Watkins. Gravitational waves from first-order cosmological phase transitions. *Physical review letters*, 69(14):2026, 1992.
- [44] R. R. Caldwell and B. Allen. Cosmological constraints on cosmic-string gravitational radiation. *Physical Review D*, 45(10):3447, 1992.

- [45] T. Regimbau and J. A. de Freitas Pacheco. Stochastic background from coalescences of neutron star-neutron star binaries. *The Astrophysical Journal*, 642(1):455, 2006.
- [46] T. Regimbau and J. A. de Freitas Pacheco. Gravitational wave background from magnetars. *Astronomy & Astrophysics*, 447(1):1–7, 2006.
- [47] V. Ferrari, S. Matarrese, and R. Schneider. Gravitational wave background from a cosmological population of core-collapse supernovae. *Monthly Notices of the Royal Astronomical Society*, 303(2):247–257, 1999.
- [48] J. Abadie et al. Upper limits on a stochastic gravitational-wave background using ligo and virgo interferometers at 600-1000 hz. *Physical Review D*, 85(12):122001, 2012.
- [49] B. P. Abbott et al. Gw170817: Implications for the stochastic gravitational-wave background from compact binary coalescences. *Physical review letters*, 120(9):091101, 2018.
- [50] K. A. Olive et al. Review of particle physics. *Chinese physics C*, 38(9):090001, 2014.
- [51] T. Callister et al. Polarization-based tests of gravity with the stochastic gravitational-wave background. *Physical Review X*, 7(4):041058, 2017.
- [52] B. P. Abbott et al. Upper limits on the stochastic gravitational-wave background from advanced ligo’s first observing run. *Physical review letters*, 118(12):121101, 2017.
- [53] B. P. Abbott et al. Gw170817: Implications for the stochastic gravitational-wave background from compact binary coalescences. *Physical review letters*, 120(9):091101, 2018.
- [54] J. Aasi et al. The characterization of virgo data and its impact on gravitational-wave searches. *Classical and Quantum Gravity*, 29(15):155002, 2012.
- [55] E. Thrane, N. Christensen, and R. M. S. Schofield. Correlated magnetic noise in global networks of gravitational-wave detectors: Observations and implications. *Physical Review D*, 87(12):123009, 6 2013.
- [56] E. Thrane, N. Christensen, R. M. S. Schofield, and A. Effler. Correlated noise in networks of gravitational-wave detectors: subtraction and mitigation. *Physical Review D*, 90(2):023013, 2014.
- [57] M. W. Coughlin et al. Subtraction of correlated noise in global networks of gravitational-wave interferometers. *Classical and Quantum Gravity*, 33(22):224003, 2016.
- [58] M. W. Coughlin, A. Cirone, P. Meyers, et al. Measurement and subtraction of Schumann resonances at gravitational-wave interferometers. *Physical Review D*, 97(10):102007, 5 2018.

- [59] J. Abadie et al. Search for gravitational waves from compact binary coalescence in ligo and virgo data from s5 and vsr1. *Physical Review D*, 82(10):102001, 2010.
- [60] J Abadie et al. All-sky search for gravitational-wave bursts in the first joint ligo-geo-virgo run. *Physical Review D*, 81(10):102001, 2010.
- [61] J. Abadie et al. Search for gravitational waves from low mass compact binary coalescence in ligo’s sixth science run and virgo’s science runs 2 and 3. *Physical Review D*, 85(8):082002, 2012.
- [62] S. Babak, A. Taracchini, and A. Buonanno. Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity. *Physical Review D*, 95(2):024010, 2017.
- [63] T. Hinderer et al. Effects of neutron-star dynamic tides on gravitational waveforms within the effective-one-body approach. *Physical review letters*, 116(18):181101, 2016.
- [64] A. Nagar et al. Time-domain effective-one-body gravitational waveforms for coalescing compact binaries with nonprecessing spins, tides, and self-spin effects. *Physical Review D*, 98(10):104052, 2018.
- [65] L. Blanchet, B. R. Iyer, C. M. Will, and A. G. Wiseman. Gravitational waveforms from inspiralling compact binaries to second-post-newtonian order. *Classical and Quantum Gravity*, 13(4):575, 1996.
- [66] A. Buonanno, B. R. Iyer, E. Ochsner, Y. Pan, and B. S. Sathyaprakash. Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Physical Review D*, 80(8):084043, 2009.
- [67] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Physical Review D*, 59(8):084006, 1999.
- [68] T. Damour and A. Nagar. The effective-one-body approach to the general relativistic two body problem. In *Astrophysical Black Holes*, pages 273–312. Springer, 2016.
- [69] S. Husa et al. Frequency-domain gravitational waves from nonprecessing black-hole binaries. i. new numerical waveforms and anatomy of the signal. *Physical Review D*, 93(4):044006, 2016.
- [70] L. Blanchet. Gravitational radiation from post-newtonian sources and inspiralling compact binaries. *Living Reviews in Relativity*, 17(1):2, 2014.
- [71] L. Santamaria et al. Matching post-newtonian and numerical relativity waveforms: Systematic errors and a new phenomenological model for nonprecessing black hole binaries. *Physical Review D*, 82(6):064016, 2010.
- [72] J. Aasi et al. Narrow-band search of continuous gravitational-wave signals from crab and vela pulsars in virgo vsr4 data. *Physical Review D*, 91(2):022004, 2015.

