

CONTENTS

I	Gravitational waves	3
1	WAVE SOLUTION OF EINSTEIN FIELD EQUATIONS	4
1.1	An Introduction to General Relativity	4
1.1.1	The Principle of Equivalence	5
1.1.2	General Relativity and Einstein field equations	5
1.2	Wave solution	6
1.3	The Effects of Gravitational Waves	8
2	THE DETECTION OF GRAVITATIONAL WAVES	10
2.1	Hulse–Taylor binary	10
2.2	Weber bars	12
2.3	Interferometric detectors	13
2.4	GW150914: the first detection	15
2.4.1	The other events	15
2.5	GW170814: the first Virgo detection	17
2.6	GW170817: the first BNS detection	19
2.7	The future	21
3	THE VIRGO DETECTOR	24
3.1	Operating principle	24
3.2	Implementation	25
3.3	Sensitivity and noise sources	27
3.3.1	Seismic and gravity gradient noises	28
3.3.2	Thermal noise	29
3.3.3	Quantum noise	30
3.4	The Advanced Virgo project	30
II	The Software Supervisor	33
4	SEISMIC ISOLATION SYSTEM IN ADVANCED VIRGO	34
4.1	The Superattenuators	34
4.1.1	Horizontal Attenuation	34
4.1.2	Vertical attenuation	39
4.2	Digital Control System	40
4.2.1	Top Stage Control	40
4.2.2	Payload Local Control	40
4.3	Suspension Control System Hardware	41
4.3.1	UDSPT boards	42
4.3.2	Other Devices	42

4.4	Deployment	43	
4.4.1	Top Stage Control	44	
4.4.2	Payload Local Control	45	
4.4.3	Other systems	45	
4.5	Performances during O2	45	
5	DSP SOFTWARE	47	
5.1	Boot and operating system	47	
5.2	Interrupt Service Routine	47	
5.2.1	Typical usage	49	
5.2.2	Code development	49	
5.2.3	Parameters Definition	50	
5.3	Services	51	
6	SOFTWARE SUPERVISOR	52	
6.1	Requirements	52	
6.2	TANGO toolkit	52	
6.3	Devices	53	
6.4	Deployment	57	
6.5	Performances	59	
7	APPLICATIONS	61	
7.1	GUI Clients	61	
7.2	Automation	62	
7.2.1	Interferometer automation	62	
7.2.2	Suspension automation	63	
7.2.3	F_0 centering tool	64	
7.3	Python scripting	68	
III	Low-latency searches		69
8	LOW LATENCY SEARCHES	70	
8.1	Burst searches	70	
8.1.1	coherent WaveBurst (cWB)	71	
8.2	Matched-filter searches	72	
8.2.1	GstLAL	74	
9	ON-LINE SEARCH WITH DSP	78	
9.1	Reconstruction of the strain $h(t)$	78	
9.1.1	Optical response	79	
9.1.2	Global control loop	80	
9.1.3	Reconstruction	82	
9.1.4	The Advanced LIGO approach	84	
9.2	Matched filter search	84	
9.3	Performances	86	
9.3.1	Computational cost	86	

9.3.2	Latency	88
9.4	Template bank	89
10	IMPROVEMENTS OF THE EXISTING PIPELINES	92
10.1	Hardware	92
10.2	Computational cost	94
10.2.1	Conclusions	94
10.3	Continuous waves pipelines	95
10.4	Competitors	95
11	CONCLUSIONS	97

INTRODUCTION

Since February 11, 2016 we are aware to live in a universe where compact objects like black holes and neutron stars use to merge in catastrophic events, releasing the equivalent of some solar masses of energy as gravitational radiation in a fraction of second. The announcement of the first detection of gravitational waves [1] came 100 years after the prediction of their existence by Albert Einstein.

The detection of gravitational waves has been done by very complex detectors based on laser interferometry. In the first part of this thesis we introduce the gravitational wave astronomy, up to the detection of a Binary Neutron Star merger on August 17, 2017 by the LIGO and Virgo interferometers.

Thanks to sophisticated suspension systems, that insulate the optical components of the interferometers from seismic noise, it has been possible to extend the bandwidth of the detectors down to few tens of hertz. This has been one of the most important piece in the discovery puzzle, as the gravitational wave signals detected so far carried most of the information below 100 Hz. From the beginning, the Virgo experiment started to develop the Superattenuators, complex mechanical structures used to improve the sensitivity in this bandwidth. Together with mechanics, electronics has been developed to implement digital control systems on the Superattenuators based on Digital Signal Processors (DSP). Even if Superattenuator performances, in its initial design, were already sufficient to detect gravitational waves, several upgrades have been applied in the last twenty years to improve their operations. The last version of the Superattenuators is the result of the *Advanced Virgo* project, whose one of the most important upgrade with respect to the *Virgo+* design concerned the new control electronics.

In August 2017, during the scientific run O2, 131 DSP-based boards were running on the experiment, to control the 10 Virgo Superattenuators. A large number of device requires a supervisory system able to perform automatic tasks, as well as to provide tools that allow users to monitor the state of the system. The first original work of this thesis is presented in the second part: we describe the new supervisory system, the so called *Software Supervisor*, that has been designed, developed and deployed on the Advanced Virgo detector. It actively assisted the installation and the commissioning of the new electronics, and now is integral part of experiment.

Even if the DSP-based boards have been designed specifically for the digital control of the Superattenuators, they are suitable also to perform on-line data analysis on the Virgo data. A key feature is that they are deeply integrated into the experiment, as they are part of the sensing and control system of the interferometer. This means that they already acquire the Virgo data with almost zero latency. In the third and last part of the thesis we will

describe the second original work, a feasibility study about how the large computational power of a DSP-based system can be exploited to perform an on-line low-latency search for gravitational waves from compact binary mergers. We also explored possible improvements to the existing searches using such a system.

Part I

Gravitational waves

1

WAVE SOLUTION OF EINSTEIN FIELD EQUATIONS

In 1916, the year after the final formulation of the field equations of general relativity, Albert Einstein predicted the existence of Gravitational Waves (GWs). He found that the linearized weak-field equations had wave solutions: transverse waves of spatial strain that travel at the speed of light, generated by time variations of the mass quadrupole moment of the source [2, 3]. Einstein understood that gravitational-wave amplitudes would be remarkably small; moreover, until the Chapel Hill conference in 1957 there was significant debate about the physical reality of gravitational waves. At the same time, in 1916, Karl Schwarzschild published a solution for the field equations that was later understood to describe a black hole [4].

This is how began the history of the GW physics. Exactly 100 years later occurred the first detection of the gravitational waves generated from a Binary Black Hole (BBH) system. In this chapter we introduce the theory that describes the wave solution of the Einstein field equations.

1.1 AN INTRODUCTION TO GENERAL RELATIVITY

The gravitational force dominates the universe on the large scale, binding matter into stars, stars into galaxies, and galaxies into cluster of galaxies. The classic theory of gravitation is based on Newton's law of gravity which states that two masses m_1 and m_2 separated by a distance r feel a mutual gravitational attraction

$$F = G \frac{m_1 m_2}{r^2} \quad (1)$$

where G is a constant of proportionality called "universal gravitational constant"; its value is known to be $G = 6.67408(31) \times 10^{-11} \text{ m}^3\text{kg}^{-1}\text{s}^{-2}$ [5]. This equation describes the motion of the planets around the Sun with great accuracy. However, there are several features that cannot be explained by the Newton's law. The most significant one is a tiny component in the precession of the perihelion of the orbit of Mercury, first observed by Urbain J. Le Verrier in 1843 [6].

The main problem is that equation 1 is time independent, which would mean that the gravitational force could act instantaneously at all distances. Such behavior is in flat contradiction to the "Special Theory of Relativity" or Special Relativity which requires that no signal should travel faster than the speed of light c [7].

The problem is shared with electromagnetism and Coulomb's law: in this case it was solved with Maxwell's equations, which are consistent with Spe-

cial Relativity. In November 1915, Albert Einstein published his geometric theory of gravitation, a solution to the problem called “General Theory of Relativity” or General Relativity (GR), that is a description of gravitation consistent with Special Relativity.

1.1.1 The Principle of Equivalence

GR is based on the Principle of Equivalence. This establishes the equality of gravitational and inertial mass, demonstrated at first by Galileo and Newton. Einstein interpreted this result to postulate the “weak equivalence principle”: *the motion of a neutral test body released at a given point in space-time is independent of its composition* [7].

Furthermore, Einstein reflected that, as a consequence, no external static homogeneous gravitational field could be detected in a freely falling elevator, because the observers, their test bodies, and the elevator itself would respond to the field with the same acceleration. Although inertial forces do not exactly cancel gravitational forces for freely falling systems in an inhomogeneous or time-dependent gravitational field, we can still expect an approximate cancellation if we restrict our attention to such a small region of space and time that the field changes very little over the region. Therefore, the “strong equivalence principle” was postulated by Einstein and it states that *at every space-time point in an arbitrary gravitational field is possible to choose a “locally inertial coordinate system” such that, within a sufficiently small region of the point in question, the laws of nature take the same form as in unaccelerated Cartesian coordinate systems in the absence of gravitation.* [8]

1.1.2 General Relativity and Einstein field equations

According to GR, the universe consists of an active space-time continuum that is distorted by matter and energy passing through it. GR allows to describe the metric of space-time as directly related to the energy and momentum of whatever matter and radiation are present. The relation is specified by the Einstein field equations, a set of 10 partial differential equations:

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

$g_{\mu\nu}$ is the metric tensor, that contains information about the intensity of the gravitational field. $R_{\mu\nu}$ represents the curvature tensor of the space-time and contains second derivatives of $g_{\mu\nu}$, while $R = g^{\mu\nu}R_{\mu\nu}$ is called scalar curvature. Λ was introduced by Einstein and is called the *cosmological constant*¹. Finally, $T_{\mu\nu}$ is the energy–momentum tensor and contains the distribution of energy and momentum in the space-time.

A first effect predicted by general relativity was detected by Arthur Stanley Eddington in 1919. The theory suggests that starlight which passes the limb

¹ The cosmological constant is related to the expansion of the universe. For several reasons it has to be very small. For our purposes, we can assume $\Lambda = 0$.

of the Sun on its way to the Earth should be deflected by $1.750''$. Eddington organized an expedition to the Island of Príncipe which photographed the star field around the Sun during a solar eclipse occurred on May 29 1919. When comparison was made with night photographs of the same star field, the predicted general relativistic deflection was confirmed and published in 1920 [9].

Since that day the predictions of general relativity have been confirmed in all observations and experiments. Among the other results, in the limit of low velocities and small gravitational effects, GR reduces to Newton's law with small corrections: in the case of Mercury, these corrections account precisely for the residual advance of perihelion.

1.2 WAVE SOLUTION

There are many similarities between gravitation and electromagnetism. It should therefore come as no surprise that Einstein's equations, like Maxwell's equations, have radiative solutions.

We know that electromagnetic propagation is described by d'Alembert's equations ($c = 1$)

$$\square A_\mu = -J_\mu/\epsilon_0 \quad (3)$$

deriving from Maxwell's equation, where $\square \equiv \partial^\mu \partial_\mu$ is the d'Alembert operator, $A_\mu = (\phi, \mathbf{A})$ describes the electromagnetic potentials and $J_\mu = (\rho, \mathbf{j})$ describes the source of the electromagnetic field [10]. A particular solution to this equation is represented by the retarded potentials:

$$A_\mu(\mathbf{x}, t) = \frac{1}{4\pi\epsilon_0} \int d^3\mathbf{x}' \frac{J_\mu(\mathbf{x}', t - |\mathbf{x}' - \mathbf{x}|)}{|\mathbf{x}' - \mathbf{x}|} \quad (4)$$

They show that the state of the field, in a certain point of the space-time, depends on that of the source at a previous time $t - |\mathbf{x}' - \mathbf{x}|$: the information propagates at speed c into the electromagnetic waves.

The derivation of gravitational waves from Einstein field equations (2) is more complicated than that of electromagnetic waves from Maxwell's equations, due to its non-linearity. Maxwell's equations are linear because the electromagnetic field does not itself carry charge; on the other hand, we may say that any gravitational wave is itself a distribution of energy and momentum that contributes to the gravitational field of the wave: it is impossible to separate the contributions of gravitational waves to the curvature from the contributions of the Earth, the Sun, the galaxy, or anything else. Thus, there is no way to find general radiative solutions of the exact Einstein's equations. Here we present only the weak-field radiative solutions, which describe waves carrying not enough energy to affect their own propagation.

If we suppose to be far from the source of the fields, the space-time will be nearly flat and the metric will be close to the Minkowski metric $\eta_{\mu\nu}$:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (5)$$

with $|h_{\mu\nu}| \ll 1$. Now Einstein field equations 2 can be written to first order in h [8],

$$R_{\mu\nu}^{(1)} = -8\pi G S_{\mu\nu} \quad (6)$$

with

$$R_{\mu\nu} \simeq R_{\mu\nu}^{(1)} \equiv \frac{1}{2} \left(\square h_{\mu\nu} - \frac{\partial^2}{\partial x^\lambda \partial x^\mu} h^\lambda_\nu - \frac{\partial^2}{\partial x^\lambda \partial x^\nu} h^\lambda_\mu + \frac{\partial^2}{\partial x^\mu \partial x^\nu} h^\lambda_\lambda \right)$$

and

$$S_{\mu\nu} \equiv T_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} T^\lambda_\lambda$$

We can also perform a further simplification to these equations and write equations 6 as [8]

$$\square h_{\mu\nu} = -16\pi G S_{\mu\nu} \quad (7)$$

These equations are actually very similar to 3 and naturally we can write retarded solutions

$$h_{\mu\nu}(\mathbf{x}, t) = 4G \int d^3\mathbf{x}' \frac{S_{\mu\nu}(\mathbf{x}', t - |\mathbf{x}' - \mathbf{x}|)}{|\mathbf{x}' - \mathbf{x}|} \quad (8)$$

These solutions describe the physical phenomenon of the gravitational waves produced by the source $S_{\mu\nu}$. They are transverse waves traveling with the same finite speed of propagation c of the electromagnetic waves and the same intensity decrease as function of distance from the source.

Far from the source, as $|\mathbf{x}' - \mathbf{x}| \rightarrow \infty$, the retarded solution approaches a plane wave, and the equations 7 are reduced to the homogeneous ones

$$\square h_{\mu\nu} = 0$$

Here, a solution is [8]

$$h_{\mu\nu} = e_{\mu\nu} \exp(ik_\lambda x^\lambda) + e_{\mu\nu}^* \exp(-ik_\lambda x^\lambda) \quad (9)$$

with

$$k_\mu k^\mu = 0 \quad (10)$$

and

$$k_\mu e^\mu_\nu = \frac{1}{2} k_\nu e^\mu_\mu \quad (11)$$

where $e_{\mu\nu} = e_{\nu\mu}$ is a 4×4 symmetric tensor and is called the *polarization tensor*. In general, a 4×4 matrix have 10 independent components; the gauge invariance 11 reduce them to only 6, but it can be shown [8] that of these six there are only two physically significant degrees of freedom, i.e. only two independent physical polarizations. The commonly used couple of independent polarization is

$$e_{\mu\nu}^+ = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad e_{\mu\nu}^\times = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (12)$$

and it forms a basis for the polarization space. The two elements are pronounced respectively *plus* and *cross* polarizations. We can obtain any other polarization (for example the circular polarizations) by a suitable linear superposition of these two.

Choosing the z -axis as the direction of propagation of the wave, we can write equation 9 as:

$$h_{\mu\nu} = e_{\mu\nu} \cos(k_\lambda x^\lambda) = e_{\mu\nu} \cos(\omega t - kz) \quad (13)$$

Using the polarization in 12, the general form of 13 is made up of a linear combination of the two orthogonal states

$$h_{\mu\nu}^+ = h^+ e_{\mu\nu}^+ \cos(\omega t - kz) \quad (14)$$

and

$$h_{\mu\nu}^\times = h^\times e_{\mu\nu}^\times \cos(\omega t - kz + \varphi) \quad (15)$$

where φ is an arbitrary phase angle, and h^+ and h^\times are the amplitudes of the components. A graphical explanation of these two polarization is shown in 1.

1.3 THE EFFECTS OF GRAVITATIONAL WAVES

To understand what happens to the space-time when it is perturbed by gravitational waves, one can consider two nearby bodies located in the same xy -plane at $A = (\xi, 0)$ and $B = (0, 0)$. If we suppose the weak field condition 5 to be true, their proper separation is

$$ds = \xi' \approx \sqrt{|g_{11}(t, 0)|} \xi \approx \left[1 + \frac{h_{11}(t, 0)}{2} \right] \xi$$

Applying the wave in equation 14, the proper space-time interval between A and B undergoes a strain of amplitude

$$\epsilon_x^+ = \frac{\xi' - \xi}{\xi} = \frac{h_{11}(0, 0)}{2} = \frac{h^+}{2} \quad (16)$$

Thus $\frac{h^+}{2}$ is the amplitude of the differential change in lengths between nearby points along the x -axis. A similar reasoning suggests that the same wave would produce a tidal effect along y -axis of

$$\epsilon_y^+ = -\frac{h^+}{2} \quad (17)$$


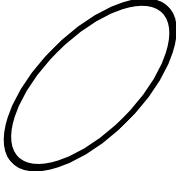
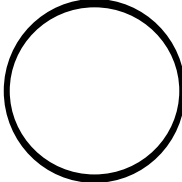
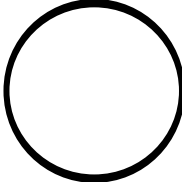
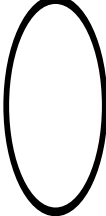
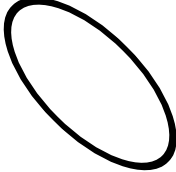
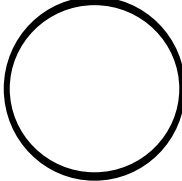
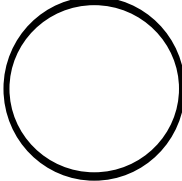
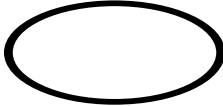
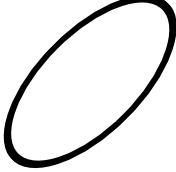
	$h_{\mu\nu}^+ = h^+ e_{\mu\nu}^+ \cos(\omega t)$	$h_{\mu\nu}^\times = h^\times e_{\mu\nu}^\times \cos(\omega t)$
$\omega t = 0$		
$\omega t = \frac{\pi}{2}$		
$\omega t = \pi$		
$\omega t = \frac{3\pi}{2}$		
$\omega t = 2\pi$		

Figure 1: The effect of gravitational waves in two different polarization ($e_{\mu\nu}^+$ and $e_{\mu\nu}^\times$) on a circle of test masses followed over one cycle. The wave is traveling along the z -axis, the paper is the xy -plane at $z = 0$, and the observer is looking towards the source.

2

THE DETECTION OF GRAVITATIONAL WAVES

The search for GW has been one of the most exciting and long challenges in the history of science. Hundreds of people have contributed, and two Nobel Prizes have been awarded in the last decades to the fundamental steps of a story that ended on September 15, 2016 at 09:50:45 UTC, when the two LIGO detectors simultaneously observed the transient gravitational-wave signal generated by two black holes. In this chapter we narrate the key events of this challenge.

2.1 HULSE–TAYLOR BINARY

The discovery of the binary pulsar system *PSR B1913+16* by Russell Alan Hulse and Joseph Hooton Taylor Jr. in 1974 [11] and subsequent observations of its energy loss by Taylor and Weisberg [12] demonstrated the existence of gravitational waves.

The system is 6400 pc far from the Earth, in the Milky Way. Currently the orbit period of the binary system is $P_b = 7.75$ h and the projected orbital velocity is $v \sim c/1000$: this suggest that there can be some measurable relativistic effects. Among the best known results are measurement of the advance of periastron at a rate $\sim 35 \times 10^3$ times that of Mercury in the solar system and, above all, the effect of gravitational waves damping, that causes a measurable rate of orbital decay [13, 14].

Few years earlier, Peters and Matthews [15] had shown that, according to GR, the resulting rate of change in orbital period of a binary system due to the radiated energy in the form of gravitational waves, measured in the orbiting system reference frame, should be

$$\dot{P}_b^{\text{GR}} = (-1.699\,451 \pm 0.000\,008) \times 10^{-12} \left[\frac{m_1 m_2 (m_1 + m_2)^{-1/3}}{M_\odot^{5/3}} \right].$$

Using this expression, we obtain for *PSR B1913+16*, by replacing the measured masses and propagating the uncertainties, a predicted value [14] of

$$\dot{P}_b^{\text{GR}} = (-2.402\,63 \pm 0.000\,05) \times 10^{-12}$$

This value is being measured since the discovery of the binary system, as shown in 2. To properly compare the values, we have to take into account a small additional contribution $\Delta\dot{P}_b = (-0.025 \pm 0.004) \times 10^{-12}$ to the observed \dot{P}_b due to the relative acceleration of that frame with respect to the




Figure 2: Orbital decay caused by the loss of energy by gravitational waves of PSR B1913+16 system. The parabola depicts the expected shift of periastron time relative to an unchanging orbit, according to general relativity. Data points represent Weisberg's measurements, with error bars mostly too small to see [14].

solar system barycenter. The most recent measurement of \dot{P}_b was made in 2016 [14] and the agreement with the predicted value is extraordinary:

$$\frac{\dot{P}_b - \Delta\dot{P}_b}{\dot{P}_b^{\text{GR}}} = 0.9983 \pm 0.0016 \quad (18)$$

This result provided conclusive evidence for the existence of gravitational waves, as predicted by Einstein's theory. In 1993 Hulse and Taylor were awarded the Nobel Prize in Physics "for the discovery of a new type of pulsar, a discovery that has opened up new possibilities for the study of gravitation".

2.2 WEBER BARS

Experiments to detect gravitational waves began with Weber and his resonant mass detectors. According to his first paper published in 1960 [16], his method made use of the fact that relative motion of mass points, or strains in a crystal, can be produced by second derivatives of the gravitational fields. The strains in a crystal may result in electric polarization in consequence of the piezoelectric effect. Measurement of voltages then enables certain components of the Riemann tensor to be determined. If the Riemann tensor has Fourier components in the vicinity of elastic normal modes of quadrupole symmetry, these modes may be observed to have greater than thermal energy. In other words, strains in space due to an incident gravitational wave excite the resonant frequency of the bar, and these vibrations can be amplified to detectable levels.

Weber initially instrumented two aluminum cylinders to record the Fourier transform of the Riemann tensor in the vicinity of the frequency $f = 1.5$ kHz, located in different places.

Since those years, many attempts have been done exploiting Weber's technique. An international network of cryogenic resonant detectors, IGEC2, has been taking data until few years ago. The detectors taking part in the IGEC2 network were: ALLEGRO, located at Louisiana State University (Louisiana, U.S.A.) and operated by the local ALLEGRO group, AURIGA, located in the Legnaro National Laboratories (Padova, Italy) of INFN and operated by the AURIGA Collaboration, EXPLORER, located at CERN (Geneva, Switzerland) and NAUTILUS, located in the Frascati National Laboratories (Frascati, Italy) of INFN. Both latter detectors were operated by the ROG Collaboration.

Their sensitivity curve is reported in figure 3. The latest results have been presented in 2010 [17], covering 515 days of data from 2005 to 2007. They found no quadruple coincidence and 20 events coincident in the three detectors, well within the expected occurrence of accidentals (~ 22). Since the False Alarm Rate (FAR) was $\sim 14 \text{ yr}^{-1}$, this analysis cannot identify single GW candidates with a reasonable significance.

In 2017, all the Weber bars have been decommissioned.




Figure 3: Typical strain noise spectral densities (single-sided) of IGEC2 detectors [17]. All detectors are sensitive in a region around 900 Hz, and the minimum level of noise is comparable in all spectra.

2.3 INTERFEROMETRIC DETECTORS

Interferometric detectors were first suggested in the 1962 by Gertsenshtein and Pustovoi: in the abstract of their first article, they stated that “*the sensitivity of the electromechanical experiments for detecting gravitational waves by means of piezocrystals is ten orders of magnitude worse than that estimated by Weber; in the low frequency range it should be possible to detect gravitational waves by the shift of the bands in an optical interferometer*” [18]. Also Weber mentioned the idea of the interferometric detection in an unpublished laboratory notebook.

Rainer Weiss first described in detail a practical solution with an analysis of realistic limitations to the technique in 1972 [19]. This led to proposals for long-baseline broadband laser interferometers with the potential for significantly increased sensitivity. In 1989 both the Virgo project by Adalberto Giazotto and Alain Brillet [20], and the LIGO project [21] were described in detailed technical reports.

After almost 30 years, interferometric detectors are still the most interesting devices ever conceived to detect gravitational waves. The most important detectors built so far are the almost identical two LIGO detectors, one in Hanford, WA (usually referred as H1) and the other in Livingston, LA (referred as L1) and the Virgo detector in Cascina, Italy (V1). Minor experiments, less sensitive mostly because of their shorter arms, are GEO 600 (Germany), and TAMA 300 (Japan).




