Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors.

S. Biscans[1],[*] J. Warner[2], R. Mittleman[1], C. Buchanan[3], M. Coughlin[4], M. Evans[1], H. Gabbard[5], J. Harms[6], B. Lantz[7], N. Mukund[8], A. Pele[9], C. Pezerat[10], P. Picart[10], H. Radkins[2], T. Shaffer[2]

[1]LIGO, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
[2]LIGO Hanford Observatory, Richland, Washington 99352, USA
[3]Department of Physics and Astronomy, Louisiana State University, Baton Rouge, LA 70803, USA
[4]Department of Physics, Harvard University, Cambridge, MA 02138, USA
[5]Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany
[6]Università degli Studi di Urbino "Carlo Bo", I-61029 Urbino, Italy
[7]Stanford University, Stanford, California 94305, USA
[8]Inter-University Centre for Astronomy and Astrophysics (IUCAA),
Post Bag 4, Ganeshkhind, Pune 411 007, India
[9]LIGO Livingston Observatory, Livingston, Louisiana 70754, USA
[10]Université du Maine, CNRS UMR 6613, 72085 Le Mans, France

Advanced gravitational-wave detectors such as the Laser Interferometer Gravitational-Wave Observatories (LIGO) require an unprecedented level of isolation from the ground. When in operation, they are expected to observe changes in the space-time continuum of less than one thousandth of the diameter of a proton. Strong teleseismic events like earthquakes disrupt the proper functioning of the detectors, and result in a loss of data until the detectors can be returned to their operating states. An earthquake early-warning system, as well as a prediction model have been developed to help understand the impact of earthquakes on LIGO. This paper describes a control strategy to use this early-warning system to reduce the LIGO downtime by ∼30%. It also presents a plan to implement this new earthquake configuration in the LIGO automation system.

INTRODUCTION

The Laser Interferometer Gravitational-wave Observatory (LIGO) consists of two identical, 4 kilometers long interferometric detectors installed at the Hanford, Washington (H1) and Livingston, Louisiana (L1) sites in the United States. The detectors are Michelson interferometers with Fabry-Perot resonant cavity arms [1]. A prerequisite for the detectors to collect scientific data is that the cavities are held in optical resonance.

LIGO completed its most recent observation run (referred as O1) from September 18, 2015 to January 12, 2016. During this period, the detectors were kept in an operating mode and the commissioning activities reduced to minimum. Keeping the detectors in operating mode is a complex task and large environmental disturbances can disrupt this process. These disturbances reduce the duty cycle and observing time, as summarized in table I.

Environmental hazards such as earthquakes were one of the primary sources of disturbance during the first observation run. Earthquakes with a Richter magnitude above 4.5 generated a non-negligible increase of the ground motion in vertical and horizontal directions, from a few mHz to 100mHz, as illustrated by the seismic vibration spectra in figure 1. This increase of the ground motion overwhelmed the LIGO seismic isolation system and prevented the detectors from operating. Once the cavities of the interferometer are out of resonance, it can take many hours to return to normal operation, as demonstrated in figure 10. Overall, H1 experienced 265 earthquakes and L1 243 earthquakes (magnitude 4.5 and above) while observing. Operation was disrupted 62 times at Hanford and 83 times at Livingston due to these earthquakes. This study will only focus on Hanford, knowing that similar results are expected for Livingston.

There is a direct correlation between the operating status of the interferometers and the ground velocity, as shown in figure 2a: the interferometer becomes unstable at higher velocities. An Earthquake mitigation scheme is thus developed to limit the extra disturbance induced by earthquakes. The goal of this configuration is to maintain operation at the expense of sensitivity. For this reason, it cannot be permanently activated and needs to be part of a smart automation system.

	H1	L1
Observation	66 %	59 %
Commissioning	9 %	7 %
Environmental disturbances	**17 %**	**24 %**
Other	8 %	10 %

TABLE I: Detectors' status over the O1 period. Commissioning time represents the vital maintenance tasks needed to keep the interferometers running. Environmental disturbances encompass earthquakes, high-wind and distant storms.