- [73] B. P. Abbott et al. First narrow-band search for continuous gravitational waves from known pulsars in advanced detector data. *Physical Review D*, 96(12):122006, 2017.
- [74] B. P. Abbott et al. Narrow-band search for gravitational waves from known pulsars using the second ligo observing run. *Physical Review D*, 99(12):122002, 2019.
- [75] B. P. Abbott et al. All-sky search for periodic gravitational waves in the o1 ligo data. *Physical Review D*, 96(6):062002, 2017.
- [76] B. P. Abbott et al. First low-frequency einstein@ home all-sky search for continuous gravitational waves in advanced ligo data. *Physical Review D*, 96(12):122004, 2017.
- [77] B. P. Abbott et al. Full band all-sky search for periodic gravitational waves in the o1 ligo data. *Physical Review D*, 97(10):102003, 2018.
- [78] B. P. Abbott et al. First search for nontensorial gravitational waves from known pulsars. *Physical review letters*, 120(3):031104, 2018.
- [79] N. Christensen. Measuring the stochastic gravitational-radiation background with laser-interferometric antennas. *Physical Review D*, 46(12):5250, 1992.
- [80] B. Allen and J. D. Romano. Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities. *Physical Review D*, 59(10):102001, 1999.
- [81] N. Leroy et al. Data quality and vetoes for the gravitational-wave burst and inspiral analyses in virgo’s first science run. *Classical and Quantum Gravity*, 26(20):204007, 2009.
- [82] F. Robinet et al. Data quality in gravitational wave bursts and inspiral searches in the second virgo science run. *Classical and Quantum Gravity*, 27(19):194012, 2010.
- [83] A. Giazotto. Interferometric detection of gravitational waves. *Physics Reports*, 182(6):365–424, 1989.
- [84] A. Abramovici et al. Ligo: The laser interferometer gravitational-wave observatory. *science*, 256(5055):325–333, 1992.
- [85] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Physical review letters*, 116(6):061102, 2016.
- [86] The Virgo Collaboration. Advanced Virgo Technical Design Report. *Technical Report: VIR-0128A-12*, 2012.
- [87] F. Acernese et al. Increasing the astrophysical reach of the advanced virgo detector via the application of squeezed vacuum states of light. *Physical Review Letters*, 123(23):231108, 2019.

- [88] L. Aiello, E. Cesarini, V. Fafone, M. Lorenzini, Y. Minenkov, I. Nardecchia, A. Rocchi, and V. Sequino. Thermal compensation system in advanced and third generation gravitational wave interferometric detectors. In *Journal of Physics: Conference Series*, volume 1226, page 012019. IOP Publishing, 2019.
- [89] T. Kelly, P. J. Veitch, A. F. Brooks, and J. Munch. Accurate and precise optical testing with a differential hartmann wavefront sensor. *Applied optics*, 46(6):861–866, 2007.
- [90] S. Hild et al. Demonstration and comparison of tuned and detuned signal recycling in a large-scale gravitational wave detector. *Classical and Quantum Gravity*, 24(6):1513, 2007.
- [91] T. Accadia et al. Performance of the virgo interferometer longitudinal control system during the second science run. *Astroparticle Physics*, 34(7):521–527, 2011.
- [92] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, 1 2015.
- [93] G. Cella, V. Sannibale, R. DeSalvo, S. Marka, and A. Takamori. Monolithic geometric anti-spring blades. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 540(2-3):502–519, 2005.
- [94] G. Ballardini et al. Measurement of the VIRGO superattenuator performance for seismic noise suppression. *Review of Scientific Instruments*, 72(9):3643–3652, 9 2001.
- [95] S. Braccini et al. Measurement of the seismic attenuation performance of the virgo superattenuator. *Astroparticle Physics*, 23(6):557–565, 2005.
- [96] F. Acernese et al. First locking of the virgo central area interferometer with suspension hierarchical control. *Astroparticle Physics*, 20(6):629–640, 2004.
- [97] G. Losurdo et al. An inverted pendulum preisolator stage for the virgo suspension system. *Review of scientific instruments*, 70(5):2507–2515, 1999.
- [98] A. Bernardini, E. Majorana, P. Puppo, P. Rapagnani, F. Ricci, and G. Testi. Suspension last stages for the mirrors of the Virgo interferometric gravitational wave antenna. *Review of Scientific Instruments*, 70(8):3463–3472, 1999.
- [99] Y. Levin. Internal thermal noise in the ligo test masses: A direct approach. *Physical Review D*, 57(2):659, 1998.
- [100] G. Cagnoli, L. Gammaitoni, J. Kovalik, F. Marchesoni, and M. Punturo. Suspension losses in low-frequency mechanical pendulums. *Physics Letters A*, 213(5-6):245–252, 1996.
- [101] M. Lorenzini et al. The monolithic suspension for the virgo interferometer. *Classical and Quantum Gravity*, 27(8):084021, 2010.
- [102] <https://git.ligo.org/gwinc>.