Figure 4: Amplitude spectral density of strain sensitivity of the Advanced LIGO–Advanced Virgo network, estimated around the time of GW₁₇₀₈₁₄ [22].

During several years, LIGO and Virgo have been continuously upgraded in order to improve their sensitivity. The latest major upgrades, named Advanced LIGO and Advanced Virgo, have been completed respectively in 2015 and 2017. The latest sensitivity curve, measured in August 2017, is shown in figure 4.

2.4 GW150914: THE FIRST DETECTION

On September 14, 2015 at 09:50:45 UTC the two LIGO detectors simultaneously observed the transient gravitational-wave signal, named GW150914 [1]. At that time they were just beginning the first observation run, named O1, after the Advanced LIGO upgrade.

The signal is shown in figure 5. The initial detection was made by low-latency searches for generic gravitational-wave transients, *cWB*, and was reported within three minutes of data acquisition. Subsequently, it has been confirmed by matched-filter analyses that use relativistic models of compact binary waveforms.

GW150914's source corresponds to a stellar-mass BBH with individual source-frame masses $m_1 = 36_{-4}^{+5} M_\odot$ and $m_2 = 29_{-4}^{+4} M_\odot$, with a final black hole mass $M_f = 62_{-4}^{+4} M_\odot$. The total energy radiated as gravitational waves has been $E_{\text{rad}} = 3.0_{-0.4}^{+0.5} M_\odot c^2$, or equivalently $E_{\text{rad}} = 5.3_{-0.8}^{+0.9} \times 10^{47} \text{ J}$, with a peak luminosity of $100_{-20}^{+30} M_\odot c^2/\text{s}$ [23].

Only the LIGO detectors were observing at the time of GW150914, because the Virgo detector was still being upgraded. The signal was observed with a matched-filter Signal-to-Noise Ratio (SNR) of 24 and a FAR estimated to be less than 1 event per 203 000 yr, equivalent to a significance greater than 5.1σ .

GW150914 has been the first direct detection of gravitational waves and the first observation of a binary black hole merger. In 2017 this discovery was rewarded by the Royal Swedish Academy of Sciences, with the Nobel Prize in Physics assigned to Rainer Weiss, Kip Thorne and Barry Barish of the LIGO/Virgo Collaboration “for decisive contributions to the LIGO detector and the observation of gravitational waves”.

2.4.1 The other events

The observation run O1 ended on January 19, 2016. During O1, two high-mass BBH events were identified with high confidence ($> 5\sigma$): GW150914 and GW151226 [24]. A third signal, LVT151012, was also identified with 1.7σ confidence. Although LVT151012 is not significant enough to claim an unambiguous detection, it is more likely to have resulted from a gravitational-wave signal than from an instrumental or environmental noise transient [25].

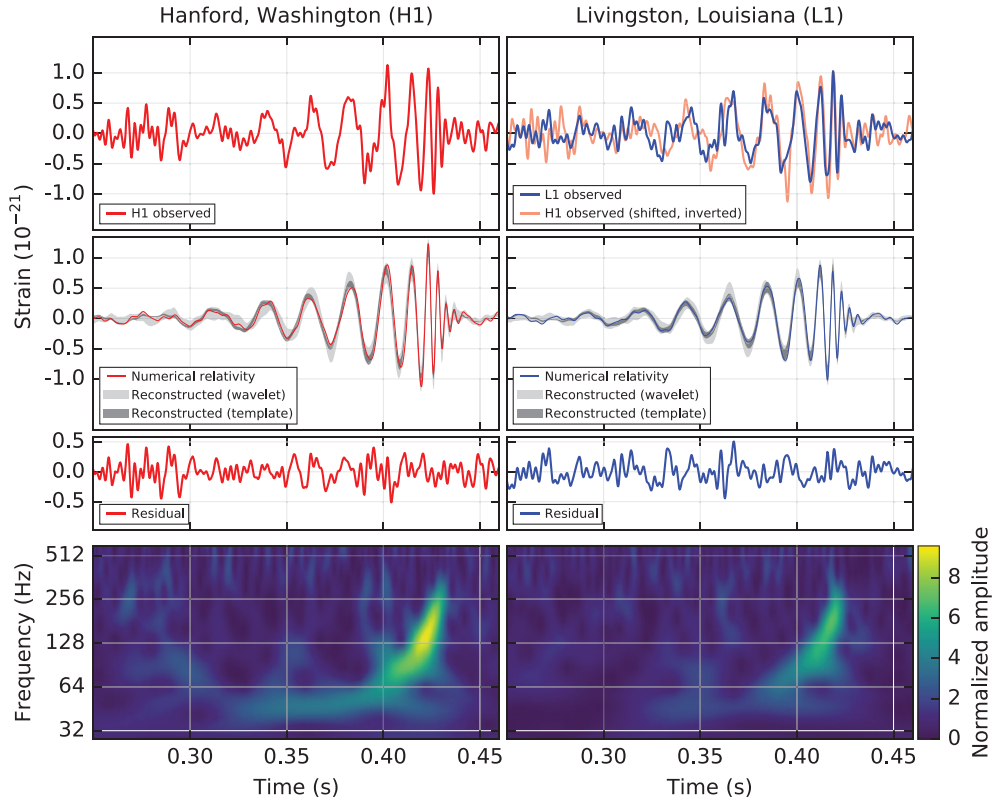


Figure 5: The gravitational-wave event GW150914 observed by the H1 (left column panels) and L1 (right column panels) detectors [1]. Times are shown relative to September 14, 2015 at 09:50:45 UTC. Top row, left: H1 strain. Top row, right: L1 strain. Second row: Gravitational-wave strain projected onto each detector in the 35 – 350 Hz band. Solid lines show a numerical relativity waveform for a system with parameters consistent with those recovered from GW150914. Shaded areas show 90% credible regions for two independent waveform reconstructions. One (dark gray) models the signal using binary black hole template waveforms. The other (light gray) does not use an astrophysical model, but instead calculates the strain signal as a linear combination of sine-Gaussian wavelets. Third row: Residuals after subtracting the filtered numerical relativity waveform from the filtered detector time series. Bottom row: A time-frequency representation of the strain data, showing the signal frequency increasing over time.

Advanced LIGO's second observing run began on November 30, 2016. On January 4, 2017, a gravitational wave signal was detected with high statistical significance, named GW170104 [26].

2.5 GW170814: THE FIRST VIRGO DETECTION

On August 1, 2017 the upgraded Virgo detector joined LIGO for a month of joint observation, until August 25. On August 14, 2017, GWs from the coalescence of two black holes with masses of $m_1 = 30.5^{+5.7}_{-3.0} M_\odot$ and $m_2 = 25.3^{+2.8}_{-4.2} M_\odot$, were observed in all three detectors, as reported in figure 6. The signal was first observed at the LIGO Livingston detector at 10:30:43 UTC, and at the LIGO Hanford and Virgo detectors with a delay of ~ 8 ms and ~ 14 ms, respectively [22].

GW170814 was first identified with high confidence ~ 30 s after its arrival by two independent low-latency matched-filter pipelines, *GstLAL* and *PyCBC*. In a off-line analysis, they reported a FAR respectively of 1 in 140 000 yr and 1 in 27 000 yr, using only the LIGO data. The difference in significance is due to the different techniques used to rank candidate events and measure the noise background in these searches. The significance of GW170814 was confirmed on the full network of three detectors by *cWB*. This more generic search reports a false-alarm rate < 1 in 5900 yr. By comparison, when we limit this analysis to the two LIGO detectors only, the false-alarm rate is approximately 1 in 300 yr; the use of the data from Virgo improves the significance by more than an order of magnitude.

The biggest Virgo contribution to the analysis of the event is in the sky localization: the addition of a third detector to the network allows to reduce dramatically the uncertainty on this localization. For the rapid localization using only H1 and L1, the 90% credible area on the sky is 1160 deg^2 and shrinks to 100 deg^2 when including Virgo data. The full parameter estimation further constrains the position to a 90% credible area of 60 deg^2 centered at the maximum *a posteriori* position of right ascension $RA = 03^{\text{h}}11^{\text{m}}$ and declination $dec = -44^\circ 57^{\text{m}}$ (J2000).

Incorporating Virgo data also reduces the luminosity distance uncertainty from $570^{+300}_{-230} \text{ Mpc}$ to $540^{+130}_{-210} \text{ Mpc}$, and consequently the three-dimensional credible volume also decreases by an order of magnitude, from $71 \times 10^6 \text{ Mpc}^3$, to $2.1 \times 10^6 \text{ Mpc}^3$. If we assume that the number density of galaxies is in the order of $\sim 0.01 \text{ Mpc}^{-3}$ [27], the number of galaxies in the fiducial volume is however pretty large, $\mathcal{O}(10^4)$.

A new interesting result from this event, that can be achieved only with at least three detectors, is the measurement of the polarization of the wave. One of the key predictions of GR is that metric perturbations possess two tensor degrees of freedom. These two are only a subset of the six independent modes allowed by generic metric theories of gravity, which may in principle predict any combination of tensor (spin-2), vector (spin-1), or scalar (spin-0) polarizations: the result is that our estimations strongly favor the purely

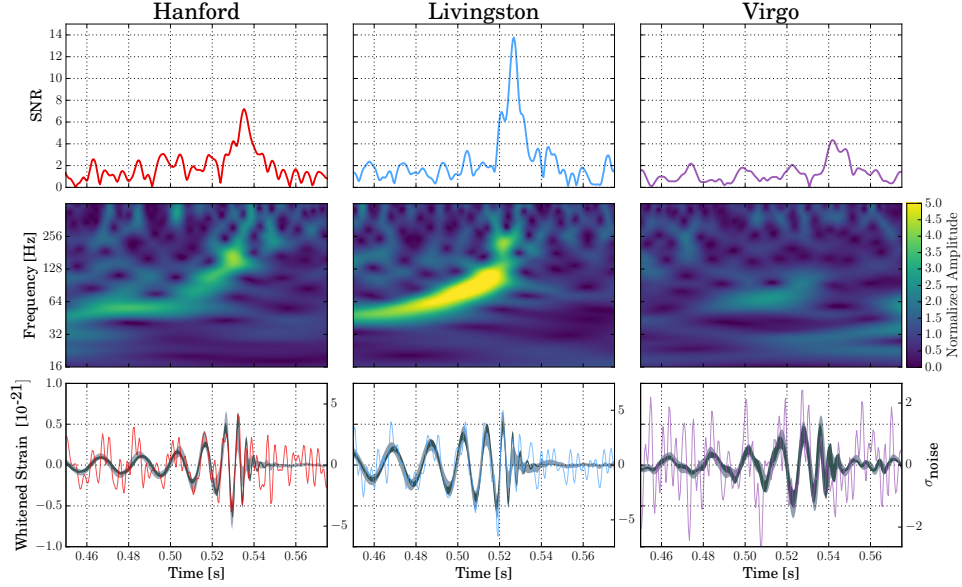


Figure 6: The event GW170814 observed by H1, L1, and V1 [22]. Times are shown from August 14, 2017, 10:30:43 UTC. Top row: SNR time series produced in low latency and used by the low-latency localization pipeline on August 14, 2017. The time series were produced by time shifting the best-match template from the on-line analysis and computing the integrated SNR at each point in time. The single-detector SNRs in H1, L1, and V1 are 7.3, 13.7, and 4.4, respectively. Second row: Time-frequency representation of the strain data around the time of GW170814. Bottom row: Time-domain detector data (in color), and 90% confidence intervals for waveforms reconstructed from a morphology-independent wavelet analysis (light gray) and BBH models, whitened by each instrument's noise amplitude spectral density between 20 Hz and 1024 Hz. For this figure the data were also low passed with a 380 Hz cutoff to eliminate out-of-band noise. *Whitening* is the process to get a flat PSD, and it corresponds to divide the data by their PSD in the frequency domain. The whitening emphasizes different frequency bands for each detector, which is why the reconstructed waveform amplitude evolution looks different in each column. The left ordinate axes are normalized such that the physical strain of the wave form is accurate at 130 Hz. The right ordinate axes are in units of whitened strain, divided by the square root of the effective bandwidth (360 Hz), resulting in units of noise standard deviations.




Figure 7: Localization of GW₁₇₀₈₁₄ [22]. The rapid localization using data from the two LIGO sites is shown in yellow, with the inclusion of data from Virgo shown in green. The full Bayesian localization is shown in purple. The contours represent the 90% credible regions. The left panel is an orthographic projection and the inset in the center is a gnomonic projection; both are in equatorial coordinates. The inset on the right shows the posterior probability distribution for the luminosity distance, marginalized over the whole sky.

tensor polarization against purely vector and purely scalar, as expected from GR [8].

2.6 GW₁₇₀₈₁₇: THE FIRST BNS DETECTION

Just three days after the first triple-coincidence GW₁₇₀₈₁₄, on August 17, 2017 the LIGO-Virgo detector network observed a gravitational-wave signal, named GW₁₇₀₈₁₇, from the inspiral of two low-mass compact objects consistent with a Binary Neutron Star (BNS) merger. This discovery comes four decades after Hulse and Taylor discovered the first neutron star binary. GW₁₇₀₈₁₇ was identified by matched filtering low-latency pipelines. This gravitational-wave signal was the loudest observed, with a combined SNR of 32.4 with values 18.8, 26.4, and 2.0 in the H₁, L₁ and V₁ respectively, and a FAR estimate of less than one per 8.0×10^4 yr.

After ~ 100 s (calculated starting from 24 Hz) in the sensitive band, the inspiral signal ended at 12:41:04.4 UTC. In addition, a Short Gamma Ray Burst (SGRB) named GRB_{170817A} was observed (1.74 ± 0.05) s after the coalescence time by space-based Fermi GBM detector and INTEGRAL detector [29], as reported in figure 8.

A first key result comes from this delay. If we conservatively assume that the peak of the GW signal and the first photons were emitted simultaneously, attributing the entire lag to faster travel by the GW signal, this time

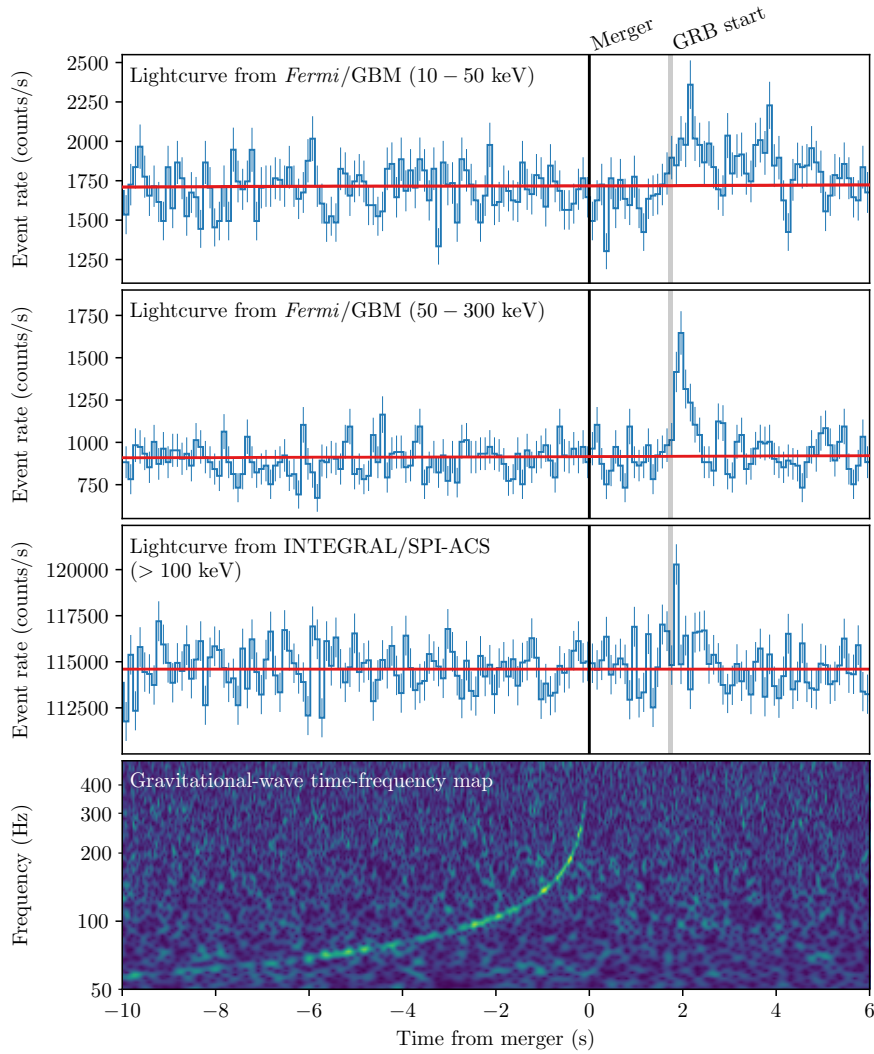


Figure 8: The multi-messenger detection of GW₁₇₀₈₁₇ and GRB_{170817A} done by Fermi GBM (first two rows), INTEGRAL (third row) and LIGO/Virgo (last row) [28]. The GW signal has been obtained by coherently combining H₁ and L₁ data.

difference provides an upper bound on $\Delta c = c_{\text{GW}} - c$, the difference between the speed of gravity c_{GW} and the speed of light c . To obtain a lower bound on Δc , one can assume that the two signals were emitted at times differing by more than $(1.74 \pm 0.05) \text{ s}$ with the faster electromagnetic signal making up some of the difference. As a conservative bound¹ we assume the SGRB signal was emitted 10 s after the GW signal. The intergalactic medium dispersion has negligible impact on the gamma-ray photon speed, with an expected propagation delay many orders of magnitude smaller than our errors on c_{GW} . The resulting constraint

$$-3 \times 10^{-15} < \frac{\Delta c}{c} < 7 \times 10^{-16}$$

is the first direct measure ever. Future joint GW–SGRB detection should resolve this problem, as the lower bound depends on the distance, while the upper bound is independent on it [28].

The combination of data from the LIGO and Virgo detectors allowed a precise sky position localization to an area of 28 deg^2 (90% confidence level). It is important to remark that the Virgo contribution allowed to reduce the localization uncertainty from 190 deg^2 using only the two LIGO detectors to 31 deg^2 , more 6 times less. Even if the signal in the Virgo detector was not sufficient to generate a trigger, it was enough to compute the delays between time of arrival in the three detectors allowing the triangulation.

This measurement enabled an electromagnetic follow-up campaign that identified a counterpart near the galaxy NGC 4993, consistent with the localization and distance inferred from gravitational-wave data [31]. The localization of the event reconstructed by LIGO–Virgo and the position of the host galaxy NGC 4993 is reported in figure 9.

It allowed also a completely independent measurement of the Hubble constant,

$$H_0 = 70.0_{-8.0}^{+12.0} \text{ km s}^{-1} \text{ Mpc}^{-1}$$

still with large uncertainty and consistent with the current existing estimations [32].

2.7 THE FUTURE

LIGO and Virgo concluded O2 on until August 25, 2017. The third joint observing run O3 is going to start within the end of 2018, leaving to commissioners about a year to perform important upgrades on the detectors and improve the sensitivity. The next years are expected to be exciting, with almost an event per day when the sensitivities will approach the design ones.

¹ The onset of gamma-ray emission from a BNS merger progenitor is predicted to be within a few seconds after the merger, given that the central engine is expected to form within a few seconds and that the jet propagation delays are at most of the order of the SGRB duration. [28]

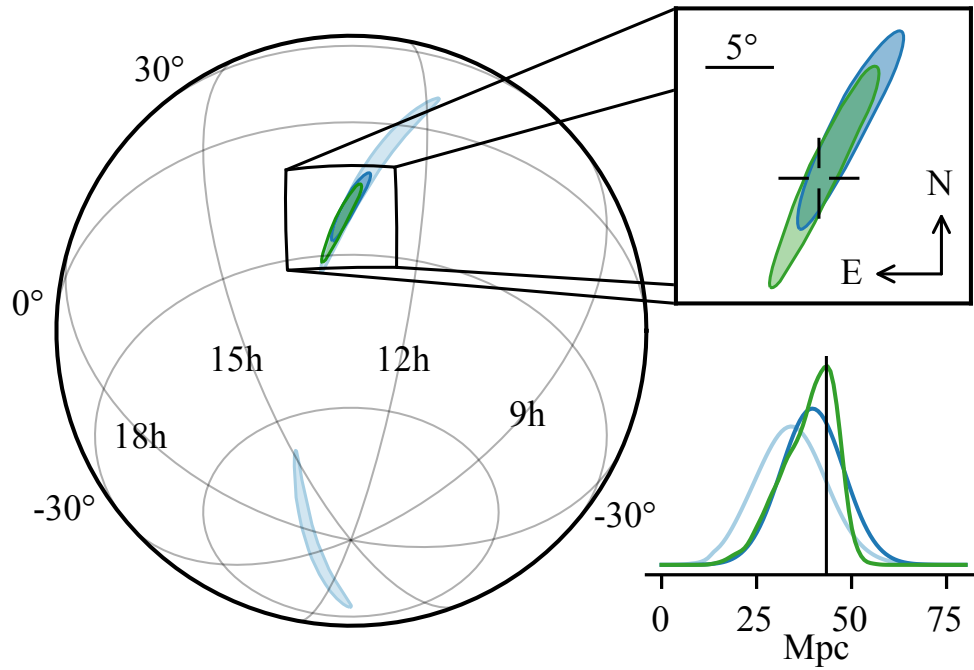


Figure 9: Sky location reconstructed for GW₁₇₀₈₁₇ by a rapid localization algorithm from a H₁-L₁ (190 deg², light blue contours) and H₁-L₁-V₁ (31 deg², dark blue contours) analysis [30]. A higher latency H₁-L₁-V₁ analysis improved the localization (28 deg², green contours). In the top-right inset panel, the reticle marks the position of the apparent host galaxy NGC 4993. The bottom-right panel shows the a posteriori luminosity distance distribution from the three gravitational-wave localization analyses. The distance of NGC 4993 is shown with a vertical line.

In the meanwhile, the scientific community is continuing to discuss about the next generation detectors. The main improvements to the sensitivity will come from cryogenics and a location underground, as well as squeezing techniques.

The KAGRA interferometric detector will join the search for gravitational waves in the next years. It will take advantage of an underground location at the Kamioka Observatory (Japan), and most of all of cryogenic mirrors.

With the BNS event from NGC 4993, the multi-messenger astronomy has just begun: in the next future the joint observation of electromagnetic radiation, gravitational waves, neutrinos, and cosmic rays from astrophysical events is expected to produce new spectacular results.

Another interesting project is also LISA: it will be the first dedicated space-based gravitational wave detector, using laser interferometry to monitor the fluctuations in the relative distances between three spacecrafts, arranged in an equilateral triangle with 2.5×10^9 m arms.

3

THE VIRGO DETECTOR

Virgo is a Michelson gravitational wave interferometer with a 3 km Fabry-Perot cavity in each arm, located in Cascina, Italy. The experiment has been developed and built by a French-Italian collaboration founded in 1989 by Adalberto Giazotto and Alain Brillet.

The initial Virgo detector recorded scientific data from 2007 to 2011 during four science runs. Below ~ 100 Hz, Virgo was by far the most sensitive detector on the Earth, with the sensitivity curve close to the design [33]. Virgo has undergone a major upgrade since 2011, within a project named Advanced Virgo (AdV), and joined LIGO for a month of joint observation in August 2017.

In 2017, the collaboration counts about 350 members from several French, Italian, Dutch, Polish, Hungarian and Spanish institutions.

3.1 OPERATING PRINCIPLE

The operating principle of an interferometric detector for gravitational waves is very simple: we can determine the distance between two test masses (nothing more than mirrors) by measuring the round trip travel time of light beams sent over large distances, and thus it is natural to aim at Michelson interferometer. The key difference with the 1887 interferometer is that here we do not connect the mirrors in a single rigid structure, but each mass is left in free fall, so that it responds in an independent way to gravitational effects.

We can calculate the time that it takes to a ray to travel in each arm when it is crossed by a gravitational wave. In equations 16 and 17 we have already seen the effect on the distance between two points; in a Michelson interferometer the arms are orthogonal and their lengths L equal each other ($L = x_0 = y_0$), so that the relative variation is

$$\begin{aligned}\Delta L(t) &= x(t) - y(t) \\ &= L \cdot h(t)\end{aligned}\tag{19}$$

This corresponds to a difference in time of arrival

$$\begin{aligned}\Delta\tau(t) &= 2\frac{\Delta L(t)}{c} = 2\frac{L \cdot h(t)}{c} \\ &= \tau_0 \cdot h(t)\end{aligned}\tag{20}$$

where the additional factor 2 takes into account for the round trip in the interferometer, and $\tau_0 = 2L/c$ is the travel time in absence of gravitational

waves¹. Corrections due to the effect of the gravitational wave itself are negligible². We can also express 20 as a phase shift:

$$\Delta\varphi(t) = \frac{2\pi c}{\lambda_L} \tau_0 \cdot h(t) = \frac{4\pi L}{\lambda_L} \cdot h(t) \quad (21)$$

where λ_L is the wavelength of the light used in the interferometer. It is clear that the effect is directly proportional to h : this immediately says that the longer is the optical path in the apparatus, the larger will be the phase shift due to the gravitational wave.

