Search for Cosmic Ray Induced Transient Noise at LIGO Hanford
LIGO-DCC-T2400290

Benjamin Mannix and Raymond Frey
University of Oregon

Abstract

We analyze LIGO data from the O4a observing run to search for excess transient noise in LIGO's gravitational wave channel (DARM) induced by cosmic ray showers. The cosmic ray detector at LIGO Hanford and its characterization with data is described. Several methods are discussed and implemented in attempts to uncover induced noise. We find no evidence for an association of transient noise in DARM with cosmic rays, including dense showers of ∼1000 cosmic ray muons incident near a suspended test mass.

1 Introduction

Cosmic ray particles, primarily muons, reach the earth's surface at a rate of approximately $60/\mathrm{hr}/\mathrm{cm}^2$ at sea level. Several potential mechanisms have been considered by which cosmic ray particles might produce observable noise transients in detectors such as LIGO through interactions of cosmic rays with the suspended optics, especially the final test masses. Cosmic ray muons pass through an optic, depositing energy via ionization of the medium in the process. This results in direct transfer of momentum and heating to the optic mass. Also, the addition of charge might create transient disruptions of the test mass with the electrostatic drive. As discussed briefly below and by Braginsky, et al.[1], an isolated cosmic ray muon will not impart an impactful momentum transfer to a test mass. A similar conclusion was reached for transient modes rung up by heating of the optic[2]. On the other hand, a different calculation of transients induced by cosmic ray heating[5] implicated cosmic rays as the source of blip glitches. An analogous, detailed calculation has not been carried out for potential charge-related transients, however the potential for noise due to charge hopping was considered by Weiss[3]. The Weiss model was used to estimate potential charging effects for aLIGO by Harry[4], who found such effects to be negligible.

Even very energetic cosmic ray showers, which originate high in the atmosphere and become diffuse at sea level, would not be considered problematic sources of transient noise. However, a high energy primary cosmic ray particle (typically protons) can interact near sea level and initiate a large cosmic ray shower which remains relatively collimated as it interacts with the LIGO detectors. Such a possibility was also considered in Ref.[2] and found to be vanishingly unlikely. Nevertheless, as near-detector showers have the greatest potential for being observable in DARM, and given the uncertainties in predicting the effects, especially for charge-related mechanisms, we designed the cosmic ray detector to be tightly proximate to a test mass. For example, taking the hypothesis that blip glitches are due to large cosmic ray showers, it would be expected that approximately 1/4 of the glitches observed at LHO – those affecting ITMX – would have an associated signal in the cosmic ray detector.

As shown in Fig. 1, a cosmic ray detector is installed on the LVEA floor, directly below a test mass (ITMX) at LHO. It was originally installed for iLIGO[9] and remains in the same location for aLIGO. The iLIGO system employed a readout which produced event records which were independent of LIGO frame data. Nevertheless, the system worked well[10]. No correlation between (large) cosmic ray signals and DARM transients was observed. Had an effect been observed, more detectors would have been built to provide full coverage of test massses at both sites.[1] The same instrumentation, but with different readout electronics, was employed at LHO for aLIGO. This was used successfully for initial GW discoveries to demonstrate that cosmic rays were not responsible for the observed signals. However, as the GW detectors become increasingly sensitive, it is important to continue experimental interrogation, especially since modeling the potential effects have uncertainties.

In this note, we present results from an evaluation of cosmic ray and DARM data during the O4a run. This was aided by an enhancement to the electronic readout during the O4a observing run. The following section describes the cosmic ray instrumentation at LHO. This is followed in Section 3 with results for the response of the detector to cosmic rays. Section 4 provides the main results searching for transient noise in DARM coinciding with cosmic ray showers. Finally, Section 5 provides a summary and perspective.

