Dated: 29 December 2021
The Virgo Collaboration + additional authors
. To be submitted to CQG
Virgo INTERNAL DOCUMENT – NOT FOR PUBLIC DISTRIBUTION

The Virgo O3 run and the impact of the environment

The Virgo Collaboration + a few additional authors

Abstract. Sources of geophysical noise (such as wind, sea waves and earthquakes) or of anthropogenic noise impact ground-based gravitational-wave interferometric detectors, causing transient sensitivity worsening and gaps in data taking. During the one year-long third Observing Run (O3: from April 01, 2019 to March 27, 2020), the Virgo Collaboration collected a statistically significant dataset, used in this article to study the response of the detector to a variety of environmental conditions. We correlated environmental parameters to global detector performance, such as observation range, duty cycle and control losses. Where possible, we identified weaknesses in the detector that will be used to elaborate strategies in order to improve Virgo robustness against external disturbances for the next data taking period, O4, currently planned to start at the end of 2022. The lessons learned could also provide useful insights for the design of the next generation of ground-based interferometers.

Contents

1. Introduction

The past decade has seen the ramp-up of the second-generation ('Advanced') earth-based gravitational-wave (GW) detectors. Design improvements and technological upgrades have paved the way to the first direct detections of GWs by the global network made up of the two aLIGO instruments [1] (located in the USA: Hanford, WA and Livingston, LA) and of the Advanced Virgo detector [2] (located in Cascina, Italy). The main results achieved by the LIGO Scientific Collaboration and the Virgo Collaboration – recently joined by the KAGRA collaboration whose detector [3], located under the Kamioka mountain in Japan, is nearing completion – include the first detection of a binary black hole merger (GW150914 [4]); the first detection of a binary neutron star merger (GW170817 [5]) that lead to the birth of multi-messenger astronomy with GW [6]; and now dozens of detections of compact binary mergers that add up in a GW Transient Catalogue regularly updated [7, 8, 9]. These detections contribute to opening a new window onto the Universe by providing insights to the populations of compact objects and the binary merger rates [10]; they also allow scientists to perform stringent tests of general relativity [11] in a new regime of gravitation never probed before.

The operation of ground-based GW detectors is organized into successive steps forming a recurring sequence over the years: upgrades; commissioning and sensitivity improvement (the so-called *noise hunting* phase); data-taking periods called observing runs (or simply runs and labelled On). So far there have been three runs for the global network of advanced detectors.

- O1 (09/2015 – 01/2016) with only the two LIGO detectors taking data;
- O2 (11/2016 – 08/2017) with Virgo joining LIGO on August 01, 2017;
- finally O3 (04/2019 – 03/2020), that saw the three detectors take data jointly during 11 months in total: 6 months first (a run called O3a), followed by a 1-month break (October 2019) and then another period of 5 months of data taking (O3b), interrupted about a month earlier than expected due to the worldwide COVID-19 pandemic.

The above listing shows that the O3 run was the first *long* data-taking period for the Advanced Virgo detector. Therefore, we have used the wealth of unprecedented data collected during this year to make an in-depth analysis of the instrument performance. In this article, we study the impact of the environment on Advanced Virgo, along the lines of previous publications from Virgo [12], LIGO [13] or KAGRA [14]. We focus on various types of seismic noises, on earthquakes and on bad weather periods. We also briefly investigate the effect of other possible disturbances: magnetic noise, lightning and cosmic muons. Our goal is threefold: to quantify how the Virgo sensitivity and duty cycle depend on these external parameters; to use this knowledge to prepare the next run, O4, scheduled to start in the second semester 2022; finally, to build experience for future GW detectors, in particular for the Einstein Telescope project [15].

The Virgo detector is located in Italy at EGO, the European Gravitational Observatory, in the municipality of Cascina. The EGO site is in the countryside, about 12 km south-east of Pisa and about 17 km east of the Tyrrhenian coast. Virgo is not far from some industrial and commercial sites that can generate noise. Within 7 km from EGO there are: elevated highways, railway tracks, wind turbines, earth quarries, electroducts and the Pisa airport. To avoid pressure waves potentially shaking the ground, a no-fly zone has been enforced in a cylindrical volume (600 m radius and height) above each of the Virgo experimental buildings.

Advanced Virgo is a power-recycled Michelson interferometer with Fabry-Perot cavities in its 3 km-long arms. All core optics are suspended to long suspensions, called the superattenuators [16], that have a twofold use: first, to isolate as much as possible the mirrors from seismic motions (both vertical and longitudinal), and then to control very accurately their positions in all six degrees of freedom. Many feedback systems are used to bring the detector to its working point and maintain it there [17, 18]. This state – the same for O2 and O3: the Michelson interferometer on a dark fringe, the Fabry-Perot and power recycling cavities in resonance – is the only one in which the detector is sensitive to the passing of GWs.

During a run, priority is obviously given to taking data of quality good enough to be included in physics analysis. In that case, Virgo is said to be in *Science mode*. During O3, the average duty cycle in Science mode has been around 76% [19], with the remaining time almost equally divided into three categories.

- Control acquisition and adjustment phases, to restore the working point and restart taking data in Science mode;
- Recurring controlled actions on the detector: maintenance (usually a few hours on Tuesday mornings local time), calibration (usually every Wednesday evening) or commissioning (measurements, working point tuning or tailored improvements: sessions organized when the need arises);
- Problems preventing a smooth running of the detector.

The article is organized as follows. Section 2 describes the environmental monitoring of the Virgo detector during the O3 run. Section 3 is dedicated to the different seismic noise contributions (either natural or human-related): how to disentangle them, how to monitor them and what their impacts on the detector are in terms of sensitivity and duty cycle. Section 4 provides an analysis of the impact of earthquakes on the detector. Section 5 studies the impact of bad weather on data quality and duty cycle, disentangling contributions from sea activity and wind. Section 6 goes through other environment impacts: magnetic noise, lightning and a study of the cosmic muon rate on the Virgo central building. Then, Section 7 concludes this article by opening outlooks to the future O4 run. Finally, Appendix A provides a detailed and quite complete classification of the control losses during the O3 run. Although that study has a scope

broader than the present article, it is included here for reference and also because its results were used, in particular to find out which control losses were due to earthquakes.

2. The Virgo environmental monitoring during O3

The Virgo detector is equipped with a large set of probes used to monitor the conditions of the surrounding environment. Since these conditions can influence the detector response, or even mimic a GW event, it is very important to track their evolution, to assess the right working condition of the detector or to use them as veto against possible fake signals.

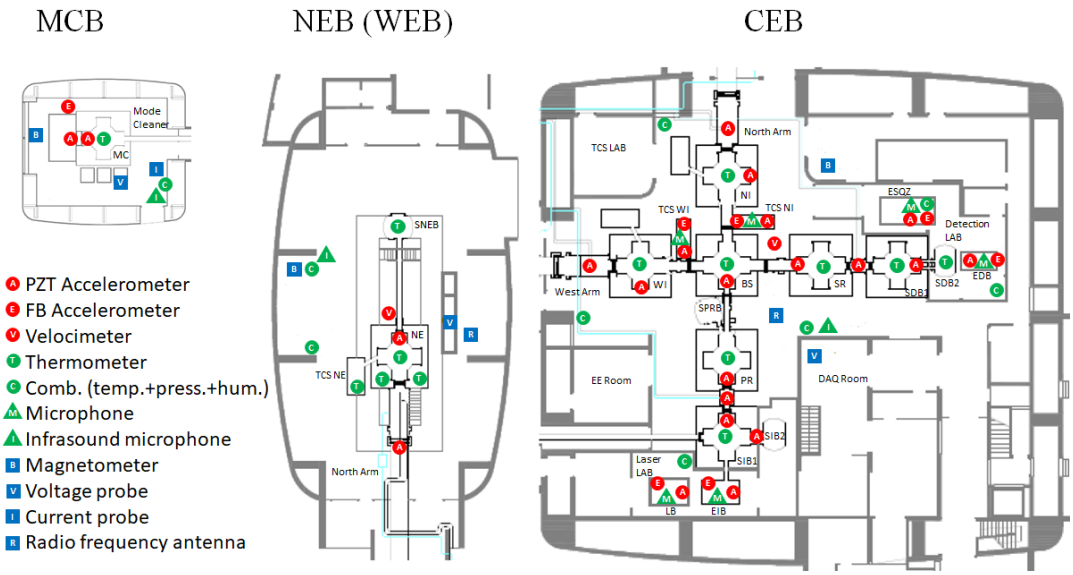

Figure 1: Location of the probes used for the Virgo environmental monitoring system. Maps of most relevant building are shown: left MCB, middle NEB, right CEB. The WEB is very similar to NEB and is not shown.

The set of probes and their conditioning electronics constitute the Environmental Monitoring System (EMS). The EMS was initially composed by a few tens of environmental probes (EPs) [20] and then improved during the detector upgrades that occurred in the past years. During O3, the total number of channels belonging to EMS was about 420.

The EMS is also helpful to understand the origin of some noise sources affecting the detector sensitivity. Indeed it was largely used during the commissioning phase following each detector upgrade, to recover and improve the Virgo performance in terms of sensitivity and duty cycle [12].

Data acquired for EMS can be grouped in two classes depending on the sample rate used for the different EPs. High-rate, or fast class, includes those EPs acquired at rate up to 20 kHz like seismometers, piezoelectric (PZT) accelerometers, force balance

(FB) accelerometer, magnetometers, microphones, voltage and current sensors, radio-frequency (RF) antennas, while low-rate or slow class includes temperature, pressure, humidity, weather and lightning probes, acquired at 1 Hz rate.

Type	Model	Frequency Band
Seismometer	Guralp CMG-40T	0.01 – 50 Hz
FB Accelerometer	Kinemetrics FBA ES-T	0.1 – 200 Hz
PZT Accelerometer	Wilcoxon 731-207 or PCB 393B12	1 Hz – 1 kHz
Magnetometer	Metronics MFS-06 or MFS-06e	0.1 mHz – 10 kHz
Microphone	Brüel & Kjær 4190 or 4193	0.1 – 10 kHz
RF antenna	AAS STA 5 A/D/0.01-100	10 kHz – 100 MHz
Voltage probe	Talema 0015P1-2-009	DC – 10 kHz
Current probe	IME 0015P1-2-009	DC – 10 kHz
Temperature probe	Analog Device AD590	DC – 0.5 Hz
Humidity probe	Honeywell HIH-5031-001	DC – 0.5 Hz
Pressure probe	NXP MPXA4115A6U	DC – 0.5 Hz
Weather station	Davis Advantage Pro 2	DC – 0.3 Hz
Lightning detector	Boltek LD 250	DC – 0.5 Hz

Table 1: Characteristics of the Virgo environmental probes used during O3.

The main characteristics (type, model and frequency band) of the EPs in use during O3 are listed in Table 1. Figure 1 shows the arrangement of the EPs inside the main Virgo buildings. Most probes are located in the experimental halls of the relevant buildings of the detector: Central Building (CEB), North and West End Buildings (NEB and WEB) and Mode Cleaner Building (MCB). Usually, the probes are in contact with critical elements of the detector, like the walls of the vacuum chambers containing the test mass suspensions, or the optical benches hosting the laser injection and GW detection systems. Figure 2 shows a bird eye's view of the Virgo detector at EGO, with an emphasis on the location of the buildings that are identified in this article.

Few probes are placed outside the buildings and are not shown in the schematics, namely the weather station, the lightning detector and two additional magnetometers. These two low-noise induction coil magnetometers are deployed at 0.5 m depth in the soil, at about 100 m from the CEB, oriented along the geographic North and West directions. Their data are shared in real time with the EM antenna network "Radio waves below 22 kHz" [21].

3. Seismic noise

In this section we introduce the main sources of seismic noise at EGO. They are disentangled and monitored by examining seismic probes in specific frequency bands.

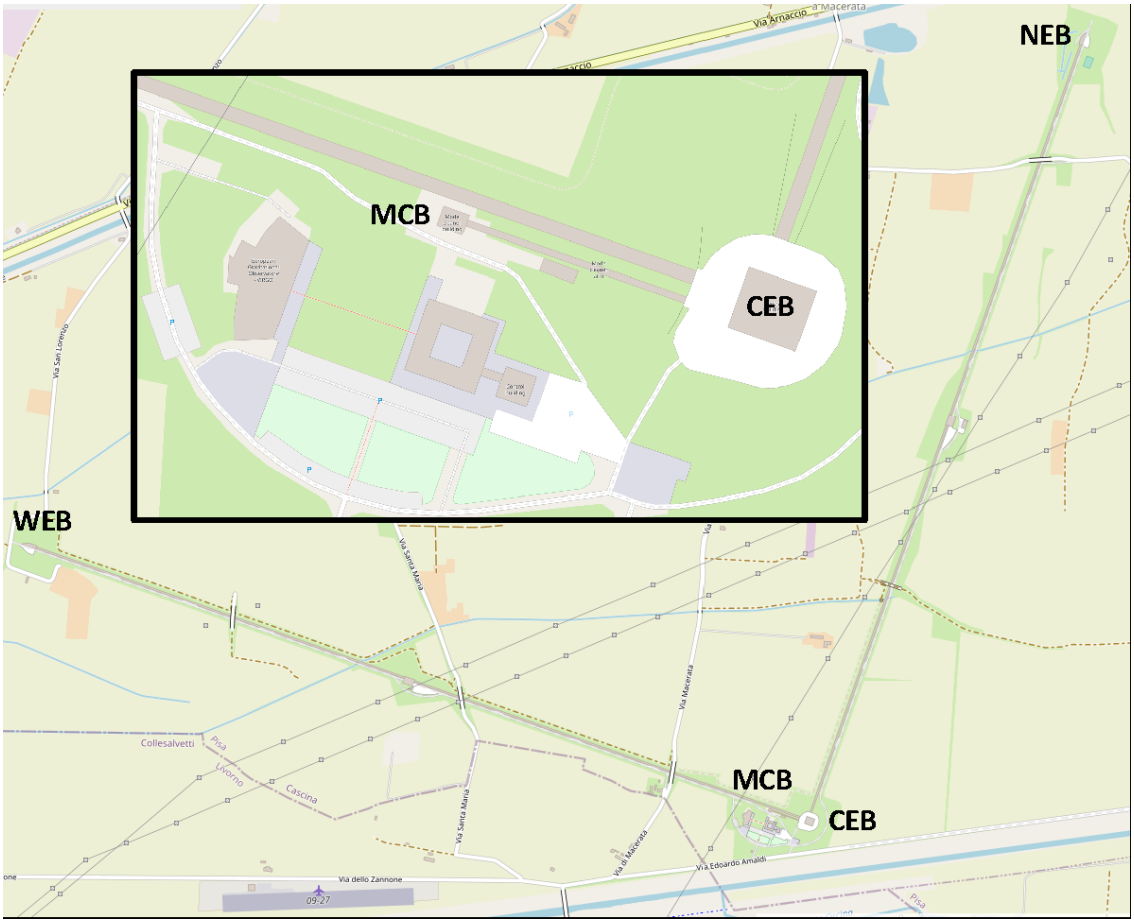

Figure 2: Map [22] of the EGO site showing the Virgo detector and in particular the location of the main buildings identified in the text. The central insert shows a zoom around the interferometer vertex, with the CEB and MCB highlighted. The Mode-Cleaner cavity is 144 m-long, while the Virgo arms are 3-km long.

We provide a statistical description of the noise and evidence its main recurring features. Then, we describe how they impacted on the detector during the O3 run.

3.1. The seismic frequency bands and their evolution during the O3 run

The seismic wavefield at EGO, the site of the Virgo detector, is the sum of several sources [24]. Seismic spectrum variability during the O3 run is illustrated in Fig. 3. The largest contribution to seismic ground motion in the frequency range between 0.1 Hz and 1 Hz, referred to as *microseism*, is due to the interaction between shallow water sea waves and the bottom of the sea [25, 26]. At EGO, the prevailing microseimic peak is around 0.35 Hz.