This scaling law doesn't hold for arbitrarily long arms because, on the other hand, large values of L reduce the bandwidth of our experiment [34]. Indeed, for $\lambda = L$, where λ is the wavelength of the gravitational wave, the effects of the gravitational wave on the detector are averaged on a entire wavelength, and no signal can be detected. We can get a graph of the sensitivity of a Michelson interferometer, expressed in $\Delta\varphi$ per h unit as function of f , where $f = c/\lambda$ is the frequency of the gravitational wave, that is shown in 10 for several values of L : looking at the behavior at high frequencies, given a certain f there is a L beyond which there are no further gains.

Of course, the feasibility of a earth-based interferometer with L larger then few kilometers is almost null, at least for the cost. Luckily an optical arrangement could come to the aid of us: implementing two Fabry-Pérot cavities in both Michelson interferometer arms, we achieve the same performance of a longer interferometer. In brief, photons are trapped in the cavities for an average time [34]

$$\tau_s = \frac{2L}{c} \frac{\mathcal{F}}{2\pi} = \tau_0 \cdot \frac{\mathcal{F}}{2\pi} \quad (22)$$

or, in other words, travel averagely $\mathcal{F}/2\pi$ times through the cavity before to go out. The quantity \mathcal{F} is a quality index of the cavity, called *finesse*.

For example, the sensitivity of a 3 km long interferometer with $\mathcal{F} = 150$ Fabry-Pérot cavities is roughly equivalent to a 400 km standard interferometer, as shown in 10.

3.2 IMPLEMENTATION

The Advanced Virgo is a Dual-Recycled Michelson Interferometer with Fabry-Pérot arm cavities. Its optical layout with the naming convention of the

-
- ¹ Note that the analytic solution of the equation 20 should be obtained by calculating the interval between the neighboring space-time events linked by the a light beam corresponding to the laser, with the usual relation $ds^2 = 0$, where $ds^2 = g_{\mu\nu} dx^\mu dx^\nu = (\eta_{\mu\nu} + h_{\mu\nu}) dx^\mu dx^\nu$, and then obtaining the time by an opportune integration. This accounts the fact that the light beam wavelength itself is strained by the gravitational wave passing through the interferometer. Eventually, we would find the same result in 20, if the approximation $L \ll \lambda$, where λ is the wavelength of the gravitational wave, applies.
 - ² It is relative to the fact that the gravitational waves period is much greater than the average storage time of the light into the detector, so that the metric tensor can be instantaneously assumed constant on the detector.

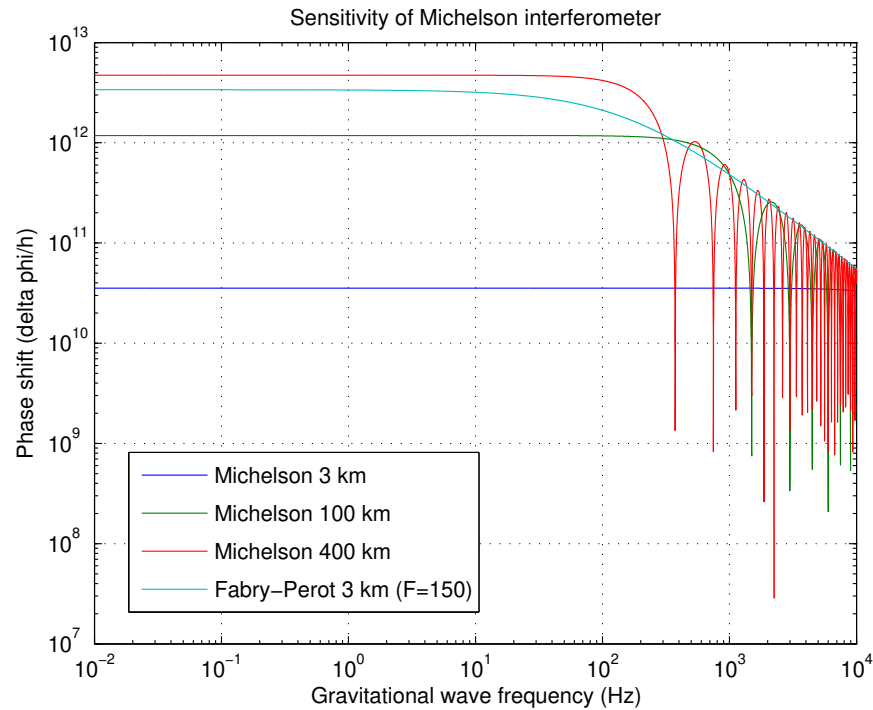


Figure 10: Phase shift per h unit $\Delta\phi/h$ in a L -long Michelson interferometer as function of the frequency of a gravitational wave f [35]. The light has wavelength $\lambda_L = 1.064 \mu\text{m}$. The peaks correspond to $\lambda = L/n$ with $n \in \mathbb{N}$, when an certain number of wavelengths fits into the interferometer, canceling the effects.

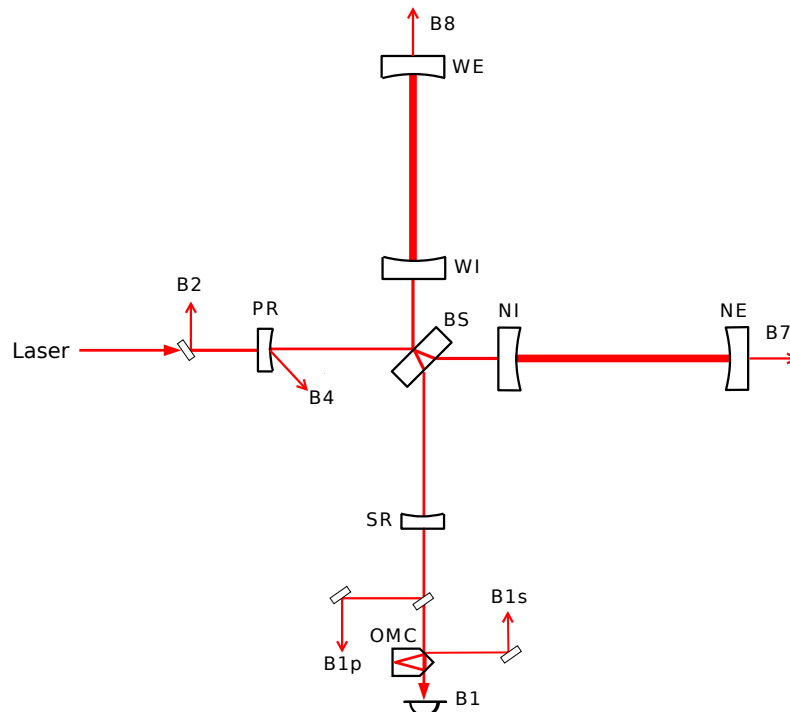


Figure 11: Advanced Virgo optical layout with mirrors and photodiodes naming convention [36].

various components and beams is reported in figure 11. All mirrors are suspended to multistage passive isolation systems designed to filter the transmission of ground micro-seismic vibration to the test masses.

The input laser beam is provided by a Nd:YAG laser amplifier with a wavelength of $\lambda_L = 1064 \mu\text{m}$. Before entering the main part of the interferometer, the beam passes through an input mode cleaner to filter its transverse mode and reduce its jitter.

In first approximation, the interferometer is operated at the *dark fringe* mode: the destructive interference guarantees that no light reaches the photodiode B1 at the *asymmetric port*³, while in the presence of a gravitational wave signal the light is phase-modulated by the gravitational wave signal itself, and signal will form. A Power Recycling (PR) mirror serves the purpose of reflecting back into the interferometer the light that would otherwise travel back to the *symmetric port* and, therefore, be lost. A Signal Recycling (SR) mirror can be placed in front of the *asymmetric port* to recycle the differential signal coming from the arms, in order to increase the storage time of the signal inside the detector and, as a consequence, increase the sensitivity.

In order to maintain the correct resonance condition inside the interferometer and the best possible detector sensitivity, distances between mirrors must be controlled by active feedback systems with typical accuracy better than 10^{-12} m. The longitudinal control system is also called the locking system and it deals with four main degrees of freedom: the differential change in length of the two Fabry-Pérot cavities (DARM) which corresponds to the main degree of freedom sensitive to the effect of a gravitational wave signal; the mean change in length of the two cavities (CARM) which is also equivalent to a change in the main laser frequency; the length of the power recycling cavity (PRCL) composed of the power recycling mirror and the two arm input mirrors; the differential change in length of the two short Michelson arms (MICH) between beam splitter and the two input mirrors. [37]

3.3 SENSITIVITY AND NOISE SOURCES

Due to local gravity fluctuations ground-based interferometers are sensitive to gravitational waves with frequencies higher than a few Hz. On the high side of the frequency band the gravitational wave signal strongly decreases due to source dynamics. The Virgo interferometer has been designed to cover a frequency range starting from 10 Hz up to a few kHz. At the time of the first proposals, expected sources in this range were coalescing binary systems like neutron star or black hole binaries, stellar collapses, rotating neutron stars, and cosmological background radiation.

The Virgo sensitivity is limited by noise sources that can be grouped into different categories. The position of the test masses can fluctuate due to local

³ The photodiode B1 location is usually called *asymmetric port*, while the location near the input laser (left side of the figure 11) is called *symmetric port*.




Figure 12: Reference AdV sensitivity and expected noise contributions [36].

perturbations, such as residual seismic noise, local gravity fluctuations and thermal motion. Other noise sources affect the detected signal without a real mirror displacement, the main source being photon shot noise. Now we briefly describe which are the noises that are limiting the detector sensitivity. Their contributions in the Advanced Virgo layout are represented in figure 12.

The sensitivity reached by Virgo and LIGO during O2 has already been shown in figure 4.

3.3.1 Seismic and gravity gradient noises

The sensitivity of interferometric antennas for gravitational waves is limited at low frequency by seismic noise. This term indicates the stochastic movements of the soil, due to a multitude of causes that include nearby anthropogenic activities (such as traffic or heavy machinery), as well as natural phenomena like wind, sea waves, earthquakes and the position of the moon. Its PSD at ground level is not flat, and strongly depends on environmental conditions. An empiric estimation (also known as *standard seismic noise*) that applies to the Virgo site is

$$S_{x_0}(f) \sim \left| 10^{-7} \left(\frac{1 \text{ Hz}}{f} \right)^2 \frac{\text{m}}{\sqrt{\text{Hz}}} \right|^2 \quad (23)$$

and it corresponds to a white noise in ground acceleration.

Seismic noise enters in the Virgo noise budget through both the residual vibration transmitted by the mechanical structures to the mirrors and the direct coupling due to the Newtonian attraction force of the moving soil to the suspended test masses. This second interaction is known as gravity gradient noise, or Newtonian noise. While the direct transmission of the seismic noise to the mirrors can be kept under control using adequate mechanical suspensions, the second cannot be suppressed.

These noises do not limit the detector sensitivity in their detection bandwidth. However future GW detectors, the so-called 3rd generation interferometers, are expected to extend their detection band to lower frequencies ($f < 10$ Hz), where the gravity gradient noise will represent one of their main limitations. For this reason several solutions to reduce the Newtonian noise, like the possibility to build future interferometers in deep underground sites and the development of active noise subtraction techniques, are being studied.

3.3.2 Thermal noise

The interferometer mirrors are in radiative thermal equilibrium with the vacuum chambers that are at room temperature. The energy exchange generates the Brownian motion of the particles of mirror glass, coating and suspension system and consequently induces a fluctuation in the measured cavity length.

The main limit to the Virgo sensitivity between 5 and 500 Hz is the thermal noise, due to the pendulum thermal fluctuation (dominant in the 5 to 50 Hz range) and the mirror internal mode vibration (dominant in the 50 to 500 Hz range).

The pendulum thermal fluctuations are so called because, as we will see, harmonic oscillators are the solution used to suspend free fall test masses in an GW detector. In fact there are two forms of dissipation in a pendulum, one due to viscous friction with gas, the other caused by the elasticity of the wire. While the first contribution can be eliminated using vacuum, the second can only be reduced and depends on the geometry and the material of the wire: monolithic fused silica has been proven to be the best material for the last stage suspension [38].

For what concerns the mirror internal mode vibration, there are three different dissipation processes involved: the bulk Brownian noise, the thermoelastic noise and the mirror coating noise. The first contribution is caused by bistable states distributed homogeneously in the substrate that convert the oscillating energy of the beam into heat, perturbing the mirror surface position. The thermoelastic noise is produced instead by the oscillating temperature distribution generated by the squeeze and stretch of the mirror substrate. Finally the last dissipation process is caused by imperfections in the mirror coating material. In fact, in order to reach the high reflectivity required for interferometer mirrors, several layers of dielectric materials are deposited on the substrate. Although the amount of material used for coat-

ing is small compared to the total mass of the mirror, its mechanical losses can constitute a large contribution to the thermal noise [35].

All these contributions will be reduced by cryogenic technologies in the next generation detectors.

3.3.3 Quantum noise

A fundamental limit to the sensitivity of a Michelson interferometer comes from the quantum nature of light, thus from the Heisenberg's uncertainty principle. It enters the detector noise budget through two fundamental mechanisms: shot noise, arising from statistical fluctuations in the arrival time of photons at the interferometer output, and radiation pressure noise, which is the recoil of the mirrors due to the radiation pressure arising from quantum fluctuations in the photon flux.

Quantum noise is the sum of the quantum radiation pressure noise and shot noise. At high frequencies, the sensitivity of a GW detector is limited by shot noise, as it generates a noise $\propto f$ at frequency $f > 1/(4\pi\tau_s)$ [35]. It can be reduced by increasing the laser power. On the other hand, the quantum radiation pressure dominates at lower frequencies, and increases if we increase the laser power. It can be adequately reduced by larger and heavier test masses.

To decrease the shot noise while increasing the radiation-pressure noise, or vice-versa, the idea of the *squeezed-state* technique has been proposed in 1981 [39]. The LIGO collaboration demonstrated for the first time in 2011 that the injection of a squeezed vacuum state into the interferometer can reduce the shot noise due to the quantum nature of light. This result was achieved with the GEO600 interferometer, and then replicated in 2013 with the Livingston interferometer [40]. In 2013, also the Virgo collaboration started developing the squeezed-state technique [41].

3.4 THE ADVANCED VIRGO PROJECT

The aim of Advanced Virgo has been to achieve a sensitivity that is improved by one order of magnitude with respect to the original Virgo [36], which corresponds to an increase of the detection rate by three orders of magnitude. Therefore, most of the detector subsystems have been improved to be compatible with the design sensitivity, to tune the effects of the noises described in the previous section accordingly to the technological and industrial improvements of the last years. The most important upgrades are described below.

INTERFEROMETER OPTICAL CONFIGURATION The Signal Recycling cavity has been added, to make dual recycled interferometer. The tuning of the cavity allows to change the shape of the sensitivity curve and to optimize the

detector most sensitive frequency band for different astrophysical sources. Furthermore, the cavity finesse is higher with respect to Virgo: a reference value of $\mathcal{F} = 443$ has been chosen.

LASER Improving the sensitivity at high frequency requires higher power, as more power in the arm cavities reduces the shot noise. The AdV reference sensitivity is computed assuming 125 W entering the interferometer (after the input mode cleaner). Therefore, considering the losses of the injection system, the laser must provide a power of at least 175 W. However, this upgrade has been postponed after O₂, and will go through a first step of 80 W during O₃.

MIRRORS Given the much larger optical power in the cavities, due to the new laser and the higher finesse, radiation pressure noise becomes a limit in the low frequency range and heavier test masses are needed to reduce it. The AdV test masses have the same diameter as the Virgo ones (35 cm) but are twice as thick (20 cm) and heavy (42 kg).

THERMAL COMPENSATION A thermal compensation system is designed to manage beam aberrations induced by thermal effects as well as mirror “cold” defects (such as refraction index inhomogeneities).

DETECTION AdV uses a DC detection scheme, which improves the quantum noise and reduces the impact of some technical noises. Furthermore, the main photodiodes are placed on the suspended optical bench, in vacuum.

SUSPENSIONS The Virgo suspensions provided a seismic isolation already compliant with the AdV requirements. However, some upgrades were necessary to improve the robustness in bad weather conditions:

- new monolithic inverted pendulum legs have been mounted: the higher resonance frequency allows to increase the bandwidth of the inertial damping servos;
- piezo actuators have been installed in the inverted pendulum feet, allowing to perform a control of the tilt.

Moreover, the control electronics has been completely redesigned. The new electronics will be described in the next chapter.

PAYLOADS The payloads, that are the last stage of the suspensions to which the mirrors are attached, have been modified to be compliant with the thicker and heavier mirrors and to suspend the new components. A new layout has been developed, getting rid of the recoil mass (see 17).

VACUUM The Virgo vacuum have been improved by a factor of about 100 in order to be compliant with the AdV sensitivity. The value is related to

the vacuum in the arms, where is most of the optical path of the laser and where the scattering happens. The vacuum in the towers is less important, since the viscous damping of the mirror is not dominant.

Part II

The Software Supervisor

4

SEISMIC ISOLATION SYSTEM IN ADVANCED VIRGO

As seen in the previous chapter, below a few tens of hertz interferometric detection of gravitational waves is masked by seismic vibrations of the optical components. On the other hand many astrophysical sources such as coalescing binaries are expected to emit mainly low-frequency gravitational waves, usually at $f < 100$ Hz. For these reasons it is very important to lower the frequency detection threshold as much as possible.

In order to isolate the mirrors of the Virgo interferometer, a sophisticated suspension system, called *Superattenuator*, has been developed. Its working principle is based on a multistage pendulum acting on seismic vibrations as a chain of second order mechanical low-pass filters.

At the beginning of the development of the interferometric GW detectors, different technical solutions had been proposed by other collaborations. In figure 13 we show the performances of the various suspensions developed for the past, existing and future detectors. The Virgo Superattenuators were by far the most performing tools developed, and they still hold this record: from the sensitivity curves of LIGO and Virgo in the 2009-2010 science runs in figure 14 it is easy to see how the Virgo suspensions enhanced the detector bandwidth at low frequencies.

The attenuation performance of the Superattenuators was considered to be compliant with the AdV requirements since the time of project design. Only relatively minor upgrades have been done on the system, with the major contribution brought by a completely new digital control system.

In this chapter we describe the Superattenuators, including their working principles and the control electronics.

4.1 THE SUPERATTENUATORS

Superattenuators, often referred to simply as *suspensions*, and represented in a rendering in figure 15, are the mechanical support of each optical component of the interferometer. They have been designed to suppress the seismic noise transmission to the last stage, the mirror, by more than ten orders of magnitude starting from about 4 Hz, well below the thermal noise. In the low-frequency region, indeed, thermal noise is the dominant noise mechanism acting on the optical components.

4.1.1 Horizontal Attenuation

The Superattenuators are composed by:




Figure 13: Vibration isolation for the initial LIGO, Virgo, TAMA, GEO600, Adv. LIGO, KAGRA and the future Einstein Telescope [42].




Figure 14: Best strain noise spectra from the LIGO and Virgo detectors during the 2009–2010 science runs [43].

- a pre-isolation stage, a ~ 10 m high three-leg inverted pendulum attached to the ground;
- six standard filters in cascade, (16) where the higher standard filter, named F_0 , is attached to the top of the inverted pendulum, and the others are named, from top to bottom, F_1, F_2, F_3, F_4 and F_7 ;
- the mirror suspension, attached at F_7 , and composed by a marionette and the mirror itself (17).

A chain of mechanical filters in cascade acts as a multi-stage pendulum. Each stage behaves as a simple pendulum, and acts a second order low-pass filter for the horizontal motion of the suspension point transmitted to the mass suspended at the last stage. Overall, a n -stage pendulum acts as a $2n$ -order low pass filter at a frequency f higher than the frequencies of the chain normal modes ($f > f_0 > f_1 > \dots > f_n$). In particular, the ratio between the linear spectral density of the last mass horizontal displacement (the optical component) and the linear spectral density of the suspension point horizontal displacement (where the excitation is applied) decreases as $(f_0^2 \cdot f_1^2 \cdot \dots \cdot f_n^2) / f^{2n}$. With this system a very large attenuation of the seismic noise horizontal component can be obtained at frequencies above the highest pendulum resonance.




Figure 15: A rendering of Advanced Virgo Superattenuator.




Figure 16: Drawing of a mechanical seismic filter of the initial Virgo [44]. All the main components are indicated.

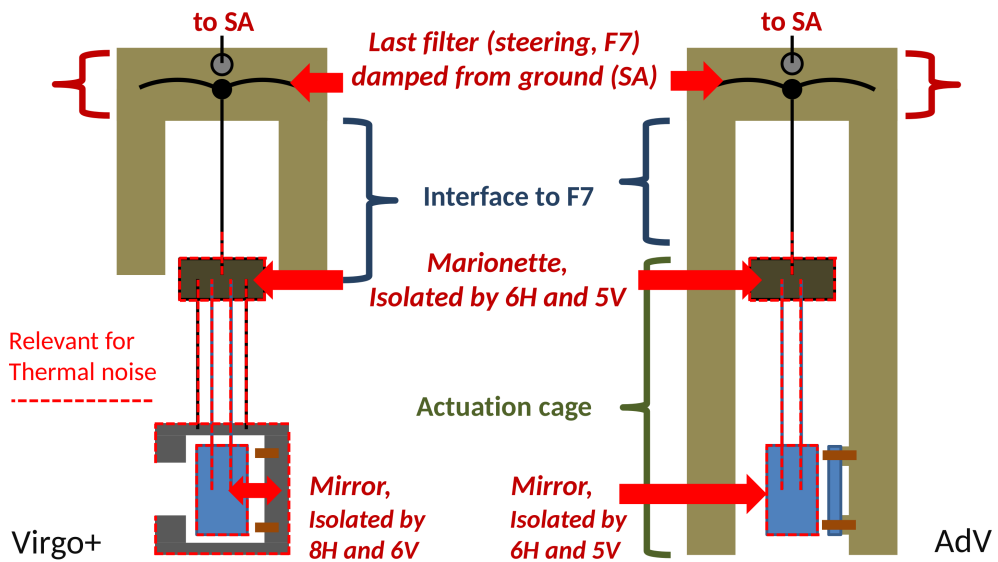


Figure 17: Different layouts of the payload in Virgo+ and Advanced Virgo. The main difference is the lack of the recoil mass in the new layout. [45]

Overall, the horizontal displacement is filtered by a 8-stage pendulum (F_0 , F_1 , F_2 , F_3 , F_4 , F_7 , marionette, mirror wires). The theoretical transfer function should be proportional to f^{-16} . However, due to other resonances of the mechanical structures and to not null mechanical impedances, the modulus of the transfer function has been measured to be proportional to f^{-15} up to ~ 10 Hz.

4.1.2 Vertical attenuation

The end mirrors, suspended 3 km away in our interferometer, are misaligned with respect to one another by about 0.3 mrad because of the different directions of the plumb line on the curved Earth surface. This means that they are not perfectly parallel to the plumb line¹: any vertical vibration will be partially transmitted to the interferometer horizontal axis, the laser beam direction, because of an unavoidable coupling among different degrees of freedom. Thus vertical motion will cause a phase change of the laser beam. For this reason, a vertical attenuation of seismic noise comparable to the horizontal one is fundamental keep low frequency performances. In this case, the passive attenuation is performed by the blades that suspend the mechanical filters to the upper one. In this case the mechanical filters are only six, because the F_0 lies at the top of the inverted pendulum and is not suspended by blades, and there are no blades also between the marionette and the mirror. A system of anti-springs in parallel to the blades allows to keep the vertical resonances below ~ 1 Hz.

Horizontal displacements of the mirror along the beam direction can also be induced by rotations of the pendulum chain around the vertical axis. To confine these rotational mode frequencies below the detection band each pendulum mass has to be replaced by a structure having a high momentum of inertia. In addition, the diameter of the suspension wire, connecting two consecutive stages, has to be small enough to reduce its restoring torque which opposes the rotation of the chain determining its rotational frequencies. An interconnection of the stages at small distance and as close as possible to their centers of mass, guarantees low-frequency tilt modes and negligible coupling effects on the horizontal displacement of the suspended mass. [44]

¹ Actually, with spherical mirrors the only requirement is that their center of curvature lies on the laser beam direction, and that the beam gaussian spot is contained within the mirror (preferably at its center). This is holds for both input and end mirrors. Another useful constraint is that the direction of the actuation along the beam axis is shared by the two mirrors, so that one can easily control the cavity length with both the mirror actuator systems, in a single degree of freedom. In practice, input mirrors is aligned with the plumb line, while the end mirrors are slightly tilted (the two rear silica fibers are shorter).

4.2 DIGITAL CONTROL SYSTEM

These mechanical suspensions introduce ~ 80 resonances in the system in the low frequency range, between 30 mHz and 1 Hz, that amplify the seismic noise.