2 Instrumentation

As shown in Fig. 1, the LHO cosmic ray detector consists of two 1-inch thick, 80 cm by 80 cm plastic scintillator slabs, which were rescued from a particle physics project at the SLAC National Lab, cut and polished and each affixed with two photomultiplier tubes. For each detector, one photomultiplier tube (PMT) is low gain (10 stages) and one high gain (11 stages). This combination provides a very large dynamic range, so that they are sensitive to single cosmic ray muons, but will not saturate for the largest imaginable low-altitude showers, i.e. thousands of muons). Having the PMTs coupled directly to the (cut and polished) scintillator corners without any additional light collection (e.g. a wavelength shifting edge panel) meant that the photon efficiency was sub-optimal and non-uniform over the detector area. It was expected that the overall efficiency for detecting single cosmic ray muons is 10-20%. However, since the focus of the design was for multi-muon showers, this was not considered problematic and was out-weighed by the simple, robust design. Cosmic ray showers are easily distinguished from any noise artifacts by demanding time-coincident signals in the two separate scintillator slabs, and less robustly by time coincidences between the two PMTs of each slab. In fact, coincident signals could be easily seen using an oscilloscope.

The intrinsic timing resolution of the cosmic ray detector (1 ns) far exceeds what is needed for LIGO. However, in the iLIGO readout, the timing resolution was preserved by using a readout path which did not write into frame data. The rising edge of the signals were used for timing, while the amplitude was taken from the integrated signals. For aLIGO, the decision was made that it would be more valuable to include the data in standard frame

[1] A cosmic ray detector was also built and sent to LLO, but was never installed.

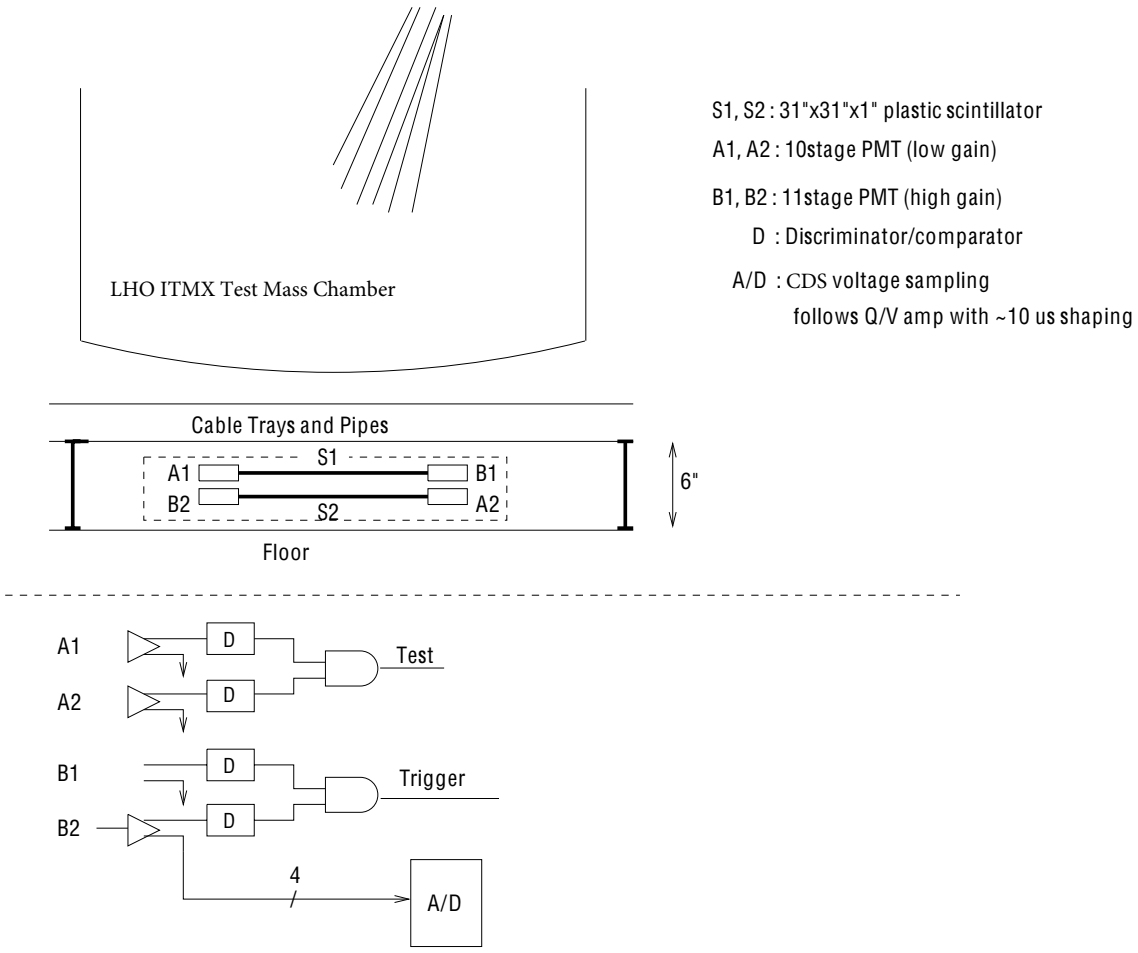

Figure 1: Schematic of cosmic ray detector at LHO. Above the dashed line is a sketch of the physical layout. Below the dashed line is a diagram of the front end electronics.