Figure 4 shows the time evolution of microseism during the O3 run, while Fig. 5 shows the corresponding cumulative distribution, split by season. Microseism intensity follows seasonal variations, being larger in fall and winter, due to the stronger wind and sea

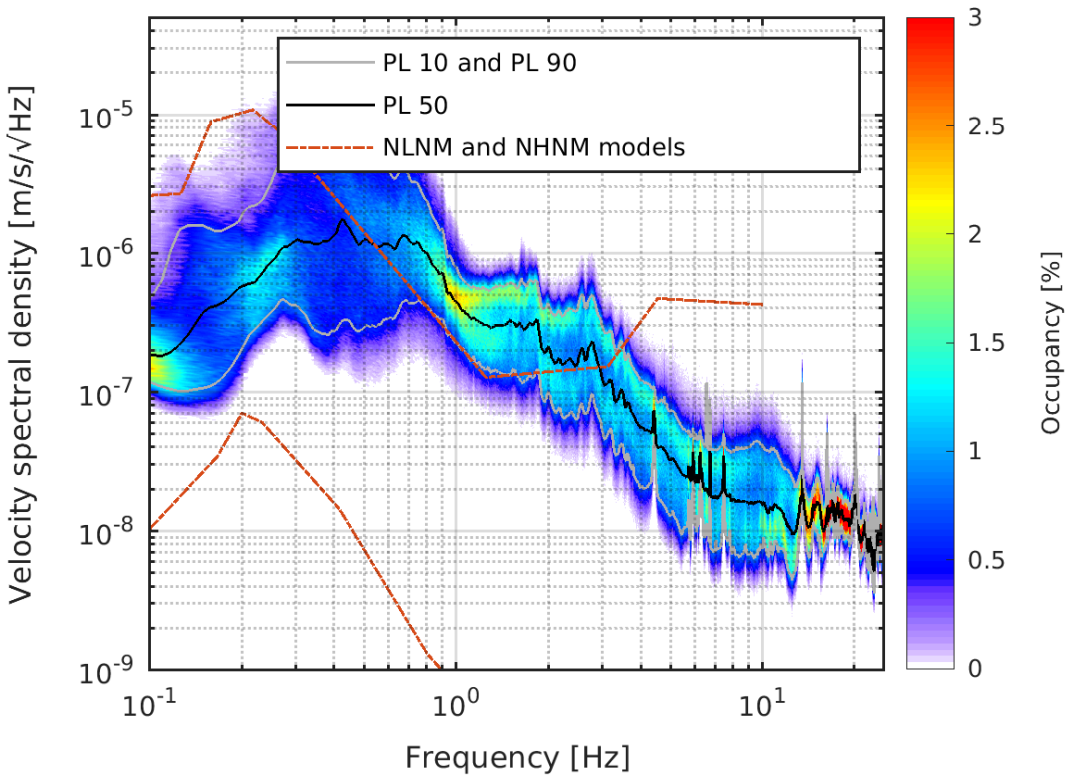

Figure 3: Variability of horizontal velocity of the Virgo NEB ground floor during O3. The quantity shown is the 2D histogram of the E-W velocity amplitude spectral density computed for the whole dataset recorded during O3 divided into 128 second long chunks. All intereferometer maintenance periods are excluded from the computation. The color scale indicates the percent occupancy of histogram bins. The superposed continuous curves show different percentile levels (labelled *PL* on the plot): 10% (gray), 50% (black) and 90% (gray as well). The two red dashed curves correspond to the Peterson low-noise ('NLNM') and high-noise ('NHNM') models [23].

activity.

Above 1 Hz, anthropogenic sources dominate the spectrum. Heavy vehicles (trucks and alike) on ∼1 km distant elevated roads are the prevailing source of seismic noise in the 1-10 Hz band [24].

As illustrated in Fig. 6, the RMS of seismic noise in the 1-5 Hz band follows a working day/night cycle with higher levels during working hours (from 8:00 to 17:00 local time – LT), with small reduction during lunch break (12:00-14:00 LT) and minima during weekends and holidays. The blue curve, used as reference, covers the whole O3 run. The green curve is based on a 4-week period, from Monday 16 December, 2019 to Sunday 12 January, 2020: the noise reduction during the two consecutive Wednesdays, Christmas 2019 and the New Year's Day 2020, is quite impressive. A significant reduction of the anthropogenic noise is also visible during the Spring 2020 lockdown in Italy, due to the

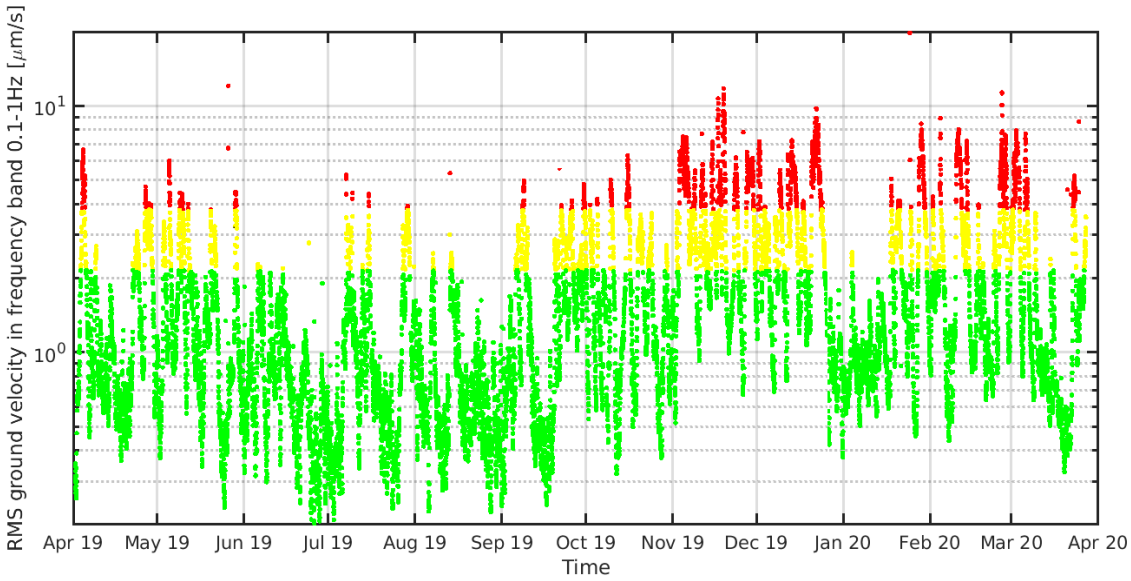

Figure 4: Evolution during O3 of seismic RMS in the 0.1 to 1 Hz frequency band. Data colored in yellow and red exceed the 75^{th} and 90^{th} percentile, respectively.

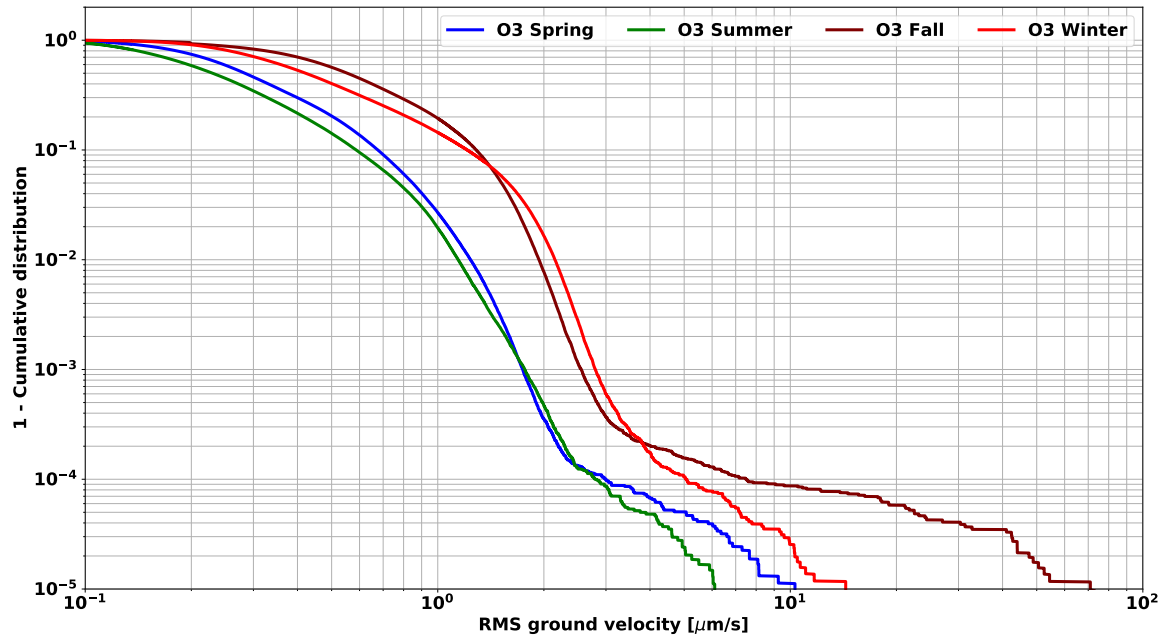

Figure 5: Cumulative distribution of microseism in the frequency band 0.1-1 Hz (dominated by sea activity), measured at EGO during each season in 2019-2020.

COVID-19 pandemic (red curve, covering a 8-week period from 09 March to 03 May). That decrease is smaller than for the Christmas and New Year holidays but it is more global as it is visible for all days of the week.

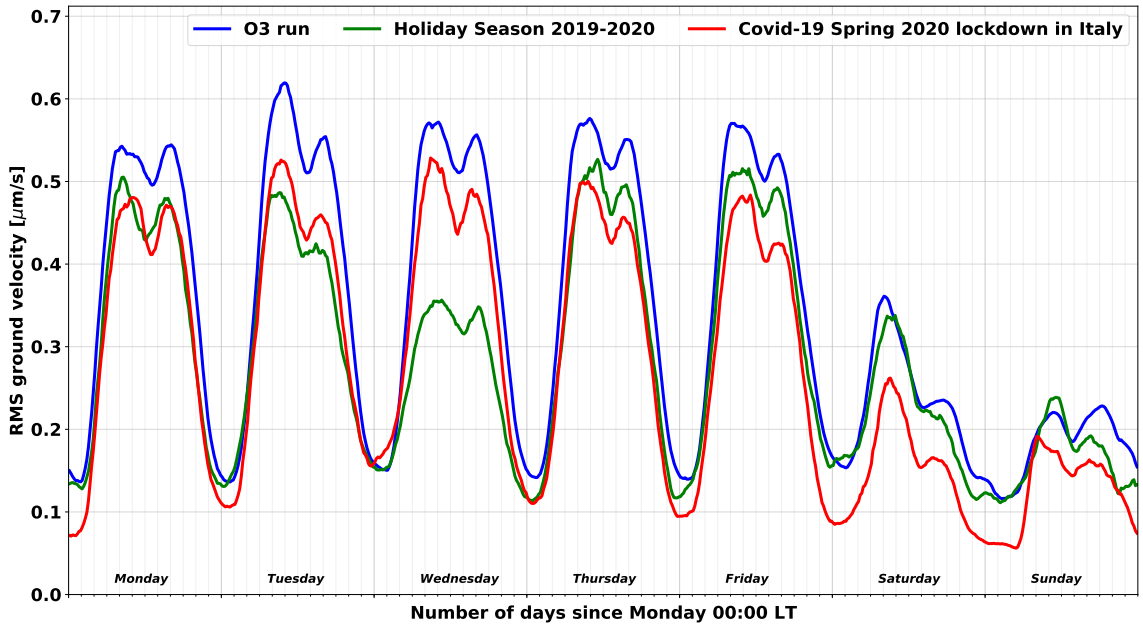

Figure 6: Average evolution on a weekly basis of the seismic anthropogenic noise (frequency band: 1-5 Hz) measured at EGO during different times in 2019-2020.

Finally, above 10 Hz, the dominant seismic contribution is generated locally: vehicles in nearby and on-site roads, agricultural work on neighbouring land, etc. Figure 7 shows the average day-night variations, computed in the 10-40 Hz band on a weekly basis: in blue during the O3 run; in magenta during the 1-month commissioning break (October 2019) separating the two halves of O3; finally in orange for the second semester of 2020, during which hardware upgrades and construction or infrastructure works for the Advanced Virgo+ project [27] took place.

The common feature between the three curves is the dominant peak on Tuesday mornings, the usual slot used for the weekly maintenance of the Virgo detector. This activity includes in particular the refilling of Nitrogen‡ tanks by heavy trucks coming on-site, and the possibility to have people moving around and working inside experimental areas whose access is forbidden during data taking periods. The on-site seismic noise level was slightly higher during the commissioning break compared to the O3 run, but not by much: that 1-month shutdown was not long enough to allow for invasive works that could have jeopardized the restart of data taking on November 01, 2019, alongside the two LIGO detectors. On the other hand, on-site activites are more evenly distributed over working days during the post-O3 upgrade. Though, activities were the lowest on

‡ Liquid Nitrogen is used to cool down the Advanced Virgo cryotraps [2].

weekends during that period because of site access restrictions, enforced because of the pandemic.

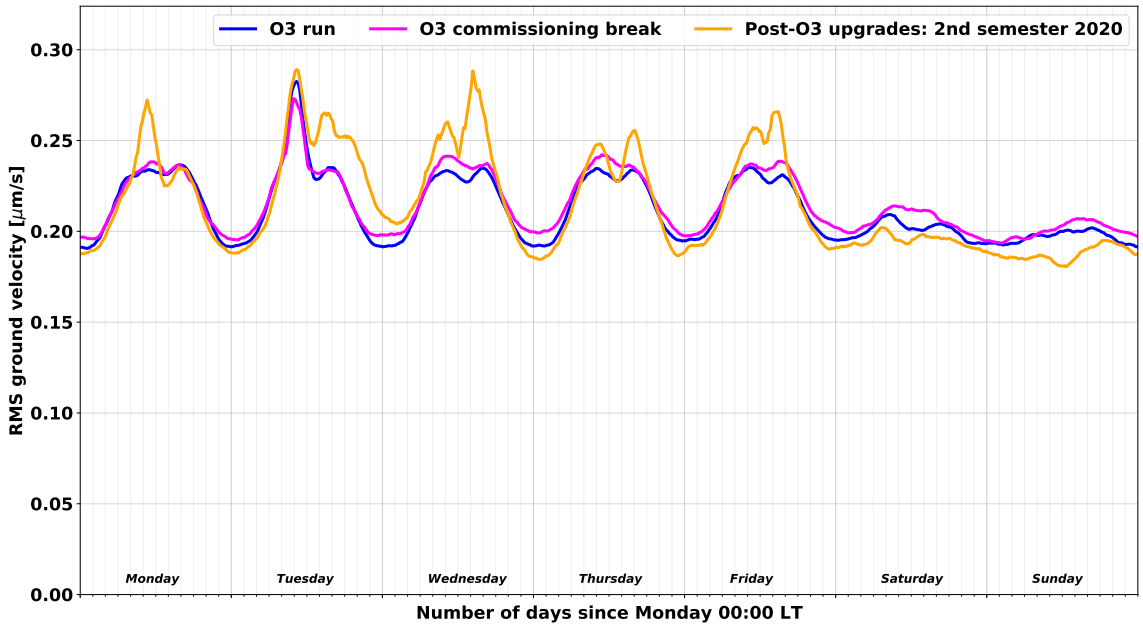

Figure 7: Average evolution on a weekly basis of the seismic on-site noise (frequency band: 10-40 Hz) measured at EGO during different times in 2019-2020.

3.2. Impact on the Virgo detector

The previous sections have demonstrated that the Virgo collaboration is accurately monitoring the seismic environment at EGO and that the recorded data show significant variations over time, in agreement with expectations from known noise sources. It is then interesting to see how these noises impact the performance of the Virgo detector, namely its sensitivity and duty cycle.

3.2.1. Sensitivity A convenient way to monitor the sensitivity of a gravitational-wave detector like Virgo is to study the evolution of the *BNS range*, that is the average distance up to which the merger of a standard binary neutron star system (BNS) can be detected with a signal-to-noise ratio (SNR) set to 8, roughly corresponding to one false alarm per year with purely Gaussian noise. The average is taken over the position of the BNS in the sky and over the orientation of its orbital plane. Broadly speaking, the lower (higher) the noise in the frequency band of interest –from a few tens of Hz to a few hundreds of Hz depending on the actual sensitivity curve –, the larger (smaller) the BNS range.

In addition to its potential dependence on the surrounding environment, the BNS range can fluctuate significantly due to changes in the control accuracy of the detector. Therefore, averaging raw BNS range values, especially over long timescales, is not expected to provide meaningful information as one would mix together too many effects that cause the BNS range to vary. Therefore, the method used in the following consists in computing a moving daily average of the BNS range and to focus on the local fluctuations around this level. Figures 8 and 9 show these variations, averaged over the whole O3 run, and projected over a weekly or daily time range, respectively. On both plots, the red dots show daily variations while the blue curve is a moving median profile of the scatter plot. The variations seen are clearly of anthropogenic origin, with a day-night pattern and a reduced spread during the weekend. Although they are significant, they are also limited in size: ∼1 Mpc compared to an average BNS range of about 50 Mpc during the O3 run, hence a ∼2% fluctuation. This shows the robustness of the Virgo detector.