A digital control system, named Suspension Control System (SCS), takes care of setting and maintaining the correct operation point for Superattenuators, and reducing the mirror motion due to the chain resonances and tidal strains. It uses 20 local sensors, plus 3 global sensors available when the full interferometer is locked. ~ 15 Digital Signal Processor (DSP)-based boards per suspension use these sensors to compute a real-time feedback action, that eventually drives coil-magnet actuators [46]. Each suspension has two main multi-variable control loops: Top Stage Control and Payload Local Control.

4.2.1 Top Stage Control

The Top Stage Control (often referred as “Inertial Damping”) is the digital control loop in charge of the reduction of suspended payload free motion. During normal operation of the Superattenuator, the Top Stage Control loop is always active.

Top stage displacement is monitored using accelerometers working in the DC-100 Hz frequency range. Together with those accelerometers, the actual position with respect to the surrounding supporting structure is monitored using Linear Variable Differential Transformer (LVDT) displacement sensors. On the Fo stage, there are 3 horizontal accelerometers placed at 120° , 2 vertical accelerometers, 3 horizontal and 1 vertical LVDTs. For the actuation, there are 3 horizontal and 2 vertical coil magnet actuators.

The digital control loop operates with a 10 kHz sampling frequency and a loop unity gain frequency of about 5 Hz^2 . Four degrees of freedom are currently controlled: two translations in the horizontal plane (x and z directions, where z is along the beam direction for that suspension and x is its orthogonal direction), the rotation along the vertical axis y (θ_y) and the vertical translation.

4.2.2 Payload Local Control

The Payload Control is another digital control loop used for the positioning of the suspended payload. It acts on the payload itself, while the Iner-

² In practice the control loop could be executed at few tens of hertz because of the low bandwidth of the control. On the other hand, currently 10 kHz is the frequency of the Virgo global control, and it is much easier to use the same frequency avoiding resamplings, since the computational power is not a problem with the current hardware. Additionally, there are some historical reasons: in the DSP-boards used in the older upgrades of Virgo it was impossible to handle a resampling of all the channels, as the operation is computationally expensive. Now this could be done, but in practice is not needed.

tial Damping acts on the top stage. It has two different operational modes. When the interferometer is unlocked, payload control makes use of measurements made with respect with a local reference frame with an system based on optical levers.

When the Virgo interferometer is in locked state, payload control makes use of set points distributed by the Global Control (GC) [47], a dedicated hardware and software in charge of the locking and the alignment loops. The GC is based on processing units named Real Time PC (RTPC) [36]. Actuation is performed by coil magnet actuators both on marionette and mirror.

Mirror position is corrected along three degrees of freedom: z , θ_x and θ_y : the first is used for the locking of the interferometer, while the others are used for the automatic alignment of the optical elements.

The low frequency components of the z error signals are sent to the actuators on the marionette ($< 4 - 5$ Hz) and on the Fo (< 10 mHz), where we can achieve bigger displacements.

4.3 SUSPENSION CONTROL SYSTEM HARDWARE

The SCS hardware is composed of several devices.

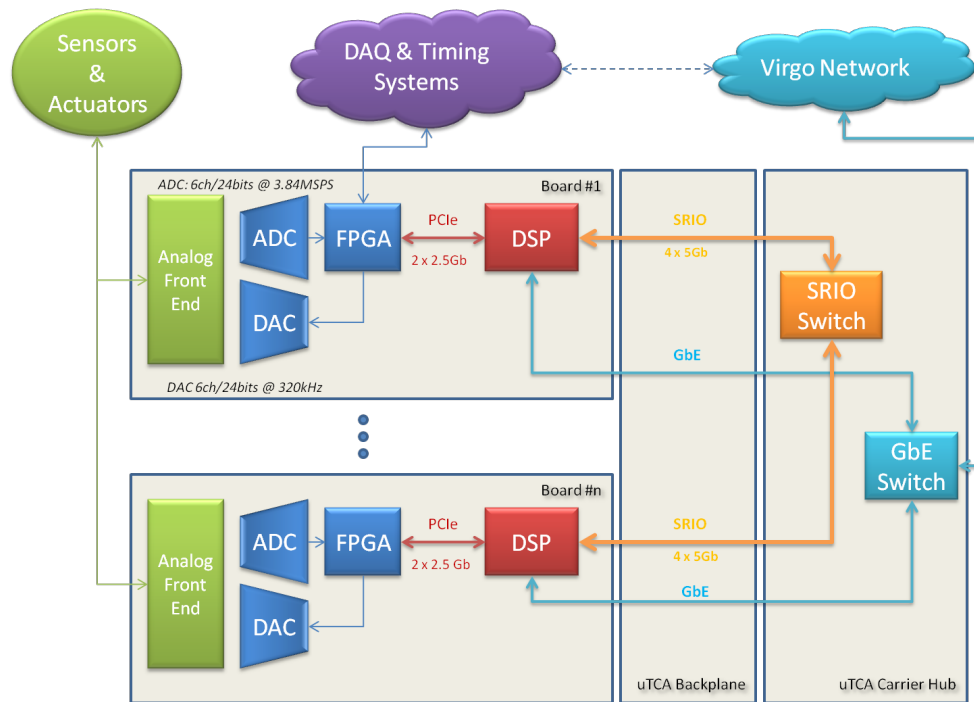


Figure 18: Block diagram of the hardware inside a MicroTCA crate [46]. On the left, there are n UDSPT boards, and on the right the MicroTCA Carrier Hub.

4.3.1 UDSPT boards

The core of the SCS is represented by the UDSPT boards: integrating 6 Analog-to-Digital Converters (ADCs), a DSP and 6 Digital-to-Analog Converters (DACs), they are used to acquire the data of the local sensors, to compute the feedback action and to drive the coil-magnet actuators.

These boards have been designed by INFN Pisa Virgo group. Key features of the UDSPT board include [36, 46]:

- Texas Instruments multi-core DSP TMS320C6678, that can perform up to 60 GFLOPS in double precision;
- 512 MB of DDR3-1333 memory;
- 170 pin B+ style AMC Interface containing RapidIO and Gigabit Ethernet;
- IRIG-B timing input;
- 1 Gbps Optical fiber link to send signals to the Virgo Data Acquisition (DAQ) and to receive signals from the interferometer global control;
- 2 Field Programmable Gate Arrays (FPGAs): one Xilinx Spartan3 dedicated to processing unit and one Altera Cyclone IV interfacing DSP with data converters and DAQ;
- 6 ADCs converters, 24 bit @ 3.84 MHz;
- 3 DACs stereo converters (6 channels individually addressable), 24 bit @ 320 kHz;
- Converters sampling frequency DSP interrupts synchronous with IRIG-B signal;
- Fully differential input channels and balanced output channels.

UDSPT boards are installed in MicroTCA crates, up to 12 for each crate. A crate with the boards plugged inside is called Local Control Unit (LCU), and is schematically represented in figure 18. The 20 Gbaud RapidIO link is used to communicate with the other boards of the same crate. A 5 Gbaud PCIe link is used to communicate with the Altera FPGA, while a Gigabit Ethernet link is used to control the DSP from remote.

4.3.2 Other Devices

MICROTCA CARRIER HUBS Each crate is equipped with a MicroTCA Carrier Hub (MCH), a MicroTCA controller produced by NAT Europe [48]. It provides the central management and data switching entity for a MicroTCA system, both for RapidIO and Gigabit Ethernet. The MCH can also provide centralized clock distribution to all the boards in the system. Every crate has a MCH, so there are ~ 25 of them installed in Advanced Virgo, including some spares.

STEPPER MOTOR CONTROLLERS Each suspension is equipped with a dozen stepper motors, used to set and balance working points of the mechanical components. They are used only during maintenance periods, due to their high noisiness. All the stepper motors are driven by Galil DMC-4183 controllers, three for each suspension, that make them remotely controllable.

POWER DISTRIBUTION UNITS The devices are plugged to remotely controlled Power Distribution Units (PDUs). The model is Switched Rack PDU AP7921 produced by APC. There is a PDU for each suspension. Each PDU has 8 slots, individually controllable: two are used for the MicroTCA chassis, three for the Stepper Motor Controllers, and another one for a piezo driver used for the actuators at the suspension base ring.

This is fundamental, both for turn off completely the UDSPT boards when needed, but also to keep the Stepper Motor Controllers turned off during the data taking, as they could introduce electromagnetic noise into the mechanical system.

4.4 DEPLOYMENT

In Advanced Virgo, there are 10 Superattenuators all operating in Ultra-High Vacuum (UHV) chambers. 7 Superattenuators are installed in the central building while the additional 3 are located in north end, west end buildings (3 km far from central area) and mode cleaner building. Two classes of Superattenuators are present, short and long, depending on actual chain length and number of seismic filters used in the chain. The short ones are

Suspension	Type	Chassis	UDSPT Boards
NE	Long	2	6+9
WE	Long	2	6+9
NI	Long	2	6+8
WI	Long	2	6+8
BS	Long	2	6+8
PR	Long	2	6+8
SR	Long	2	7+8
INJ	Short	2	5+5
MC	Short	2	6+7
DET	Short	1	5
BPC	n.a.	1	2

Table 1: List of suspensions with the relative type and the number of chassis and UDSPT boards plugged in the chassis. Then there are two chassis, the number of boards is reported as sum of the boards in each chassis. The Beam Pointing Control (BPC) is not a suspension, but it is reported in the table because it is controlled using the same electronics.

composed by a single stage pendulum chain with the F7 directly suspended to a Fo placed over an about 2.5 m-high Inverted Pendulum.

The Long Suspensions are:

- NE - north end mirror suspension;
- WE - west end mirror suspension;
- NI - north input mirror suspension;
- WI - west input mirror suspension;
- BS - beam splitter mirror suspension;
- PR - power recycling mirror suspension;
- SR - signal recycling mirror suspension.

The Short Suspensions are:

- INJ - injection bench suspension;
- MC - mode cleaner mirror suspension;
- DET - detection bench suspension.

The electronics installed in each suspension is reported in table 1. Typically, each suspension is equipped with two chassis, one for the Inertial Damping and the other for the Local Control. The only exception is represented by the short DET suspension is controlled by USDPT boards only at the base ring and Fo level, while the bench is controlled using a different system: this is why only one chassis is installed there.

Each chassis has its MCH, and each suspension has a PDU. Globally, in AdV 131 USDPT boards have been installed in 20 chassis, and powered by 10 PDUs.

4.4.1 Top Stage Control

In the Inertial Damping chassis, by default, there are six boards:

- Sensing of base ring LVDTs
- Sensing of Fo horizontal LVDTs and vertical LVDT
- Sensing of vertical LVDTs of F1, F2, F3, F4 and F7
- Sensing of horizontal accelerometers
- Sensing of vertical accelerometers
- Control loops and driving of Fo coil drivers

In addition a seventh board has been installed at SR and MC to drive the piezoelectric actuators at the base ring. The third board (vertical LVDTs) is not installed in the short suspensions.

4.4.2 Payload Local Control

In the Local Control chassis, by default, there are eight boards:

- Sensing of the F7 LVDTs
- Driving of the F7 coil drivers
- Sensing of the optical levers of the local control, which maintains the angular position of the mirrors with respect to local references (four boards)
- Driving of the marionette coil drivers
- Control loops and driving of the mirror coil drivers

At the end mirror suspensions there is an additional board for the actuation at the mirror level, because the coil drivers there have two distinct parts: one high power section for lock acquisition and one low noise section for linear regime. The two sections are driven by two independent digital to analog converter channels, with different values of resistance in series with the coil drivers.

The two suspensions part of the injection system are slightly different. The lower part of MC has a different geometry of its payload, while at the injection bench suspension there is a optical bench suspended instead of a mirror, and also there the logic and the number of board is different.

4.4.3 Other systems

An additional chassis is installed at the Beam Pointing Control (BPC), still part of the injection system. The fundamental idea of this system is to use beam pointing sensors as error signals to control beam pointing: this allows to lower the beam jitter at low frequency and keep the beam in a given reference position [36].

4.5 PERFORMANCES DURING O2

The performances of the Superattenuators and their SCS during the Virgo O2 in August 2017 are reported in figure 19. The official Virgo sensitivity (red line) is reconstructed only for $f > 8$ Hz (see section 9.1). The dark green is a reconstruction of the Virgo sensitivity at $f < 8$ Hz using the low frequency part of the z error signal sent to the end test masses (NE and WE), with a technique similar to the one that will be described in the section 9.1. We see how the contribution of the horizontal and vertical seismic noise (light blue and dark blue lines) is not limiting the Virgo sensitivity at frequencies $f > 2$ Hz [49].

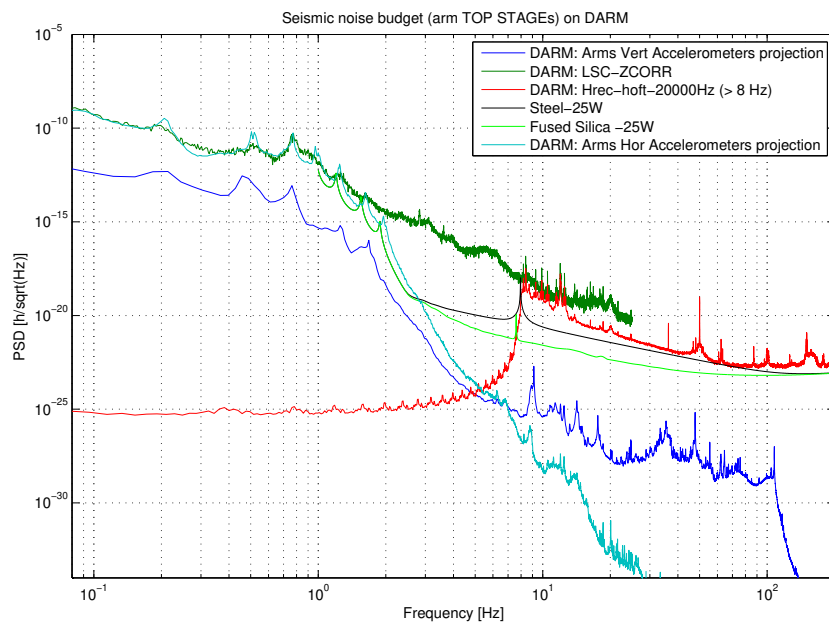


Figure 19: Projected contribution of the horizontal and vertical residual seismic noise on the official Virgo sensitivity (red) measured on August 14, 2017 [49].

5 | DSP SOFTWARE

The Texas Instruments multi-core DSP TMS320C6678 is a 8-core processor [50]. Cores are numbered from 0 to 7, and we'll refer to the single cores as *core0*, *core1*, ..., *core7*.

In this chapter we describe the software that is executed by the DSPs of the UDSPT boards used for the Suspension Control System.

5.1 BOOT AND OPERATING SYSTEM

The operating system is loaded at start time using a standard network booting system based on BOOTP/DHCP/TFTP protocols. The *core0* is responsible for the boot: in our configuration, at startup it executes the so-called *Intermediate Boot Loader* (IBL). IBL is flashed into an EEPROM connected to the DSP using I2C. The boot loader is programmed to download a bootable image from a TFTP server discovered using BOOTP. The boot image file is then loaded on memory, and eventually the IBL transfers control to the applications: a separate application is provided for each core, and together these applications constitute the operating system.

The operating system is based on the real time operating system *SYS/BIOS*. The kernel is made up of a number of discrete components, called modules. Each module can provide services via an API and is individually configurable.

It is completely written in C, with few assembly lines for some time critical operations.

A block diagram of the usage of each core is provided in figure 20. The system is composed of two parts, described in the next sections.

5.2 INTERRUPT SERVICE ROUTINE

Two cores of the DSP (*core 1* and *core 4*) are used to execute respectively two independent Interrupt Service Routines (ISRs). This is the "programmable" part of the DSP code. ISRs are always activated periodically with sampling frequencies f_{core1} and f_{core4} and consist of single functions.

ISRs are activated by polling transitions on dedicated clock lines generated by the Altera FPGA and synchronized with the GPS time. The approach is changed with respect to the old Virgo system, where the routines were activated by Interrupt Request (IRQ) generated by the analogous timing system:

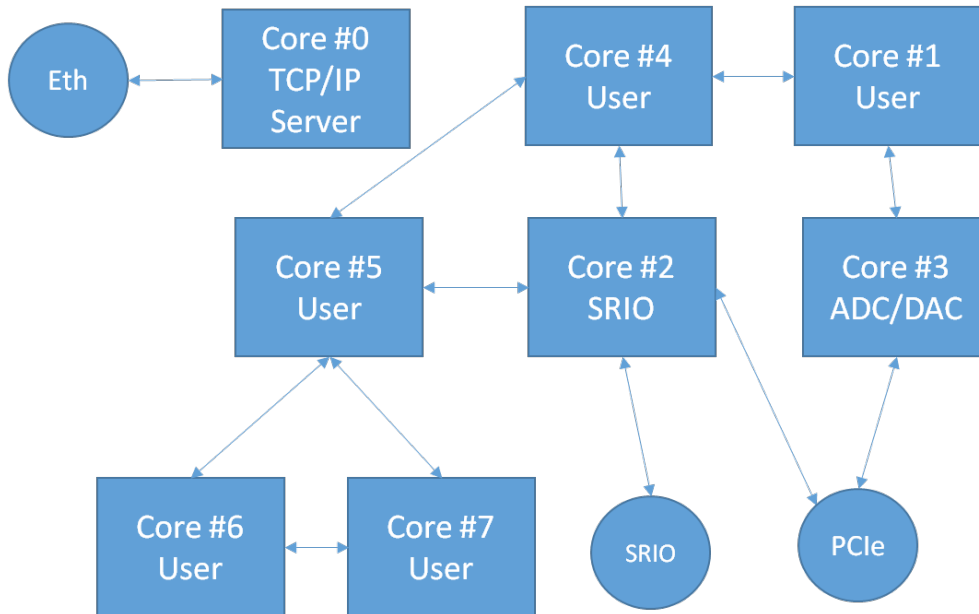


Figure 20: Usage of the eight cores of TMS320C6678.

this allow a faster activation of the routines, avoiding the time required by the processor for the context switch.

The sampling frequencies f_{core1} and f_{core4} are also defined by the user, among these:

- 10 kHz
- 20 kHz
- 40 kHz
- 80 kHz
- 160 kHz
- 320 kHz

By software design, the only constrain is that $f_{core1} \geq f_{core4}$. Both cores can read and write the 6 ADC and the 6 DAC channels on the board appropriately resampled. Signal can be sent from a routine to the other, with an opportune resampling. By design, only the routine running on *core 4* can access to the optical fiber and the RapidIO links to send and receive signals from other devices.

5.2.1 Typical usage

Usually, the routine on *core 1* is activated at $f_{\text{core1}} = 320 \text{ kHz}$ and is used only to demodulate the LVDT displacement sensors¹. During O₂, it was used only on the boards plugged to LVDTs and accelerometers.

On the other hand, *core 4* runs slowly at $f_{\text{core4}} = 10 \text{ kHz}$ and is used for the controls. This is because a larger number of computations can be done in $100 \mu\text{s}$, allowing to insert high-order digital filters and large matrix by column products, and also because the global control signals are produced and received at that frequency.

5.2.2 Code development

The ISRs are completely programmable by the user: from a source code written in a simple object-oriented language we generate assembly code for the cross-compiler that produces DSP binary code. The programming is done using a specific cross-compiler: the *Damping* application. Actually Damping offers something more than a cross-compiler since it is possible to use it as editor and generic user interface for the DSP program.

Cross-Compiler

The Damping cross-compiler produces the assembly code for the DSP starting from a description of the algorithm to be implemented using a functional blocks library[51]. The main output of the compiler is an assembly file with extension **.asm**. The **.asm** file is then passed to the Texas Instruments *c6000* proprietary assembler to produce the binary file for the DSP with extension **.out**.

In parallel to the **.out** file, the compiler produces two files: the **.map** file contains a set of information used at run time, including the full path of the binary code and variables location in DSP data memory. The compiler takes as input a **.net** file. The **.net** file contains the algorithm “netlist” that is the full list of used functional blocks with all connections, variables and parameters. This is the only file containing all the information required to rebuild the executable code for the DSP. The **.net** file is created by a pre-compiler which takes as input the source file, with extension **.hrd** as well as the other files referenced in the source.

¹ In Virgo, a typical LVDT has three coils placed around a common axis. The center coil is the primary, and the two outer coils are the secondaries. The primary coil is attached to a structure, while the secondaries are fixed each-other attached to another structure. The sensor measure the distance between the two structures: an alternating current drives the primary and causes a voltage to be induced in each secondary, proportional to the distance between the primary coil and its electrical center, typically at the middle between the secondary coils. The $f_{\text{core1}} = 320 \text{ kHz}$ sampling frequency is used because LVDT resonant circuits attached to the primary coils are tuned at $50 - 80 \text{ kHz}$, depending on the sensor.

Editor

The `.hrd` file is the file associated to the editor display. The Damping application in fact has a dedicated editor allowing a simple introduction of input and output channels, variables, functional blocks and functional blocks parameters. The editor creates a `.hrd` file containing a full description of the algorithm but only at high level. All parameters displayed in sub windows of the editor, filters poles and zeros, matrices coefficients and state space representations are not included in the `.hrd` file but in the referenced `.flt`, `.mat` and `.ssp` files respectively. The Damping editor allows to modify those parameters by editing the corresponding file.

5.2.3 Parameters Definition

ISR parameters are rigidly divided in two classes: *configuration parameters* and *monitoring parameters*. Configuration parameters are data that can be changed at runtime (“on the fly”), while monitoring parameters are data that can be only read at runtime.

By design choice, parameters of the filters, as poles and zeros location or damping factors, cannot be changed at runtime², and do not belong to any of the two classes. This choice comes from two different considerations. First of all, is not possible, in general, to change arbitrarily the position of a pole or zero in a filter without introducing heavy discontinuities in the filter output. The second consideration comes from the fact that DSP code usually implements feedback controls and therefore a change in filters parameters produces a big change in performances. Thus, two ISRs differing only in filters poles and zeros cannot be considered as a same application with two different configurations.

CONFIGURATION PARAMETERS Once excluded filters, it should be clear what configuration parameters are: gains and offsets. Gains and offsets are in fact the only input to change at runtime the operation of a given DSP program. Configuration parameters can be defined only on the *core4*. By default, the configuration parameters are sent and stored at $f_s = 0.5$ Hz in the Virgo Data Acquisition (DAQ) through the Software Supervisor that will be described in the next chapter.

MONITORING PARAMETERS Monitoring parameters are read only data. This class includes all named variables used in the DSP code, including ADC and DAC channels, RapidIO and optical fiber channels. It is up to the developer to select which are the parameters to be monitored in order to certify the quality of the service ensured by the application. By default, the selected monitoring parameters, historically named *probes*, are sent and stored at $f_s = 10$ kHz in the Virgo DAQ through the optical fiber link [52]. *Probes*

² Actually two operation are allowed: the change of the overall filter gain and the reset of filter internal state.

are available in *core4* only. In case the *core4* is running at $f_{\text{core4}} > 10\text{ kHz}$, data are downsampled in the DSP before to be sent out through the optical fiber. All the monitoring channels can also be acquired from both *core1* and *core4*, directly at f_{core1} and f_{core4} without downsampling, via a TCP/IP dedicated streaming.

5.3 SERVICES

The other cores of the DSP are used for services, as shown in figure 20. In particular:

- *core0* handles the Ethernet link, and executes a TCP/IP server that allows to control the status of the ISRs from remote;
- *core2* handles the RapidIO link, used to communicate with the other boards of the same chassis;
- *core3* handles the PCIe link, used to communicate with the on board Altera FPGA, that provides the interface with the analog frontend as well as the optical fiber link;
- *core5* is used to manage the streaming of monitoring parameters to a client via TCP/IP;

The last two cores are currently unused. In the next future, *core5*, *core6* and *core7* will provide an on board programmable spectrum analyzer. This will greatly simplify spectral analyses of the suspension controls, avoiding the current necessity to send a lot of monitor channels in the time-domain either to a client or via optical fiber to the DAQ.

6

SOFTWARE SUPERVISOR

Due to the large number of devices installed in the experiment, and in particular for the SCS, a Supervisory Control And Data Acquisition (SCADA) system is needed. A SCADA is a control system architecture that uses computers, networked data communications and graphical user interfaces for high-level process supervisory management, but uses other peripheral devices to interface to the process plant or machinery.

Historically in Virgo such a system has been called Software Supervisor, or simply *Supervisor*. A new Supervisor was required, because the Advanced Virgo hardware architecture is completely changed with respect to Virgo+, but also because the old supervisor was based on a discontinued toolkit. The new system is based on TANGO, an open source toolkit actively maintained and largely used in the scientific community.

In this chapter we describe the software supervisor developed for Advanced Virgo, which has been used during the scientific run O2. In the next chapter we'll describe the clients developed on the top of this system. The content of this two chapters is the first personal contribution described in this thesis.

6.1 REQUIREMENTS

The requirements of the Supervisor are:

- to provide a layer for the Virgo DAQ to read and store low frequencies variables ($f < 1$ Hz), like the configuration parameters;
- to provide tools that allow users (operators and experts) to monitor the state of the system;
- to provide tools that allow users to perform actions on the system or in a part of it, based on a permission control system, and to log them;
- to provide an automation layer to perform tasks on the system, for example to load the default ISR on the DSP.