The first section of this paper describes the LIGO interferometers and its seismic isolation systems in more detail. We then present the earthquake strategy and how it will be integrated into the LIGO infrastructure.
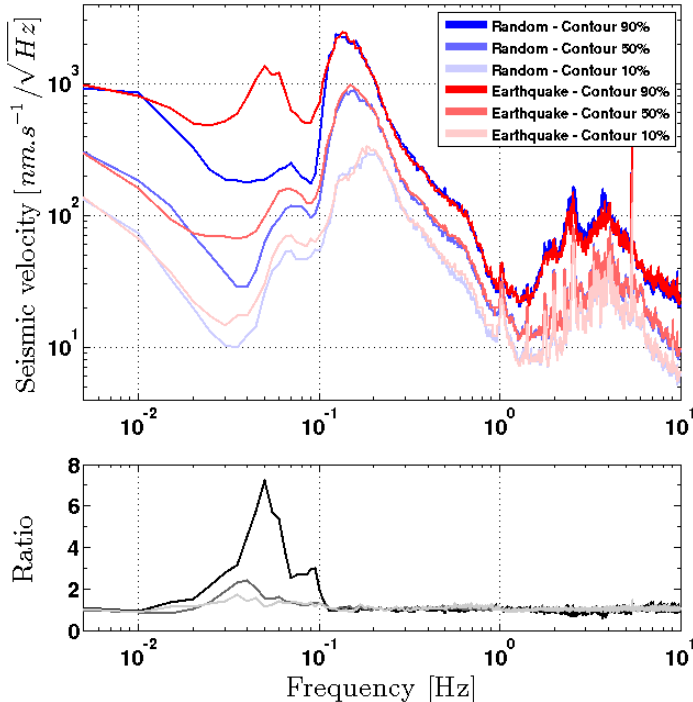
FIG. 1: 1000-second long data stretches were selected over the total time span of the O1 period at H1. Specifically, the blue curves correspond to stretches selected at random times, while the red curves correspond to stretches selected during Earthquakes of Richter magnitude 4.5 and above. For each frequency bin, the data were histogrammed and a set of probabilities was found. The different shades of colour indicate different probability values (10%, 50%, 90%), the darker tone being a 90% probability. The bottom part of this figure represents the ratio between the red and the blue curves for each set. We observe an amplification of the ground motion up to a factor of 7 in the [30mHz-100mHz] bandwidth. Only the horizontal direction along the Y-arm of the interferometer is represented here, but we obtain similar results in the X-horizontal direction and the vertical direction.

THE LIGO DETECTORS

The LIGO detectors are Michelson interferometers coupled with Fabry-Perot cavities in the arms. A beamsplitter is used to separate the input light into the two arms. A simplified optical layout of the LIGO detector is shown in figure 3. Each arm is comprised of an input test mass (ITM) and an output test mass (ETM) forming the Fabry-Perot arm cavity. Other cavities and auxiliary optics such as the power recycling mirror and the signal recycling mirror are present to enhance the signal quality [2].

All the LIGO's optics are attached to seismic isolation platforms, which seek to decouple the optics from the ground in the observational bandwidth (above 10Hz). A Hydraulic External Pre-Isolator (HEPI) supports the vacuum chamber. The optics are mounted inside the chambers on an active Internal Seismic Isolation (ISI)

system. There are two types of ISI systems: a Single-Stage platform and a Two-Stages platform. The Single-Stage platforms are used for the auxiliary optics in the small vacuum chambers called HAMs. The Two-Stages platforms are used for the core optics of the interferometers, in the large vacuum tanks called BSCs. In total, a detector has 11 vacuum tanks (six HAM chambers, five BSC chambers) each with a seismic isolation platform. Despite some mechanical differences between the different vibration isolation systems, the general concept is identical for all of them. A combination of active and passive isolation is provided to bring the BSC-ISI platform motion down to $1 \cdot 10^{-12}$ m/$\sqrt{\text{Hz}}$ at 10Hz and the HAM-ISI platform motion down to $2 \cdot 10^{-11}$ m/$\sqrt{\text{Hz}}$ at 10Hz.

The mechanics and functioning of these platforms have been extensively studied in previous works ([3], [4], [5]). This paper will focus only on the active isolation configuration and performance of the BSC-ISI Stage 1 platform. The BSC-ISI Stage 1 platform is of interest for this study as it provides low frequency active isolation for the core optics. All the BSC-ISI platforms being identical, we chose to look at the ITM chamber of the Y arm (called ITMY) in the horizontal Y-direction. Figure 4 illustrates the layout of the BSC chamber.