format, so the slower, integrated signals were to be used. The responsibility for carrying out the new readout was taken by LIGO Lab. But unfortunately, this was not carried out. The only signal which appeared in the frames for runs O1-O3 was a discriminator output (labeled 'Test' in the figure). This was good enough to make checks for GW candidate events. For example, for GW150914, the distribution of these signals in time around the GW signal was found to be consistent with the Poissonian expectation for the nearby off-source data. However, lacking an amplitude channel(s), this readout was not sufficient to look for subtle transients, since such effects would be expected to be correlated with larger cosmic ray showers.

In Summer 2023, the analog signals from the old electronics, which included a front-end enalog integrator with $\sim$ 10 μs time constant, were connected directly to standard CDS A/D converters (labeled 'A/D' in the figure). There are two issues with the configuration as implemented. First, the A/D digitization rate of 16 kHz is a bit too slow compared to the 10 μs shaping time of the analog front end integrator. And because there is no peak sensing, this means that on average we miss about 20% of the peak amplitude, with an estimated RMS error of $\sim$ 10%. (**To be checked!**) Second, the analog electronics output appears as a differential pair, to be consumed as such by the A/D converter. However, as seem in Fig. 3 both polarities of the analog output are digitized, resulting in bi-polar signals. We did check that the negative and positive peak amplitudes are tightly correlated. In the analyses below, we use only the positive polarity [**check**] signals. Nevertheless, these limitations were found to not limit the analyses presented here in any significant manner.

The analysis software can easily take the place of any coincidence logic indicated in the figure. Happily, the analog signals were still healthy, although one (of the four) PMT channels showed indications of a degradation, most likely in the front end analog electronics. However, we found this to not be crucial to our analysis. The subsequently recorded data from the O4a run from Fall 2023 is used for the analyses reported in this note.

3 Observing Cosmic Rays

We tested the efficiency of the cosmic ray detector using the predicted rate of 60 muons/hr/cm^2 and the area of the scintillator plates ($\sim$6,400 cm^2). Thus, we would predict that roughly 384,000 muons/hr or 100 muons/s are passing through the cosmic ray detector. We can compare that to the rate of events seen in our detector in Figure 2. Figure 2 shows the rate of cosmic ray events taken from a week of data. We can test the efficiency of our detector by focusing on the low amplitude peak with finer binning, which is shown in Figure 2b. The shape of this low amplitude peak aligns with what we expect of the rate of individual muons. Integrating under this peak gives a rate of 40 muons/s and thus a detector efficiency of $\sim$40%. In this study we are interested in studying the highest amplitude events on the right side of Figure 2a which correspond to cosmic ray showers with the largest number of particles passing through our detector.

When many cosmic rays are passing through our detector, this will likely trigger multiple PMTs. By looking for times when all of our PMTs are detecting high amplitude events, we