- [103] K. S. Thorne. Gravitational radiation. *Three hundred years of gravitation*, pages 330–458, 1987.
- [104] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton. Findchirp: An algorithm for detection of gravitational waves from inspiraling compact binaries. *Physical Review D*, 85(12):122006, 2012.
- [105] M. Bassan. Advanced interferometers and the search for gravitational waves. *Astrophysics and Space Science Library*, 404, 2014.
- [106] P. R. Saulson. Fundamentals of Interferometric Gravitational Wave Detectors. 1994.
- [107] A. Buonanno and Y. Chen. Quantum noise in second generation, signal-recycled laser interferometric gravitational-wave detectors. *Physical Review D*, 64(4):042006, 2001.
- [108] Roman Schnabel. Squeezed states of light and their applications in laser interferometers. *Physics Reports*, 684:1–51, 2017.
- [109] G. G. Sorrells, J. A. McDonald, Z. A. Der, and E. Herrin. Earth motion caused by local atmospheric pressure changes. *Geophysical Journal International*, 26(1-4):83–98, 1971.
- [110] T. Akutsu et al. Kagra: 2.5 generation interferometric gravitational wave detector. *Nat Astron*, 3(35-40), 2019.
- [111] M. Punturo et al. The einstein telescope: a third-generation gravitational wave observatory. *Classical and Quantum Gravity*, 27(19):194002, 2010.
- [112] L. Naticchioni et al. Microseismic studies of an underground site for a new interferometric gravitational wave detector. *Classical and Quantum Gravity*, 31(10):105016, 2014.
- [113] J. Aasi et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7), 2015.
- [114] T Accadia et al. Virgo: a laser interferometer to detect gravitational waves. *Journal of Instrumentation*, 7(03):P03012–P03012, 3 2012.
- [115] J. Y. Vinet et al. Scattered light noise in gravitational wave interferometric detectors: a statistical approach. *Physical Review D*, 56(10):6085, 1997.
- [116] R. Weiss. Electromagnetically coupled broadband gravitational antenna. *Quarterly Progress Report, Research Laboratory of Electronics*, 105, 1972.
- [117] P. R. Saulson. Terrestrial gravitational noise on a gravitational wave antenna. *Physical Review D*, 30(4):732–736, 1984.
- [118] Scott A. Hughes and Kip S. Thorne. Seismic gravity-gradient noise in interferometric gravitational-wave detectors. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 58(12):1–27, 1998.