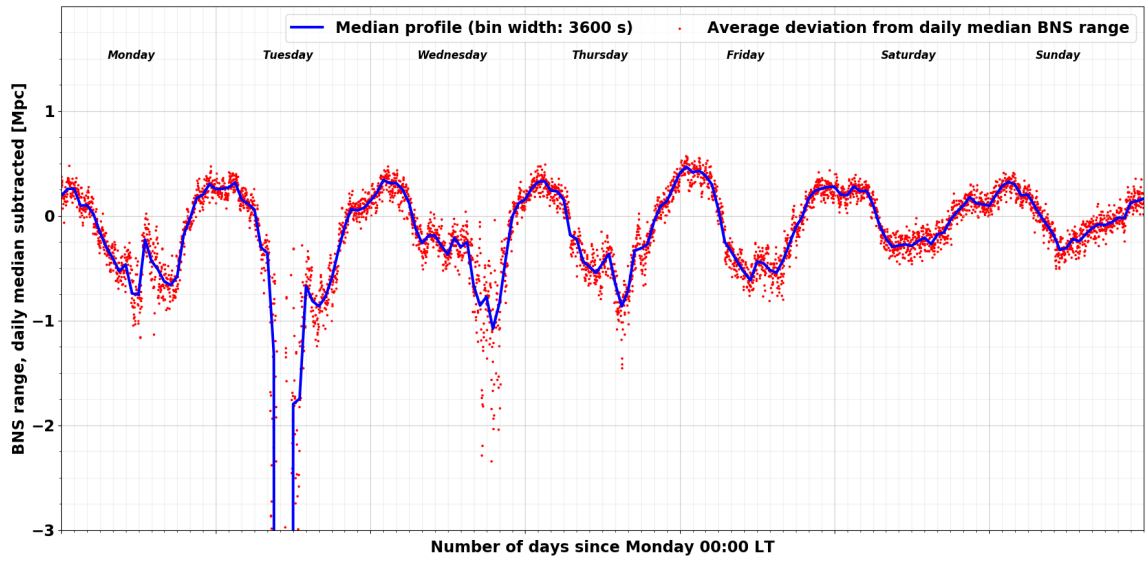

Figure 8: Average variation of the BNS range around its local average, computed on a weekly basis. The blue trace is a moving median profile of the red scatter plot, each dot showing the fluctuation at a particular weekday and time. The lack of available data on Tuesday morning corresponds to the weekly maintenance period of the Virgo detector, while the sharper variations on Wednesday and Thursday afternoons are due to the fact that these times have often been used for calibration or detector activities. Therefore, the BNS range is less stable than usually when nominal data taking gets restored.

3.2.2. Duty cycle Figure 10 shows the average duty cycle of the Virgo detector during the O3 run. The top plot displays its average variation over a week, while the bottom one focuses on a day. The red curve normalizes the Science mode data taking by the elapsed

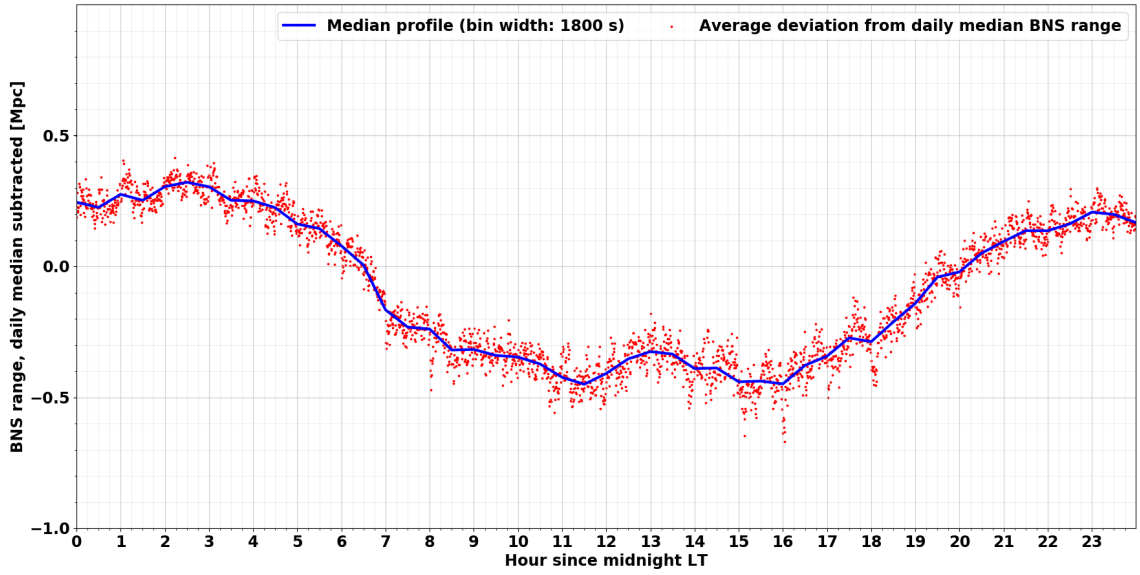

Figure 9: Average variation of the BNS range around its local average, computed on a daily basis. The blue trace is a moving median profile of the red scatter plot, each dot showing the fluctuation at a particular time of the day.

real time, while the green one is computed by excluding the calibration, commissioning and maintenance periods. Thus, the latter curve shows the fraction of the time available for data taking that is actually used for that. Activities on the detector are concentrated during working hours as expected, with maintenance on Tuesday morning, calibrations on Wednesday evenings and commissioning slots from Monday to Friday depending on the needs. There is a non-negligible recovery time from maintenance, while the transition from calibration back to data taking is smoother and quicker on average. During the quietest hours of the night, when no work takes place on the interferometer except in case of an emergency, the average duty cycle reaches a plateau around 85%.

4. Earthquakes

Earthquakes radiate energy through different types of seismic waves that are commonly divided in "body" and "surface" waves, depending on the path followed from the source to the receiver. Body waves that travel through the Earth are usually detected first. The fastest are named P-waves and are compressional longitudinal waves whose speed can reach 8 km/s. Then come the S-waves, transverse shear waves whose velocity scales by a factor of $\sqrt{2}$ with respect to P-waves. Surface waves are slower and their size dominates at large epicentral distance since their amplitude scaling factor is 1/distance while body waves scale with 1/distance^2. Most relevant surface waves are Rayleigh waves, that originate from P-wave and S-wave (with vertical polarization) coupling at the Earth surface. The result is a wave with both longitudinal and transversal components and a

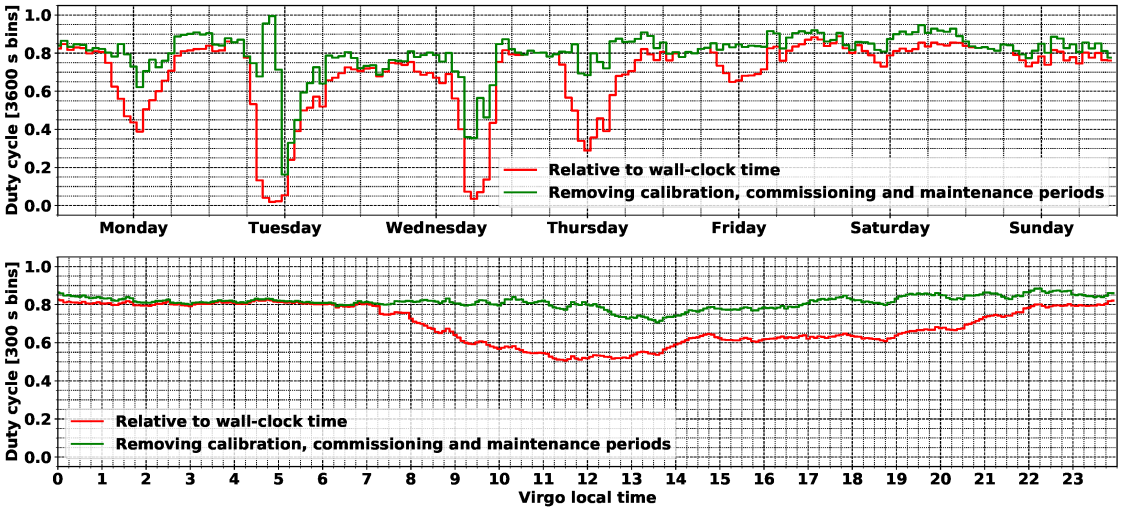

Figure 10: Average weekly (top) and daily (bottom) duty cycle of the Virgo detector during the O3 run. The red curve uses the elapsed real time as normalization, while the green one is produced excluding the times spent doing calibration, commissioning or maintenance, three activities that are incompatible with Science-mode data taking.

propagation speed up to a few km/s.

Since seismic waves excite buildings even at great distance, the Virgo monitoring set at EGO includes local, regional and teleseismic earthquakes since it was observed that all of them can induce large motion of the interferometer elements. This can then saturate the control capability of the feedback systems that keep Virgo at its nominal working point, leading to a loss of control. Following each control loss (regardless of its origin: an earthquake or another cause), data taking stops immediately and can only restart after the completion of the semi-automated sequence that allows restoring the Virgo global working point – during the O3 run, that procedure took about 20 minutes on average [19]. But the time lost can be much longer in case of a control loss due to an earthquake, in case the suspension normal modes are excited by the seismic waves. In that case, one may have to wait up to one hour after the event that the high-quality factor modes of the suspensions are damped, before initiating the control acquisition procedure. Since each control loss reduces the Virgo duty cycle, it is therefore important to understand which fraction of these are due to earthquakes, and what are the earthquakes that induce them.

Large earthquakes at local and regional scale do not happen very often, so the type of earthquakes on which this analysis is focused is large earthquakes that occur along the boundaries of the main tectonic plates. Most of them are quite distant from EGO, meaning that a low-latency framework relying on data from a variety of seismic stations worldwide could produce early warning notices that would be received and processed *ahead* of the seismic waves arrival. In that case, one could take preventive measures to

try to mitigate the effect of the ground shaking, with the goal of avoiding the control loss. In the following, we describe the strategy implemented at Virgo during the O3 run, that relies on the Seismon framework [28, 29, 30] developed at LIGO – an example of the existing teamwork among members of the global GW detector network.

Furthermore, as explained below, the study has also highlighted another contribution from much weaker earthquakes, quite close to EGO (the majority of which occur on the Italian Apennines). Those have been more difficult to identify as they do not lead to early warnings from Seismon and the frequency of their seismic waves is much higher when they arrive at EGO: up to $\sim$1 Hz, whereas teleseism waves are in the frequency band $10 - 100$ mHz. In addition, the proximity of their epicenters makes useless the use of warnings that would always come too late. Thus, the only way to mitigate these earthquakes is to understand how they impact the Virgo control system and what could be done to strengthen it.

4.1. O3 Seismon setup at EGO

In addition to making the whole detector as robust as possible against the passing of strong seismic waves, the only other leverage one can use to mitigate the impact of earthquakes is to rely on early warnings provided by worlwide arrays of seismometers.

Following initial tests done during the O2 run and the upgrade period that followed, we ran at EGO during O3 an instance of the Seismon framework, developed by LIGO to process earthquake early warnings provided by the US Geological Survey (USGS) [31] and to compute information relevant for each site of the LIGO-Virgo network. Namely, for each earthquake, Seismon potentially predicts the arrival time of the different seismic waves (P-, S- and surface), their amplitude at site and the probability of losing the detector control in consequence of that earthquake That framework was split into four consecutive steps, each associated with a server integrated within the Virgo online data acquisition system (DAQ) used to steer and monitor the detector.

- Reception of the USGS alerts.
- Processing of each alert by the Seismon framework.
- Extraction of the subset of Seismon data pertinent to the EGO site and provision to the Virgo online framework.
- Local processing of these data.

In addition to producing a plot summarizing all information available from the early warning, a loose cut is applied on magnitude and distance to estimate whether the earthquake could be relevant, meaning that it could impact the control of the Virgo detector. In that case, and if the warning was quick enough to precede the arrival of the seismic wave on-site, an alarm would latch on the main panel of the Virgo Detector Monitoring System [32, 33], alerting the operator on duty in the control room.

In the nominal O3 control configuration, the two 3 km-long optical cavities are kept in resonance by acting on the end mirror suspensions: their actuators are the least noisy, at the price of a reduced correction range availability. Actuators located at the level of the input mirror suspensions have higher dynamics, while introducing slightly more noise. Thus, they can be used as earthquake control mode (in short *EQ mode*) to try to maintain the Virgo working point during periods of elevated seismic noise.

A smooth transition procedure, working both ways without losing the detector control, has been developed to allow switching back and forth between end-mirror and input-mirror actuations. During the O3 run, the procedure in use was the following: once alerted by Seismon, the operator on duty would monitor the optics suspension status and manually trigger the transition from nominal mode to EQ mode when the test mass suspensions would start shaking significantly. Once activated, that process would take a few tens of seconds to complete. Then, either the detector would nevertheless lose its working point (and the control acquisition procedure would have to be restarted from the beginning), or the EQ mode control would be kept until the whole seismic wave trains has passed by and the suspensions motion has been damped back to levels low enough to allow resuming the nominal control mode.

Unrelated to earthquakes, the EQ mode was also found useful during periods of high wind: gusts shake the building structures (walls and floors) and those vibrations can couple to the suspensions, potentially causing control corrections to saturate. However, since EQ mode was not validated for the production of good quality data for physics analysis, this method was used parsimoniously during most of O3 because corresponding data would have to be discarded. A few weeks before the end of the run, the EQ mode got finally qualified for regular data taking and later studies [34] showed that there was no significant degradation of the Virgo sensitivity when switching to it. Therefore, it was used more regularly from that time; the possibility to have such a backup solution for O4 as well will be studied in the coming months.

4.2. Earthquakes impact during O3

The stronger and/or the closer to EGO the earthquake, the more likely it is to impact the control of the Virgo detector. To study the impact of strong regional earthquakes or teleseisms, the USGS warnings processed by Seismon are sufficient (as they should include all such earthquakes). But it was soon realized that some moderate earthquakes occurring at local and regional distance (from few tens to few hundreds kilometers away from EGO), too weak to generate a USGS alert and thus not processed by Seismon, could cause losses of control of Virgo. To check if any of the control losses was caused by this type of earthquakes, we queried [35] the INGV (Istituto Nazionale di Geofisica e Vulcanologia) public earthquake catalogue [36] to download the list of events that occured during O3 in the Mediterranean region. This list partly overlaps with the USGS one and duplicates were removed. All results presented in the following are based

on the whole set of earthquakes, assembled by merging the USGS and INGV event lists.

The control of the Virgo detector is extremely complex. Therefore, finding out how many earthquakes induced control losses during the O3 run required a careful study of all control losses, documented below in Appendix A. An earthquake from the list of USGS warnings is associated to a recorded control loss if the loss occurs within the time range during which seismic waves were predicted to arrive on-site according to Seismon and if the seismic activity around the time of the control loss is significantly larger than its typical range of variation. In case of concurring early warnings from different earthquakes overlapping in time at EGO, the strongest is arbitrarily selected as the reason for the control loss.

Estimating the strength of an earthquake when its seismic waves arrive at EGO is not easy. Yet, this is a key point to address, first to reject quickly warnings from harmless earthquakes and then to adjust the latency and level of response for the crew in charge of steering the Virgo detector. During O3, basic rectangular cuts in the magnitude-distance plane – e.g. *if magnitude > (...) or (distance < (...) km and magnitude > (...)) or etc.* – were applied to the live earthquake warnings received from USGS and processed by Seismon. During the post-run analysis, the ranking

$$\text{ranking} = \frac{10^{\text{magnitude}/2}}{\text{distance[km]}} \tag{1}$$

was introduced. While not complete – e.g. neither the hypocenter depth nor its azimuth angle computed with respect to EGO are accounted for – this ranking appears sound: the higher its value, the more likely the control loss. Applying a (conservative) minimum cut at ranking = 0.02 allows to safely remove more than half of the earthquakes to be analyzed.

Results shown below use the largest possible earthquake statistics, meaning that one requires the Virgo detector to be fully controlled, but not necessarily in Science mode. This looser requirement enlarges the dataset of interest and hence the number of earthquake early warnings to be taken into account.

Figure 11 highlights the epicentral distance and magnitude of the earthquakes that led to a Virgo control loss. The top (bottom) row deals with the earthquake magnitude (epicentral distance) while the right column displays the ratio of the red and blue histograms shown on the left column. As expected, the larger the earthquake magnitude, the more likely the control loss, with the fraction of earthquakes leading to a control loss departing from 0 for magnitude 6 and above. That fraction saturates to 1 (meaning that all events causes a control loss) when magnitude exceeds 7.2. We also note that the fraction is not null around magnitude 3: this reflects the control loss consequence of some small local earthquakes recognizable also in the left side histogram of Figure 11. The histogram ratio is much flatter for that other variable, with the most significant bins reflecting the location of seismic regions on the globe with respect to EGO, mainly

the broad Mediterranean area and the Ring of Fire (a region covering much of the rim of the Pacific Ocean that is seismically very active).