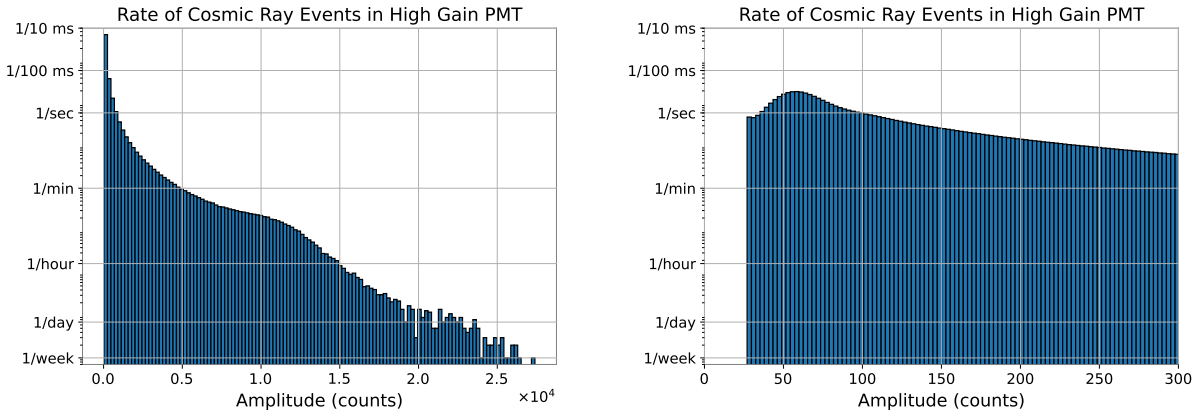

(a) (b)

Figure 2: (a) The detection rate of of cosmic ray events, categorized by amplitude, in the high-gain PMTs using a week of data. (b) focuses on the low-amplitude peak with finer binning to study the distribution of individual muons seen by the detector. The data is cut at around 30 counts because at this threshold cosmic rays cannot be confidently distinguished from detector noise.

can find the cosmic ray showers that are most likely to create an effect in LIGO, if any do. Figure 3 shows a typical cosmic ray shower recorded by the cosmic ray detector during O4a. This is a large shower, corresponding to $\sim$ 400 muons passing through the detector, and hence all four PMTs have significant signals. Similarly, Figure 4 shows an hour of cosmic ray events where the amplitude recorded by the high gain PMT is plotted against the amplitude recorded by the low gain PMT. From the slope, we can see that the high gain PMT 2 is five times more sensitive to cosmic rays as low gain PMT 2.

We now want to consider the effect a passing muon has on aLIGO test masses. The optics are made of fused silica which has a density of $\rho = 2.2$ g/cm^3, and an average muon deposits energy into fused silica through ionization at a rate of I = 2.049 MeV$\cdot$cm^2/g. We consider the scenario where the muon travels along the beam-axis and through the d = 20 cm width of the test mass. Since most cosmic rays arrive arrive at an angle that is nearly perpendicular to Earth's surface, a cosmic ray traveling in line with the beam would be a rare event. However, any cosmic ray that crosses the test mass at an angle of about 30° or greater from the normal would still travel through the full thickness of the test mass and have the same impact. The amount of energy the passing muon would deposit into the test mass in this case would be:

$$\Delta E = I\rho d = 90 \text{ MeV} \tag{1}$$

Similar to the approach Braginsky, et al. took in [1], we consider how far the test mass would move given this energy. We treat the bottom stage of the quadruple suspension system as a simple pendulum and use the momentum the muon imparts to the test mass, which for highly relativistic particles $\Delta p \approx \Delta E/c$:

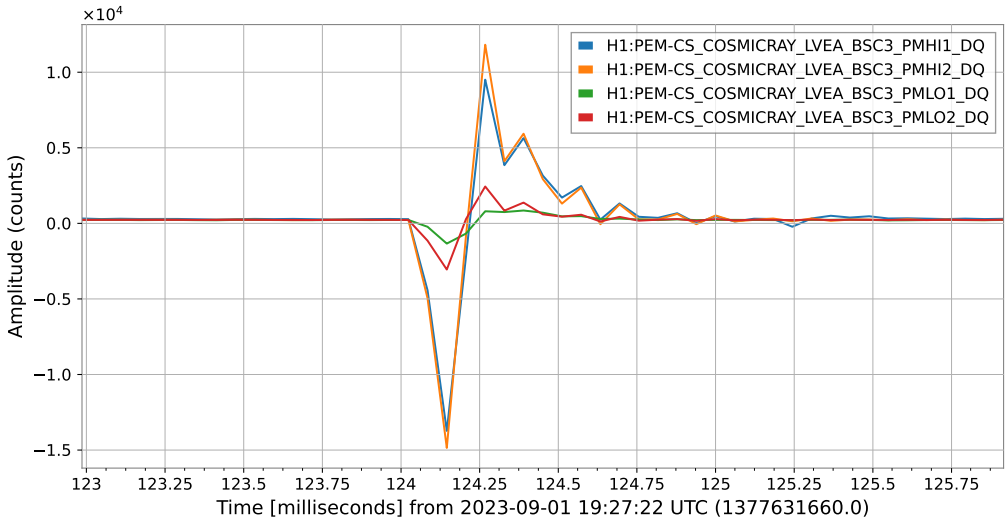

Figure 3: A typical cosmic ray shower that can be seen in all four PMTs. The two channels connected to high-gain PMTs are labelled 'PMHI' and the low-gain PMTs are labelled 'PMLO'.