- [119] M. Beccaria et al. Relevance of Newtonian seismic noise for the VIRGO interferometer sensitivity. *Classical and Quantum Gravity*, 15(11):3339–3362, 1998.
- [120] J. Harms. Terrestrial gravity fluctuations. *Living Reviews in Relativity*, 22(1):6, 2019.
- [121] G. Cella. Off-line subtraction of seismic Newtonian noise. In *Recent Developments in General Relativity*, pages 495–503. Springer, 2000.
- [122] M. G. Beker et al. Improving the sensitivity of future GW observatories in the 1-10 Hz band: Newtonian and seismic noise. *General Relativity and Gravitation*, 43(2):623–656, 2011.
- [123] J. C. Driggers, J. Harms, and R. X. Adhikari. Subtraction of newtonian noise using optimized sensor arrays. *Physical Review D*, 86(10):102001, 2012.
- [124] Donatella Fiorucci, Jan Harms, Matteo Barsuglia, Irene Fiori, and Federico Paoletti. Impact of infrasound atmospheric noise on gravity detectors used for astrophysical and geophysical applications. *Physical Review D*, 97(6):62003, 2018.
- [125] J. Harms, A. Bertolini, T. Bulik, H. J. Bulten, E. Calloni, A. Chincarini, I. Fiori, B. Idzkowski, S. Koley, A. Masserot, and M. Suchinski. Newtonian noise cancellation, "advanced virgo plus" design. *Virgo technical document: VIR-0584A-19*, 2019.
- [126] K. Pepper. Newtonian noise simulation and suppression for gravitational-wave interferometers. *Pasadena, CA: LIGO*, 2007.
- [127] K. S. Thorne and C. J. Winstein. Human gravity-gradient noise in interferometric gravitational-wave detectors. *Physical Review D*, 60(8):082001, 1999.
- [128] A. Einstein. Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der physik*, 322(8):549–560, 1905.
- [129] P. R. Saulson. Thermal noise in mechanical experiments. *Physical Review D*, 42(8):2437, 1990.
- [130] V. B. Braginsky, M. L. Gorodetsky, and S. P. Vyatchanin. Thermodynamical fluctuations and photo-thermal shot noise in gravitational wave antennae. *Physics Letters A*, 264(1):1–10, 1999.
- [131] G. Harry, T. P. Bodiya, and R. DeSalvo. *Optical coatings and thermal noise in precision measurement*. Cambridge University Press, 2012.
- [132] E. M. Gretarsson and A. M. Gretarsson. Three methods for characterizing thermo-optic noise in optical cavities. *Physical Review D*, 98(12):122004, 2018.
- [133] The Virgo Collaboration. Advanced Virgo Plus Phase I - Design Report. *Technical Report: VIR-0596A-19*, 2019.

- [134] A. Chincarini, A. Cirone, and G. Gemme. FEA Charging Simulations. *Technical Report: VIR-0827B-18*, 2018.
- [135] V. Boschi, A. Magazzu, E. Majorana, F. Paoletti, D. Passuello, and P. Ruggi. Actuation noise projections. *Technical Report: VIR-0879A-18*, 2018.
- [136] J. Aasi et al. The characterization of Virgo data and its impact on gravitational-wave searches. *Classical and Quantum Gravity*, 29(15), 2012.
- [137] Izabela Kowalska-Leszczynska et al. Globally coherent short duration magnetic field transients and their effect on ground based gravitational-wave detectors. *Classical and Quantum Gravity*, 34(7):074002, 4 2017.
- [138] I. Fiori, P. Ruggi, B. L. Swinkels, A. Chincarini, S. Farinon, M. Neri, P. Rapagnani, and P. Puppò. Tentative upper limit of magnetic field distortion factor at AdV mirror actuation magnets. *Virgo technical document: VIR-0171A-14*, 2014.
- [139] A. Cirone, A. Chincarini, M. Neri, S. Farinon, G. Gemme, I. Fiori, F. Paoletti, E. Majorana, P. Puppò, P. Rapagnani, Ruggi P., and B. L. Swinkels. Magnetic coupling to the advanced virgo payloads and its impact on the low frequency sensitivity. *Review of Scientific Instruments*, 89(11):114501, 2018.
- [140] A. Cirone, I. Fiori, F. Paoletti, M. M. Perez, A. R. Rodríguez, B. L. Swinkels, A. M. Vazquez, G. Gemme, and A. Chincarini. Investigation of magnetic noise in advanced virgo. *Classical and Quantum Gravity*, 36(22):225004, 2019.
- [141] M. Was, B. Swinkels, and R. Flaminio. Advanced Virgo Plus: Technical Noises - for IRC meeting. *Technical Report: VIR-0732A-19*, 2019.
- [142] T. Accadia et al. The noemi (noise frequency event miner) framework. In *Journal of Physics: Conference Series*, volume 363, page 012037. IOP Publishing, 2012.
- [143] M. Coughlin et al. Noise line identification in ligo s6 and virgo vsr2. In *Journal of Physics: Conference Series*, volume 243, page 012010. IOP Publishing, 2010.
- [144] F. Robinet. Omicron: An Algorithm to Detect and Characterize Transient Noise in Gravitational-Wave Detectors. *Technical Report: VIR-0545C-14*, 2015.
- [145] W. O. Schumann. Über die strahlungslosen eigenschwingungen einer leitenden kugel, die von einer luftschicht und einer ionosphärenhülle umgeben ist. *Zeitschrift für Naturforschung A*, pages 149–154, 1952.
- [146] W. O. Schumann. Über die dämpfung der elektromagnetischen eigenschwingungen des systems erde—luft—ionosphäre. *Zeitschrift für Naturforschung A*, pages 250–252, 1952.
- [147] V. B. Braginsky, O. G. Ryazhskaya, and S. P. Vyatchanin. Notes about noise in gravitational wave antennas created by cosmic rays. *Physics Letters A*, 350(1-2):1–4, 2006.