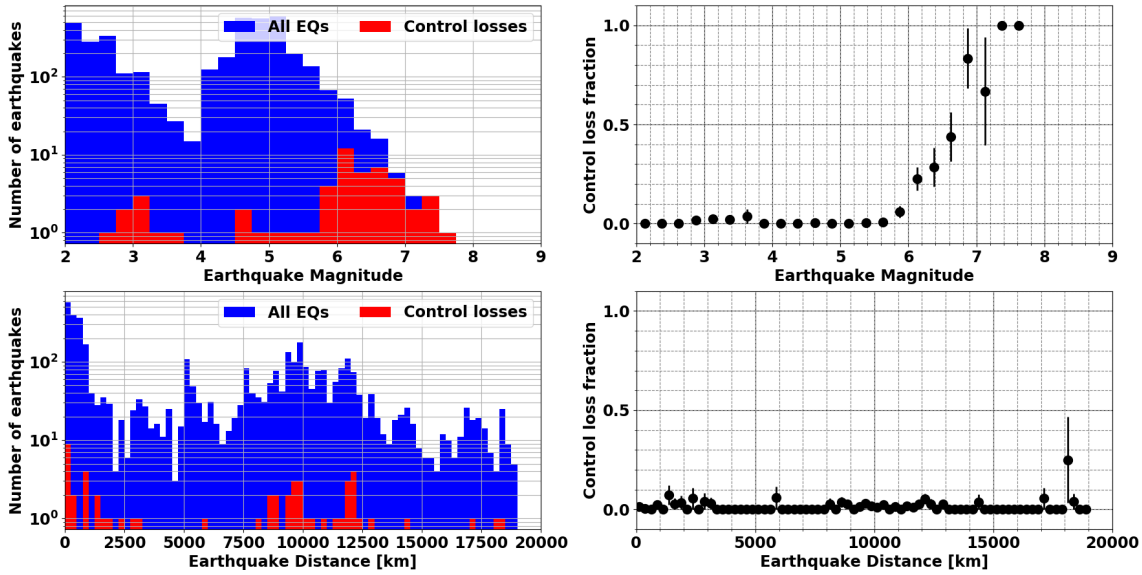

Figure 11: Summary of the impact of earthquakes on the Virgo detector during the O3 run. Left column: the blue (red) histogram shows all earthquakes (the earthquakes that have induced a control loss); top: magnitude distribution; bottom: distribution of the distance between EGO and the epicenter. Right column: corresponding fraction as a function of the earthquake magnitude (top) and distance (bottom). In all cases, the earthquakes that certainly could not impact Virgo (ranking below 0.02) were excluded.

Figure 12 shows the population of earthquakes that caused a control loss (red dots) in the two-dimensional plane epicentral distance vs. magnitude. These earthquakes form the upper envelope of the scatter plot drawn, meaning they are usually the earthquakes with largest magnitude for any distance. The separation between red and green (earthquakes that did not cause a control loss) dots is not perfect for at least two reasons. The first one is that the control of the Virgo detector is complex enough that the actual level of control (accuracy and stability) plays a role in whether or not the control is lost for earthquakes at the limits of inducing a control loss. The second reason is that our model could probably be improved by including other earthquake warning parameters: two candidates would be the hypocenter depth (the deeper the hypocenter, the lower the earthquake impact on the ground at equivalent magnitude) and the azimuthal orientation of the epicenter with respect to EGO.

Figures 13 and 14 show the location of the significant earthquakes that occurred during 03 with the same color coding used in Figure 12. Their distribution depicts the boundaries of the main tectonic plates and, as discussed above, we can observe that the most harmful earthquakes for Virgo are coming from the Mediterranean area (medium to large magnitudes but smaller distances) and part of the Pacific Ring of Fire. The mid-

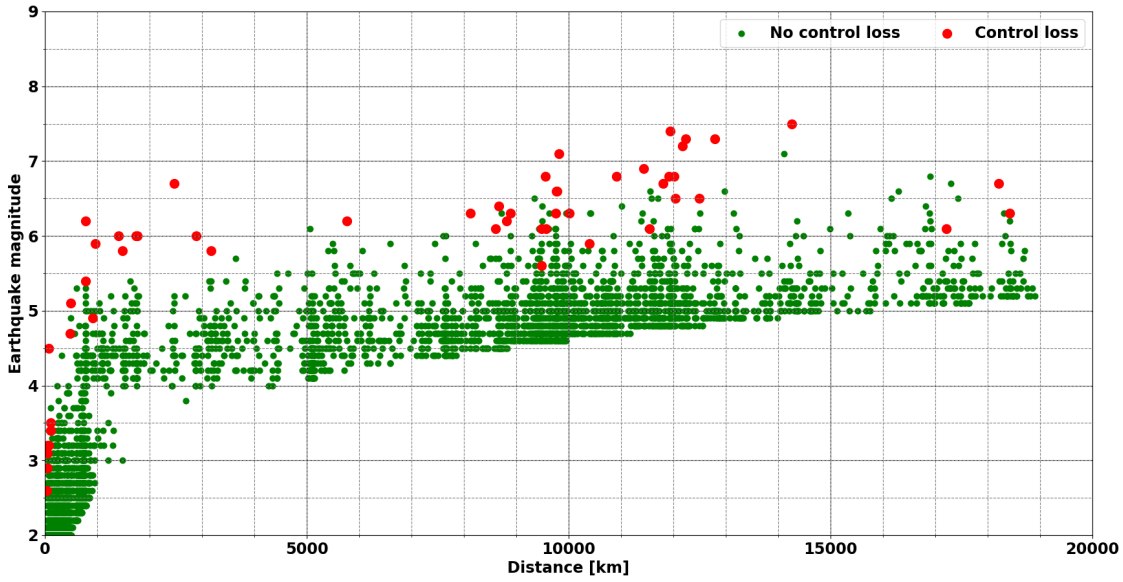

Figure 12: Distribution of earthquakes in the plane distance-magnitude during the O3 run. The earthquakes that caused a control loss (did not cause a control loss) are represented with red (green) dots. The lack of points below the main bulk of earthquakes is due to the ranking cut, set at 0.02.

Atlantic ridge and the Asian portion of the Alpide earthquake belt did not produce many earthquakes that impacted Virgo, possibly because of the limited statistics. During the O3 run, the distribution of the earthquakes leading to control loses was the following: ∼15% of close earthquakes; ∼20% from other earthquakes in the Mediterranean area; and ∼65% from distant earthquakes. We remark that this statistics has not an absolute meaning: the O3 run took place during a quiet seismic period for Italy, compared to e.g. 2009 or 2016. This analysis will be updated in the future with data from the O4 run.

Finally, Fig. 15 shows an example of the impact of a strong and distant earthquake on the Virgo detector and how the early warning information was used to change the control mode prior to the arrival of the strongest seismic waves. This allowed the crew on duty to keep the working point of the instrument by preventing the correction force (applied on mirror suspensions to maintain resonance in the arm cavities) from saturating. Should that action not have been performed, the control loss would have been unavoidable – as the correction would have saturated around 22:27 UTC. The description of the different stripcharts displayed is given below.

- Top plot: variation of the index labelling the Virgo data taking configuration: the Science mode corresponds to the value 1; other indices shown here (-1, -7, -9) indicate different control configurations that are not nominal and that were used to wait for the right moment to switch back to Science data taking mode.

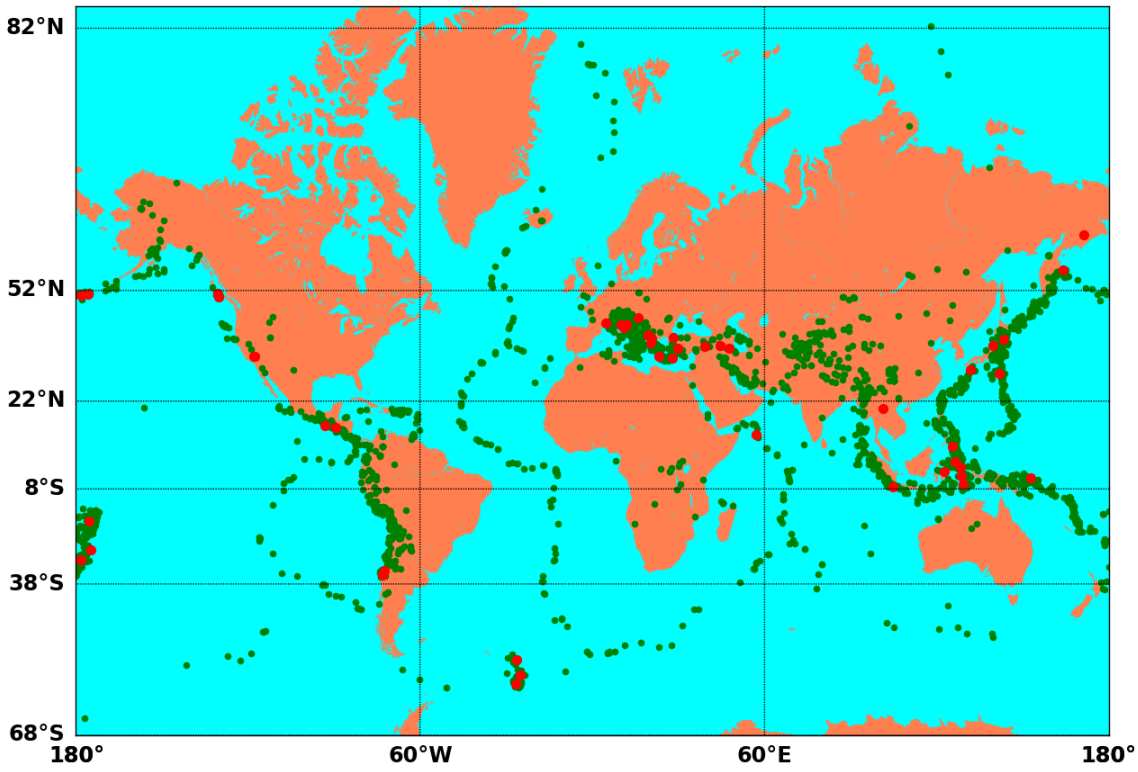

Figure 13: Location of the O3 earthquakes used in this study (ranking greater than 0.02) The earthquakes that caused a Virgo control loss (did not cause a control loss) are represented with red (green) dots.

- Second plot: stripchart of the BNS range versus time; the seismic waves clearly make the BNS range go down and fluctuate more while they are passing (see seismic activity variations recorded in the bottom plot, described below); the BNS range recovers its steady value at the end of the plotted time when the earthquake effect fades away.
- Third plot: switch showing the times when the earthquake-resilient control mode ('EQ-mode') is turned on ($0 \rightarrow 1$ transition) and later on off ($1 \rightarrow 0$ transition) manually by the operator on-duty.
- Fourth plot: For each second, maximum value of the correction applied on the test masses to keep the Virgo arms in resonance. When the nominal control mode is used, a control loss happens within two seconds at most after the time for which the correction voltage§ exceeds a 9.5 V threshold. This occurs a few times close to the middle of the time range represented here but no control loss follows, as the EQ-mode allows for larger corrections.
- Bottom plot: seismic noise measured in three orthogonal directions (vertical and

§ The mirror control is done by varying the amount of current applied to actuators (pairs of coil-magnet): see Ref. [17] for details.

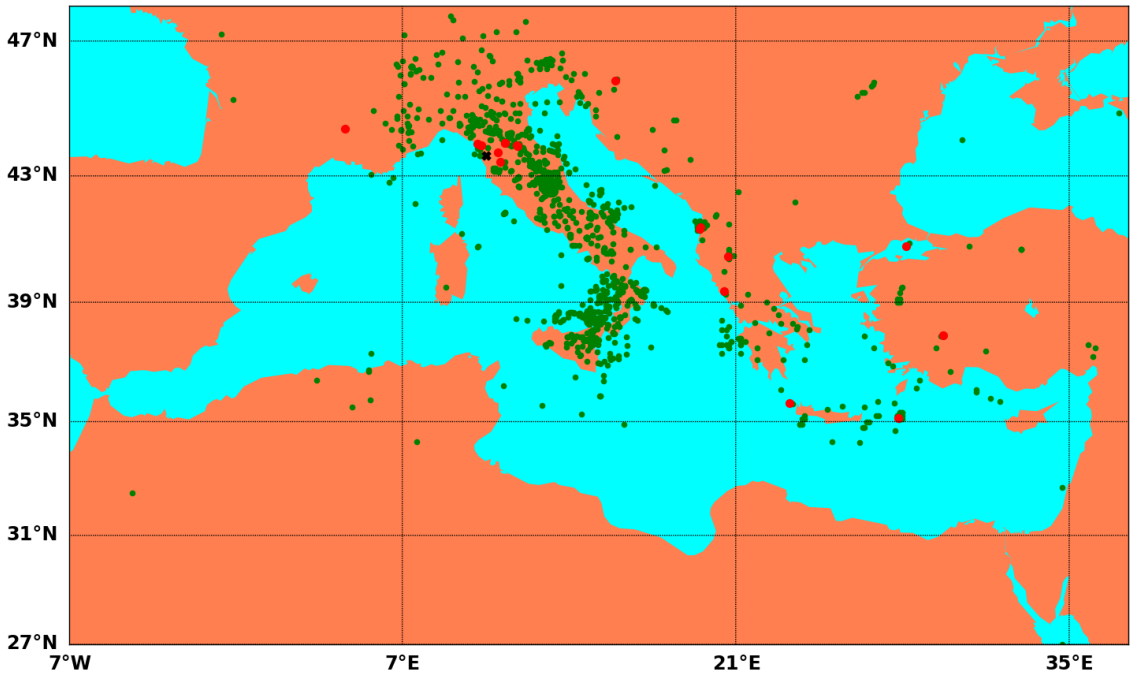

Figure 14: Zoom on the Mediterranean area of the map shown in Fig. 13 above. It shows the earthquakes nearby Virgo (whose site, EGO, is marked by a black cross) recorded during the O3 run. The earthquakes that caused a control loss (did not cause a control loss) are represented with red (green) dots.

along the two Virgo arms) using the dominant frequency range for earthquakes recorded at teleseismic distance: 10 mHz $\rightarrow$ 100 mHz.

- Finally, the vertical dashed lines common to all plots show the time of important events. From left to right: the time at which the earthquake occurred; the time at which the corresponding USGS warning had been received and processed by the Seismon framework at EGO; the expected arrival time of the seismic P-waves, S-waves and Rayleigh waves. For the latter, we use three different arrival times that stem from different assumed velocities (5, 3.5 and 2 km/s respectively).

4.3. Plans for O4

Work is in progress to build on the O3 experience and have a more performing, better integrated, earthquake early warning framework for O4 (and beyond). The plan is to run the latest version of Seismon with an improved prediction capability for EGO, achieved by means of all the data collected during the O3 run. We are also exploring the possibility to use the INGV Early-Est system (a framework for rapid location and seismic/tsunamigenic characterization of earthquakes) [37, 38] as an additional source of warnings, complementary to USGS. Tests are in progress to have this new live stream

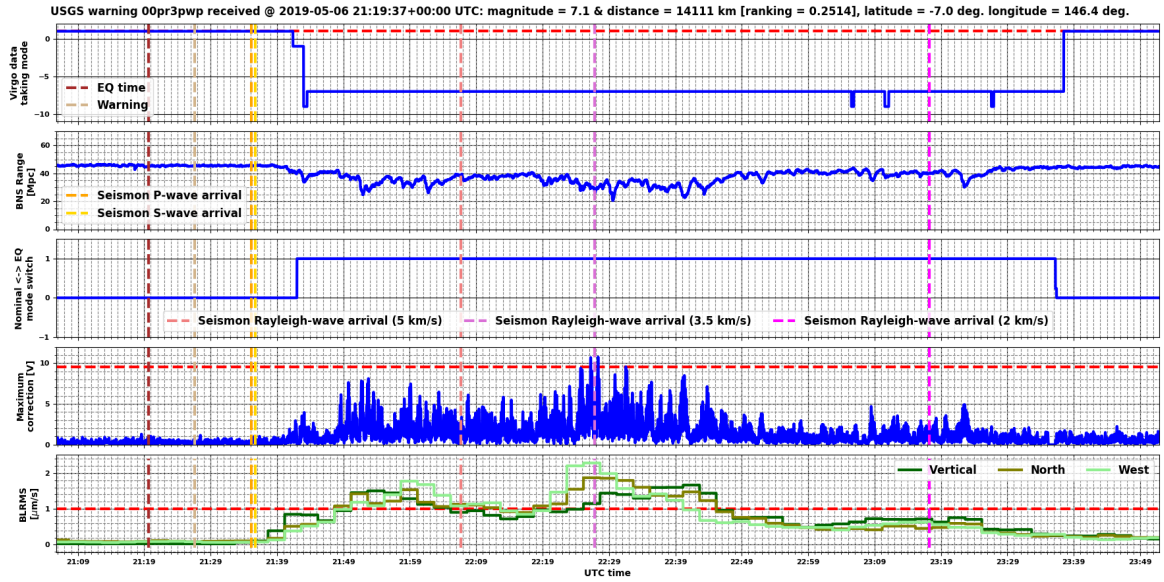

Figure 15: Example impact on the Virgo detector of a strong (magnitude greater than 7) and distant (14,000 km away from EGO) earthquake, that occurred on May 06, 2019 at 21:19:37 UTC in Eastern Papua New Guinea. The description of the different stripcharts is provided in the text.

received at EGO and integrated into the existing framework. The two sets of early warnings will then be compared, in terms of latency and accuracy.

5. Bad weather

Through O3, the Virgo interferometer performed worst during days with adverse meteorological conditions, namely high winds and intense sea activity. These periods were generally associated with increased non-stationary noise in the GW signal below about 100 Hz and with some difficulties in maintaining the interferometer in its controlled state, resulting in reduced duty cycle. In the following, we study the impact of the increased microseimic noise associated to sea waves, then the influences of wind on BNS range, as well as the effect of wind gusts on the global interferometer controls. Because of the wind action on the sea surface, high winds and rough sea often occur together. We use a statistical approach to disentangle their effects on the detector.