6.2 TANGO TOOLKIT

TANGO [53] has been chosen as toolkit for the development of SCADA system for the AdV project; nevertheless, many Virgo subsystems are still using a custom framework called *Cm* [54], developed for the initial Virgo and

no longer maintained. The LIGO experiment is using *EPICS*, a framework originally written jointly by Los Alamos National Laboratory and Argonne National Laboratory since 1994, and is now used by many large scientific facilities throughout the world.

The TANGO control system is a free open source device-oriented controls toolkit for controlling any kind of hardware or software and building SCADA systems. It is used for controlling synchrotrons, lasers, physics experiments in over 20 sites in the world. It is being actively developed by a consortium of research institutes.

It has been chosen for AdV among other SCADA systems because it is a good compromise between worldwide usage, support, and ease of use.

TANGO uses two network protocols: the *omniORB* implementation of CORBA and *ZeroMQ*. The basic communication model is the client-server model: communication between clients and servers can be synchronous, asynchronous or event driven. CORBA is used for synchronous and asynchronous communication and *ZeroMQ* is used for event-driven communication.

Three key elements are defined in TANGO: *Classes*, *Devices* and *Servers*. TANGO is based on the concept of *Devices*. Devices implement object oriented and service oriented approaches to software architecture. The Device model in TANGO implements commands/methods, attributes/data fields and properties for configuring Devices. In TANGO all control objects are Devices. Devices are objects of a certain *Class*, that is the main class that the developer has to implement. Classes can be developed in C++, Python or Java.

Devices are provided by *Servers*. Generally speaking, the server is a process whose main task is to offer one or more services to one or more clients. To do this, the server has to spend most of its time in a wait loop waiting for clients to connect to it. The TANGO Devices are hosted in the Servers. A Server is able to host several Devices of a certain Class.

Moreover, TANGO has a built-in logging system, and also provides a controlled access system, named Tango Access Control, that defines which user is allowed to execute which commands (or write which attributes) on which devices and from which hosts. By default, all the devices can be accessed in read-only mode, and only some user can write attributes on them. [55]

6.3 DEVICES

The project of a software supervisor for the suspension electronics is based on a hierarchical system, and has been designed to handle the large number of different Devices. Each kind of hardware is defined by its own Class, that describes attributes and methods to interact with it. A client can access any hardware device by a unique TANGO Device.

We have developed a Class for each device described in 4.3. The properties of these Classes are reported in 2.

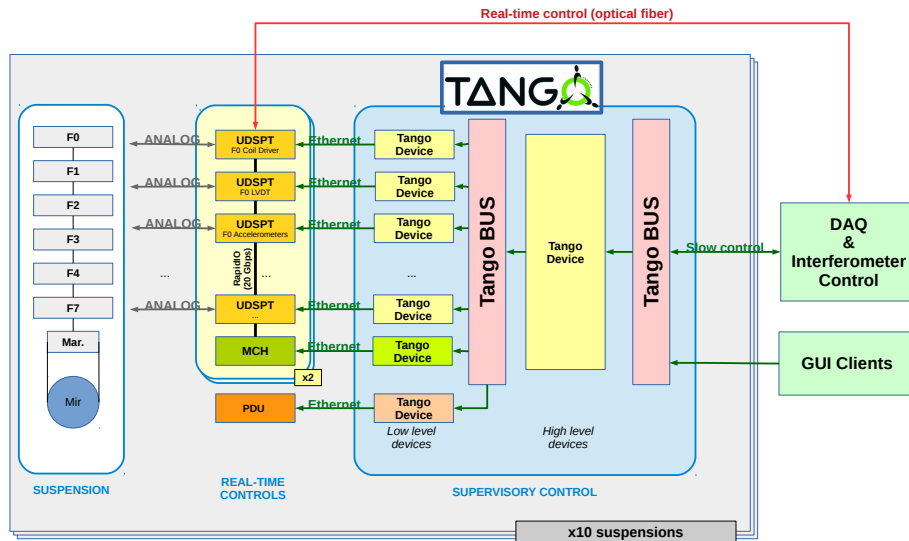


Figure 21: Hierarchy of Software Supervisor.

Class	Language	N. of Devices	TCP/IP Application layer
UDSPT	C++	~ 200	Custom
MCH	C++	25	IPMI, Telnet
PDU	Python	12	Telnet
MTR	C++	30	Telnet

Table 2: TANGO Device properties.

Class	Language	N. of Devices	TCP/IP Application layer
LCU	Python	25	TANGO
SAT	Python	12	TANGO

Table 3: TANGO Device supervisor properties.

As shown in 21, we've developed also a series of higher level Devices that interact as client with Devices of lower level, not accessing directly the hardware. This choice allows a smart management of the access permissions to methods and attributes of the Devices, reserving low level operation only to expert users, and allowing experiment operators to work in a safe environment.

Currently two high level Classes have been developed: one for the Local Control Unit (LCU), i.e. a crate consisting of a MCH, several UDSPT and a PDU slot, and another one, named "SAT", for a whole suspension, consisting of two LCU. The properties of these classes are reported in 3.

The Classes have been developed by the author with the key support of Luca Rei (INFN, Genova). The Stepper Motor class has been developed by Giulio Ballardin (EGO, Cascina). Key features of the developed Classes follows.

UDSPT CLASS The class allows to:

- load a ISR to the DSP
- get a list of all configuration parameters and monitoring parameters of the ISR
- read a parameter
- set a configuration parameter
- get uptime and GPS time
- get the version of FPGA firmware and operative system
- get temperatures of the board

MCH CLASS The class allows to:

- turn on/off a UDSPT board
- perform a reboot of the crate
- get list of UDSPT board MAC address in the crate
- get MCH firmware version
- get temperatures of the board

STEPPER MOTOR CLASS The class allows to:

- move the stepper motors of the suspension

PDU CLASS The class allows to control a PDU:

- turn on/off a device plugged in a certain slot
- get the current load

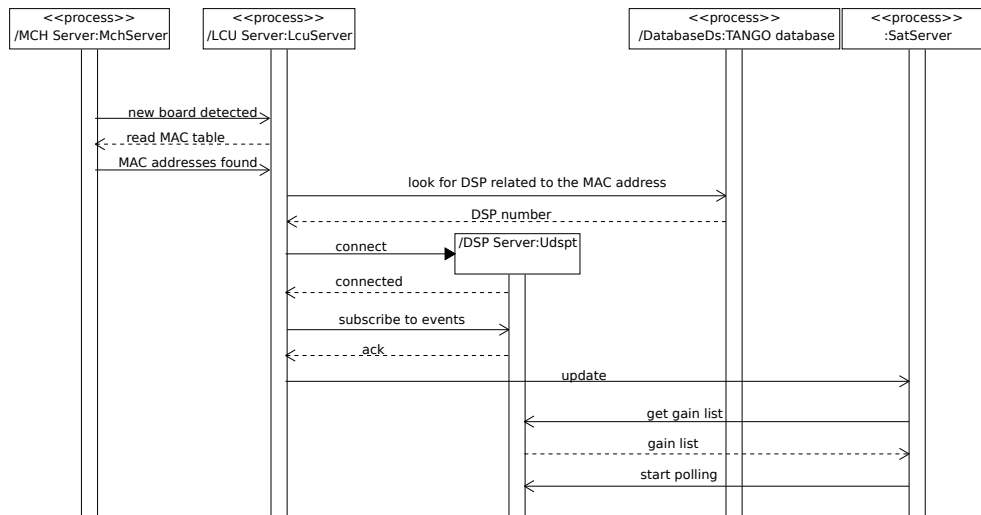


Figure 22: Sequence diagram of the board identification performed by the LCU Devices.

LCU CLASS The class is mostly used to identify the state of the UDSPT boards that are part of a LCU. It connects to three classes of lower level Devices:

- 1 MCH Device
- several UDSPT Devices
- 1 PDU Device

The procedure to identify the UDSPT board of the crate is the following. It retrieves the MAC table of the Ethernet switch embedded into the corresponding MCH, that contains the relation between the MAC address of the DSP of the UDSPT board on the crate, and the slot in which they are plugged. Then, it matches the found MAC with their serial number, querying a database. In this way, it knows which board is inserted in which slot. Then, it connects to the respective UDSPT devices: in this way it knows also their status. The process is described in figure 22.

It also monitors the PDU slot from which the crate is powered.

Once started, it monitors the state of the crate and of the boards inside it, notifying the user when something is changed. Events are sent to the users when:

- a new board is plugged or a board is removed
- a board is switched on/off
- there are failures in the UDSPT boards of a chassis
- the PDU slot to which the LCU is attached has changed its state

It allows to:

- turn on/off a UDSPT board
- perform a reboot of the crate
- get list of UDSPT board in the crate

SAT CLASS It is the highest level class. It allows to get and set the values of configuration parameters defined in the ISR running in any UDSPT of the suspension, despite of their position in the crates. There parameters are mainly switches for control loops, gains and set points. Indeed, the parameters that can be changed at runtime on a suspension are displaced in several DSPs, so the software supervisor has to know where is a board, and which routine it is running. It connects to three classes of lower level Devices:

- 2 LCU Devices ¹

Every instance of this class takes care of read and write some variables within the DSP ISR, usually used as switch and set points. For the reading, the device polls them every 2 s: when a client asks for the value of a parameter, it returns the last value read instead of sending a message to the DSP every time. The polling is needed because currently the configuration parameters can be changed also from the Damping Editor, that communicates directly with the DSP using the same TCP/IP protocol and bypassing the Software Supervisor. The 2 s polling period has been chosen as the configuration parameters are supposed to be changed on longer time scale (minutes). This allows to significantly reduce the network request to the DSPs, especially when many users or processes are asking for the same parameter.

6.4 DEPLOYMENT

TANGO runs in the EGO/Cascina environment, based on Scientific Linux 6 [56]. TANGO needs a MySQL database, where it stores the properties of each devices (IP addresses, serial numbers, suspension names, ...). As explained in the deployment diagram reported in 23, this database is installed in the local cluster and is shared with other subsystems of the experiment.

As described in the TANGO documentation [55], the minimal configuration consists of three TANGO Devices, one for each of these Classes, that have to be started manually:

- *DataBaseds*, used by any other TANGO device to access the MySQL database;
- *TangoAccessControl*, used to control the user permissions;

¹ Only 1 LCU Device for the detection suspension (B1 in 11) and for the beam pointing control system.

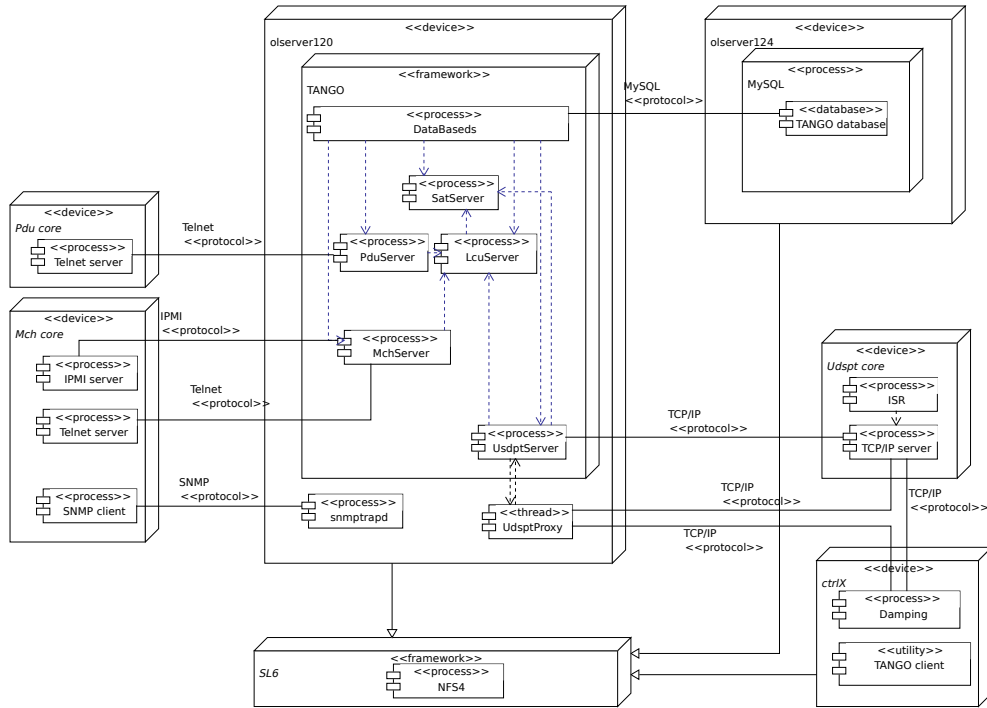


Figure 23: Deployment diagram of Suspension Software Supervisor.

Class	N. of Server	N. Devices per Server
UDSPT	11	2 to 15
MCH	11	1 or 2
PDU	1	10
MTR	10	3
LCU	11	1 or 2
SAT	1	11

Table 4: TANGO Servers providing the Advanced Virgo Software Supervisor. There are 45 TANGO Servers running on *olserver120*. In addition, there are the 3 system Servers, for *DataBases*, *TangoAccessControl* and *Starter*.

- *Starter*, used to automatically execute the user Devices defined in the database to run on that machine, and to control their status.

They don't need to run on the same machine, and for example several instances of *Starter* can be executed on different machines in case of a distributed system. However, in our case, all the TANGO Devices (the three system devices and our ~ 300 custom devices described above in tables 2 and 3) run on a machine named *olserver120*. The number of TANGO Devices is slightly larger than the number of devices actually installed in the experiment because there are devices also for the electronics installed in test facilities at the Virgo site.

Since a TANGO Server is able to host more Devices of a certain Class, there are only 45 Servers, as shown in table 4. We decided to group the Devices of a Superattenuators into the same server, in order to help maintenance and

reliability. Only the PDU Devices and the SAT Devices are running all into single instances of Servers. Similarly, in order to assure reliability of the system, Devices related to hardware installed in the test facility are executed by Servers not on *olserver120*, but on a different machine.

Permissions have been defined for 57 users. The highest level (represented by the SAT Class) is used by three types of clients:

- the Virgo DAQ (read only), through a Cm based Python server that polls the attributes on the TANGO devices and sends them to the storage servers [52];
- the interferometer automation (read and write);
- by users that need to adjust manually the value of the configuration parameters (read and write).

The lower levels are used both by the higher level devices and by expert only tools (for example the *Fo centering tool* that will be in 7.2.3, or the one in 27).

6.5 PERFORMANCES

The CPU and RAM usages of the machine *olserver120* in August 2017, i.e. during the Virgo observation run O2 are reported in 24: they do not show any relevant memory leak or strange behavior. The average network usage of that machine on the same period was

~ 180 kB/s

both in input and output.

In figure 25 we report the number of times the routines to adjust configuration parameters in the ISR have been called before and during O2. Totally,




Figure 24: Memory and CPU load of *olserver120* on O2, recorded by Virgo Ganglia system.




Figure 25: Number of calls to the SAT Server methods that set the values of configuration parameters in the ISR in the period June to August 2017. The shadowed period corresponds to Virgo O2.

there have been 273537 calls in 92 days, that means an average number of calls of

$$N_{\text{write}} \sim 3 \times 10^3 \text{ day}^{-1}$$

calls. Outside the O2 period the number of calls is higher due to the commissioning operations. Indeed, most of the calls have been done by the interferometer automation system, that will be discussed in the next section: the transition between interferometer states is performed also by changing a large number of configuration parameters. Since during a scientific run the goal is to state as much as possible in the science mode state, the number of calls is limited.

Nevertheless, only a small fraction of the calls to the supervisor are shown in figure 25: more than 400 configuration parameters were defined during O2, and each of them was read in polling every 2 s to be provided to users and to be stored in the Virgo DAQ system. This feature requires

$$N_{\text{read}} \sim 17 \times 10^6 \text{ day}^{-1}$$

calls to the DSPs, that account for most of the network usage described above.

7 | APPLICATIONS

The SCADA system can be used not only to monitor the status of the suspensions, but also to develop applications that don't require low latencies. In this chapter we describe several tools developed on the top of the TANGO framework, that can be used by operators. They are integral part of the software supervisor.

7.1 GUI CLIENTS

Several Graphical User Interface (GUI) clients have been developed, to easily access to the SCADA features described in the previous chapter. All of them have been written in Python, so they can run on any computer connected to the Virgo network, both on Linux and Windows. They are based on **PyQt4**, the Python binding of Qt GUI framework. This choice allows a fast development of the interfaces, and is compliant with the software development guidelines of the experiment. [57]

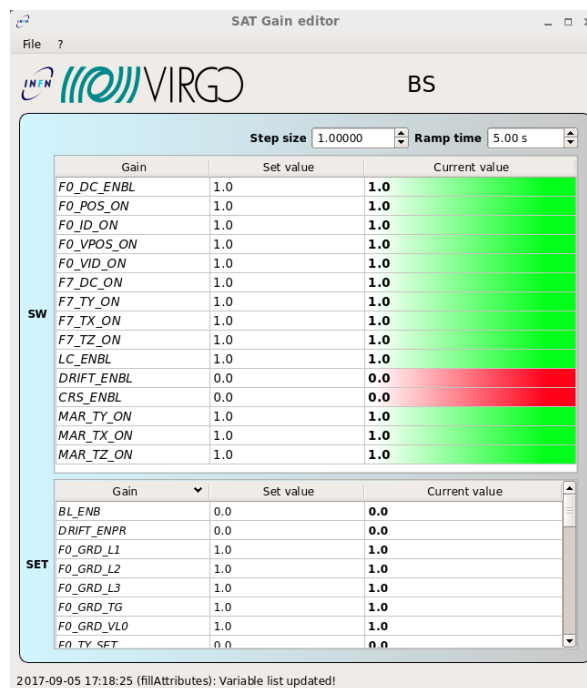


Figure 26: GUI tool to monitor and change configuration parameter of a suspension. Usually there are switches to open and close some control loops, and variable to change the work points of the suspension.

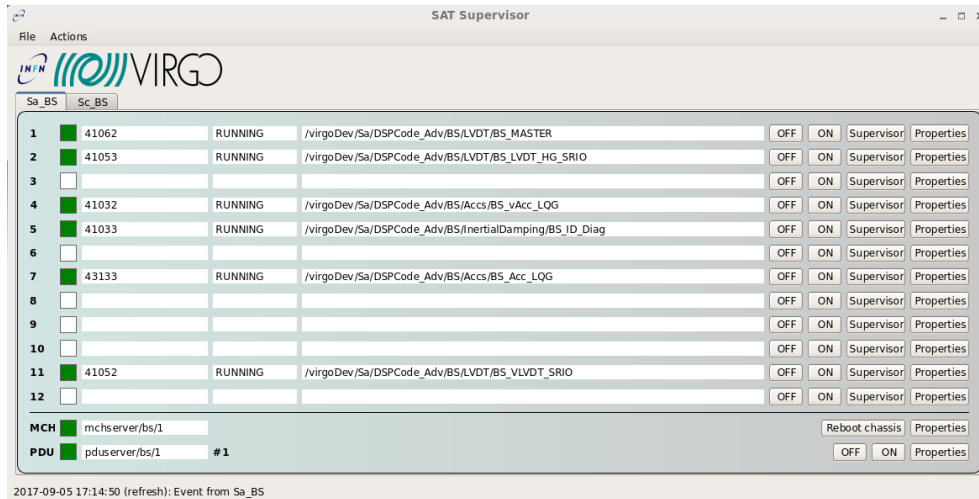


Figure 27: GUI that acts as client of the LCU Servers.

Two classes of clients exist. The first class consists of tools that can be used by non-expert users. Currently the only client of this type is shown in 26, that allows to set configuration parameters inside the ISR of a suspension: it is a client for the SAT Servers described in the previous chapter.

The second class consists of tools for experts. An example is shown in 27, and is a client for the LCU Servers. Clients exist also for all the low level devices.

7.2 AUTOMATION

7.2.1 Interferometer automation

While SCADA systems like TANGO and EPICS provide suitable mediums for supervisory control, they don't provide much in the way of structure or functionality for the development and management of higher level supervision tasks. These systems are designed for distributed control of large numbers of independent devices and typically include network message passing infrastructures as well as sequential logic programming tools for device-level automation. As stated before, Advanced Virgo relies on both *Cm* and TANGO, and LIGO relies on EPICS as the primary communication layer for supervisory control. For this reason, additional tools have been developed in order to manage this tasks, on the top of SCADA systems.

A common model used to represent automation systems is the Finite State Machine (FSM). FSMs are naturally represented by graphs, where states corresponding to particular configurations of, or commands on, a system are represented by nodes in the graph, connected together by directed edges defining allowable transitions between states. Finite state machine representations are quite powerful and intuitive and are well suited for many automation tasks.

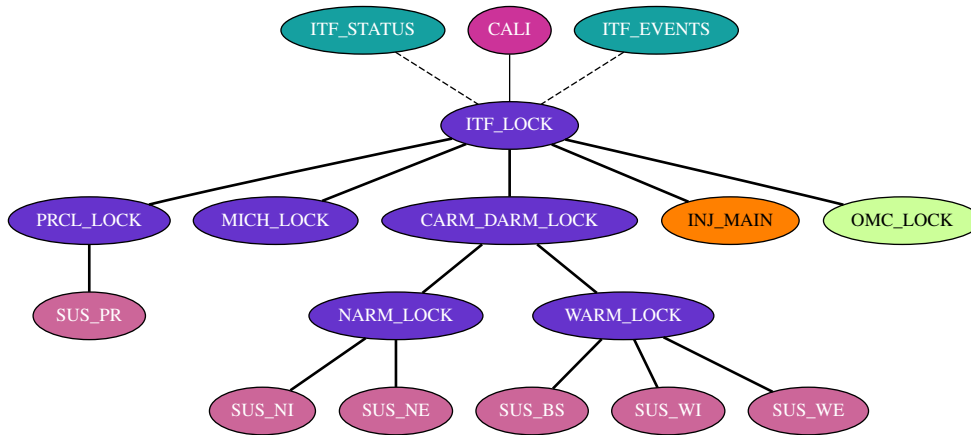


Figure 28: *Metatron* node hierarchy during O2.

The automation of the Advanced Virgo interferometer is based on a platform called *guardian*. It consists of distributed, independent, state machine automaton nodes organized hierarchically for full detector control. User code is written in standard Python and the platform is designed to facilitate the fast-paced development process associated with the commissioning of large instruments. It has been developed for Advanced LIGO, and then adopted also for Advanced Virgo, replacing the old supervisory system named *Alp* [58]. The version in use at Advanced Virgo has been renamed *metatron*. Even if the internal communication layer of *guardian* is based on EPICS, it can be expanded to use other protocols simply importing the relative libraries. In this way, in Advanced Virgo it has been expanded in order to support also *Cm* and TANGO. Indeed, TANGO framework is equipped with PyTango, a Python module that trivially exposes to Python the complete Tango C++ API. In this way, a transition between states of a generic FSM can be achieved using the attributes and the methods of an appropriate TANGO device.

Several FSMs have been hierarchically defined and deployed. The highest level FSM controls the status of the detector, and relies on the status of several lower level instances, referring to a smaller part of the experiment, e.g. a single arm of the interferometer. A figure that represents the hierarchy of the *metatron* node is represented in 28: the lowest level, represented by the suspension nodes (*SUS_XX*) use the TANGO layer to communicate with the system. The highest node (*ITF_STATUS*) is used to control the state of the whole interferometer.

7.2.2 Suspension automation

The states of the mirror suspensions are also handled by FSMs, that are part of the hierarchical automation of the experiment.

7.2.3 F_0 centering tool

The horizontal position of the top stage of a suspension, the F_0 , is controlled using 3 coil-magnet actuators driven by the UDSPT boards. Due to the finite dynamic of the DACs installed in the UDSPT boards (± 10 V), its position at rest can be set using three springs driven by stepper motors. The same mechanism applies to the vertical degree of freedom, even here there are only 2 coil-magnet actuators and a single stepper motor. So both coil-magnet actuators and stepper motors can be used as actuators to control the F_0 position. However, since the movement of the stepper motors introduces a lot of electromagnetic and mechanical noise in the system, they cannot be used when the interferometer is locked.

Problems arise in two cases:

- the reference position of the F_0 has to be changed for the alignment of the interferometer
- the position at rest of the suspension slowly drifts due to temperature changes, tidal motion and earth's subsidence.

In both cases, the mean value of the voltage sent to the coil-magnet actuators becomes not null, and could lead to a saturation of the DAC channels. For this reason, sometime it is necessary to adjust the position at rest of the suspensions using the stepper motors.

We describe here a slow control loop that makes strong usage of the TANGO Devices presented so far. The idea of the tool is simple: a reference horizontal position is chosen by the user, and a Python script takes care of moving the F_0 position at rest to that position in ~ 100 s. These are the features of the system:

- the reference position is set in the DSP ISR
- it works with the position controlled F_0 , i.e. with the loop closed using only the signal of LVDTs.
- it minimizes the corrections applied on the F_0 coil drivers in ~ 100 s
- it works in feedback

So far this procedure was done manually by an experienced user of the suspension, moving a motor per time until all the corrections were reasonably close to zero. The development of this tool allows to save a lot of time and to achieve more precise results.