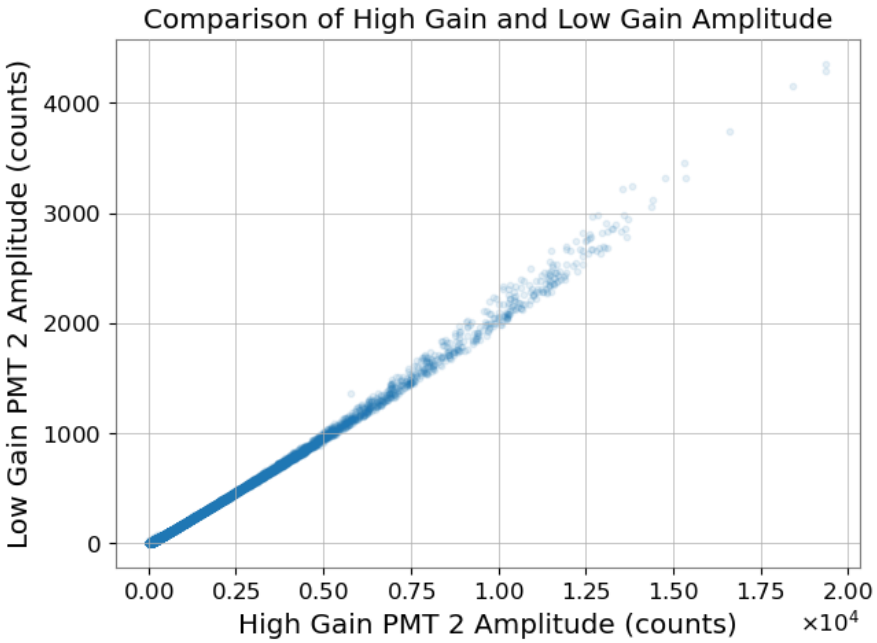

Figure 4: An hour of cosmic ray data comparing the amplitude of cosmic ray events that are witnessed in both High Gain PMT 2 and Low Gain PMT 2

$$\frac{\Delta E^2}{2mc^2} = mgh \tag{2}$$

$$\Delta x \approx \sqrt{2\ell h} = \sqrt{\frac{\ell \Delta E^2}{gm^2c^2}} \tag{3}$$

Where m is the mass of the test mass (40 kg), Δx is the horizontal movement of the test mass along the beam axis, h is vertical movement of the test mass, ℓ is the pendulum length (0.6 m), c is the speed of light, and g is the acceleration due to gravity. In our single muon example, we expect the test mass to move about $\Delta x = 3.0 \times 10^{-22}$ m and with a measured strain of $\Delta x/x_0 = 7.5 \times 10^{-26}$ in LIGO. Thus, an individual cosmic ray would be below LIGO's sensitivity curve, but when thousands of cosmic rays or more pass through the test mass, it could potentially create noise. On monthly timescales, we observe several events that will saturate the high-gain PMTs at 32,000 counts. Looking again at Figure 2, if we assume the 60 count peak is the amplitude of individual muon events and the detectors are operating at 40% efficiency, then we can estimate that the highest energy showers contain at least 1,300 muons passing through the detector.

We can model the response of the test mass as a damped pendulum, which should follow:

$$x(t) = \frac{a_0}{\sqrt{\omega_0^2 - \alpha^2}} \sin\left(t\sqrt{\omega_0^2 - \alpha^2}\right) e^{-\alpha t} \tag{4}$$

Where a_0 is the initial acceleration on the test mass from the cosmic ray, ω_0 is the natural frequency of the pendulum, and α is a damping coefficient. The test mass is damped in the pendulum mode to prevent lockloss [6] and has a Q factor of about 3. We can model the movement of the test mass, which is shown in Figure 5. Given that the effect of this impulse is at less than a Hertz, it is unlikely to observe this effect unless it up converts in some way.