5.1. Impact of sea activity

Microseism amplitude at EGO increases by more than one order of magnitude between calm and rough sea periods. For 10% of the time during O3, ground RMS velocity between 0.1 Hz and 1 Hz was above 4 μm/s, as shown in Fig. 5. This happened in particular in correspondence of the seasonal change in the first part of O3b and for

some periods of adverse weather conditions in the first months of 2020. Periods of intense sea activity were associated to larger than usual strain residual noise whose characteristics and origin require further analysis.

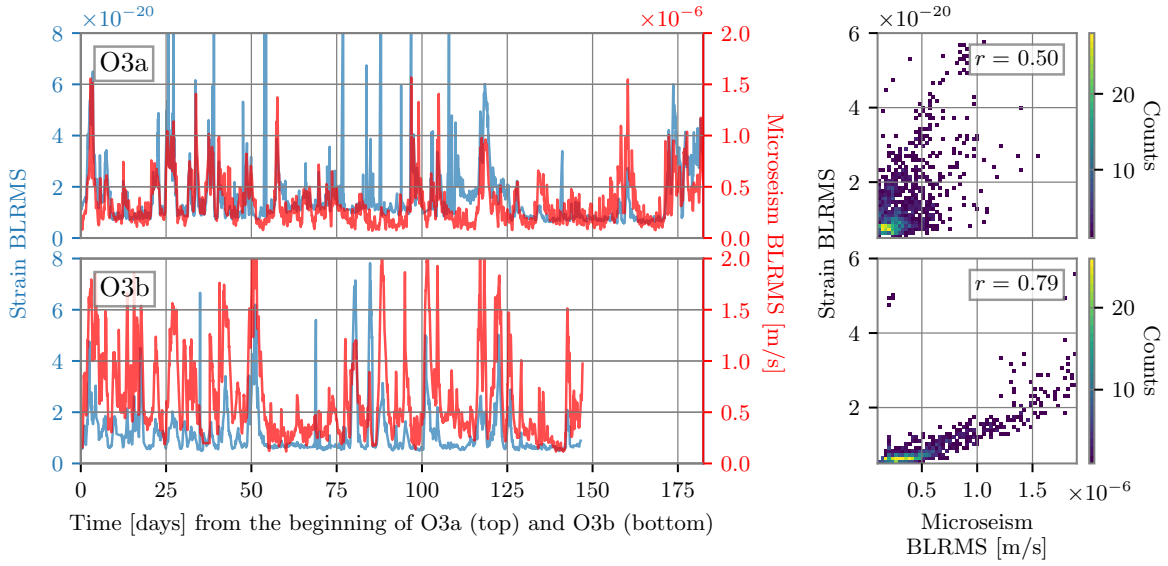

Figure 16: Correlation between the low frequency noise in Virgo GW strain and the microseism induced by the sea activity; top row for O3a and bottom for O3b. Left: in blue the time series of the strain BLRMS in band $[10, 20]$ Hz and in red that of the BLRMS in band $[0.1, 1]$ Hz of a CEB seismometer, mostly influenced by the sea activity. Right: 2D-histograms of the correlation between the two BLRMS, where the colorscale counts, for every pixel in this map, how many data points have the corresponding values of strain and microseism BLRMS. The annotation in the top-left corner reports the value of the Pearson correlation coefficient r.

5.1.1. Microseism impact on strain noise Periods of high sea activity were associated with larger strain residual noise up to about 100 Hz. To characterize this effect, we made use of the *band-limited* RMS (BLRMS), defined for a generic signal, in a certain frequency band $[f_{\min}, f_{\max}]$, as:

$$\mathrm{BLRMS}\big(t; [f_{\min}, f_{\max}]\big) := \sqrt{\int_{f_{\min}}^{f_{\max}} S(f;t) df} \tag{2}$$

where $S(f;t)$ is an estimate of the signal *power spectral density* (PSD) referred to a time t.

In Fig. 16, we report, for the entire O3 run, in blue the BLRMS of the strain in the band $[10, 20]$ Hz and, in red, the CEB seismometer BLRMS in the band $[0.1, 1]$ Hz. These have been estimated from (2), where $S(f;t)$ is computed with the Welch's method making

use of strides of 2048 seconds and FFT length of 128 seconds, overlapping by 50% [39]. The correlation between the two curves is apparent. In particular, when the microseism is intense, the peaks in the strain BLRMS are almost everywhere coincident with those in the seismometer BLRMS. This fact is also highlighted in the 2D-histograms on the right-hand side of the same figure, where the Pearson correlation coefficient has been computed for the two data taking periods, O3a (top) and O3b (bottom). In general, we observe that, despite the "spikes" in correspondence of bad weather conditions (in particular at the beginning of O3b and then during most of Winter‖), the induced strain noise at low frequency has improved during O3 and can now be mostly attributed to microseism.

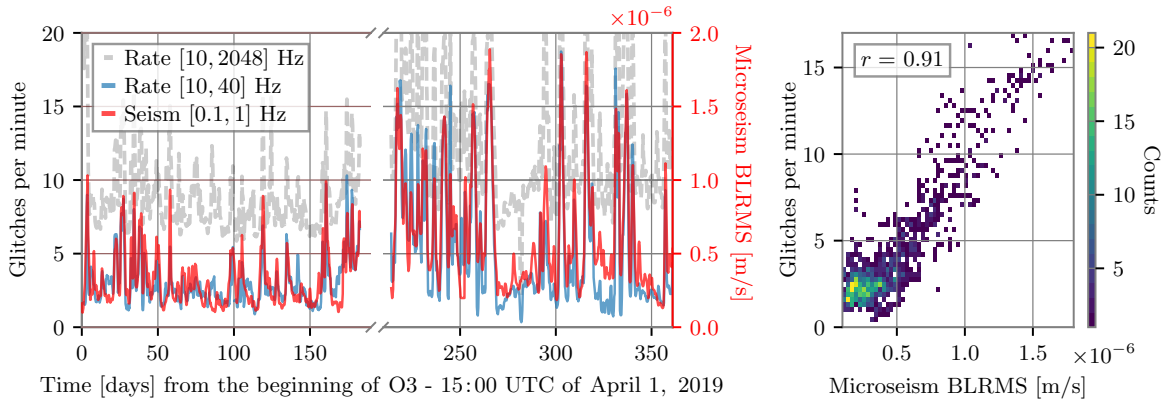

Figure 17: Correlation between Virgo glitch rate and the sea induced microseism during the O3 run. Left: the dashed gray line represents the daily moving median of the glitch rate per minute recorded by Omicron [40] for glitches with SNR > 6.5 and frequency at peak in band $[10, 2048]$ Hz, estimated over strides of 2048 seconds. The blue continuous line is the median rate referred to glitches with frequency at peak in $[10, 40]$ Hz band. The continuous red line is the BLRMS in band $[0.1, 1]$ Hz of a seismometer in the Virgo CEB. Right: 2D-histogram of the glitch rate in band $[10, 40]$ Hz and the microseism BLRMS, where the colorscale counts, for every pixel in this map, how many data points have the corresponding values of the rate and the microsiesm BLRMS. The annotations in the top-right corners report the values of their Pearson correlation coefficient r.

5.1.2. Microseism impact on glitch rates Besides an increase in the RMS value of the strain noise at low frequency, microseisms induce short transients of power excess in this channel, colloquially referred to as *glitches*. In Fig. 17 we report the minute rate of these glitches during the entire O3 run. To reduce the – usually very large – variability in their rate, we computed running daily medians. The gray dashed line represents the time evolution of daily medians for glitches with SNR > 6.5 and frequency at peak in

‖ That calendar season starts around day 50 of O3b and lasts almost until the end of the data taking.

the band $[10, 2048]$ Hz, as measured by the online Omicron pipeline [40]. The blue solid line is the median minute rate of glitches with peak frequency in the $[10, 40]$ Hz band. These glitches accounted for about 30% of the total during O3a, and for almost 40% in O3b, with peaks larger than 80% in correspondence of periods of intense sea activity. This glitch rate is highly correlated with microseism, represented in the left-hand side plot of Fig. 17 by the solid red line of the running weekly median of the BLRMS in band $[0.1, 1]$ Hz of the CEB seismometer. On the right-hand side of the same figure, we report the 2D-histogram of these two quantities and the value of their Pearson coefficient ($r = 0.91$).

5.1.3. Microseism and scattered light Glitches due to microseism often resemble arches in a time-frequency map, as illustrated for example in Fig. 18. Arches are the typical signature of scattered light (SL) noise processes, which is a major issue and topic of investigation in the second generation GW detectors [12, 13, 14, 41, 42, 43].

A stray light beam bouncing off a moving surface adds coherently to the beam main mode every time its optical path, $x(t)$, changes (increases or decreases) by an integer wavelength. It follows that the frequency of the strain noise is:

$$f_{sc}(t) = \frac{2n|\dot{x}(t)|}{\lambda} \tag{3}$$

where $\dot{x}(t)$ is the instantaneous relative velocity between the interferometer beam and the scatterer, and $\lambda = 1.064$ μm is the Virgo laser wavelength. Equation 3 is referred to as predictor. In case the scattered beam encounters a second reflective surface it can bounce back and forth n times along the same path before recombining, giving rise to higher order noise arches, reaching out n-times larger frequencies.

In O3 the main sources of scattered light affecting the sensitivity were the suspended optical benches placed beyond the end test masses in the terminal buildings (SNEB, SWEB). In this case, the noise observed in the time-frequency domain is well visible as power fluctuations in the cavity. The noise appears as a series of arches, where the typical non-stationarity and non-linearity of the noise is evident. Arch time spacing is the half-period of the oscillation of the mirror-bench relative motion, and arch amplitude (i.e. the maximum frequency extension of the induced strain noise) is $f_{max} = (4\pi/\lambda)AFn$ where A and F are the amplitude and frequency of the oscillation. If the frequency and amplitude of the oscillation are such that $f_{max} > 10$ Hz, the noise affects the GW detection frequency band.

Being those benches suspended and controlled [44], their motion induced by the microseism was supposed to be attenuated enough to push the maximum frequency of the arches below 10 Hz. Moreover, a control technique taking into account the mirror-bench differential signals was implemented in order to reduce their relative motion (BENCH-MIRROR), which is the quantity effectively responsible of the noise coupling.

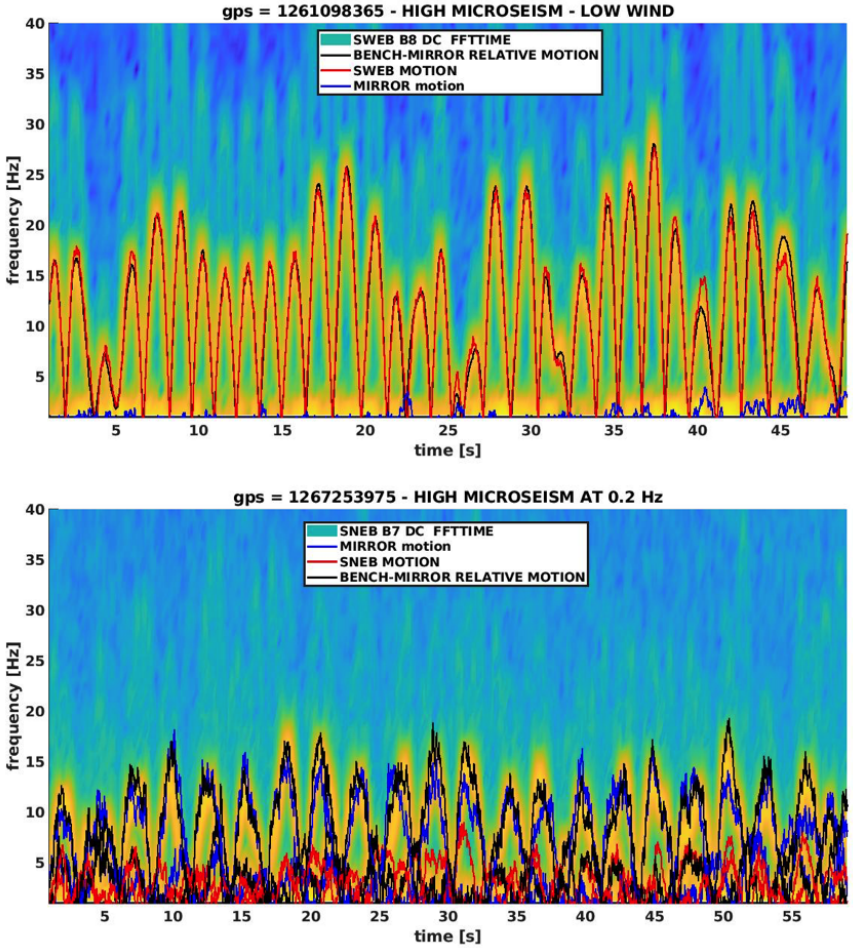

Figure 18: Spectrograms of the light transmitted at the end of the arm cavities and detected by photodiodes located behind, on suspended benches – top plot: west arm, B8 photodiode, SWEB bench; bottom plot: north arm, B7 photodiode, SNEB bench. The typical pattern of scattered light noise (arches) – both first order and second order (higher frequencies) – is visible. On the SWEB plot, arch spacing and amplitude correspond to half the period of marine microseism at Virgo (∼3 s) and a ground velocity of about 8 μm/s. The predictor for BENCH-MIRROR is shown in black, while the predictors computed from mirror and bench motions are shown in blue and red, respectively. The overlap shows that BENCH-MIRROR is the best predictor of scattered light, closely matching the observed arches.

During O3, a malfunctioning was identified in the mechanical setting of the West Bench suspension (SWEB) which caused its actual motion to be comparable to the ground motion at the frequency of the main microseismic peak. Figure 18 shows the mirror contribution and the bench contribution to the arches separately, for both North and West cavity, in two selected bad-weather conditions. In the West arm power spectrogram, the typical pattern is visible: the arches were entirely due to SWEB motion, and all the times the ground motion exceeded a certain threshold during the

run, these arches entered the detector band. In the North arm power spectrogram, the arches were normally much lower, and the contribution from the bench motion was of the same magnitude as the mirror motion. It was even possible to find some special conditions (the largest component of the ground motion centered at 0.2 Hz), in which the mirror motion was prevalent (see Figure 18, bottom panel).

The issue concerning SWEB mechanics and control has been understood and cured after O3. In O4, its residual motion is expected to be at least similar to the one observed in O3 for SNEB. Further improvements in the control strategy will be tested for both the mirror and the bench suspension.

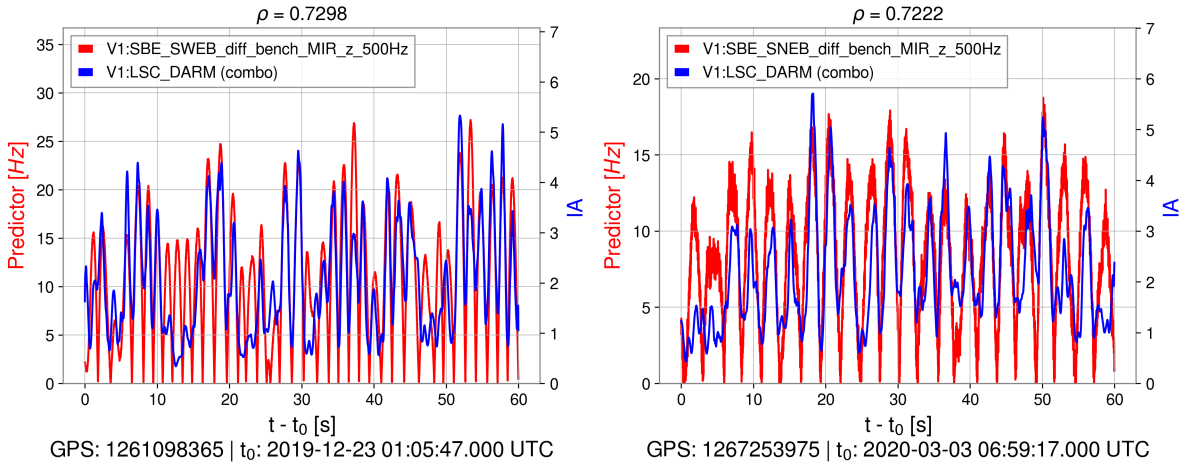

Figure 19: In red is the culprit's predictor, i.e. Equation 3 for the relative motion (diff) between the suspended end bench and the end mirror (BENCH-MIRROR) of the West end (left) and North end (right). The sum of the IA of the first two modes of DARM, extracted by tvf-EMD is shown in blue.