Theory

The diagram of the controller is described in 29. If we consider the system quiet, the voltages \vec{v} applied on the three coil drivers in the F_0 are related




Figure 29: Closed loop system. Colors represent where the parts are implemented in the system: green are in the DSP code, red in a Python script, blue parts are the mechanic system. Since we want to set the corrections close to zero, we need to set $\vec{c}_0 = [0, 0, 0]$.

to the number of steps of the three stepper motors \vec{m} , needed to set $\vec{c} = 0$ maintaining the F_0 position, according to

$$\vec{c} = M \cdot \vec{m} \quad (24)$$

where M is a 3×3 matrix.

The system is subject to nonlinearities: for example, the position range of F_0 is limited, and it could touch a safety structure. For this reason, it's better to design the controller with a diagonalized system [59], calling \vec{c}_D the new vector of the corrections on the coil drivers and S the sensing matrix:

$$\vec{c}_D = S \cdot \vec{c} \quad (25)$$

The M matrix depends on the geometry of the mechanical system, and has been measured directly in all the Advanced Virgo suspension. The sensing matrix S and the driving matrix D are obtained from M using its singular value decomposition:

$$\begin{aligned} S &= \Sigma^{-1} U^{-1} \\ D &= V'^{-1} \end{aligned} \quad (26)$$

where $M = U \Sigma V'$ according to the rules of the singular value decomposition.

The correction $\vec{m}_D = C \cdot \vec{e}$ is calculated with a PI controller C using an error signal

$$\vec{e} = \vec{c}_0 - \vec{c}_D$$

and is then reverted in the \vec{m} coordinate system with the driving matrix D :

$$\vec{m}_i = D \cdot \vec{m}_D \quad (27)$$

The stepper motors accept only an integer number of steps, so the correction \vec{m}_i is quantized by Q to the nearest integer. An integrator I_Q here is placed in feedback to the quantizer. At this point, the correction is applied by the stepper motor M_{sm} , that acts as an integrator, to the system M , as \vec{m} number of steps.

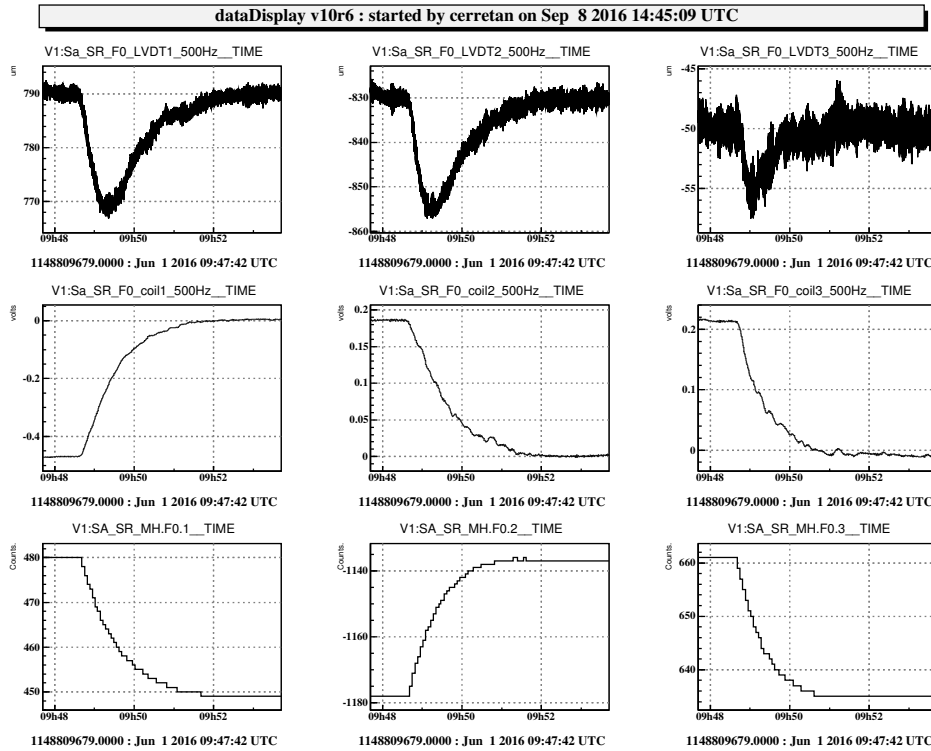


Figure 30: Application of the stepper motor centering tool on the Signal Recycling suspension.

Implementation and performances

As described in 29, the PI controller used for this tool is included in the ISR of the F_0 coil-magnet actuators DSP: they are the blocks in in green. A Python script, that corresponds to the red blocks, can be executed in any Virgo control workstation. Also a GUI client has been developed for this task, as shown in figure 31. It reads the output of the driving matrix from the DSP using a TANGO UDSPT Server, and the movements are sent to the stepper motor controller via a TANGO Stepper Motor server. The slowness of the motors must be taken into account: they takes about 3 s to perform 200 steps, including ~ 1 s overhead delay. The parts in the DSP works at 10 kHz, while the script runs at 0.25 Hz, so the number of steps is clipped at ± 200 to be sure that the motors is read at the next cycle.

An application to a Virgo suspension is reported in 30. On the first row are the signals of the LVDT displacement sensors, that read the position of the F_0 , while the second row are the correction applied on the F_0 coil driver. On the last row are the cumulative number of steps of the three motors that can move the F_0 . One minute after the begin of the time window the control loop is started, and we can see how the corrections tend to zero in ~ 100 s, while the positions remain around the reference position, except for a small overshoot at the beginning. The loop is automatically opened when

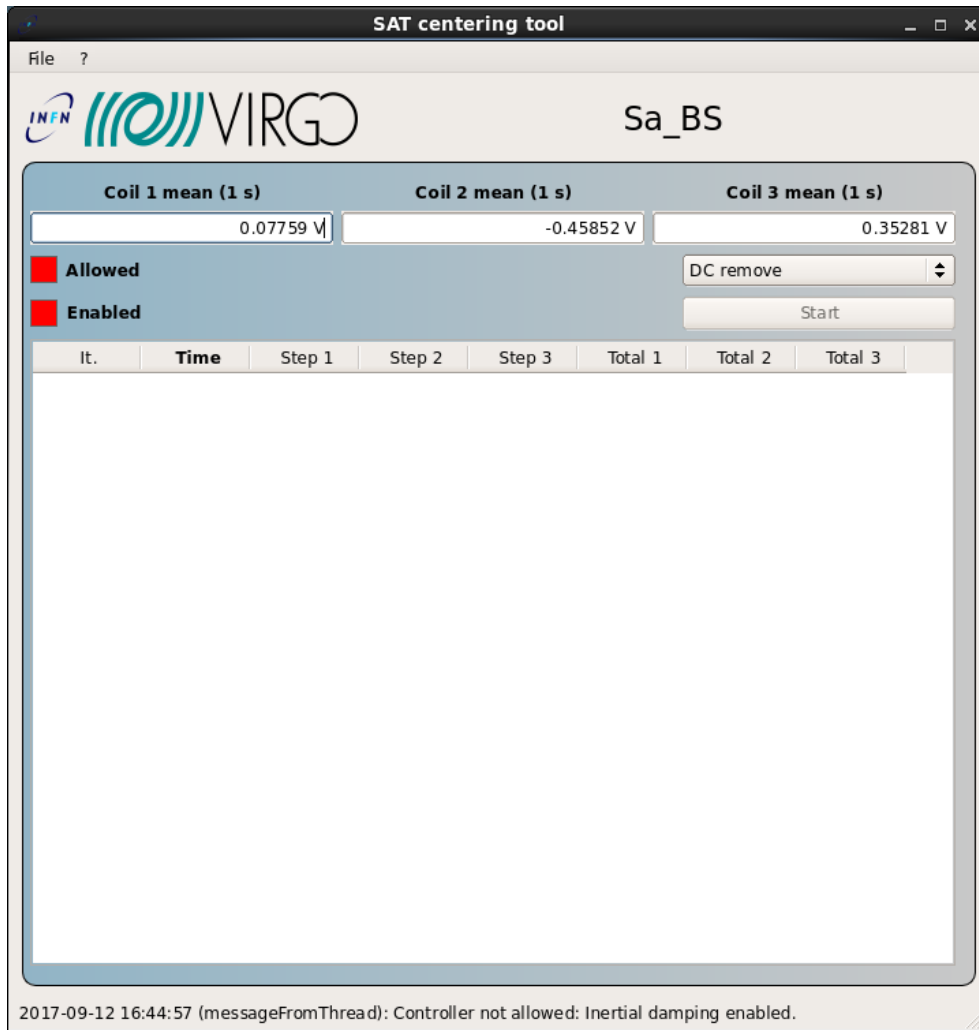


Figure 31: GUI used to control the F_0 centering tool.

the means of the last 4 corrections sent to the stepper motors are null, that usually happens in less than 2 minutes.

7.3 PYTHON SCRIPTING

Among the various clients developed for the Advanced Virgo commissioning and operations, and described in the previous sections, also other applications have been developed as simple Python scripts. For example, this Python script first reads the current value of the marionette θ_y reference position of the West End (WE) suspension, and then changes it to the position $264.8\mu\text{rad}$:

```
1 import tango
2
3 dsp = tango.DeviceProxy("satserver/we/1")
4 print(dsp.MAR_TY_SET)
5 dsp.MAR_TY_SET = 264.8
```

In this way, the TANGO devices features are used also for high level purposes, including both real-time tasks (as the interferometer automation) and the slow data acquisition, as well as scripts used to test the status of the electronics. Several tools have been developed especially to automatize the validation of the functionalities of the UDSPT boards.

Part III

Low-latency searches

8

LOW LATENCY SEARCHES

In what follows, *latency* is defined as the time a pipeline takes to send a trigger since the passage of the GWs through the detector.¹

In the search for gravitational waves, efficient and low latency on-line pipelines are fundamental if we want to be able to look for possible electromagnetic counterparts of the detected GW events: their latency must be low enough to trigger electromagnetic follow up observations by some astronomical partners.

Electromagnetic counterparts have been observed for the BNS GW₁₇₀₈₁₇, and are plausible also for Neutron Star-Black Holes (NSBHs) mergers. For what concerns BBHs mergers, there is little expectation of a detectable electromagnetic signature. In classical general relativity, a vacuum BBH merger does not produce any electromagnetic or particle emission whatsoever: even if supermassive BBHs in galactic centers may have other distinctive electromagnetic signatures due to interactions with gas or magnetic fields, stellar BBH systems are not expected to possess detectable EM counterparts [60]. As of the end of O₂, all the BBHs mergers detected have not shown any counterpart.

The initial detections of GWs have been made by two families of low-latency searches for gravitational-waves: burst searches and matched-filter searches [61]. In this chapter we describe the properties of these two families.

8.1 BURST SEARCHES

The searches belonging to this family aim at identifying time-correlated short duration transients (also known as gravitational-wave bursts) in multiple detectors, without assuming any particular signal morphology, origin, direction or time. This allows them to be sensitive to gravitational waves emitted by a wide range of sources: since burst methods do not require precise waveform models, the unmodelled search space includes also compact binary coalescences.

The first event GW₁₅₀₉₁₄ has been detected by two independent pipelines belonging to this family, coherent WaveBurst (cWB) and omicron-LALInference-Bursts (oLIB) [62].

Burst searches are performed in two operational modes: on-line and off-line. On-line, low-latency searches provide alerts within $\mathcal{O}(2 \text{ min})$, at least

¹ This definition does not apply in case of $\mathcal{O}(10 \text{ s})$ long events like BNS mergers. In this cases, we now assume the time of the passage of the GWs as the merger time. In future, we may take it as the time the cumulative SNR of a minutes long event exceeds a detection threshold. As of O₂, no low-latency pipeline provides this “pre-trigger” feature.




Figure 32: Completion time for jobs of the on-line cWB pipeline during LIGO O1.

for cWB. In the days and weeks following the data collection, burst analyses are refined using updated information on the data quality and detector calibration to perform off-line searches. These off-line searches provide improved detection confidence estimates for GW candidates, measure search sensitivity, and improve the waveform reconstruction and astrophysical interpretation.

We will describe only cWB, mainly because oLIB is slower, producing results in around thirty minutes from the arrival of the gravitational wave candidates.

8.1.1 coherent WaveBurst (cWB)

The cWB pipeline uses a method based on the likelihood ratio functional, for a coherent detection and reconstruction of burst signals.

The method is significantly different from the traditional burst detection methods: unlike coincident methods, which first identify events in individual detectors by using an excess power statistic and then require coincidence between detectors, cWB method combines data streams of two or more detectors into one coherent statistic constructed in the framework of the constrained maximum likelihood analysis.

The algorithms used in the cWB pipeline include [63]: wavelet transformation, conditioning of input data, construction of time delay filters, and generation of coherent triggers.

Latency performances

As already stated, it can work both on-line and off-line. The work is split in independent jobs. For the on-line analysis, a new job is executed every 60 s, and analyzes the last 180 s of data, with 120 s that overlap with the previous

job. On average, as shown in 32, during LIGO O1 cWB on-line pipeline took 2 min to process a job: 55 s were needed to submit a new job, and 50 s to execute them. This means that the trigger could arrive between 2 and 5 minutes after the arrival of the burst signal: in the worst case, if the signal arrives is at the beginning of a new job, then the event will be triggered 5 minutes after the arrival.

8.2 MATCHED-FILTER SEARCHES

The other family of pipelines is represented by the matched-filter searches, using relativistic models of compact binary coalescence waveforms. Since compact binary coalescence signals are well modeled by numerical solution of the GR, matched filtering is by definition the optimal method to detect these events.

In signal processing, the *matched filter* is the optimal linear filter for maximizing the SNR in the presence of additive stochastic noise. It is obtained by correlating a known signal, or template, with an unknown signal to detect the presence of the template in the unknown signal. This is equivalent to convolving the unknown signal with a conjugated time-reversed version of the template.

To perform this kind of search, a bank of search templates is defined to cover the parameter space of expected signals. This is generated before the initialization², to keep the analysis latency as low as possible. This template bank covers the parameter space we are interested in, and is dense enough to reduce problems related to the coarseness of template spacing. Differently from the burst searches, here the searches are done separately on each detector, and then triggers from the individual detectors are combined to find coincidence events [64].

Three independent matched-filter pipeline are currently used at LIGO: *PyCBC* and *GstLAL*. Another existing low-latency pipeline is Multi-Band Templated Analysis (MBTA), developed at Virgo.

PyCBC and GstLAL identify candidate events that are detected at both observatories consistent with the 10 ms inter-site propagation time. Events are assigned a detection-statistic value that ranks their likelihood of being a gravitational-wave signal. The detection statistic depends on the trigger SNR normalized by a χ^2 value. This detection statistic is compared to the estimated detector noise background to determine the probability that a candidate event is due to detector noise.

For O1, PyCBC and GstLAL used a common set of template waveforms, also known as template bank: it is shown in figure 33. For O2, each pipeline has used its own template bank. These analyses differ also in several parts: their implementations of matched filtering, their use of detector data-quality information, the techniques used to mitigate the effect of non-Gaussian noise

² An alternative could be to generate the templates at runtime instead of read them from memory.




Figure 33: The four-dimensional search parameter space covered by the template bank used by GstLAL and PyCBC during O1, shown projected into the component-mass plane, using the convention $m_1 > m_2$ [25]. The colors indicate mass regions with different limits on the dimensionless spin parameters χ_1 and χ_2 . Symbols indicate the best matching templates for GW150914, GW151226, and LVT151012. For GW150914 and GW151226, the templates were the same in the PyCBC and GstLAL searches, while for LVT151012 they differed.

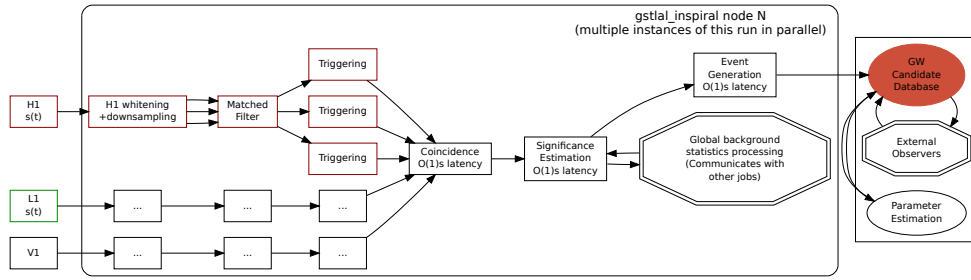


Figure 34: The `gstlal_inspiral` running topology.

transients in the detector, and the methods for estimating the noise background of the search. For what concern the generation of the templates, on the PyCBC analysis, sources with total mass less than $4 M_{\odot}$ are modeled by computing the inspiral waveform accurate to third-and-a-half post-Newtonian order. To model systems with total mass larger than $4 M_{\odot}$, they use templates based on the effective-one-body (EOB) formalism, which combines results from the Post-Newtonian approach with results from black hole perturbation theory and numerical relativity to model the complete inspiral, merger and ringdown waveform. The waveform models used assume that the spins of the merging objects are aligned with the orbital angular momentum. The GstLAL analysis uses the same waveform families, but the boundary between Post-Newtonian and EOB models is set at a total mass of $1.74 M_{\odot}$. [61]

For the O2 run, PyCBC had a latency of about 30 seconds, GstLAL of about 15 seconds, and MBTA of about 60 seconds.

8.2.1 GstLAL

GstLAL is a pipeline developed since 2011 by Chad Hanna and Kipp Cannon. It provides a suite of elements that expose gravitational-wave data analysis tools from the LALSuite³ library for use in *GStreamer*⁴ signal-processing pipelines.

The GstLAL-based inspiral search pipeline, `gstlal_inspiral`, is part of the GstLAL suite. It uses time-domain matched filtering to search for a set of plausible template waveforms in gravitational-wave data. It can perform both on-line and off-line search.

The principle of the GstLAL Low-Latency Online Inspirational Detection (LLOID) is described in figure 34. The data $s(t)$ and templates $h(t)$ are each whitened in the frequency domain by dividing them by an estimate of the power

³ The LSC Algorithm Library Suite (LALSuite) is comprised of various gravitational wave data analysis routines written in C following the C99 standard. LALSuite is the standard library used to analyze data from gravitational wave detectors.

⁴ GStreamer is a library for constructing graphs of media-handling components. The applications it supports range from simple audio playback, audio/video streaming to complex audio (mixing) and video (non-linear editing) processing. GStreamer is released under the LGPL.

spectral density of the detector noise. This procedure is applied piece-wise on overlapping Hann-windowed time-domain blocks that are subsequently summed together to yield a continuous whitened time series $s_w(t)$. The time-domain whitened template $h_w(t)$ is then convolved with the whitened data $s_w(t)$ to obtain the matched-filter SNR time series $\rho(t)$ for each template. Actually, what GstLAL does is slightly more complicated. In order to reduce the complexity of the analysis, the templates used by the pipeline are reduced to an orthonormalized basis of the original template bank, consisting of a much smaller number of templates. Furthermore, each template of this basis is split in blocks that are downsampled according to the Nyquist theorem. So, the first block of the template, far from the merger, will be processed at the lowest sampling frequency, and so on. This “partial intermediate SNR” are then upsampled and summed to get the “intermediate SNR”. Whitened data are tested against those templates, and then this “intermediate SNR” are mapped in the SNR of each original template using a matrix that contains the coefficients to reconstruct the original templates from the basis.

Peaks are then identified in the matched-filter SNR time series and these data are used to generate discrete triggers at a rate of one per second per template. If the peak is above a matched-filter SNR of 4, it is recorded as a trigger. The candidate triggers are also checked for consistency with gravitational wave signals through a χ^2 test that computes the consistency of the SNR time series surrounding the trigger. Each trigger is then checked for time coincidence with triggers from the same template in the other detector. If two triggers occur from the same template within 15 ms in both detectors, a coincident event is recorded and ranked according to a multidimensional likelihood ratio \mathcal{L} . The likelihood ratio \mathcal{L} provides a ranking of events such that larger values of \mathcal{L} are associated with a higher probability of the data containing a signal. The likelihood ratio itself is not the probability that an event is a signal, nor does it give the probability that an event was caused by noise. Instead, for each candidate event, GstLAL computes the false alarm probability \mathcal{F} . This is the probability of finding one or more noise background events with a likelihood-ratio value greater than or equal to that of the candidate event[61]. All this information is used to estimate the probability that an event is a signal.

Latency performances

During O1, its latency was above 30 s as shown in figure 35: for example, the gravitational-wave candidate GW151226 has been identified within 66 s [66].

Originally, the GstLAL low-latency analysis was designed to be used in three searches: the search for BNS, the search for NSBH binaries and the low-latency search for GRBs. Each of these three searches uses the results of one `gstlal_inspiral` analysis that cover the parameter space of compact binaries that could have an EM counterpart. At the beginning of O1, the




Figure 35: Latency of the on-line searches of GstLAL and MBTA during O1 [65]. The latency is measured as the time between the event arriving at Earth and time at which the event is uploaded to GracEDb.

on-line search parameter space was updated to include BBH mergers. During this run, the template bank contained 2.49×10^5 templates, and was the one shown in figure 33 [67]. Most of the analysis however is highly parallel: the work flow is divided into several multi-threaded jobs that run independently.

The LLOID algorithm matched filtering is the most computationally expensive portion of the analysis, requiring 64% of the computing time, while the other time is used by other services including a Python application layer (10%), the triggering and signal consistency (8%), kernel level calls (6%) and input data conditioning (4%).

The theoretical floating point operations for the LLOID method is

$$\gtrsim (1.4 \pm 0.1) \text{ MFLOPS/template}$$

to perform real-time data analysis at a 2048 Hz sample rate. [68] This roughly corresponds to ~ 350 GFLOPS, the theoretical computational power of 6 TMS320C6678 DSPs.

Since the beginning of O2, the second observation run of LIGO, some code has been rewritten, and the latency has been lowered to ~ 16 s: for example, GW₁₇₀₈₁₄ has been identified in 16.4 s⁵. Only ~ 4 s are required for the matched filter search itself, while the other time is related to the reconstruction of the strain $h(t)$ from the sensors of the interferometer, and the data transmission of it from the observatories to the computers where the on-line pipeline is executed.

⁵ This was the latency of L1 and H1 coincidence. The reconstruction of the V1 GW strain takes half a minute, so the triple coincidence was identified in 33 s

As described in the previous chapter, several pipelines have been developed to look for GW signals and have been used for the on-line search since the first Advanced LIGO observation run O1. The lowest latency pipelines are the matched filter searches. They are currently executed on standard *x86-64* servers on LIGO Scientific and Virgo Collaborations (LVC) clusters. Since the operations they perform are extremely simple and recurring, it is natural to think to implement such a kind of search in a DSP-based system. Such a system already exists, and is represented by the Advanced Virgo SCS. Several advantages are foreseen:

- The system is deeply integrated in the experiment, and most of the information used to compute the gravitational wave strain $h(t)$ measured by Virgo passes through the SCS DSP system with zero latency.
- As of the end of O2, 131 DSP-based boards are already installed in the experiment, and we are using only 6 of the 8 cores of the C6678 processors. Moreover, hundreds of spare boards are available and can be installed in the experiment for the specific task of the matched filter search.

In this chapter we describe how to implement a low-latency matched filter search for gravitational waves in DSPs: in the first part of the chapter we see how to reconstruct the gravitational wave strain $h(t)$; then, a simple implementation of a matched filter search is described, to look for waveforms contained in a specifically defined template bank up to the computation of the SNR for each template.

Then, in the next chapter we'll see how the same hardware could be used to improve the existing pipelines described in the previous chapter. The content of these chapters is the second personal contribution of this thesis.

9.1 RECONSTRUCTION OF THE STRAIN $h(t)$

Currently, the gravitational wave strain measured by the interferometer is reconstructed by an on-line process, named *Hrec*, which runs locally at Virgo in a server named *olserver52*, the main DAQ machine.

The algorithm used by *Hrec* in the initial Virgo science runs is fully described in [69]. Since the improvements of Advanced Virgo did not change significantly the optical scheme of the interferometer, most of the technique described there has been still valid for the Virgo configuration during O2. Its

latency is between 20 and 30 seconds, mostly because the Hrec reconstruction is done in the frequency domain: in case we need lower latency, the same method can be applied in the time-domain directly. This means that if we want to be able to implement lower latency search, we cannot rely on this process, but instead we need to compute the reconstructed gravitational wave strain within the DSP system.

In this section we describe the workflow of Hrec, without the feed-forward suppression of the 50 Hz power line described in [69].