- [148] G. Vajente. *Analysis of sensitivity and noise sources for the Virgo gravitational wave interferometer*. PhD thesis, Scuola Normale di Pisa, 2008.
- [149] A. Effler, R. M. S. Schofield, V. V. Frolov, G. González, K. Kawabe, J. R. Smith, J. Birch, and R. McCarthy. Environmental influences on the ligo gravitational wave detectors during the 6th science run. *Classical and Quantum Gravity*, 32(3):035017, 2015.
- [150] B. P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quantum Grav.*, 33(13), 2016.
- [151] B P Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. pages 1–48.
- [152] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters*, 119(16):161101, 10 2017.
- [153] B. P. Abbott et al. Prospects for observing and localizing gravitational-wave transients with advanced ligo, advanced virgo and kagra. *Living Reviews in Relativity*, 21(1):3, 2018.
- [154] Bernard F. Schutz. Networks of gravitational wave detectors and three figures of merit. *Classical and Quantum Gravity*, 28(12), 2011.
- [155] B. P. Abbott et al. Low-latency Gravitational-wave Alerts for Multimessenger Astronomy during the Second Advanced LIGO and Virgo Observing Run. *The Astrophysical Journal*, 875(2):161, 2019.
- [156] B. P. Abbott et al. Gw170817: observation of gravitational waves from a binary neutron star inspiral. *Physical Review Letters*, 119(16):161101, 2017.
- [157] B. P. Abbott et al. Multi-messenger observations of a binary neutron star merger. *Astrophys. J. Lett*, 848(2):L12, 2017.
- [158] Stephen Fairhurst. Source localization with an advanced gravitational wave detector network. *Classical and Quantum Gravity*, 28(10), 2011.
- [159] S. Hild et al. Sensitivity studies for third-generation gravitational wave observatories. *Classical and Quantum Gravity*, 28(9), 2011.
- [160] B. P. Abbott et al. Exploring the sensitivity of next generation gravitational wave detectors. *Classical and Quantum Gravity*, 34(4), 2017.
- [161] B. Canuel et al. Exploring gravity with the miga large scale atom interferometer. *Scientific Reports*, 8(1):14064, 2018.
- [162] H. Paik et al. Low-frequency terrestrial tensor gravitational-wave detector. *Classical and Quantum Gravity*, 33(7):075003, 2016.
- [163] A. Shoda et al. Search for a stochastic gravitational-wave background using a pair of torsion-bar antennas. *Physical Review D*, 89(2):027101, 2014.

- [164] D. J. McManus et al. Torpedo: A low frequency gravitational force sensor. In *Journal of Physics: Conference Series*, volume 716, page 012027. IOP Publishing, 2016.
- [165] H. Volland. *Atmospheric electrodynamics*, volume 11. Springer Science & Business Media: Berlin, Germany, 1984.
- [166] N. Tesla. The transmission of electrical energy without wires as a means for furthering peace. *Electrical World and Engineer*, pages 21–24, 1905.
- [167] W. O. Schumann. Über die ausbreitung sehr langer elektrischer wellen um die erde und die signale des blitzes. *Il Nuovo Cimento*, pages 1116–1138, 1952.
- [168] W. O. Schumann and H. König. Über die beobachtung von “atmospherics” bei geringsten frequenzen. *Naturwissenschaften*, pages 183–184, 1954.
- [169] M. Balser and C. A. Wagner. Observations of earth-ionosphere cavity resonances. *Nature*, 188:638–641, 1960.
- [170] M. Balser and C. A. Wagner. Effect of a high-altitude nuclear detonation on the earth-ionosphere cavity. *Journal of Geophysical Research*, 68(13):4115–4118, 1963.
- [171] H. Rowe. Extremely low frequency (elf) communication to submarines. *IEEE Transactions on Communications*, 22(4):371–385, 1974.
- [172] M. Balser and C. A. Wagner. Diurnal power variations of the earth-ionosphere cavity modes and their relationship to worldwide thunderstorm activity. *Journal of Geophysical Research*, 67(2):619–625, 1962.
- [173] A. P. Nickolaenko and L. M. Rabinowicz. Study of the annual changes of global lightning distribution and frequency variations of the first schumann resonance mode. *Journal of Atmospheric and Terrestrial Physics*, 57(11):1345–1348, 1995.
- [174] S. J. Heckman, E. Williams, and B. Boldi. Total global lightning inferred from schumann resonance measurements. *Journal of Geophysical Research: Atmospheres*, 103(D24):31775–31779, 1998.
- [175] H. Yang, V. P. Pasko, and G. Satori. Seasonal variations of global lightning activity extracted from schumann resonances using a genetic algorithm method. *Journal of Geophysical Research: Atmospheres*, 114(D1), 2009.
- [176] C. Price and D. Rind. The effect of global warming on lightning frequencies. *Proceedings of the AMS 16th Conference on Severe Storms, American Meteorological Society: Alberta, AB, Canada*, 1990.
- [177] E. R. Williams. The schumann resonance: A global tropical thermometer. *Science*, 256(5060):1184–1187, 1992.
- [178] A. P. Nickolaenko and L. M. Rabinowicz. On the possibility of existence of global electromagnetic resonances on the planets of solar system. *Space Research*, 20:82–89, 1982.