THE SEISMIC PLATFORM CONTROL SCHEME

Each stage is equipped with a set of actuators, displacement sensors and inertial sensors. They are used to actively control the stage in the three translational and three rotational degrees of freedom. The platforms have been designed to limit the cross-coupling between the different degrees of freedom, therefore, each degree of freedom can be actively controlled independently with Single-Input Single-Output compensators.

The LIGO seismic control scheme is a combination of feedback, feedforward and sensor correction. The block diagram in figure 5 shows the simplified control topology for one degree of freedom.

Feedback control

A control force F is used to reduce the inertial motion of the stage (X_{Stage}), which is induced by the ground (X_{G}). This control force is created using a combination of feedback and feedforward controllers. The feedback controller is fed by two sensors: a displacement sensor measuring the relative motion between the stage and the input motion ($X_{\text{Stage}} - X_{\text{G}}$), and an inertial sensor (seismometer) measuring the inertial motion of the stage (X_{Stage}). We wish to reduce the inertial motion measured by the seismometers, but cannot below 25mHz because of noise limitations. Therefore, a displacement sensor is
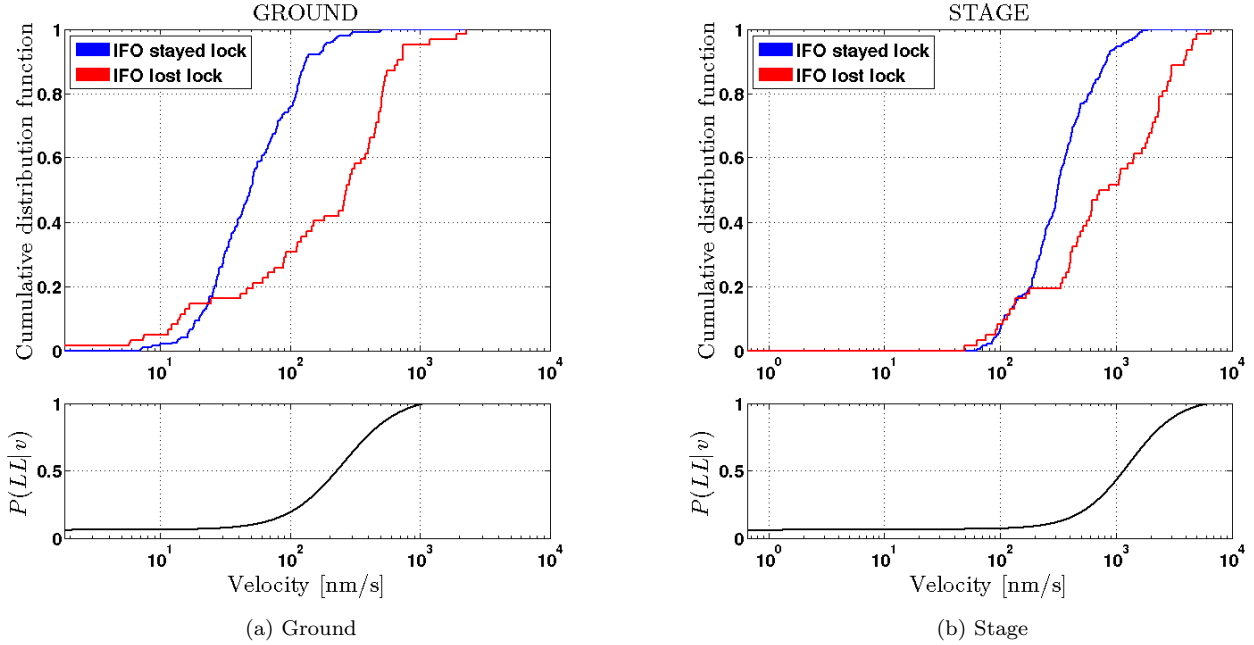
FIG. 2: Comparison of Stage 1 ITM behavior in the [30mHz-100mHz] bandwidth for different ground motions: stretches selected during earthquakes when the interferometer survived (blue curve), stretches selected during earthquakes when the interferometer stops functioning (red curve). The top part of the figure represents the cumulative distribution function for the ground and the stage respectively, as a function of velocity. The plots indicate a direct correlation between velocity and the interferometer status. We observe a net increase of the stage velocity compared to the ground, due to a self-inflicted gain peaking in this frequency band. The bottom part of the plots represents the smoothed probability of losing lock (dropping out of optical resonance) as a function of velocity.
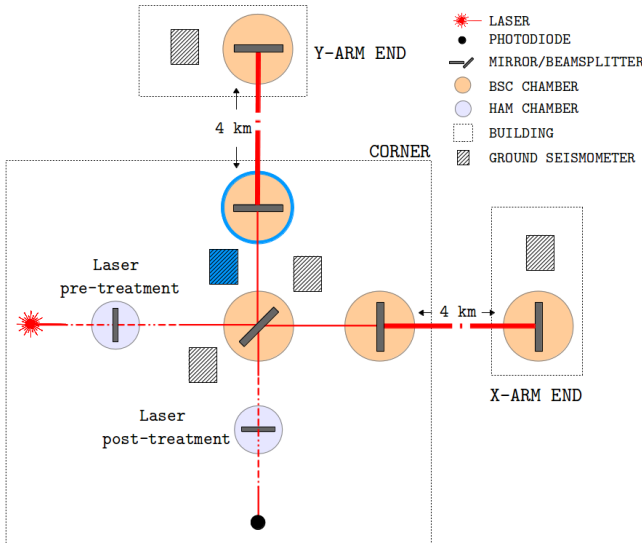