9.1.1 Optical response

In the ideal case of interferometer mirrors not moving relatively to the others, the power-recycled Michelson interferometer (ITF) B1 output power, called \mathcal{P}_{DC} ,¹ depends on the differential arm length through the so-called ITF optical response $G_{ITF}O_{ITF}(f)$ of the ITF, and through the sensing transfer function $S(f)$:

$$\mathcal{P}_{DC}(f) = S(f) \times G_{ITF}O_{ITF}(f) \times L_0 \times h(f) \quad (28)$$

where L_0 is the length of the arms and $h(f)$ the Fourier Transform of a gravitational wave passing through the detector. $O_{ITF}(f)$ describes the frequency dependence of the transfer function while G_{ITF} is the low frequency gain. The frequency dependent part is a single pole transfer function

$$O_{ITF}(f) = \frac{1}{1 + i\frac{f}{f_p}}$$

where f_p is the cavity pole frequency. The pole is given by the fact that, when the propagation time of the laser beam inside the Fabry-Perot cavities, with finesse \mathcal{F} and length L_0 , is no longer negligible with respect to the period of the length fluctuations, the effect of the fluctuations are averaged over various round-trips. The frequency can be analytically obtained as

$$f_p = \frac{c}{4\mathcal{F}L_0}$$

It was around 500 Hz in Virgo and is now around 55 Hz in Advanced Virgo due to the higher finesse. In Virgo, this frequency slowly varied with time by $\pm 3.5\%$. The main source of variation was the etalon effect in the Fabry-Perot input mirrors which had parallel flat faces. In Advanced Virgo, during O2, the effect has been similar.

The so-called optical gain, G_{ITF} , measured in W/m, is the gain of the ITF optical response. It also varies slowly, in particular with the ITF alignment.

All these parameters are measured on-line at 0.1 Hz thanks to calibration lines injected permanently in the ITF, and the computed value is used by Hrec.

¹ In the Advanced Virgo optical layout, the detection of GW is done using the DC channel of the photodiode B1 at the *asymmetric port* (see 11), while the control of ITF is done using the AC channel.

9.1.2 Global control loop

Mirrors are subject to seismic noise, and their positions are controlled to keep the detector at its operating point, i.e. in dark fringe mode. The main controlled longitudinal degree of freedom is the differential arm length $\Delta L = L_N - L_W$, also known as **DARM**, defined as the difference between the length of the north arm L_N and that of the west arm L_W . The controller is mostly a Proportional-Integral-Derivative controller (PID) with some notches, with its response cut at 1 kHz by a 3rd order low-pass Butterworth filter.

For this reason, in the ITF output we have also to take into account the relative motion of the mirrors $\delta L_i(f)$ induced by this control loop. In other words, we need to add a term

$$S(f) \times \sum_i [G_i O_i(f) \times \delta L_i(f)]$$

to the right side of equation 28.

Different optical responses $G_i O_i(f)$ are defined, associated to the responses of the ITF to variations of the positions of the i -th mirror. In particular, the responses to motions of the end mirrors (NE, WE) and of the beam-splitter have the same shape $O_{ITF}(f)$. The optical gains associated with the motion of the end mirrors are expected to be equal to G_{ITF} , while that associated with the beam splitter mirror is lower, being it outside the Fabry-Perot cavities: $G_{BS} \sim \frac{G_{ITF}}{2\mathcal{F}/\pi}$. For what concerns PR, the gain of the optical response to PR displacement is low compared to the other mirrors.

The histograms of the measured DC gain G_i during Virgo O2 are displayed in figure 36, and show that the fluctuations are relatively small:

$$\begin{aligned} O_{BS} &= (1.32 \pm 0.09) \times 10^7 \text{ W/m} \\ O_{PR} &< 1 \times 10^{-2} \text{ W/m} \\ O_{NE} &= (4.00 \pm 0.04) \times 10^9 \text{ W/m} \\ O_{WE} &= (3.89 \pm 0.04) \times 10^9 \text{ W/m} \end{aligned}$$

For what concerns $\delta L_i(f)$, they can be extracted as

$$\delta L_i(f) = A_i^{\text{mir}}(f) \times zC_i^{\text{mir}}(f) + A_i^{\text{mar}}(f) \times zC_i^{\text{mar}}(f)$$

where $zC_i^{\text{mir}}(f)$ is the voltage sent to the mirror coil actuators, measured in V, and $A_i^{\text{mir}}(f)$ is the transfer function between the voltage and the displacement, measured in m/V. The same are $A_i^{\text{mar}}(f)$ and $zC_i^{\text{mar}}(f)$, but related to the marionette.

The estimated actuation gain A_i^{mir} , perpendicular to the mirror surface, is

$$A_i^{\text{mir}}(f) = \frac{n_{\text{coils}} \gamma \alpha}{M_{\text{mir}} \omega_0^2} \times p_i^{\text{mir}}(f)$$

the first part is the response due to the electromagnetic actuator, that can be estimated from the number of coils n_{coils} , the nominal values of the conversion factors of the coil driver electronics γ , the current-force conversion

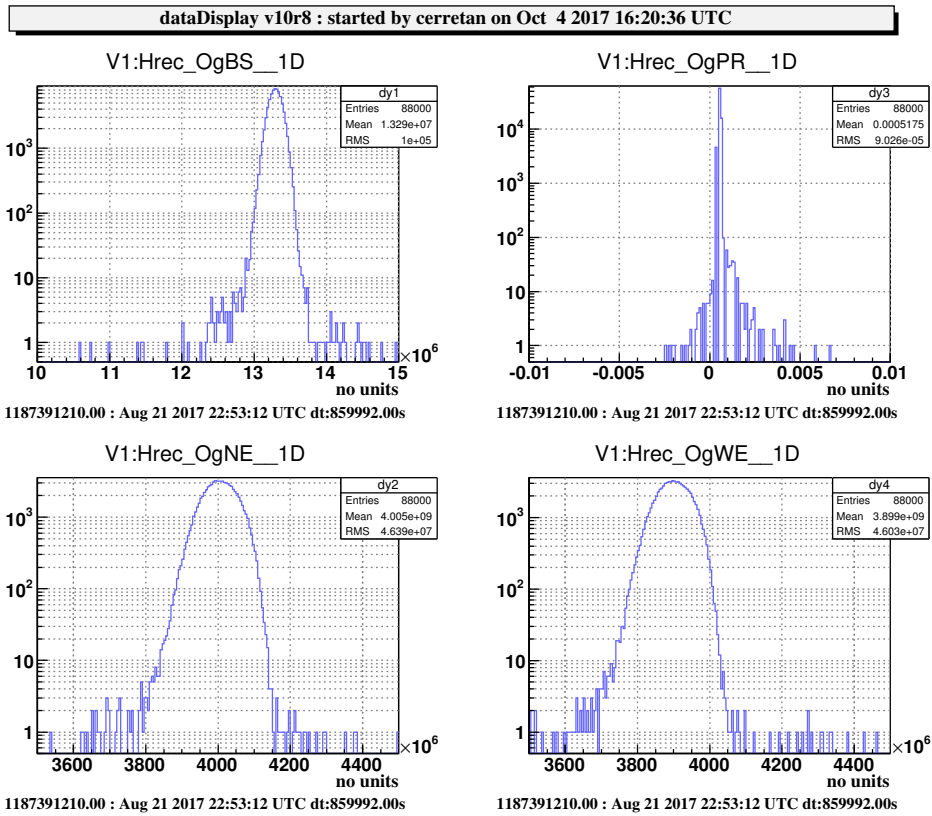


Figure 36: Optical gains of the various mirror between August 12 and August 22, 10 days of Virgo O2. Units on x-axes are W/m.

factor α and the mechanical response of the pendulum. The second term, $P_i^{\text{mir}}(f)$, is the mechanical response of a pendulum, with $f_0 = 0.6$ Hz and $Q = 1000$.

For what concerns A_i^{mar} , it is computed in a similar way

$$A_i^{\text{mar}}(f) = \frac{n_{\text{coils}}\gamma\alpha}{M_{\text{mar}}\omega_0^2} \times P_i^{\text{mar}}(f)$$

while now $P_i^{\text{mar}}(f)$ represents a double stage pendulum, with two pairs of complex poles at $f_0 = 0.6$ Hz and $Q = 1000$.

9.1.3 Reconstruction

In conclusion, the output power \mathcal{P}_{AC} , that is the main signal of the interferometer, is sensing the effective differential arm length variations which come partly from the imposed motions of the different controlled mirrors $\delta L_i(f)$, and partly from the free variations, $L_0 \times h(f)$:

$$\mathcal{P}_{DC}(f) = S(f) \times \left\{ \sum_i [G_i O_i(f) \times \delta L_i(f)] + G_{ITF} O_{ITF}(f) \times L_0 \times h(f) \right\} \quad (29)$$

Now we can extract $h(f)$:

$$h(f) = \frac{1}{L_0 \times G_{ITF} O_{ITF}(f)} \left[\frac{\mathcal{P}_{DC}(f)}{S(f)} - \sum_i [G_i O_i(f) \times \delta L_i(f)] \right] \quad (30)$$

Of course, one could state that the filters of the longitudinal control loop could be modeled to extract directly $h(t)$ from the dark fringe signal \mathcal{P}_{AC} , because the same \mathcal{P}_{AC} is the error signal used to compute all the zC_i^{mir} and zC_i^{mar} . However, the Virgo reconstruction method for $h(t)$ is based on the subtraction of the control contributions from the dark fringe signal, in order to recover the signal of a free ITF. This method makes the reconstruction independent of the global control system since the knowledge and monitoring of its filters are not needed.

All the computation so far are described in the frequency domain, and are used to compute what we call $h_{fd}(t)$. It is easy to implement such algorithm directly in the time domain to compute $h_{td}(t)$, using the time series of zC_i^{mir} , zC_i^{mar} and \mathcal{P}_{DC} , and replacing, with convolution with appropriate Infinite impulse response (IIR) filters, the products with the transfer functions.

In figure 37 we report both the transfer function and the coherence between $h_{fd}(t)$ and $h_{td}(t)$. The plots are relative to a 100 s long segment acquired on August 2011, during the Virgo VSR4: the former is the output of the Hrec process, while the second is the one obtained using the procedure described in this section in the time domain. The only main difference is that in our case we have not applied the power line suppression described in [69]. This is why there is low coherence on 50 Hz and its harmonics. For what concerns the transfer function, its module is contained between 0.8




Figure 37: Transfer function and coherence between $h_{fd}(t)$ and $h_{td}(t)$.

and 1.7. The phase shift is not a problem, and can be adjusted applying appropriate delays to signals used in the process.

In conclusion, it is very simple to obtain an estimate of $h(t)$ using the standard procedure, with zero latency and an accuracy within $\sim 50\%$.

9.1.4 The Advanced LIGO approach

For completeness, I will describe briefly how Advanced LIGO reconstructs the gravitational wave strain in a different way [70]. Calibrated data is produced in the time domain, using both a low-latency, on-line procedure and a high-latency, off-line procedure.

The low-latency $h(t)$ data stream is produced in two stages, the first of which is performed on the same computers that operate the detector's feedback control system. This stage, referred to as the *front-end* calibration, uses IIR filtering and performs all operations on 16384 Hz clock cycles. Due to several limitations, this procedure currently introduces certain systematic errors in the calibrated strain data, motivating the second stage of the low-latency procedure, known as the low-latency *GstLAL* calibration pipeline. This second stage uses the same software used by the low-latency matched filter search described in the previous chapter. The *GstLAL* calibration pipeline uses Finite impulse response (FIR) filtering to apply corrections to the output of the front-end calibration. It applies time-dependent correction factors to the sensing and actuation components of the calibrated strain to reduce systematic errors. The difference between the outputs of the two stages during O2 was $< 5\%$ in the band between 20 Hz and 2000 Hz. The latency of the first stage is ~ 0 , while the latency of the second stage is ~ 10 s.

The *GstLAL* calibration pipeline is also used in high-latency (up to several weeks) to recalibrate the data, which is necessary due mainly to on-line dropouts in the calibrated data and systematic errors in the calibration models and filters made in low latency. Additionally, data are recalibrated when significant improvements to the calibration can be identified. The difference between the low-latency and the high-latency streams during O2 was $< 2\%$ in the band between 20 Hz and 4000 Hz.

9.2 MATCHED FILTER SEARCH

We describe in detail an algorithm to perform a matched filter search on a sampled stream of $h(t)$ reconstructed using the method described in the previous section. It is based on the *FINDCHIRP* algorithm described in [71], and can be easily implemented in a DSP-based device.

Suppose that $n(t)$ is a stationary Gaussian noise process with power spectral density $S_n(f)$. Then the matched-filter output of a data stream $h(t)$,

that represents an eventual gravitational wave strain $s(t)$ added to that noise $n(t)$,

$$h(t) = n(t) + s(t)$$

with a filter template $s_m(t)$, is

$$\begin{aligned} \chi_m(t_0) &= 2 \int_{-\infty}^{\infty} \frac{\tilde{h}(f) [\tilde{s}_m^*(f)]_{t_0=0}}{S_n(f)} e^{2\pi i f t_0} df \\ &= 4\Re \int_0^{\infty} \frac{\tilde{h}(f) [\tilde{s}_m^*(f)]_{t_0=0}}{S_n(f)} e^{2\pi i f t_0} df \end{aligned}$$

where the signal $s_m(t)$ is implicitly taken to depend on a termination time t_0 . The $\tilde{}$ and the dependence on f mean that that term has been Fourier transformed, while the $*$ represents the conjugate. The waveforms $s_m(t)$ have unknown parameters:

- the amplitude,
- the coalescence phase,
- the binary companion masses and spins.

The amplitude simply sets a scale for the matched-filter output, and is unimportant for matched-filter templates (these can be normalized). The unknown phase φ_0 can be found by maximizing $\chi(t_0)$ over φ_0 . In practice, we have to compute the matched filter with two orthogonal templates, phase-shifted by $\pi/2$. An analytical solution consists in computing the complex output

$$z_m(t_0) = 4 \int_0^{\infty} \frac{\tilde{h}(f) [\tilde{s}_m^*(f)]_{t_0=0, \varphi_0=0}}{S_n(f)} e^{2\pi i f t_0} df \quad (31)$$

where $z = x_{re} + ix_{im}$, and x_{re} and x_{im} are the values of the equation 31 with $\varphi_0 = 0$, $t_0 = 0$ and $\Re \rightarrow \Im$. The output is given by the modulus of z ; the phase $\varphi_0 = \frac{1}{2} \arg z$.

For what concerns masses and spins, since they change significantly the frequency evolution of the merger, we have to construct a template bank to cover our parameter space. The construction of the template bank will be discussed next in this chapter.

By convention, the waveform templates are constructed for systems with an effective distance of $D_{eff} = 1$ Mpc. To produce SNR series as output, a normalization constant for each template m is computed

$$\sigma_m^2 = 4 \int_0^{\infty} \frac{|\tilde{s}_m(f)|^2}{S_n(f)} df$$

that is a measure of the sensitivity of the instrument for the m -th template. Thus the quantity

$$\rho_m(t) = \frac{|z_m(t)|}{\sigma_m} \quad (32)$$

is the amplitude SNR ratio of the (complex modulus) matched filter.

In the discrete world, if we use the sampled version of h and s_m , 31 can be written as

$$z_m [j] = 4\Delta f \sum_{k=0}^{N-1} \frac{\tilde{h} [k] \tilde{s}_m^* [k]}{S_n [k]} e^{2\pi i j k / N}$$

where also here in \tilde{s}_m both t_0 and φ_0 are set to zero and where the template is made of N points. In this case, the $\tilde{\cdot}$ represent the Fast Fourier Transform (FFT) applied to the signal, and the index k is related to a frequency bin. Of course, the output of this algorithm can be improved windowing and overlapping the segments of sampled stream.

Now S_n is a discrete estimation of the power spectral density, that can be done for example using the Welch method with an appropriate window [72]. Once known S_n , it is straightforward to compute σ_m as

$$\sigma_m^2 = 4 \sum_{k=1}^{N-1} \frac{|\tilde{s}_m [k]|^2}{S_n [k]} \Delta f$$

In conclusion, the matched filter algorithm corresponds to compute the value

$$\tilde{z}_m [k] \Delta f = 4\Delta f \frac{\tilde{h} [k] \tilde{s}_m^* [k]}{S_n [k]} \quad (33)$$

and then to obtain $z_m [j]$ through the inverse Discrete Fourier Transform (DFT). The sampled SNR is given by

$$\rho_m [j] = \frac{|z_m [j]|}{\sigma_m} \quad (34)$$

9.3 PERFORMANCES

Now that we've described how to implement a matched filter search on a sampled stream of data, let's see which is the computational cost and the latency of this algorithm.

9.3.1 Computational cost

For each template m , the computational cost of the matched filter output is mainly represented by the inverse DFT of the N element long series \tilde{z}_m . If we use the FFT algorithm, the complexity goes as $N \cdot \log N$, while the products in 33 goes just as N . Indeed there is just one product for each element in \tilde{z}_m , because S_n changes much slower than the length of a template², so that we compute the amount \tilde{s}_m^*/S_n in advance in a separated process. The

² This is not always true. For a BNS template, the length can be hundreds of seconds and, in some frequency bands, S_n can have variations below this time scale.

output of the inverse FFT gives simultaneously the matched filter output for N samples in the time domain.

We can compute how many templates can be analyzed simultaneously by a single core of the TMS320C6678 DSP. First of all, the N -samples long FFT of the strain, \tilde{h} , have to be computed just once for all the templates, so it can be neglected. Then, each template requires

$$C_{\text{conv}} = N$$

clock cycles for the convolution product in 33. The computation of the inverse FFT, assuming N is a power of 4, the fastest solution is to use the *radix-4*³ base block. Each radix-4 block requires 12 products and 22 sums, that can be computed in

$$C_{\text{radix-4}} = 10$$

clock cycles by the DSP, according to the specifications of the processor [50]. The number of radix-4 to be computed for template is

$$N_{\text{radix-4}} = N/4$$

so that each inverse FFT requires

$$C_{\text{IFFT}} = C_{\text{radix-4}} \cdot N_{\text{radix-4}} \cdot \log_4 N$$

where the term $\log_4 N$ is given by the Cooley–Tukey FFT algorithm.

Then, according to equation 34, we must compute the modulus of the complex value $z_m[j]$, and to multiply it by σ_m^{-1} . This adds other 2 clock cycles

$$C_{\text{norm}} = 2N$$

It means that the clock cycles required to get N samples of the SNR time series for the m -th template are

$$C_{N\text{-samples}} = C_{\text{conv}} + C_{\text{IFFT}} + C_{\text{norm}}$$

and, for each sample of the m -th template,

$$\begin{aligned} C_{\text{sample}} &= \frac{C_{N\text{-samples}}}{N} \\ &= \frac{C_{\text{conv}} + C_{\text{IFFT}} + C_{\text{norm}}}{N} \\ &= \frac{N + C_{\text{radix-4}} \cdot N_{\text{radix-4}} \cdot \log_4 N + 2N}{N} \\ &= \frac{N + 10 \cdot N/4 \cdot \log_4 N + 2N}{N} \\ &= 3 + \frac{10}{4} \log_4 N \end{aligned} \quad (35)$$

³ In *radix-4* is a possible *butterfly* of the FFT. The *butterfly* term appears in the context of the Cooley-Tukey FFT algorithm, and represents a portion of the computation that combines the results of smaller DFTs into a larger DFTs.

It is reasonable to choose $N = 4^7 = 16384$, so that at $f_s = 10$ kHz we have templates 1.6 s long, because almost all the BBH detected so far had a duration less ~ 1.6 s in the interferometers detection bandwidth, as we'll see in the next section in figure 39. Then

$$C_{\text{sample}} = 20.5 \quad (36)$$

The DSP are clocked at $f_{\text{DSP}} = 1$ GHz⁴ so that each instruction is performed in $T_{\text{DSP}} = 1$ ns. So, in order to have a real time computation, on average we must compute all the templates within $f_s^{-1} = 100$ μ s, that corresponds to $f_s^{-1}/T_{\text{DSP}} = 10^5$ clock cycles of the DSP. So, the theoretical maximum number of templates that can be computed by a single core of the TMS320C6678 DSP is

$$N_{\text{max}} = \frac{f_s^{-1}/T_{\text{DSP}}}{C_{\text{sample}}} \simeq 4800 \quad (37)$$

Several simplifications have been done in this estimation. First of all, we've neglected the time required to read the templates from the memory and the cache efficiency, that can reduce our performance by $\sim 80 - 90\%$. We also neglected the 50% overlap between two segments, that brings another factor 2 to the computational cost since each sample in $s[j]$ is actually computed twice.

On the other hand, we could think to downsample the original $h(t)$ stream from $f_s = 10$ kHz to a frequency at least twice the maximum frequency of our templates, according to the Nyquist–Shannon sampling theorem. This is the approach of GstLAL. Since data has to be downsampled only once, reducing the sampling frequency by a factor 10 increases by 10 times N_{max} . Moreover, in this case also N is reduced by a factor 10, and we get $C_{\text{sample}} = 24.7$, increasing N_{max} by an additional 25%.

9.3.2 Latency

There is an intrinsic delay in the algorithm presented in the previous section, that is given by the length of the segment, that could be up to 1.6 s as just described. This is actually a limitation of the frequency domain analysis, since we need to wait for N new samples in order to start the matched filter algorithm. So, the sample at the beginning of the segment will be processed after a time $N \cdot f_s$, while the last sample of the segment will have only a delay equal to the computation time. A possible solution to overcome this problem is to perform the matched filter search in the time domain.

Exploiting the convolution theorem, the matched filter output could also be computed in this way, with no intrinsic delay in the algorithm other than the computation itself:

$$z_m[j] = \sum_{l=0}^{N-1} s_m^*[l] h[j-l]$$

⁴ The newest version of the processor can be clocked up to $f_{\text{DSP}} = 1.4$ GHz.

It means that for each new sample in the strain signal $s[j]$ we have to compute

$$C_{\text{sample}} = N \quad (38)$$

multiply and add operations for each template m , instead of the logarithmic behavior of the frequency domain algorithm presented in 35. For this reason, the number of templates that can be computed gets greatly reduced, making this approach rather unsuitable.

9.4 TEMPLATE BANK

The last item we need to define now is which are the expected signals to look for in this matched filter search.

The shape of the incoming GW signals depend on various parameters, which are not known a priori (e.g. the masses of the binary's components). Thus, we have no choice but to filter the data through a set of templates, known as a template bank. Analysts fix an acceptable minimal match between any signal and its nearest template, so that the number of templates is minimal (to reduce computational cost) and the loss of ideal event rate is as low as possible. [73]

In the LIGO O1, the templates were placed using a combination of geometric and stochastic methods such that the loss in matched-filter SNR caused by its discrete nature was $\lesssim 3\%$. 249 077 templates were used to cover this parameter space, both by PyCBC and GstLAL [61]. In LIGO O2, the size of the GstLAL bank has been increased to 611 336 waveforms, mostly to extend the mass range up to $400 M_{\odot}$ and to be compliant with the $\lesssim 3\%$ loss constrain, after the detector sensitivity improvements; for this run, each pipeline used its own template bank.

The length of a template depends both on the sampling rate and on the starting frequency. The starting frequency f_{start} depends on the detector sensitivity: it is reasonable to choose the lowest frequency of the so-called detector sensitivity. In other words, there is very little information about the gravitational wave on the detector under that frequency, and applying matched filter for lower frequency signal does not help our search. For O2, GstLAL used $f_{\text{start}} = 15$ Hz, while PyCBC $f_{\text{start}} = 20$ Hz. In figure 38 we see how $> 80\%$ of the reconstructed SNR is provided by frequencies $f > 30$ Hz.

Under this assumptions, we can think to compute our templates starting by $f_{\text{start}} \leq 30$ Hz, and to fix a maximum length to 16,384 samples, or 1.6384 s, also to be consistent with the assumption in 36: this greatly reduce the length of the templates, but does not degrade significantly the SNR.

In figure 39 we show the bank used by GstLAL low-latency search during O2, reduced to those templates that matches this property. This way, the template bank is reduced from 611 336 to $N = 7375$ waveforms, but still contains most of the BBH events detected so far.




Figure 38: Cumulative SNR for GW170814 on the three detectors for the GstLAL low-latency search.

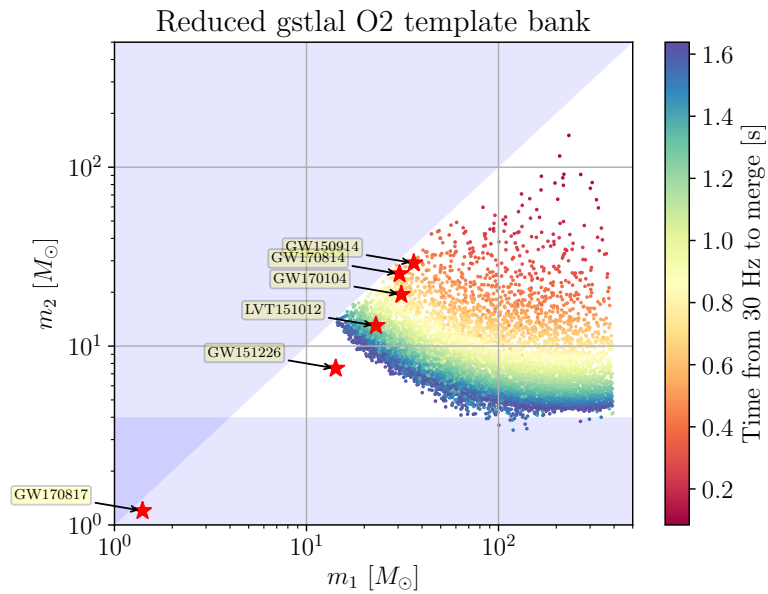


Figure 39: Template bank used by GstLAL low-latency search during O2, reduced to those templates that matches the property described in this section. The duration of the template is the time the signals enter the detection bandwidth (assumed at 30 Hz). The signals detected so far are reported as \star on the best fit values (error bars not provided here).