- [179] O. Pechony and C. Price. Schumann resonance parameters calculated with a partially uniform knee model on earth, venus, mars, and titan. *Radio science*, 39, 2004.
- [180] C. Price. Elf electromagnetic waves from lightning: the schumann resonances. *Atmosphere*, 7(9):116, 2016.
- [181] H. J. Christian et al. Global frequency and distribution of lightning as observed from space by the optical transient detector. *Journal of Geophysical Research: Atmospheres*, 108, 2003.
- [182] E. Greenberg and C. Price. Diurnal variations of elf transients and background noise in the schumann resonance band. *Radio Science*, 42(02), 2007.
- [183] A. P. Nickolaenko and M. Hayakawa. *Resonances in the Earth-ionosphere cavity*, volume 19. 2002.
- [184] Nanometrics. *Centaur Digital Recorder User's Guide*.
- [185] V. Boschi and A. Gennai. Gralp cmg-3td and cmg-eam quick guide. *Virgo technical document: VIR-0172A-10*, 2010.
- [186] <https://www.geo-metronix.de/mtxgeo/index.php/mfs-06e-overview>.
- [187] <http://www.lemisensors.com/?p=245>.
- [188] A. Kulak and others. Extremely low frequency electromagnetic field measurements at the hylaty station and methodology of signal analysis. *Radio Science*, 49(6):361–370, 2014.
- [189] <https://www.agh.edu.pl/en/science/info/article/agh-ust-scientist-heads-work-on-unique-radiolocation-system/>.
- [190] K. Schellbach. Probleme der variationsrechnung. *Journal fr die reine und angewandte Mathematik*, 41:293–363, 1851.
- [191] R. Courant. *Variational methods for the solution of problems of equilibrium and vibrations*. Verlag nicht ermittelbar, 1943.
- [192] O. C. Zienkiewicz. *The finite element in engineering science*. 1971.
- [193] G. Strang and G. J. Fix. *An analysis of the finite element method*, volume 212. Prentice-hall Englewood Cliffs, NJ, 1973.
- [194] William BJ. Zimmerman. *Multiphysics modeling with finite element methods*. World Scientific Publishing Company, 2006.
- [195] <https://www.simscale.com/docs/content/simwiki/fea/whatisfea.html>.
- [196] Ragnar Holm. *Electric Contacts*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1967.
- [197] Milenko Braunovic, Nikolai K. Myshkin, and Valery V. Konchits. *Electrical Contacts: Fundamentals, Applications and Technology*. CRC Press, 2007.