FIG. 3: Simplified optical layout of the LIGO detector. The data presented in this study is extracted from the blue ground seismometer and the BSC chamber circled in blue (called ITMY).
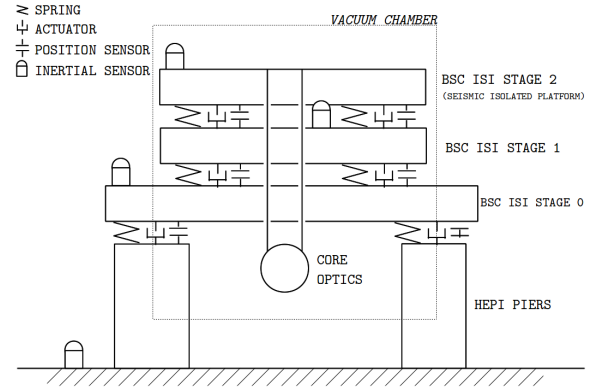

FIG. 4: Schematic of the LIGO BSC chamber. Each stage is equipped with multiple actuators, position and inertial sensors (only a few are represented here for clarity). The core optics are supported by a quadruple pendulum (not shown for clarity) which provides additional seismic isolation in all degrees of freedom.

used at low frequencies and both sensors are blended together to feed the controller. The relative motion signal is low-passed by a filter L_{disp}, and the inertial motion signal is high-passed by a filter H_{in}. L_{disp} and H_{in} are designed to be complementary, meaning $L_{disp} + H_{in} = 1$. The frequency at which the low-pass and high-pass filters cross is called the blend frequency. The controller provides isolation up to 30Hz, with high loop gain below 1Hz. The expression of the stage inertial motion when the feedback control is on and the feedforward and sensor correction are off is given by equation 1. N_{pos} and
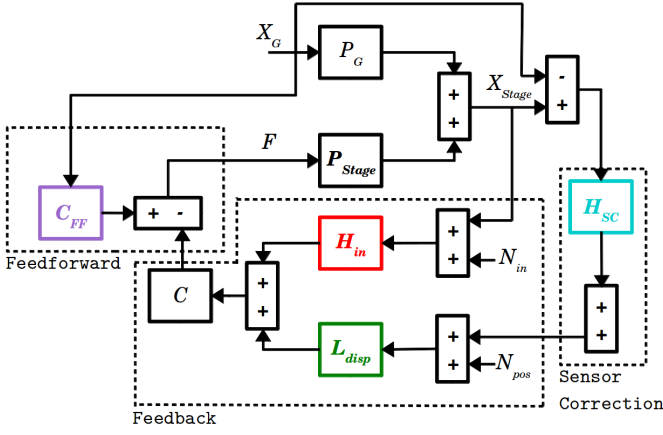
FIG. 5: Control block diagram of a LIGO seismic isolation stage for one degree of freedom. The coloured blocks are related to figure 6.