The duration of the templates has been evaluated using the *SimInspiralChirpTimeBound* function of the lalsuite: this routine actually estimates the time it will take for point-particle inspiral from a specified frequency to infinite frequency. The estimate is intended to be an over-estimate, so that the true inspiral time is always smaller (usually by few milliseconds) than the time this routine returns.

This cut roughly corresponds to excluding some regions in the (m_1, m_2) plot. The density of the remaining area has not been altered, so that the $\lesssim 3\%$ loss constrain still applies.

According to [37](#), at least two TMS320C6678 cores are needed for an on-line search based on this template bank.

10

IMPROVEMENTS OF THE EXISTING PIPELINES

Another approach to the usage of a DSP-based system for the on-line analysis of gravitational waves is to use such a dedicated hardware to improve an existing Compact Binary Coalescence (CBC) detection pipelines in use by the LVC pipeline. Among the ones at the time of writing, the GstLAL low-latency matched filter pipeline has the lowest latency, and medium computational cost, and so it is a reasonable representative choice.¹ Furthermore, because of its “streaming” architecture inherited from GStreamer, the filtering algorithm employed by the GstLAL pipeline is well-suited for conversion to a DSP-based implementation.

This is why it is reasonable to study how the usage of DSP-based boards could improve the GstLAL low-latency pipeline. In this chapter we describe in detail how we could improve the most computationally expensive parts of the pipeline with such a hardware acceleration.

10.1 HARDWARE

The computational power of the DSP used in the UDSPT boards have already been exhaustively described. Since the GstLAL is currently executed in the dedicated LVC clusters outside the observatory sites, there is no actual necessity to use a boards that hosts an analog front-end. At the time of this thesis, the TMS320C6678 is still the Texas Instruments DSP with the highest computational power, even if updated versions have been released, with the only main feature of the clock frequency f_s scalable up to 1.4 GHz with respect of the 1.0 GHz and 1.25 GHz currently used in ones shipped by the UDSPT boards.

Since 2012 Texas Instruments is producing also a new family of DSP: the most performing of this type is the model *66AK2H12*, that integrates a four-core ARM Cortex-A15 @ 1.4 GHz and an eight-core processor almost identical to the TMS320C6678. It could be useful in case we need to run a standard version of Linux on the ARM cores.

As of 2017, few commercial solutions deploying the TMS320C6678 and 66AK2H12 processors exist and are in production. For example, Advantech is producing devices as full-length PCI Express cards (like the Advantech DSP-8682 with 8 processors shown in figures 40 and 41) and as ATCA blades, integrating up to 20 DSPs per board. [74]

¹ The MBTA pipeline has lower computational cost but its detection sensitivity is not well measured due to lack of data from software injection tests.




Figure 40: Picture of Advantech DSP-8682 without the cooler unit.

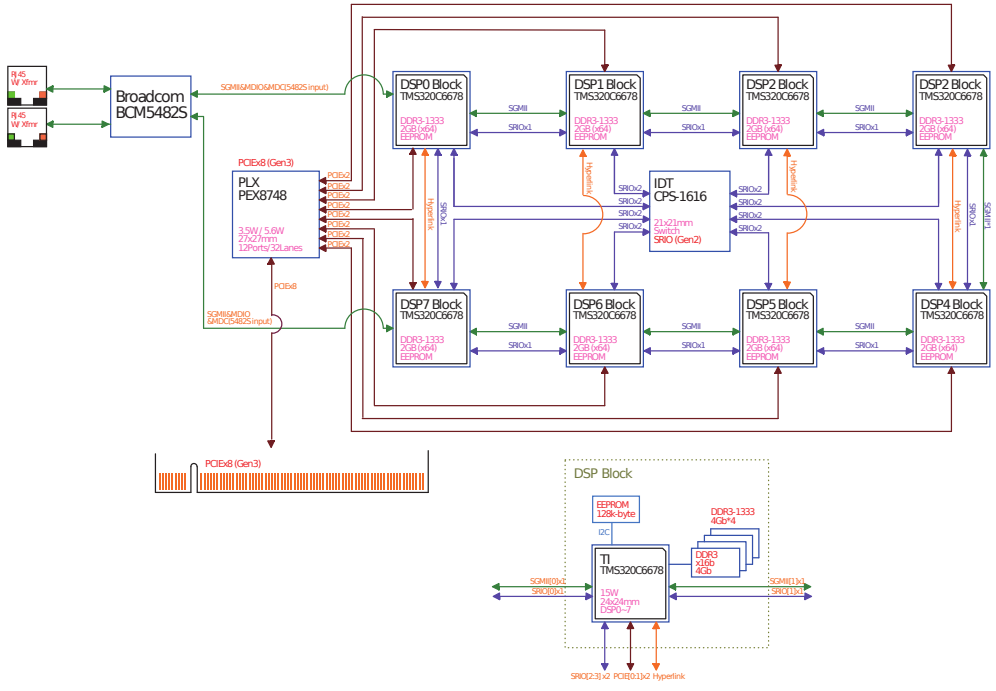


Figure 41: Board scheme of Advantech DSP-8682.

10.2 COMPUTATIONAL COST

The GstLAL detection pipeline consists of two principal subsystems. The first has been already described in the 8.2: implemented almost entirely in C using the GStreamer signal processing software framework, it transforms the strain time series data from the GW antennas into a sequence of “candidates”. Roughly speaking, these are times when the data took on the character of a CBC GW signal and are described by the time of the candidate and the physical parameters of the signal. The second subsystem, implemented almost entirely in Python, performs a statistical analysis of the ensemble of candidates to assess their significance; it estimates a False Alarm Probability (FAP) and FAR for each trigger, and reports the most statistically significant candidates to the user. The first subsystem, the “filtering” component, is the most computationally expensive.

A detailed analysis of the computational cost of the filtering component of the GstLAL-based detection pipeline is presented in [67, 68]. Some aspects of the algorithm and its configuration have changed since those documents were written, but the computational cost is still dominated by single-precision vector-vector inner products, either comprising FIR filters or the inner loops of matrix-matrix multiplications. In what follows we make two assumptions: that the total floating point operation count is dominated by multiply-add pairs and can be divided by two when considering hardware with a dedicated multiply-add operation, and that the dynamic range of the (whitened) strain and SNR time series is such that all calculations can be satisfactorily implemented with 16-bit fixed-point numbers. Given these assumptions, we can equate 1 GMAC s^{-1} (billion fixed-point multiply-accumulate operations), a standard operation count unit for DSPs, to 2 GFLOPS.

The O2 analysis configuration employed 661 335 complex-valued templates requiring a total of 949 GMAC to convert 1 s of strain data from a single antenna to 1 s of SNR time-series data. In other words, a computation rate of 949 GMAC/s/antenna was required to keep up with the data.²

The template bank is partitioned into groups of approximately 500 templates each, and each group must be processed as a unit. For the H1 instrument, but the counts are essentially the same for other instruments, the template group requiring the largest storage for its matrix and filter coefficients requires 5 691 300 coefficients while the smallest requires 179 950 coefficients. The average is 448 046 coefficients per group. In total, 591 420 280 coefficients are required to describe the entire H1 template bank.

10.2.1 Conclusions

Looking only at computational costs, the whole GstLAL real-time analysis could run in ~ 10 processors, or in only two Advantech DSP-8682 boards, if

² Computed using the `gstlal_inspiral_flopulator` tool and dividing the total operation count by 2.

we assume to be able to exploit half of the theoretical computational power of the DSPs.

We've explored also the idea to port only the most computationally expensive part to a DSP-based system, that includes the FIR, the matrix mixing and the various upsampling up to the SNR generation for each template. In this case, the bandwidth between the DSP system and a CPU running GstLAL, required to extract 661 335 16-bit stream at 2048 Hz, is 21 Gbps/antenna.

On the other hand, even if the DSP can be programmed with standard C code, it would take years of work to port the current GStreamer based code to the DSP architecture.

10.3 CONTINUOUS WAVES PIPELINES

The computational cost of the Continuous Wave (CW) search is extremely computationally expensive. The latest results for a deep all-sky search for periodic gravitational waves [75] have been produced through Einstein@Home, that was able to provide an average of ~ 2 PFLOPS [76]. This computational power is not currently available by the LVC clusters.

The DSP performances will improve in the next years as well as those of the CPUs, so it is reasonable to investigate if the usage of DSP-based board could allow to build a cluster with a computational power as high as that of Einstein@Home.

Nowadays a $O(\text{PFLOPS})$ cluster could be achieved by $O(10^4)$ TMS320C6678 boards. Of course, we have to take into account also the cost to develop the software.

Because CPUs will never be fast enough for the CW search, and DSPs will always be faster, the code for porting and developing new code could be a win, even if $O(1)$ years long.

10.4 COMPETITORS

Meanwhile, several competitors for high performance computing have been presented. The Intel is producing the Intel® Xeon Phi™ series since 2012, that consists of manycore processors³ intended for use in supercomputers, servers, and high-end workstations. The latest architecture of the series, named Knights Landing, contains a device able to run up to 288 threads in a single 72-Atom core CPU.

The existing GstLAL code is able to run on these processors without major modifications, making this choice more appropriate for future developments.

³ Manycore processors are multi-core processors designed for a high degree of parallel processing, containing a large number of simpler, independent processor cores.

We also note the appearance of hybrid devices that combines a traditional multi-core CPU with a high-end FPGA device. While there is no practical development path that can be identified at this time, a single device of this kind could provide approximately 1% of the processing power of Einstein@Home.

On August 14, 2017 the Advanced Virgo interferometer for the first time ever detected a gravitational wave event: GW₁₇₀₈₁₄, the first triple coincidence of Binary Black Hole (BBH) merger done by the two LIGO and the Virgo detectors. After three days, another event has been detected by the three interferometers: GW₁₇₀₈₁₇, the first event ever from a Binary Neutron Star (BNS) system and with a coincident electromagnetic counterpart. The signal SNR could be cumulated down to ~ 30 Hz thanks to the very good seismic isolation of the Advanced LIGO and Advanced Virgo detector

The Advanced Virgo Superattenuators are complex mechanical suspensions used to insulate the optical elements of the interferometer from the seismic noise, and to extend the bandwidth of the detector down to few tens of hertz, where most of the information of the detected events lies. They are operated by the digital Suspension Control System (SCS), running on a DSP-based distributed computing infrastructure.

Hundreds of electronic devices are controlled by a system in charge of provide an automation layer, as well as tools to monitor the status of the Superattenuators in real time. We have presented the new *Software Supervisor*, the first personal contribution in this thesis: it is a Supervisory Control And Data Acquisition (SCADA) system developed for the Advanced Virgo SCS entirely based on TANGO, an open source toolkit for controlling hardware or software widely used in the scientific community.

The Software Supervisor is based on 45 servers that provide access to more than 300 physical and abstract devices. It provides also clients that are used to monitor the status of the electronics, but also to perform maintenance tasks. We have presented here the most significant Graphical User Interface (GUI) clients developed with *Python* and with the graphics library *Qt4*: one of them, largely used during maintenance, allows to use a slow control loop to center the rest position of the Superattenuator top stage using stepper motors. This is a situation where a control system that has to interact with different types of hardware (stepper motor controllers and DSP-based boards in this case) benefits from the usage of the Software Supervisor.

The Software Supervisor is also part of the automation of the whole experiment. It is monitoring the SCS since the installation of new Advanced Virgo electronics; it supported the commissioning, and allowed to have a system ready for the LIGO-Virgo joint observation run in August 2017, when the two milestone events were detected.

The second personal contribution presented in this thesis is a feasibility study about the usage of a DSPs-based distributed computing system for data analysis. An on-line low-latency search for gravitational waves

would take advantage of computational power on a system deeply integrated within the experiment like the Advanced Virgo SCS. The lowest latency of the existing searches for signals from compact binary mergers is currently ~ 16 s, achieved by the GstLAL pipeline. Most of this delay is due to the reconstruction of the gravitational wave strains and to the data transfer from the interferometer sites to the computing farms. We have shown how a simple matched filter search performed on the SCS could dramatically reduce this time to few seconds, or even better in case of a different architecture of the pipeline.

BIBLIOGRAPHY

- [1] B. P. Abbott *et al.*, “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.*, vol. 116, no. 6, p. 061102, 2016.
- [2] A. Einstein, “Approximative Integration of the Field Equations of Gravitation,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1916, pp. 688–696, 1916.
- [3] A. Einstein, “Über Gravitationswellen,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1918, pp. 154–167, 1918.
- [4] K. Schwarzschild, “On the gravitational field of a mass point according to Einstein’s theory,” *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, vol. 1916, pp. 189–196, 1916.
- [5] C. Patrignani *et al.*, “Review of Particle Physics,” *Chin. Phys.*, vol. C40, no. 10, p. 100001, 2016.
- [6] U. Le Verrier, *Détermination nouvelle de l’orbite de Mercure et de ses perturbations*. 1843.
- [7] I. R. Kenyon, *General Relativity*. 1990.
- [8] S. Weinberg, *Gravitation and cosmology: principles and applications of the general theory of relativity*. John Wiley & Sons, Inc., 1972.
- [9] F. W. Dyson, A. S. Eddington, and C. Davidson, “A Determination of the Deflection of Light by the Sun’s Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919,” *Phil. Trans. Roy. Soc. Lond.*, vol. A220, pp. 291–333, 1920.
- [10] R. Feynman, R. Leighton, and M. Sands, *The Feynman Lectures on Physics - Mainly Electromagnetism and Matter*. Addison-Wesley Publishing Company, Inc., 1964.
- [11] R. A. Hulse and J. H. Taylor, “Discovery of a pulsar in a binary system,” *Astrophys. J.*, vol. 195, pp. L51–L53, 1975.
- [12] J. H. Taylor and J. M. Weisberg, “A new test of general relativity: Gravitational radiation and the binary pulsar PS R 1913+16,” *Astrophys. J.*, vol. 253, pp. 908–920, 1982.
- [13] J. Weisberg, D. Nice, and J. Taylor, “Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16,” *Astrophys. J.*, vol. 722, pp. 1030–1034, 2010.

- [14] J. M. Weisberg and Y. Huang, "Relativistic Measurements from Timing the Binary Pulsar PSR B1913+16," *Astrophys. J.*, vol. 829, no. 1, p. 55, 2016.
- [15] P. Peters and J. Mathews, "Gravitational radiation from point masses in a Keplerian orbit," *Phys.Rev.*, vol. 131, pp. 435-439, 1963.
- [16] J. Weber, "Detection and Generation of Gravitational Waves," *Phys. Rev.*, vol. 117, pp. 306-313, 1960.
- [17] P. Astone *et al.*, "IGEC2: A 17-month search for gravitational wave bursts in 2005-2007," *Phys. Rev.*, vol. D82, p. 022003, 2010.
- [18] M. E. Gertsenshtein and V. I. Pustovoit, "On the Detection of Low Frequency Gravitational Waves," *Sov. Phys. JETP*, vol. 16, p. 433, 1962.
- [19] R. Weiss, "Electromagnetically coupled broadband gravitational antenna," techreport MIT Report No. 105, Quarterly Report of the Research Laboratory for Electronics, 1972.
- [20] A. Brilliet and A. Giazotto, "Virgo Project Technical Report," techreport, 1989.
- [21] R. W. P. Drever, F. J. Raab, K. S. Thorne, R. Vogt, and R. Weiss, "Laser Interferometer Gravitational-wave Observatory (LIGO) Technical Report," techreport, 1989.
- [22] B. P. Abbott *et al.*, "GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence," *Phys. Rev. Lett.*, vol. 119, no. 14, p. 141101, 2017.
- [23] B. P. Abbott *et al.*, "Properties of the Binary Black Hole Merger GW150914," *Phys. Rev. Lett.*, vol. 116, no. 24, p. 241102, 2016.
- [24] B. P. Abbott *et al.*, "GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence," *Phys. Rev. Lett.*, vol. 116, no. 24, p. 241103, 2016.
- [25] B. P. Abbott *et al.*, "Binary Black Hole Mergers in the first Advanced LIGO Observing Run," *Phys. Rev.*, vol. X6, no. 4, p. 041015, 2016.
- [26] B. P. Abbott *et al.*, "GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2," *Phys. Rev. Lett.*, vol. 118, no. 22, p. 221101, 2017.
- [27] L. P. Singer *et al.*, "Going the Distance: Mapping Host Galaxies of LIGO and Virgo Sources in Three Dimensions Using Local Cosmography and Targeted Follow-up," *Astrophys. J.*, vol. 829, no. 1, p. L15, 2016.
- [28] B. P. Abbott *et al.*, "Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A," *Astrophys. J.*, vol. 848, no. 2, p. L13, 2017.

- [29] A. Goldstein *et al.*, “An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A,” *Astrophys. J.*, vol. 848, no. 2, p. L14, 2017.
- [30] B. P. Abbott *et al.*, “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral,” *Phys. Rev. Lett.*, vol. 119, no. 16, p. 161101, 2017.
- [31] “Multi-messenger Observations of a Binary Neutron Star Merger,” *Astrophys. J.*, vol. 848, p. L12, 2017.
- [32] B. P. Abbott *et al.*, “A gravitational-wave standard siren measurement of the Hubble constant,” *Nature*, 2017.
- [33] T. Accadia *et al.*, “Virgo: a laser interferometer to detect gravitational waves,” *JINST*, vol. 7, p. P03012, 2012.
- [34] P. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors*. World Scientific Publishing Company, Incorporated, 1994.
- [35] V. Boschi, *Modeling and Simulation of Seismic Attenuation Systems for Gravitational Wave Interferometers*. PhD thesis, Università di Pisa, 2010.
- [36] The Virgo Collaboration, “Advanced Virgo Technical Design Report,” Tech. Rep. VIR-0128A-12, CNRS and INFN, Apr 2012.
- [37] T. Accadia *et al.*, “Performance of the Virgo interferometer longitudinal control system during the second science run,” *Astropart. Phys.*, vol. 34, pp. 521–527, 2011.
- [38] P. Amico, L. Carbone, C. Cattuto, L. Gammaitoni, M. Punturo, F. Travasso, and H. Vocca, “Thermal noise limit in the Virgo mirror suspension,” *Nucl. Instrum. Meth.*, vol. A461, pp. 297–299, 2001.
- [39] C. M. Caves, “Quantum Mechanical Noise in an Interferometer,” *Phys. Rev.*, vol. D23, pp. 1693–1708, 1981.
- [40] J. Aasi *et al.*, “Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light,” *Nature Photon.*, vol. 7, pp. 613–619, 2013.
- [41] M. Leonardi, *Development of a squeezed light source prototype for Advanced Virgo*. PhD thesis, Università di Trento, 2016.
- [42] R. X. Adhikari, “Gravitational Radiation Detection with Laser Interferometry,” *Rev. Mod. Phys.*, vol. 86, p. 121, 2014.
- [43] J. Abadie *et al.*, “Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3,” *Astrophys. J.*, vol. 760, p. 12, 2012.

- [44] G. Ballardin *et al.*, “Measurement of the VIRGO superattenuator performance for seismic noise suppression,” *Rev. Sci. Instrum.*, vol. 72, pp. 3643–3652, 2001.
- [45] L. Naticchioni, “The payloads of Advanced Virgo: current status and upgrades,” Tech. Rep. VIR-0781A-17, 2017.
- [46] A. Gennai, “Real Time Control of Suspended Test Masses in Advanced Virgo Laser Interferometer,” in *20th IEEE-NPSS Real Time Conference*, 2016.
- [47] N. Arnaud *et al.*, “The global control of the Virgo experiment,” *Nucl. Instrum. Meth.*, vol. A550, pp. 467–489, 2005.
- [48] N.A.T. GmbH, *NAT-MCH User’s Manual*, 1.31 ed., 08 2016.
- [49] L. Trozzo. PhD thesis, Università di Siena, 2017.
- [50] Texas Instruments, “TMS320C6678 Multicore Fixed and Floating-Point Digital Signal Processor (Rev. E),” tech. rep., 03 2014.
- [51] The MathWorks, “Simulink help,” 2017.
- [52] F. Acernese, P. Amico, M. Alshourbagy, F. Antonucci, S. Aoudia, P. Astone, S. Avino, D. Babusci, G. Ballardin, F. Barone, *et al.*, “Data acquisition system of the virgo gravitational waves interferometric detector,” *IEEE Transactions on Nuclear Science*, vol. 55, no. 1, pp. 225–232, 2008.
- [53] J. M. Chaize, A. Gotz, W. D. Klotz, J. Meyer, M. Perez, and E. Taurel, “TANGO - an object oriented control system based on CORBA,” *Conf. Proc.*, vol. C991004, pp. 475–479, 1999. [475(1999)].
- [54] C. Arnault and P. Massart, “Cm: a multitask communication package,” Tech. Rep. VIR-MAN-LAL-5100-112, Virgo internal note, 1998.
- [55] The TANGO Team, *The TANGO Control System Manual*, 06 2013.
- [56] G. Cerretani, “TANGO deployment at VIRGO,” 31st TANGO Collaboration Meeting VIR-0430A-17, 2017.
- [57] F. Carbognani and B. Swinkels, “Interferometer automation with Python @ 21 Jul 2016 DIM,” Tech. Rep. VIR-0366A-16, 2016.
- [58] J. Graef Rollins, “Distributed state machine supervision for long-baseline gravitational-wave detectors,” *Rev. Sci. Instrum.*, vol. 87, no. 9, p. 094502, 2016.
- [59] J. Maciejowski, *Multivariable feedback design*. Electronic systems engineering series, Addison-Wesley, 1989.
- [60] B. P. Abbott *et al.*, “Localization and broadband follow-up of the gravitational-wave transient GW150914,” *Astrophys. J.*, vol. 826, no. 1, p. L13, 2016.

- [61] B. P. Abbott *et al.*, “GW150914: First results from the search for binary black hole coalescence with Advanced LIGO,” *Phys. Rev.*, vol. D93, no. 12, p. 122003, 2016.
- [62] B. P. Abbott *et al.*, “Observing gravitational-wave transient GW150914 with minimal assumptions,” *Phys. Rev.*, vol. D93, no. 12, p. 122004, 2016.
- [63] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, “A coherent method for detection of gravitational wave bursts,” *Classical and Quantum Gravity*, vol. 25, no. 11, p. 114029, 2008.
- [64] T. Adams, “Low latency search for compact binary coalescences using MBTA,” in *Proceedings, 50th Rencontres de Moriond Gravitation : 100 years after GR: La Thuile, Italy, March 21-28, 2015*, pp. 327–330, 2015.
- [65] B. P. Abbott *et al.*, “Upper Limits on the Rates of Binary Neutron Star and Neutron Star–black Hole Mergers From Advanced Ligo’s First Observing run,” *Astrophys. J.*, vol. 832, no. 2, p. L21, 2016.
- [66] S. Sachdev, “Status of Advanced LIGO gstlal search for compact binary coalescences,” Tech. Rep. G1601025-v2, Jun 2016.
- [67] C. Hanna and A. Pace, “gstlal_inspiral: code performance and scaling June 2016 Update,” tech. rep., 06 2016. LIGO-T1600193.
- [68] C. Hanna and K. Cannon, “gstlal_inspiral: Code performance and scaling,” tech. rep., 04 2015. LIGO-T1400542-v13.
- [69] T. Accadia *et al.*, “Reconstruction of the gravitational wave signal $h(t)$ during the Virgo science runs and independent validation with a photon calibrator,” *Class. Quant. Grav.*, vol. 31, p. 165013, 2014.
- [70] A. Viets, M. Wade, A. Urban, and S. Kandhasamy, “Making $h(t)$ for Advanced LIGO,” Tech. Rep. LIGO-P1700236, 2017. LIGO-P1700236-v5.
- [71] 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,” *Phys. Rev.*, vol. D85, p. 122006, 2012.
- [72] P. Welch, “The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms,” *IEEE Transactions on audio and electroacoustics*, vol. 15, no. 2, pp. 70–73, 1967.
- [73] T. Cokelaer, “Gravitational waves from inspiralling compact binaries: Hexagonal template placement and its efficiency in detecting physical signals,” *Phys. Rev.*, vol. D76, p. 102004, 2007.
- [74] Advantech, “Dsp-8682 datasheet,” tech. rep., 2014.

- [75] B. P. Abbott *et al.*, “Results of the deepest all-sky survey for continuous gravitational waves on LIGO S6 data running on the Einstein@Home volunteer distributed computing project,” *Phys. Rev.*, vol. D94, no. 10, p. 102002, 2016.
- [76] “BOINCStats webpage.”.