5.1.4. Identification of scattered light culprits Part of the effort regarding scattered light noise mitigation consists in the localisation of scattered light sources, referred to as culprit, through data analysis. This can be a difficult and time consuming operation in a km-long detector with many possible sources of SL. Adaptive algorithms for time series analysis can be used to this end, due to their ability to decompose non-linear and non-stationary data into a set of oscillatory modes [45, 46]. The methodology described in [46] and based on the time varying filter empirical mode decomposition (tvf-EMD) [47] adaptive algorithm is applied to the two data segments shown in Fig. 18. SL noise couples with the differential motion of the arm cavities (*DARM*, the Virgo longitudinal degree of freedom sensitive to GW) time series, which is first low-passed and then decomposed using tvf-EMD to extract its oscillatory modes, from which the instantaneous amplitude (IA) is obtained using the Hilbert transform. Computing Equation 3 for a broad list of position sensors and correlating with the IA of DARM's

oscillatory modes allows to quickly identify the most correlated channel, i.e. the culprit. The two data segment considered are

- GPS: 1261098365 UTC - 2019/12/23 01:05:47 + 60s,
- GPS: 1267253975 UTC - 2020/03/03 06:59:17 + 60s.

Obtained results are reported in Fig. 19, showing the predictors of the culprit for the end benches, based on Equation 3, correlated with the IA of DARM. The culprits are related to the BENCH-MIRROR channel in both cases. The resulting values of correlation are $\rho = 0.73$ for SWEB and $\rho = 0.72$ for SNEB. Since after low-passing the data the first two oscillatory modes of DARM were found to be the most correlated with the same predictor, the sum of their IA is considered and is shown in Fig. 19 for both cases, referred to as *combo*. As a counter proof, in Fig. 18 the predictors of the culprits are overlapped on the spectrograms of the WEB and NEB photodiodes. It can be seen that they closely match the scattered light arches. In particular, for the SWEB case, the mirror motion is small and the bench motion is mainly responsible for the observed scattered light. For SNEB case, while the mirror motion is significant the BENCH-MIRROR predictor, identified with adaptive analysis, better matches the arches also in this case.

5.2. Impact of wind

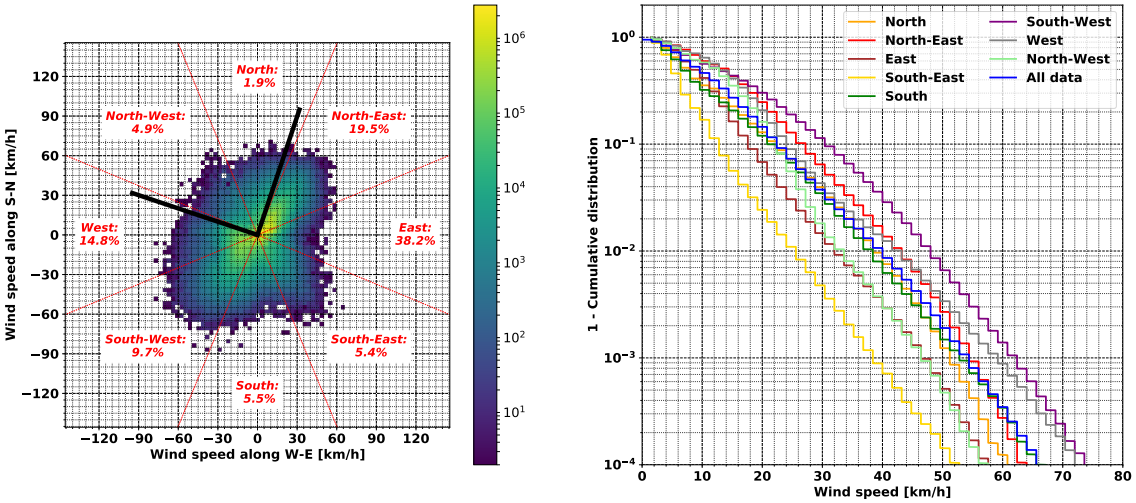

Figure 20: Wind statistics as measured by the EGO weather station during the O3 run. The left plot shows the joint distribution of the wind speed and orientation, with the two black bars showing the directions of the two arms of the Virgo detector. The right plot shows the complementary cumulative distribution of the wind speed for each of the eight quadrants of the wind rose.

Figure 20 summarizes the wind statistics recorded at EGO during the O3 run. Wind is blowing more often from the East while the stronger winds are predominantly coming from the West – the sea shore. The method described in Sec. 3.2.1 can be applied to quantify the impact of the instantaneous wind speed on the sensitivity. Figure 21 shows that the sensitivity is pretty much unaffected until a wind speed of $\sim 20 - 25$ km/h, while the detector gets sensitive to larger speeds: the BNS range decrease exceeds $\sim$4 Mpc for a wind speed of 50 km/h or above. Yet this variation is limited (about 10% of nominal BNS range values during O3), meaning that the detector is quite robust against wind. Another consequence of high-wind conditions is the need for the Virgo global control system to use larger corrections to keep the instrument at its nominal working point. And the larger these corrections, the more the detector is vulnerable to additional disturbances that could make the corrections saturate and lead to an almost immediate control loss.

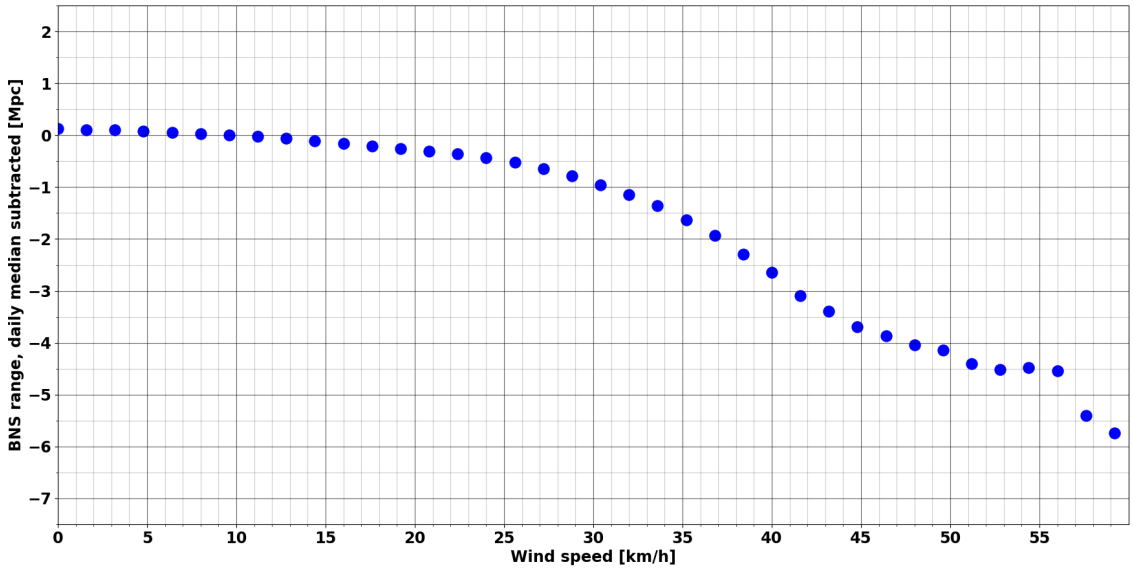

Figure 21: Average variation of the BNS range around its local average, as a function of the wind speed. In the Virgo DAQ, the BNS range and the wind speed are updated every 4 and 2 seconds, respectively.

The effect of the wind speed is clearly visible on Fig. 22 that compares the complementary cumulative distribution functions of the kilometric Fabry-Perot cavity longitudinal corrections for different ranges of wind speed. Clearly, the larger the wind speed, the higher the correction. On this plot, the average wind speed and the maximum correction have been computed using non-overlapping time windows of 30 seconds each. The largest displayed correction range stops on purpose at 9 V because the actual physical correction saturates at 9.5 V, a value that can be reached or even exceeded when there is a control loss. As the control system has some small but non-zero internal latency, it is not always clear whether the observed saturation is the cause of the control

loss or a consequence of it. Therefore, for a cumulative plot like the one shown on Fig. 22, corrections above 9 V have been cut away to avoid contamination from correction signals posterior to control losses.

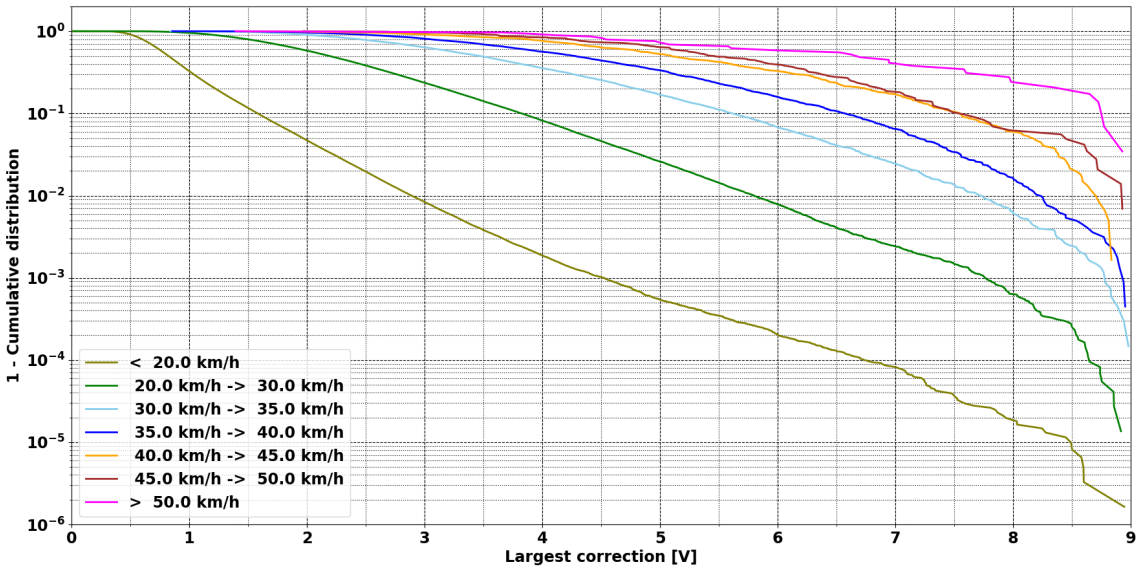

Figure 22: O3 complementary cumulative distribution functions of the maximum longitudinal corrections (in volts) keeping the Virgo arm cavities resonant for different wind speed ranges. The mean wind speed and the maximal corrections have been computed over 30 s time windows. The x-axis ends at 9 V, a bit below the saturation level of 9.5 V for that particular correction.

5.3. Disentangling sea activity and wind

Fig. 23 attempts to disentangle the impact of high microseism levels (due to the nearby rough sea) and high wind, by looking at the O3 Virgo duty cycle as a function of the microseism level for three different wind conditions: no cut on wind speed (blue histogram); low wind speed (below 25 km/h, green); high wind speed (above 25 km/h, red). One can see that in low wind conditions the duty cycle is pretty much independent from microseismicity, whereas it is lower and decreases more quickly when the wind level increases. Therefore, the Virgo detector appears robust against microseism but more sensitive to wind. Note that the extreme bins on the histograms plotted on Fig. 23 may have low statistics compared to others (low wind and high microseism, or high wind and low microseism are rare conditions): this explains why the duty cycles reported there fluctuate significantly compared to neighboring bins.

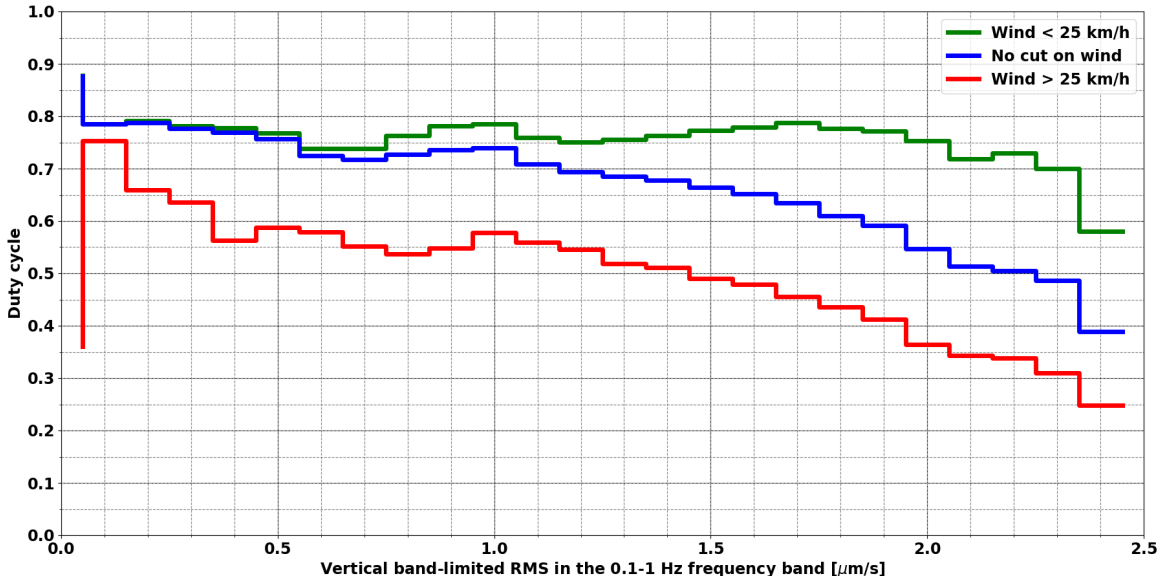

Figure 23: Virgo duty cycle during the O3 run versus microseism activity, for three different wind conditions: blue ↔ no cut on wind speed; green ↔ low wind (speed below 25 km/h); red ↔ high wind (speed above 25 km/h).

6. Other environment impacts

Additional sources of external noise have potential impact on the interferometer. Hereafter we describe those sources that we have further investigated during O3, namely: Schumann's resonance magnetic fields, lightning strikes and cosmic ray muons.

6.1. Magnetic noise

Ambient magnetic fields can couple to GW interferometers, for example through the magnetic actuators used for the control of the seismic isolation platforms of optical components and of the test masses [12, 48]. Like gravitational waves, electromagnetic (EM) waves travel at the speed of light, and, due to their strength, could affect multiple detectors with time differences compatible with those expected from some GW.

Magnetic fields that extend over the entire planet, such as the Schumann resonances [49] (SR), or large-current lightning strikes, can limit the sensitivity to GW signals correlated over multiple detectors [50, 51]. One purpose of the EGO external magnetometers (see Sec. 2) is to monitor the level of these global magnetic fields.

At Virgo, the external magnetic environment is much quieter than inside experimental halls where stray magnetic fields are radiated by electric loads and cables where large currents are circulating. Figure 24 compares inside and outside magnetometer spectra recorded at Virgo during O3 and in the very quiet environment inside the Sos Enattos mine in Sardinia [52]. The most intense spectral noise features are narrow lines at the

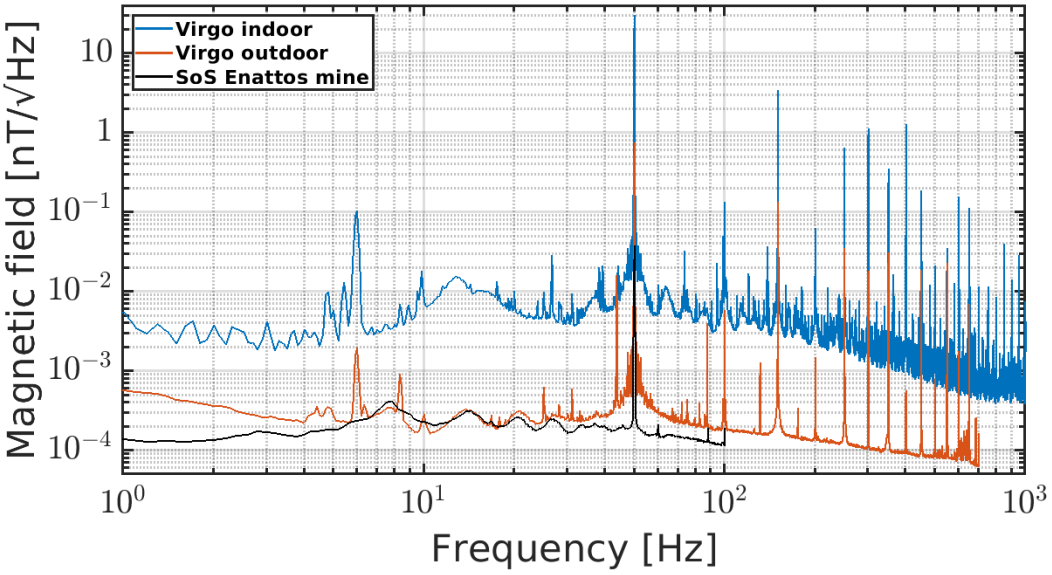

Figure 24: Amplitude spectral densities of indoor (blue curve) and outdoor (red curve) magnetometers at EGO and at Sos Enattos mine in Sardinia (black curve). The quiet Sos Enattos location shows evidence of Schumann resonances peaked at approximately 8, 14, 21, 27 and 33 Hz.

50 Hz electric mains frequency and its odd harmonics. The RMS amplitude of the 50 Hz line measured at Virgo is of the order of 0.1 nT in the external location, while it is at least 50 times larger in any inside location.