N_{in} represent the noise associated with the displacement sensors and inertial sensors respectively. P_G and P_{Stage} represent the plant model of the ground and stage respectively.

$$|X_{Stage}|^2 = |\frac{P_G + L_{disp}CP_{Stage}}{1 + CP_{Stage}}X_G|^2$$
$$+ |\frac{H_{in}CP_{Stage}}{1 + CP_{Stage}}N_{in}|^2 \qquad (1)$$
$$+ |\frac{L_{disp}CP_{Stage}}{1 + CP_{Stage}}N_{pos}|^2$$

Sensor Correction

Sensor correction is a feedforward technique using a seismometer from the ground. The seismometer signal X_G is filtered by H_{SC} and added to the position sensor signal $(X_{Stage} - X_G)$ to create a virtual inertial sensor [6]. Our sensor correction is designed to increase performance around 100mHz. By adding the sensor correction to the feedback loop, the stage inertial motion becomes:

$$|X_{Stage}|^2 = |\frac{P_G + L_{disp}CP_{Stage}(1 - H_{SC})}{1 + CP_{Stage}}X_G|^2$$
$$+ |\frac{H_{in}CP_{Stage}}{1 + CP_{Stage}}N_{in}|^2 \qquad (2)$$
$$+ |\frac{L_{disp}CP_{Stage}}{1 + CP_{Stage}}N_{pos}|^2$$

Feedforward control

A standard feedforward controller is added from the ground in parallel with the feedback and sensor correction loops. This operates where the coherence with the ground is high (above 1Hz). The stage absolute motion becomes:

$$|X_{Stage}|^2 = |\frac{P_G + L_{disp}CP_{Stage}(1 - H_{SC}) + C_{FF}P_{Stage}}{1 + CP_{Stage}}X_G|^2$$
$$+ |\frac{H_{in}CP_{Stage}}{1 + CP_{Stage}}N_{in}|^2$$
$$+ |\frac{L_{disp}CP_{Stage}}{1 + CP_{Stage}}N_{pos}|^2$$
$$(3)$$

Equations 1, 2 and 3 are plotted in an example in figure 6 to illustrate the performance of each loop and the combined overall performance. In this example, we used the nominal BSC-ISI stage 1 filters on a damped single degree of freedom system with a 2Hz resonance (which is similar to the BSC-ISI Stage 1). The figure shows the transfer function between the stage and ground motion $\frac{X_{Stage}}{X_G}$ with the feedback loop only (solid orange curve), with the feedback loop and sensor correction on (solid brown curve) and with feedback, sensor correction and feedforward on (solid black curve). The motion associated with sensor noise is not represented on this figure for clarity. The open loop (not represented) has a 30Hz unity gain frequency. The Sensor correction filter (dash cyan curve) is designed to provide extra isolation between 50mHz and 200mHz, whereas the feed forward filter (dash purple curve) provides isolation at 1Hz and above. At low frequencies where the loop gain is effectively infinite, the performance is limited by the low pass filter (dash green curve), and limited by the finite loop gain at higher frequencies. Typically, the low pass filter is tuned to provide as much isolation as possible in the control bandwidth at the cost of some gain peaking around the blend frequency (in this case ~ 45mHz). Some sharp notches are also present in this filter to target known payload resonances.

NOMINAL FILTERS

The nominal filters are designed to meet LIGO requirements and maximize the seismic isolation above 100mHz, at the expense of some gain peaking at lower frequencies. The blue curve in figure 7 demonstrates the performance of this configuration at H1 during a typical ground motion period. The stage provides a factor 38 of isolation at 200mHz at the expense of a gain peaking of 4 at 54mHz. While this gain peaking is not a limitation during typical ground behavior, it becomes problematic during an earthquake.