- [198] <https://www.comsol.it/multiphysics/fea-software>.
- [199] Douglas Montgomery. *Design and Analysis of Experiments*. Hoboken, NJ: John Wiley & Sons, Inc, 8th edition, 2013.
- [200] O. Kempthorne. *The Design and Analysis of Experiments*. Wiley eastern private limited, 1952.
- [201] R.L. Plackett and J.P. Burman. The design of optimum multifactorial experiments. *Biometrika*, 33, 1946.
- [202] <https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=39199>.
- [203] L. Naticchioni. The payloads of Advanced Virgo: current status and upgrades. Technical report, 2017.
- [204] J. C. Driggers, M. Evans, K. Pepper, and R. Adhikari. Active noise cancellation in a suspended interferometer. *Review of Scientific Instruments*, 83(2):024501, 2012.
- [205] M. Coughlin, N. Mukund, J. Harms, J. Driggers, R. Adhikari, and S. Mitra. Towards a first design of a newtonian-noise cancellation system for advanced ligo. *Classical and Quantum Gravity*, 33(24):244001, 2016.
- [206] Aki K. and Richards P. G. Quantitative seismology, theory and methods.
- [207] G. F. Miller, H. Pursey, and E. C. Bullard. On the partition of energy between elastic waves in a semi-infinite solid. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 233(1192):55–69, 1955.
- [208] M. W. Coughlin, J. Harms, J. Driggers, D. J. McManus, N. Mukund, M. P. Ross, B. J. J. Slagmolen, and K. Venkateswara. Implications of dedicated seismometer measurements on newtonian-noise cancellation for advanced ligo. *Physical review letters*, 121(22):221104, 2018.
- [209] B. A. Boom, A. Bertolini, E. Hennes, R. A. Brookhuis, R. J. Wiegerink, J. F. J. Van den Brand, M. G. Beker, A. Oner, and D. Van Wees. Nano-g accelerometer using geometric anti-springs. In *2017 IEEE 30th International Conference on Micro Electro Mechanical Systems (MEMS)*, pages 33–36. IEEE, 2017.
- [210] M. Beker et al. Innovations in seismic sensors driven by the search for gravitational waves. *The Leading Edge*, 35(7):590–593, 2016.
- [211] J. Harms and K. Venkateswara. Newtonian-noise cancellation in large-scale interferometric gw detectors using seismic tiltmeters. *Classical and Quantum Gravity*, 33(23):234001, 2016.
- [212] L. Errico. High sensitivity tiltmeters for Seismic and Newtonian Noise studies. *Technical Report: VIR-0850A-19*, 2019.
- [213] M. C. Tringali et al. Seismic array measurements at virgo’s west end building for the configuration of a newtonian-noise cancellation system. *Classical and Quantum Gravity*, 2019.

- [214] J. C. Driggers et al. Improving astrophysical parameter estimation via offline noise subtraction for advanced ligo. *Physical Review D*, 99(4):042001, 2019.
- [215] S. J. Orfanidis. Optimum signal processing an introduction. *Macmillan, New York*, 2007.
- [216] J. Benesty, Y. A. Huang, and J. Chen. Wiener and adaptive filters. In *Springer Handbook of Speech Processing*, pages 103–120. Springer, 2008.
- [217] L. R. Vega and H. Rey. Wiener filtering. In *A Rapid Introduction to Adaptive Filtering*, pages 7–17. Springer, 2013.
- [218] F. A. Fortin, F. M. De-Rainville, M. A. Gardner, M. Parizeau, and C. Gagné. DEAP: Evolutionary algorithms made easy. In *Proceedings of the 14th annual conference companion on Genetic and evolutionary computation*, volume 13, pages 2171–2175. ACM, 2012.
- [219] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [220] M. D. Zeiler. Adadelata: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [221] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [222] J. H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. University of Michigan press Ann Arbor, 1975.
- [223] T. Bäck, D. B. Fogel, and Z. Michalewicz. *Evolutionary computation 1: Basic algorithms and operators*. CRC press, 2018.
- [224] F. Badaracco and J. Harms. Optimization of seismometer arrays for the cancellation of newtonian noise from seismic body waves. *Classical and Quantum Gravity*, 36(14):145006, 2019.
- [225] H. Lamb. On the propagation of tremors over the surface of an elastic solid. *Phil. Trans. R. Soc. Lond. A*, 203(359-371):1–42, 1904.
- [226] C. L. Pekeris. The seismic surface pulse. *Proc. Nat. Acad. Sci.*, 41:469–480, 1955.
- [227] C. C. Chao. Dynamical response of an elastic half-space to tangential surface loadings. *Journal of Applied Mechanics*, 27(3):559–567, 1960.
- [228] H. M. Mooney. Some numerical solutions for lamb’s problem. *Bulletin of the Seismological Society of America*, 64(2):473–491, 1974.
- [229] P. G. Richards. Elementary solutions to lamb’s problem for a point source and their relevance to three-dimensional studies of spontaneous crack propagation. *Bulletin of the Seismological Society of America*, 69(4):947–956, 1979.