Virgo external magnetometers detect the SR field. This consists of steady EM waves that resonate inside the waveguide formed by the Earth surface and the ionosphere, and which are excited by globe-wide lightning activity. The second and third SR modes (peak frequency around 14 Hz and 21 Hz, respectively) are visible above noise at almost any time, their median amplitude during O3 is a few tenth of pT, their intensity follows a 24-hour modulation. The measured daily modulation of the third SR mode is shown in Fig. 25. This modulation is thought to be associated to temperature-driven variations in the height of the ionosphere EM waveguide [53]. The first SR mode and those of order greater than three, are often covered by anthropogenic magnetic noise. Figure 25 shows that during the COVID-19 lockdown period from March to May 2020, the external magnetic field median RMS in the low frequency region from 1 to 6 Hz reduced by about 50% with respect to the reference period between December 2019 and February 2020. At the same time, the magnetic field RMS amplitude between 18 Hz and 24 Hz around the 3^{rd} Schumann mode, did not change appreciably.

At EGO, anthropogenic external magnetic noise follows a daily modulation: broad maxima during working hours and minima around 01:00 LT. This noise has the form of short transients with intensity of $\approx$ 10 pT extending from DC up to approximately 20 Hz. We believe this noise is associated to train transits along railway tracks at about

666 6 km distance from the site. The sudden trunk-line change when a train passes from an
667 electro-duct section to another one creates stray currents and magnetic fields that are
668 observed as magnetic glitches at EGO. According to the measured coupling of ambient
669 fields [12] we estimate a negligible impact of Schumann's and anthropogenic magnetic
670 noise on the sensitivity of the future Virgo upgrades. More relevant might be the impact
671 of the correlated Schumann noise on multiple interferometers, which is under evaluation.

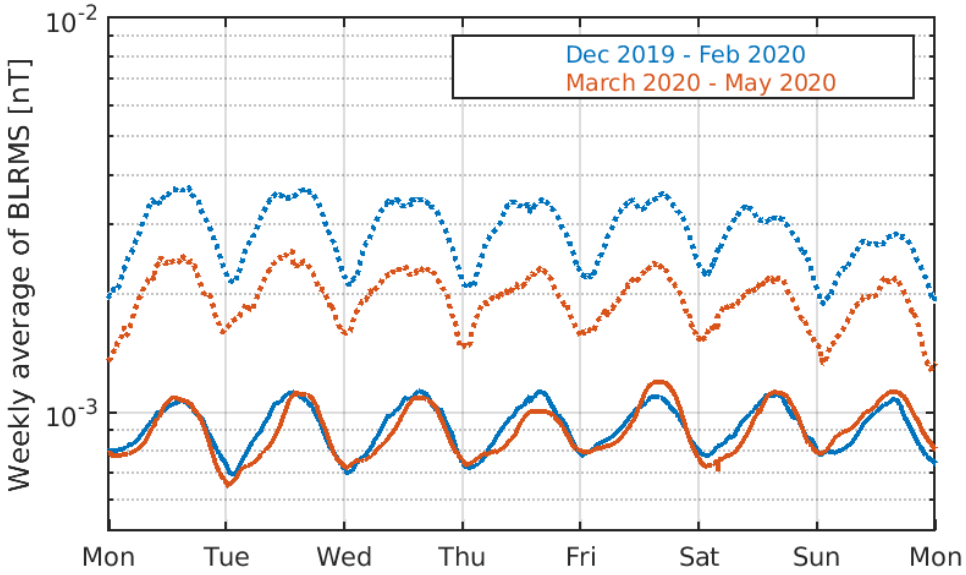

Figure 25: Weekly averaged magnetic field band-limited RMS values computed in two frequency bands: 1 to 6 Hz (dashed) and 18 to 24 Hz (solid). Magnetic field intensity is measured externally of Virgo experimental buildings, in the reference period between December 2019 and February 2020 (blue curves) and in the period between March 15 and May 15 (red curves) which corresponds to reduced anthropogenic activity within and outside of EGO because of the COVID-19 pandemics.

672 *6.2. Lightnings*

673 Lightning strikes produce prompt EM waves and much slower air pressure waves which
674 induce vibrations of the ground and of the detector mechanical components. There are
675 studies of correlated lightnings noise between the Virgo and LIGO sites [51] and at the
676 KAGRA underground observatory [54].

677 The typical effect of the impact of a lightning strike occurring at approximately 10
678 kilometers from the Virgo detector during O3 is illustrated in Fig. 26. A distinctive
679 feature of lightning strikes is a coincident short transient noise in magnetometers located
680 inside the 3 km-distant Virgo experimental buildings (top graph of Fig. 26). The
681 magnetic impulse is followed by the slower sound shock wave detected by seismometers

(middle graph of Fig. 26). The bulk of displacement noise reaching the buildings is below 10 Hz.

The bottom graph of Fig. 26 illustrates the effect of the lightning in the GW strain signal. In coincidence with the spike in magnetometers, we observe a prompt broadband low-frequency noise and the onset of a 48 Hz narrow spectral noise, with a minute-long decay time, leading to a ∼30% drop of the live BNS range. This latter noise has been associated to one structural mode of the West end test mass suspension, which gets excited because of the coupling of ambient magnetic fields with the magnetic actuators located along the suspension. Moreover, associated with the delayed acoustic and seismic bursts of ambient noise reaching the experimental buildings, a broadband strain noise shows up, extending up to about 100 Hz. This is likely due to scattered light processes within the interferometer.

Data quality flags triggered by lightning strikes were produced during the O3 run; they proved useful in a test aiming at filtering out part of the false-alarm triggers found by a real-time transient GW search [19]. Further studies are planned during the O4 run preparation.

6.3. Cosmic muons

Ground-based GW detectors are constantly passed through by *muons*, produced by the interaction of cosmic rays with Earth's atmosphere [55]. These energetic charged particles can interact with the detector test masses and constitute an additional source of noise, as addressed in the literature since the first prototypes of resonant mass GW detectors [56, 57, 58, 59].

We report here the preliminary results on the first measurement of potential effects of these muons on the Virgo detector noise. Further result can be found in [60]. This study has been carried out by means of about 17 days, at the end of the O3b run, of joint data acquisition of Virgo and a muon telescope designed by the IP2I laboratory [61], installed in the CEB close to the beam splitter mirror. Two kind of tests have been performed. In the first one, we have evaluated whether the rate of muons in the correspondence of GW candidate events was larger than the reference values of the period: we have found no statistical evidence of an excess of muons in correspondence of these triggers. In the second test, we have estimated the correlation of this rate with the rate of glitches in Virgo noise. Figure 27 shows the time series corresponding to the rates of glitches and muons, averaged on strides of 30 minutes. Here, a correlation is clearly evident. This is actually not surprising, for the number of the muons arriving at ground being highly dependent on air density and ultimately on parameters like atmospheric pressure and temperature. These quantities are also witnesses of the weather conditions, which in turn can determine an increase of the detector noise, as we have commented in Sec. 5.1. Therefore, both the variations of these rates share the same main cause, which explain

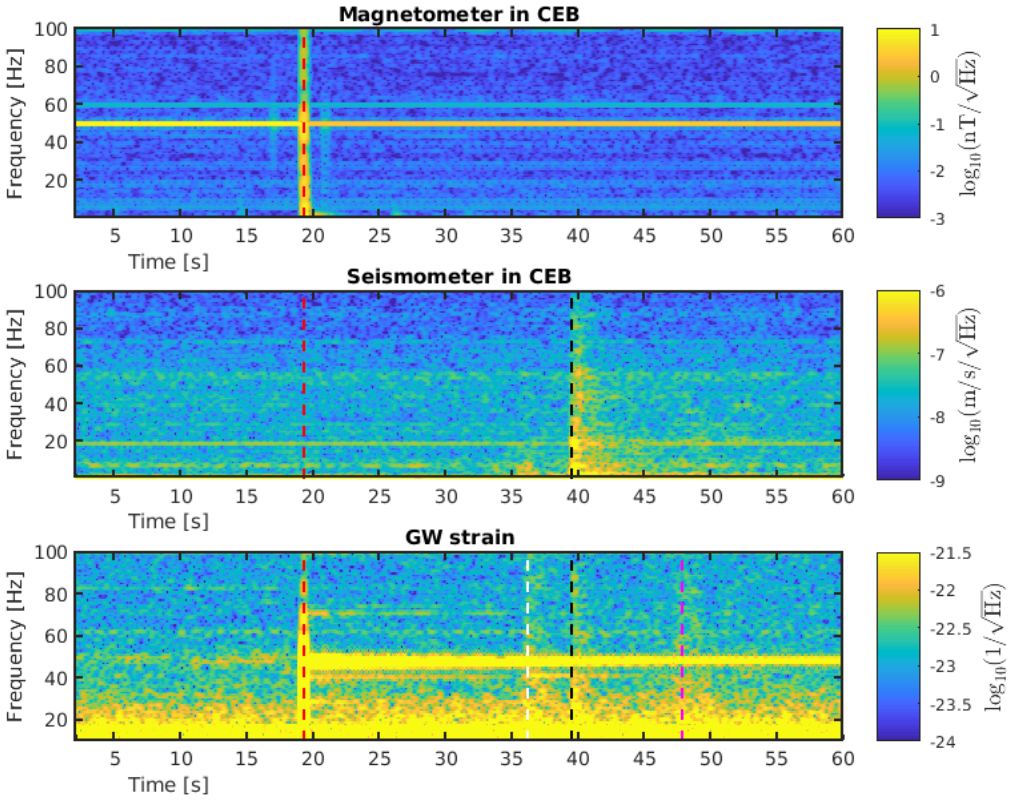

Figure 26: Impact on the Virgo environment and detector of one lightning strike which occurred 6 to 10 km away from Virgo buidings on November 15, 2019 at 23:25:51 UTC. The spectrograms of a few relevant signals are shown. (Top) A prompt magnetic transient is detected by magnetometers at the time of the event, marked by the red vertical line. (Middle) A few seconds later, a seismic (and acoustic, not shown) transient is detected in the central experimental area, marked by the black vertical line. The bottom spectrogram shows the reconstructed GW strain during the same time interval. The red vertical line marks the lightning strike occurence, the black, magenta and white vertical lines mark the occurrence of seismic transients detected in the CEB, NEB and WEB, respectively.

their large correlation. Once the effects of the atmospheric conditions are removed via a regression analysis, the residuals exhibit no significant correlation.

7. Outlook and prospects for O4

The Virgo detector performances are affected by external environment conditions; in particular, seismic noise, earthquakes, bad weather, magnetic noise and lightnings have an impact on the detector sensitivity or duty cycle. The main coupling mechanisms are:

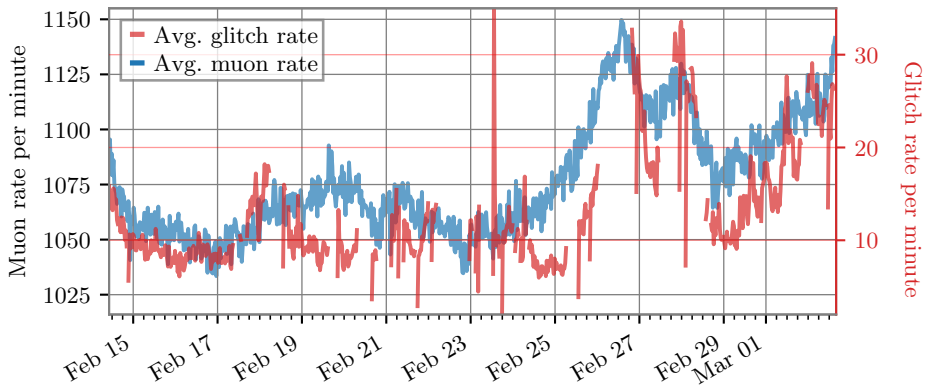

Figure 27: Correlation between muon and glitch rates. The blue line represents the time series of the average rate per minute of muons while the red one is the time series of the rate per minute of glitches with SNR > 4.5 and frequency at peak in $[10, 4096]$ Hz as identified by the Omicron pipeline [40]. Gaps in the latter correspond to periods when the detector was not taking data.

direct excitation of suspended mirrors, vibration of experimental buildings, shaking of benches hosting auxiliary optical systems, disturbances on critical electronic equipment, scattered light.

If the detector control system is able to manage the effect of a disturbance, the interferometer can remain at its working point with a reduced sensitivity. Otherwise the global control gets lost and the procedure to recover it has to be started again from the beginning, thus impacting on the duty cycle – see Appendix A.

In this work we reported the results of the analysis of such events during the O3 run. Thanks to the large amount of data collected, we were able to perform a careful statistical analysis of the impact of several kinds of external disturbances.

The results confirm that the Virgo detector is a very robust apparatus. The sensitivity reduction due to anthropogenic seismic noise is very low: less than 2% in terms of BNS range. Also the degradation due to the wind is limited: it appears only for wind speeds larger than 25 km/h, reaching a sensitivity reduction as large as 10% only for very high speed (larger than 50 km/h).

In these cases, the sensitivity reduction is due to an increased noise at low frequency as well as to the appearance of short high frequency glitches. In few cases, such kind of noise was indirectly originated by lightnings.

Both microseism and wind have an impact on the detector duty cycle, since the increasing correction signals acting on the mirror during bad weather can saturate, finally resulting to a control loss. It results that the Virgo detector global control is more robust against microseism while it is less effective against strong wind.

The analysis of control losses during O3 confirms that earthquakes are a relevant source

of these. The Seismon framework, useful to keep the detector in a safe state to try to avoid loosing its control during such events, was used during the whole O3 run and it is now being upgraded for the next scientific run.

An upgrade of the environmental monitoring system is in progress to better face the influence of external disturbances: installation of a new lightning detector in the central area; installation of two additional weather stations at the end buildings to monitor local wind gusts; and the installation of more sensitive accelerometers on locations prone to light scattering (viewports, external optical benches, etc.).

These actions, together with several other upgrades of the Virgo detector, already performed or presently in progress, will have a crucial role for the success of the next scientific run O4, which is expected to start at the end of 2022.

Acknowledgements

October 2020 version – `https://tds.virgo-gw.eu/ql/?c=15940`

The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Netherlands Organization for Scientific Research, for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Spanish Agencia Estatal de Investigación, the Consellera d'Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain, the National Science Centre of Poland and the Foundation for Polish Science (FNP), the European Commission, the Hungarian Scientific Research Fund (OTKA), the French Lyon Institute of Origins (LIO), the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek – Vlaanderen (FWO), Belgium. The authors gratefully acknowledge the support of the NSF, STFC, INFN, CNRS and Nikhef for provision of computational resources.

We would like to thank all of the essential workers who put their health at risk during the COVID-19 pandemic, without whom we would not have been able to complete this work.

Appendix A. Study of the control losses during O3

The Virgo detector needs to be controlled accurately in order to be sensitive to gravitational-wave signals [17, 18]. Schematically, there is an automated procedure [19] that brings the instrument from an initial state where the optics and the laser are controlled independently one from another, to the nominal state where the different

optical cavities are jointly resonant and the interferometer itself is used as a length etalon to control further the laser frequency. That procedure typically takes about 15-20 minutes and requires 1-2 attempts to complete. Then, the global control of the detector is kept as long as possible, with feedback loops maintaining Virgo at its nominal working point. When that control is lost for whatever reason, data taking stops and the control acquisition procedure has to be started again. This leads to a decrease of the instrument duty cycle and can cause transient gravitational waves to be missed. Therefore, it is important to find out the causes of the control losses and to use this information to improve the feedback systems and make them more robust.

As explained in Sec. 4 above, a global study of the control losses was needed to be able to extract those likely due to earthquakes. It was decided to focus on the 601 control losses that occured during O3 while the detector was taking data in nominal conditions (Science mode), to be sure that no particular human action was happening on the instrument at any of these times. Related to the duration of the O3 run (about 11 months) and to the duty cycle of the Virgo detector (about 75%), this corresponds to about 1 control loss every 10 hours of data taking on average. And, in reality, uninterrupted data taking stretches could be much longer as control losses usually cluster in time when a particular problem impacts the detector.

The first part of the study was to define the time the control loss occurred for each of these events. For that, we have used three different Virgo DAQ channels.

- Two fast channels, sampled at 10 kHz: `ARM_POWER`, latching when the power stored in the kilometric arm cavities goes below some threshold, meaning that they are not resonant anymore; `DARK_FRINGE_SHUTTER`, triggered when the fast shutter protecting the dark fringe photodiodes from an excess of light [2] closes.
- One slow channel, sampled at 1 Hz: `AUTOMATION_STATUS`, monitoring the global status of the detector, as seen by the automation process that steers the instrument.

`ARM_POWER`	`DARK_FRINGE_SHUTTER`	`AUTOMATION_STATUS`	Total
14	559	28	601

Table A1: Number of control losses in Science mode witnessed first by each DAQ channel used to time accurately control losses. As expected, the two fast channels are by far those that detect a control loss first. Most of the time the fast shutter protecting the dark fringe photodiodes closes before the arm power loss gets large enough to trigger the other fast channel.