Figure 2 compares the performance of the stage during earthquakes. It shows a clear amplification of the ground motion by the stage over the bandwidth of interest (figure 2a versus figure 2b). With this configuration,
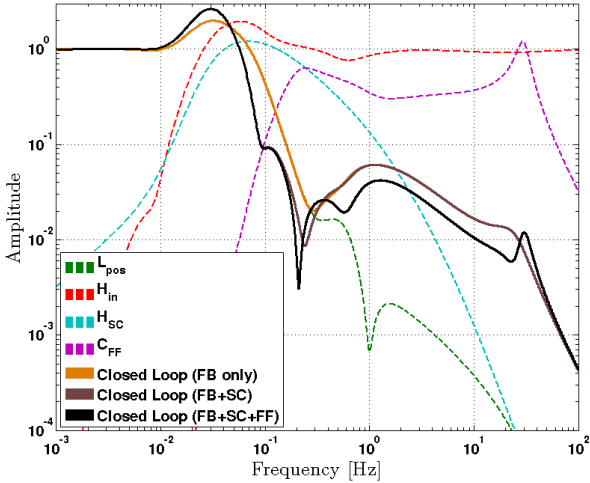
FIG. 6: Example of the LIGO seismic control scheme on a damped single degree of freedom system. The dash curves shows the different filters used, as opposed to the solid curves showing the transfer functions with different loops engaged.
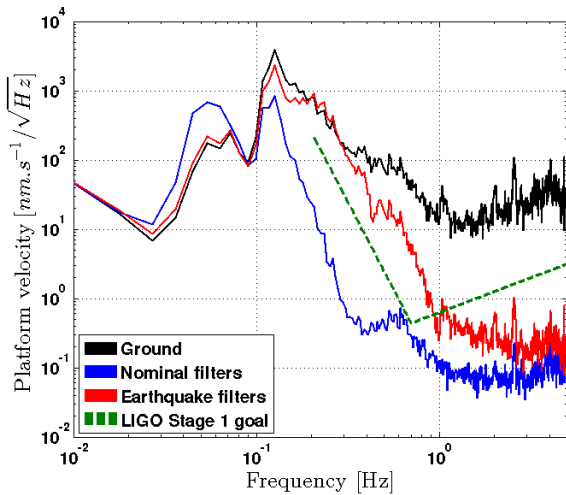

FIG. 7: Seismic isolation provided by BSC-ISI Stage 1 in the horizontal Y-direction at H1. The black curve represents a typical ground motion, and the blue curve the measured motion of the stage. The red curve is a simulation of the stage motion with the Earthquake filters in place. The dotted curve indicates the LIGO goal to obtain from 200mHz to higher frequencies for Stage 1.

the interferometer has a 50% chance of losing lock for a stage velocity higher than 1000 nm/s.

EARTHQUAKE FILTERS

The goal of the earthquake control configuration is to find a good balance between limited gain peaking at low frequencies (because the nominal gain peaking is at the maximum earthquake power) and performance at higher frequencies. The gain peaking observed is directly depen-
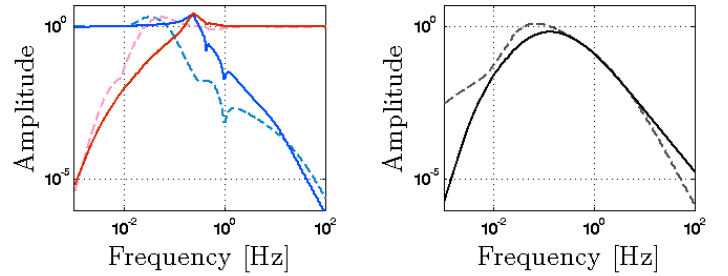

FIG. 8: Comparison of the filters used during O1 (dash lines) and the new designed filters for Earthquakes (solid lines). The left part of the figure show the complementary low-pass and high-pass filters. The right part shows the sensor correction filters.

dent on the low-pass filter L_{disp} and the sensor correction filter H_{SC}, as demonstrated in equation 4. The feedback and feedforward controllers stay untouched.

$$\lim_{CP_{Stage}\to\infty} |X_{Stage}|^2 = |(L_{disp}(H_{SC}-1)X_G)|^2$$
$$+ |H_{in}N_{in}|^2 \qquad (4)$$
$$+ |L_{disp}N_{pos}|^2$$

Figure 8 shows a comparison between the newly designed low, high and sensor correction filters with the nominal filters. To move the gain peaking out of the earthquake band, the blend frequency between the low-pass and the high-pass filter has been increased from 45mHz to 250mHz, and the sensor correction has been slightly modified to be less aggressive. The overall gain peaking is reduced by a factor of 3.1 with this configuration. However, there is no longer any isolation at 200mHz, which greatly impacts the sensitivity of the detectors. The stage does not provide the required seismic isolation to observe gravitational waves when in operation, but will deliver enough isolation to keep the interferometer locked most of the time. It will not provide enough isolation to keep the detector locked if other disturbances, such as high-wind or distant storms, increase the ground motion around 200mHz. Figure 7 shows the simulated stage 1 motion with the new earthquake configuration (red curve).