- [230] E. Kausel. Lamb's problem at its simplest. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 469(2149), 2013.
- [231] L. Cagniard. *Réflexion et réfraction des ondes séismiques progressives*, volume 6. Gauthier-Villars Paris, 1939.
- [232] A. T. De Hoop. A modification of cagniard's method for solving seismic pulse problems. *Applied Scientific Research, Section B*, 8(1):349–356, 1960.
- [233] L. R. Johnson. Green's function for lamb's problem. *Geophysical Journal International*, 37(1):99–131, 1974.
- [234] J. McKenna, M. McKenna, S. Yushanov, J. Crompton, and K. Koppenhoefer. Computational modeling of wave propagation in a geophysical domain. Technical report, 2008.
- [235] M. G. Beker, J. F. J. van den Brand, E. Hennes, and D. S. Rabeling. Towards time domain finite element analysis of gravity gradient noise. In *Journal of Physics: Conference Series*, volume 228, page 012034. IOP Publishing, 2010.
- [236] M. G. Beker, J. F. J. van den Brand, E. Hennes, and D. S. Rabeling. Newtonian noise and ambient ground motion for gravitational wave detectors. In *Journal of Physics: Conference Series*, volume 363, page 012004. IOP Publishing, 2012.
- [237] P. Danguang. An optimization solution for rayleigh damping coefficients in seismic response analysis. *Engineering Mechanics*, 30(11):15–20, 2013.
- [238] E. L. Wilson. *Static and Dynamic Analysis of Structures*. Computers and Structures, Inc. Berkeley, CA (4th ed.), 2004.
- [239] S. Youzhi. Contribution of rayleigh damping parameters to site response under influence of rayleigh wave. *Rev. Téc. Ing. Univ. Zulia*, 39:48–59, 2016.
- [240] M. Ilankatharan and B. Kutter. Modeling input motion boundary conditions for simulations of geotechnical shaking table tests. *Earthquake spectra*, 26(2):349–369, 2010.
- [241] A. Singha, S. Hild, and J. Harms. Newtonian-noise reassessment for the virgo gravitational-wave observatory including local recess structures. *Virgo technical document: VIR-1254A-19*, 2019.
- [242] S. Nikolić and G. Nikolić. Analysis of financial time series in frequency domain using neural networks. In *Fourier Transforms-Century of Digitalization and Increasing Expectations*. IntechOpen, 2019.
- [243] G. Musser. Immaginazione artificiale. *Le Scienze*, 612, 2019.
- [244] A. Gulli and S. Pal. *Deep learning with Keras*. Packt Publishing Ltd, 2017.
- [245] <https://blog.goodaudience.com/introduction-to-1d-convolutional-neural-networks-in-keras-for-time-sequences-3a7ff801a2cf>.

- [246] <https://machinelearningmastery.com/gentle-introduction-mini-batch-gradient-descent-configure-batch-size/>.
- [247] D. M. Hawkins. The problem of overfitting. *Journal of chemical information and computer sciences*, 44(1):1–12, 2004.

Acknowledgements

During this long journey of three years, I had the opportunity to learn how to become a researcher. The merit of this goes entirely to Andrea Chincarini, my PhD supervisor and teacher of life, but also my friend. I give credit to him and also to Gianluca, Fiodor and to Prof. Maurizio Canepa for having created a wonderful place to work, even if with tiny problems of space. This has allowed generations of young scientists to grow within a thriving connection of minds, between brains and gravitational waves. I am afraid I will be long, but my room buddies deserve personalized thanks. Luca, thank you for your vital IT help, for the coffee breaks (not useful for the coffee itself) and for knowing everything. Diego, thank you for all your insights (and complain) into the Virgo collaboration life. Francesco, thank you for your incredible selflessness and senseless willingness (wow, that's poetry!). Enrico, thank you for reviving the lab with your presence alone, for the long and instructive (and not) conversations and for being a friendly thorn in my side (to say the least XD). I greatly thank all the aquarium with Beatrice, Martina, Barbara, Gloria, Valeria, Gabriella and Nicola, blessed among the women (or maybe not so blessed); I hope they don't complain about the order of appearance, I have a lawyer, a good one. Among the new acquaintances, I want to thank all the beautiful people of the Virgo and LIGO collaborations with whom I had the opportunity to work and travel. A list would in no way be exhaustive, but I still want to mention Irene and Federico, the ideal guides for all young Virgo researchers. Moving away from gravitational people, how to not hugely thank all the friends who accompany my life from longer times before my PhD: Amedeo, Giulia, Carlo, Betta, Ronny, Andrea, and the ones met more recently: Alessio and Alessandro. When I think about roots and friends, lifetime friends, the people you grew up with, a sincere thanks goes to Fabrizio and Luca. My family has been the support that every man in the world would wish to have, always excited for all my new achievements. *Dulcis in fundo*, a thanks as big as the size of all the multiverse plus one (hope it is enough) to Francesca, the love of my life.



Figure 3.4: Characteristic clean room outfit.



Figure 3.5: Characteristic electric vehicle to reach the Virgo end buildings.



Figure 3.6: ... some are less fortunate.