The time of a control loss is defined as the earliest time one of these three switches flips from its nominal value to the value corresponding to an uncontrolled detector. Most of the time, as expected, the fast channels are the first ones to latch. And they do almost simultaneously, given that the cavity resonance losses are all connected. Though, in

practice, the dark fringe shutter closes almost always before the cavity arm power has decreased below its nominal threshold. In addition there are a few cases for which the central automation system triggers first a shutdown of the detector global control, either because it has detected an issue or because it has received a manual abort request from the operator on duty. Table A1 shows the breakout of witnesses for the O3 control losses that occurred while taking Science data.

Then, the selected strategy consists in testing several hypothesis in parallel for each of these events – the main hypothesis investigated are listed in Tabs. A2 and A3 and documented in the neighbouring text.

Various algorithms scanning the data around the control loss have thus been developed, with the twofold goals of being

- *complete*: to have as many control losses as possible tagged by at least one control loss hypothesis;
- *selective*: to find the right control loss origin as often as possible.

Achieving (close to) completeness requires testing many hypothesis, while a profusion of algorithms could be detrimental to the selectivity of the method. Therefore, the classification starts with a subset of hypothesis, those that, when identified, certainly cause a control loss and are also very likely to be the root cause of that particular event. Obvious examples in that category – called *sure* in the following – are control losses induced manually by the operator on duty, or hardware problems unambiguously identified by the real-time monitoring system of the Virgo detector. These control loss hypothesis are independent by definition and the associated algorithms should be selective. This has been checked by processing the 601 O3 control losses studied. All these events have been associated with at most one control loss hypothesis belonging to the sure category: 24% with one, 76% with none.

Error	Manual	Hardware	Control software	PI	Earthquakes	Total
2	10	92	7	2	30	143 (24%)

Table A2: Sure causes for 143 O3 control losses in Science mode – see text for details.

Table A2 provides details about the 143 control losses whose cause has been tagged as sure, as described above. The dominant class is hardware problems, mainly transient interruptions of the data flow coming from some suspensions and causing feedback control systems to fail. The faulty components have been identified and replaced during the post-O3 shutdown and upgrade phase. Therefore, these problems are not expected to reoccur during the O4 run. Then, earthquakes are the second most common source of control losses in the sure category; about three times a month on average. Manual control losses induced by the operator on shift follow: they are due to the need to switch from nominal data taking to another task: weekly maintenance, regular calibration or

commissioning activity. In O4 and beyond, such control losses should no longer occur as the procedure will be updated to require leaving Science mode before manually aborting the control. In 7 cases (only 1% of the total control losses) the source of the event could be traced to some software problem; 2 more cases were due to human errors.

Finally, two control losses are labelled as *PI* for parametric instabilities, an optomechanical phenomenon due to the interaction between optical and mechanical modes of the detector and that had been observed at LIGO in 2015 before finally being seen in Virgo as well in January 2020 [62]. If not mitigated, a PI can make control systems saturate in a deterministic way (meaning that the saturation will consistently reoccur as long as the detector remains in a configuration favourable for its appearance and growth), thus impacting the detector duty cycle. Moreover, it is impossible to predict exactly what combinations of the instrument parameters will lead to a PI. Therefore, a dedicated simulation framework has been developed to estimate the susceptibility of Virgo to PIs during O3, for O4, and beyond [63].

Fast unlocks	Actuation saturation	DARM control inaccuracy	Power loss in sidebands	Arm power asymmetry	Likely missing data	Automation decision	Others	Total
173	85	77	22	4	10	23	64 (11%)	458 (76%)

Table A3: Breakout by category of control losses not tagged as sure. 64 (about 11% of the total number of control losses recorded in Science mode during the O3 Virgo run) control losses have not been accurately classified, either because none of the tested hypothesis seemed to match the recorded data or because too many hypothesis were found matching, making their classification unconclusive. Further studies will be done when pre-O4 control losses data become available, in order to make the current classification more complete.

Table A3 describes how the remaining control losses (∼76%) have been classified. 11% of the total remain unclassified, either because none of the hypothesis tested matched, or because too many did and there was no clear way to find out which one was the root cause (if identified).

The largest category by far (29%) are the so-called *fast unlocks*, events that are almost instantaneous and occur within the laser injection system, upstream of the interferometer. Such control losses have been present for years, at rates that strongly vary over time, ranging from crisis periods lasting some hours to very quiet times. This past Summer, following detailled investigations of the fast unlock characteristics, this problem has been finally solved by installing [64] a pigtailed Electro-Optic-Modulator (EOM) with a larger dynamic, and thus able to compensate the laser frequency glitches that were found to be the cause of fast unlocks. They have not reoccured since this new EOM has been operated.

The next five categories are all related to the variety of feedback control systems that are running in parallel to keep the whole detector at its nominal working point. Improving the accuracy and the robustness of these systems while making the instrument more complex and thus more sensitive to the passing of gravitational wave is a permanent challenge, taken up during each upgrade or commissioning phase.

The analysis of the O3 control losses has been made using two independent software frameworks whose results have been compared: they have been found in good agreement, in particular for the dominant control loss categories. With the experience gained during O3, the goals for O4 are to improve the monitoring of the control losses and to reduce the latency of their analysis. A software framework similar to the Data Quality Reports (DQR) [19, 65, 66] used to vet in real time the gravitational-wave transient candidates that are significant enough to trigger a public alert is under development:here control losses play the role of GW event candidates and the set of checks run to assess the data quality is replaced by the test of the various hypotheses for control loss. This upgraded tool will be available to improve the overall performance of the instrument during the commissioning phase and associated noise-hunting for the new, double-recycled Advanced Virgo interferometer.

References

[1] Aasi J *et al.* (LIGO Scientific) 2015 *Class. Quant. Grav.* **32** 074001 (*Preprint* 1411.4547)

[2] Acernese F *et al.* (Virgo Collaboration) 2015 *Class. Quant. Grav.* **32** 024001 (*Preprint* 1408.3978)

[3] Akutsu T *et al.* 2021 *Progress of Theoretical and Experimental Physics* **2021** ISSN 2050-3911 05A102 (*Preprint* https://academic.oup.com/ptep/article-pdf/2021/5/05A102/38109702/ptab018.pdf) URL https://doi.org/10.1093/ptep/ptab018

[4] Abbott B *et al.* (LIGO Scientific Collaboration, Virgo Collaboration) 2016 *Phys. Rev. Lett.* **116** 061102 (*Preprint* 1602.03837)

[5] Abbott B *et al.* (LIGO Scientific Collaboration, Virgo Collaboration) 2017 *Phys. Rev. Lett.* **119** 161101 (*Preprint* 1710.05832)

[6] Abbott B *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, Fermi GBM, INTEGRAL, IceCube Collaboration, AstroSat Cadmium Zinc Telluride Imager Team, IPN Collaboration, Insight-HXMT Collaboration, ANTARES Collaboration, Swift Collaboration, AGILE Team, 1M2H Team, Dark Energy Camera GW-EM Collaboration, DES Collaboration, DLT40, GRAWITA, Fermi-LAT Collaboration, ATCA, ASKAP, Las Cumbres Observatory Group, OzGrav, DWF (Deeper Wider Faster Program), AST3 and CAASTRO Collaborations, VINROUGE Collaboration, MASTER Collaboration, J-GEM, GROWTH, JAGWAR, CaltechNRAO, TTU-NRAO and NuSTAR Collaborations, Pan-STARRS, MAXI Team, TZAC Consortium, KU Collaboration, Nordic Optical Telescope, ePESSTO, GROND, Texas Tech University, SALT Group, TOROS Collaboration, BOOTES Collaboration, MWA, CALET Collaboration, IKI-GW Follow-up Collaboration, H.E.S.S. Collaboration, LOFAR Collaboration, LWA, HAWC Collaboration, Pierre Auger Collaboration, ALMA Collaboration, Euro VLBI Team, Pi of Sky Collaboration, Chandra Team at McGill University, DFN, ATLAS Telescopes, High Time Resolution Universe Survey, RIMAS, RATIR, SKA South Africa/MeerKAT) 2017 *Astrophys. J. Lett.* **848** L12 (*Preprint* 1710.05833)

[7] Abbott B *et al.* (LIGO Scientific Collaboration, Virgo Collaboration) 2019 *Phys. Rev. X* **9** 031040 (*Preprint* 1811.12907)

[8] Abbott R *et al.* (LIGO Scientific Collaboration, Virgo Collaboration) 2021 *Phys. Rev. X* **11** 021053 (*Preprint* 2010.14527)
[9] Abbott R *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration and the KAGRA Collaboration) 2021 *arXiv e-prints* (*Preprint* 2111.03606)
[10] Abbott R *et al.* 2021 *The Astrophysical Journal Letters* **913** L7 URL https://doi.org/10.3847/2041-8213/abe949
[11] Abbott R *et al.* (LIGO Scientific Collaboration and Virgo Collaboration) 2021 *Phys. Rev. D* **103**(12) 122002 URL https://link.aps.org/doi/10.1103/PhysRevD.103.122002
[12] Fiori I *et al.* 2020 *Galaxies* **8** ISSN 2075-4434 URL https://www.mdpi.com/2075-4434/8/4/82
[13] Nguyen P *et al.* 2021 *Classical and Quantum Gravity* URL http://iopscience.iop.org/article/10.1088/1361-6382/ac011a
[14] Washimi T *et al.* 2021 *Classical and Quantum Gravity* **38** 125005 URL https://doi.org/10.1088/1361-6382/abf89a
[15] Punturo M *et al.* 2010 *Classical and Quantum Gravity* **27** 194002 URL https://doi.org/10.1088/0264-9381/27/19/194002
[16] Acernese F *et al.* 2004 *Astroparticle Physics* **20** 629–640 ISSN 0927-6505 URL https://www.sciencedirect.com/science/article/pii/S0927650503002603
[17] Acernese F *et al.* 2020 *Astroparticle Physics* **116** 102386 ISSN 0927-6505 URL https://www.sciencedirect.com/science/article/pii/S0927650519301835
[18] Allocca A *et al.* 2020 *Galaxies* **8** ISSN 2075-4434 URL https://www.mdpi.com/2075-4434/8/4/85
[19] The Virgo Collaboration 2021 *In preparation*
[20] Barone F, De Rosa R, Eleuteri A, Milano L and Qipiani K 2002 *IEEE Transactions on Nuclear Science* **49** 405–410
[21] Romero R *et al.* RADIO WAVES below 22 kHz URL http://www.vlf.it
[22] The OpenStreetMap contributors OpenStreetMap URL https://www.openstreetmap.org
[23] Peterson J R 1993 *Open-File Report* URL http://pubs.er.usgs.gov/publication/ofr93322
[24] Koley S *et al.* 2017 *SEG Technical Program Expanded Abstracts* 2946–2950 URL https://doi.org/10.1190/segam2017-17681951.1
[25] Longuet-Higgins M S 1950 *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* **243** 1–35
[26] Cessaro R K 1994 *Bulletin of the Seismological Society of America* **84** 142–148
[27] Flaminio R 2020 Status and plans of the Virgo gravitational wave detector *Ground-based and Airborne Telescopes VIII* vol 11445 ed Marshall H K, Spyromilio J and Usuda T International Society for Optics and Photonics (SPIE) pp 205 – 214 URL https://doi.org/10.1117/12.2565418
[28] Coughlin M *et al.* 2017 *Classical and Quantum Gravity* **34** 044004 URL https://doi.org/10.1088/1361-6382/aa5a60
[29] Biscans S *et al.* 2018 *Classical and Quantum Gravity* **35** 055004 URL https://doi.org/10.1088/1361-6382/aaa4aa
[30] Mukund N *et al.* 2019 *Classical and Quantum Gravity* **36** 085005 URL https://doi.org/10.1088/1361-6382/ab0d2c
[31] Product distribution layer git repository URL https://github.com/usgs/pdl
[32] Berni F *et al.* 2012 The Detector Monitoring System https://tds.virgo-gw.eu/ql/?c=9005
[33] F Berni 2020 DMS help manual https://tds.virgo-gw.eu/ql/?c=15469
[34] 2020 Virgo logbook entry validating the use of the eq mode control configuration to take science-quality data URL https://logbook.virgo-gw.eu/virgo/?r=48612
[35] 2020 Query to the public ingv website URL http://webservices.ingv.it/fdsnws/event/1/query?starttime=2019-04-01T15%3A00%3A00&endtime=2020-03-27T17%3A00%3A00&minmag=2&maxmag=10&mindepth=-10&maxdepth=1000&minlat=27.0&maxlat=48.0&minlon=-7.0&maxlon=37.5&minversion=100&orderby=time-asc&timezone=UTC&format=text&limit=10000

[36] INGV seismic surveillance center public website URL http://terremoti.ingv.it
[37] Bernardi F *et al.* 2015 *Natural Hazards and Earth System Sciences* **15** 2019–2036 URL https://nhess.copernicus.org/articles/15/2019/2015/
[38] Early-est: Earthquake rapid location system with estimation of tsunamigenesis URL http://early-est.rm.ingv.it/warning.html
[39] Welch P 1967 *IEEE Transactions on audio and electroacoustics* **15** 70–73
[40] Robinet F *et al.* 2020 *SoftwareX* **12** 100620 ISSN 2352-7110 URL https://www.sciencedirect.com/science/article/pii/S2352711020303332
[41] Canuel B, Genin E, Vajente G and Marque J 2013 *Opt. Express* **21** 10546–10562 URL http://www.opticsexpress.org/abstract.cfm?URI=oe-21-9-10546
[42] Wąs M, Gouaty R and Bonnand R 2021 *Classical and Quantum Gravity* **38** 075020 URL https://doi.org/10.1088/1361-6382/abe759
[43] Accadia T *et al.* 2010 *Classical and Quantum Gravity* **27** 194011 URL https://doi.org/10.1088/0264-9381/27/19/194011
[44] Van Heijningen J V *et al.* 2019 *Class. Quant. Grav.* **36** 7 URL https://iopscience.iop.org/article/10.1088/1361-6382/ab075e
[45] Valdes G, O'Reilly B and Diaz M 2017 *Class. Quant. Grav.* **34** 235009
[46] Longo A *et al.* 2020 *Class. Quant. Grav.* **37** 145011 (*Preprint* 2002.10529)
[47] Li H, Li Z and Mo W 2017 *Signal Processing* **138** 146–158 URL https://www.sciencedirect.com/science/article/pii/S0165168417301135
[48] Cirone A *et al.* 2018 *Rev. Sci.* **89** 114501 URL https://doi.org/10.1063/1.5045397
[49] Schumann W 1952 *Zeitschrift Naturforschung Teil A* **7** 149
[50] Coughlin M W *et al.* 2018 *Phys. Rev. D* **97**(10) 102007 URL https://journals.aps.org/prd/abstract/10.1103/PhysRevD.97.102007
[51] Kowalska-Leszczynska I *et al.* 2017 *Classical and Quantum Gravity* **34** 074002 URL https://doi.org/10.1088%2F1361-6382%2Faa60eb
[52] Naticchioni L *et al.* 2020 *Journal of Physics: Conference Series* **1468** 012242 URL https://doi.org/10.1088/1742-6596/1468/1/012242
[53] Sentman, D D 1995 *Schumann Resonances* (CRC Press) chap 11
[54] Washimi T *et al.* 2021 *Journal of Instrumentation* **16** P07033 URL https://doi.org/10.1088/1748-0221/16/07/p07033
[55] Maurin D, Melot F and Taillet R 2014 *Astronomy & Astrophysics* **569** A32
[56] Beron B L and Hofstadter R 1969 *Phys. Rev. Lett.* **23**(4) 184–186 URL https://link.aps.org/doi/10.1103/PhysRevLett.23.184
[57] Amaldi E and Pizzella G 1986 *Il Nuovo Cimento C* **9** 612–620
[58] Giazotto A 1988 *Physics Letters A* **128** 241–244 ISSN 0375-9601
[59] Chiang J, Michelson P and Price J 1992 *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **311** 603–612 ISSN 0168-9002
[60] The ENV team 2021 Future publication about muons separate paper about Virgo and cosmic muons
[61] Lesparre N *et al.* 2012 *Geoscientific Instrumentation, Methods and Data Systems* **1** 33–42 URL https://gi.copernicus.org/articles/1/33/2012/
[62] Puppo, P for the Virgo Collaboration 2021 Parametric Instability Observation in Advanced Virgo second European Physical Society Conference on Gravitation: measuring gravity URL https://agenda.infn.it/event/26098/contributions/132480/attachments/83185/109525/EPS_Online2021_PI.pdf
[63] Cohen D *et al.* 2021 Towards optomechanical parametric instabilities prediction in ground-based gravitational wave detectors (*Preprint* 2102.11070)
[64] 2021 Virgo logbook entry reporting the installation of the new pigtailed eom URL https://logbook.virgo-gw.eu/virgo/?r=52813

[65] The LIGO Scientific Collaboration and The Virgo Collaboration 2018 Data Quality Report User Documentation https://docs.ligo.org/detchar/data-quality-report
[66] Davis D *et al.* 2021 *Classical and Quantum Gravity* **38** 135014 URL https://doi.org/10.1088/1361-6382/abfd85