IMPROVEMENT OF LIGO ROBUSTNESS AND DUTY CYCLE

In this section we estimate the improvement in duty cycle due to reducing the sensitivity to earthquakes. Earthquake data collecting during O1 is used to simulate the effect of the new earthquake filters on Stage 1 velocity. We compare the averaged gain peaking induced by O1 earthquakes in the [30mHz-100mHz] band between the nominal and the earthquake filters. The velocity distribution
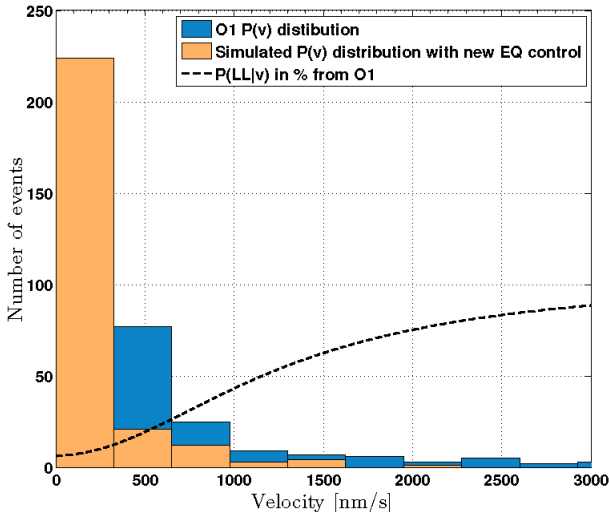
FIG. 9: New P(v_{new}) distribution based on O1 data with P(v_{new}) = $\frac{P(v)}{1.5}$. As a reminder from figure 2, we plotted the probability of loosing lock as a function of stage velocity P(LL|v) in black.
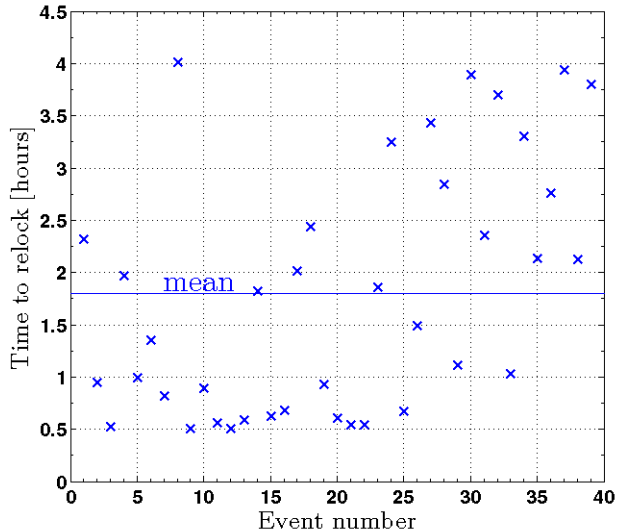

FIG. 10: Time spent returning to science mode after a problematic earthquake. If the time-span to re-acquire resonance was longer than 5 hours, it was assumed that the interferometer was put in stand-by and the event was not taken in account for the statistics. It takes 1.8 hours on average to return the instrument to its previous state.

will change from P(v) to P(v_{new}), with P(v_{new}) = $\frac{P(v)}{1.5}$. Based on this new distribution (plotted in figure 9), we can re-calculate the probability of losing resonance P(LL) as written in equation 5. We expect a 32.3% reduction.

$$P(LL) = \sum_{\substack{velocity \\ bin}} P(LL|v)P(v_{new}) \qquad (5)$$

Each time the interferometer goes out of operation there is an associated downtime. Due to the complexity of the instrument and the duration of an earthquake, it can take several hours to re-acquire resonance. During O1, it took an average of 1.8 hours to go back to science mode after an earthquake, as explained in figure 10: H1 was down more than 111 hours because of earthquakes. If the number of times the interferometer goes out of resonance is reduced as expected, the downtime will be also be reduced from 111 to 75 hours.

IMPLEMENTATION AT THE SITES

As the earthquake control configuration described in this paper doesn't meet LIGO observing requirements, it is important to switch to this configuration only when an problematic earthquake arrives at the site. An early alert system called Seismon has been developed and is currently installed at the observatories [7]. Seismon predicts the arrival time of the primary, secondary and surface waves of an earthquake with a few seconds accuracy, and with a notification latency of less than 10 minutes. This prediction is using data from a world-wide network of seismic observatories that has been compiled by the United States Geological Service (USGS) and can be accessed on the World-Wide Web [8]. This gives enough time to switch the seismic configuration as needed.

The second role of Seismon is to predict, based on a logistic regression algorithm explained in [7], the likelihood of the interferometer to go out of science mode. This gives us the information to assess if an upcoming event would be problematic. If there is a high likelihood of losing resonance, we would trigger a change in the control configuration using the LIGO automation system called Guardian.

Each interferometer is supervised by a state machine called Guardian [9]. It consists of state machine automation nodes capable of handling control changes automatically. It is composed of multiple nodes, organised in a hierarchical fashion for each system and subsystem. In the case of the BSC-ISI, multiple intermediate states are required to bring the platform from the initial state to a fully isolated chamber. More states can be added to enable the switch between low-pass, high-pass and sensor correction filters. This change of state will only require a few minutes, as it is possible to switch these filers in either direction without having to turn off the isolation loops. This is due to a filter switching system which is already part of the LIGO infrastructure [10].

CONCLUSION

In this paper, we have discussed the problem induced by earthquakes on the LIGO gravitational-wave detectors and presented a seismic control strategy to minimize their impact. This new design represents a 32.3% improvement of the interferometers' robustness to earthquakes, which implies a direct increase in the number of detections for LIGO. We have described that the LIGO automation infrastructure is capable of switching to this new configuration based on early-warning predictions. Further effort will be spent on implementing this strategy for future observation runs.

ACKNOWLEDGEMENTS

* Corresponding author: sbiscans@ligo.mit.edu

[1] BP Abbott, R Abbott, TD Abbott, MR Abernathy, F Acernese, K Ackley, C Adams, T Adams, P Addesso, RX Adhikari, et al. Gw150914: The advanced ligo detectors in the era of first discoveries. *Physical review letters*, 116(13):131103, 2016.

[2] J Aasi, BP Abbott, R Abbott, T Abbott, MR Abernathy, K Ackley, C Adams, T Adams, P Addesso, RX Adhikari, et al. Advanced ligo. *Classical and Quantum Gravity*, 32(7):074001, 2015.

[3] F Matichard, B Lantz, R Mittleman, K Mason, J Kissel, B Abbott, S Biscans, J McIver, R Abbott, S Abbott, et al. Seismic isolation of advanced ligo: Review of strategy, instrumentation and performance. *Classical and Quantum Gravity*, 32(18):185003, 2015.

[4] R Abbott, R Adhikari, G Allen, S Cowley, E Daw, D DeBra, J Giaime, G Hammond, M Hammond, C Hardham, et al. Seismic isolation for advanced ligo. *Classical and Quantum Gravity*, 19(7):1591, 2002.

[5] R Abbott, R Adhikari, G Allen, D Baglino, C Campbell, D Coyne, E Daw, D DeBra, J Faludi, P Fritschel, et al. Seismic isolation enhancements for initial and advanced ligo. *Classical and Quantum Gravity*, 21(5):S915, 2004.

[6] B Lantz. *Description of the Sensor Correction FIR and IIR Filter Components*, 2012. `https://dcc.ligo.org/LIGO-T1200285/public`.

[7] M Coughlin, P Earle, J Harms, S Biscans, C Buchanan, E Coughlin, F Donovan, J Fee, H Gabbard, M Guy, et al. Limiting the effects of earthquakes on gravitational-wave interferometers. *Classical and Quantum Gravity*, 34(4):044004, 2017.

[8] *Data from the IRIS network of the USGS can be found at.* `https://earthquake.usgs.gov/earthquakes/`.

[9] JG Rollins. Distributed state machine supervision for long-baseline gravitational-wave detectors. *arXiv preprint arXiv:1604.01456*, 2016.

[10] R Kurdyumov, C Kucharczyk, and B Lantz. *Blend Switching User Guide*, 2012. `https://dcc.ligo.org/LIGO-T1200126/public`.