We also consider the effect proposed by Yamamoto, et al. [2] where vibrations are rung up in the test mass due to heated tracks cosmic rays leave as they pass through. In this paper, they conclude that the fundamental mode of these vibrations would be dominant and the frequency could be described as:

$$\omega_1 \sim \frac{\pi}{H}\sqrt{\frac{Y}{\rho}} \tag{5}$$

where H is the thickness of the test mass, and Y and ρ are the Young's modulus and density of the optic material respectively. Using the properties of fused silica, we would expect this effect to be around 14,000 Hz, making it unlikely to be a concern in LIGO.

The final effect we consider is the charge hopping model suggested by Weiss in [3]. In later work, Ugolini, et al. [7] predicted this electrostatic effect would create noise that follows a $1/f^3$ shape. Although the exact amplitude for this effect is difficult to predict, we compare how $1/f^3$ relates to the slope of LHO's sensitivty curve in Figure 6. At lower frequencies, the sensitivity curve is falling more rapidly than $1/f^3$, and thus this noise could be relevant at frequencies near the bottom of the sensitivity curve.

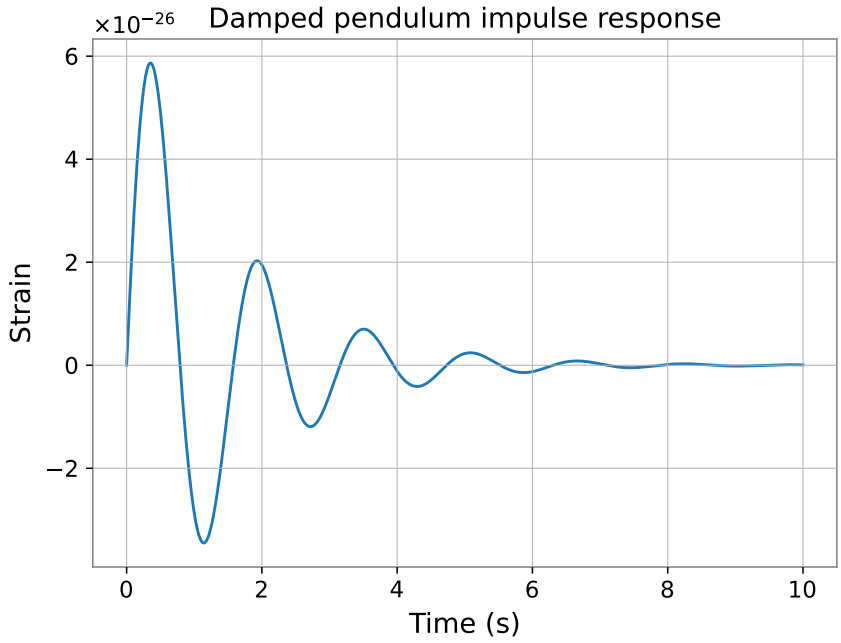

Figure 5: A model of the measured strain resulting from the energy a single muon deposits could deposit on the test mass.

Thinking forward to Cosmic Explorer, we can make predictions for the same potential effects. We assume the expectation that Cosmic Explorer will 40 km laser arms and 320 kg test masse [8]. For simplicity, we assume these test masses are still made of fused silica. In the case of momentum transfer, the measured strain of an individual muon would be 3.0×10^{-27}, which is 25 times less than what is predicted for aLIGO. Similar to LIGO, if cosmic rays created noise in this mechanism, we would still expect it to take potentially thousands or more muons in a single cosmic ray shower to create an observable effect. For vibrations induced by heating of cosmic ray tracks, the test masses would be twice as thick, so the predicted frequency would be around 7000 Hz. Cosmic explorer is typically predicted to be an order of magnitude more sensitive at most frequencies but there could be even greater improvement in lower frequencies, which would be more applicable to the momentum transfer mechanism and the charge hopping effect.

4 Cosmic Rays and DARM

We now discuss our search for excess noise in DARM associated with cosmic rays. We first investigated if cosmic rays could be a source of glitches where the cause is unkown. Then, we analyzed Amplitude Spectral Density (ASD) and Power Spectral Density (PSD) data of DARM near large cosmic rays showers to search for evidence of noise in our gravitational wave channel.

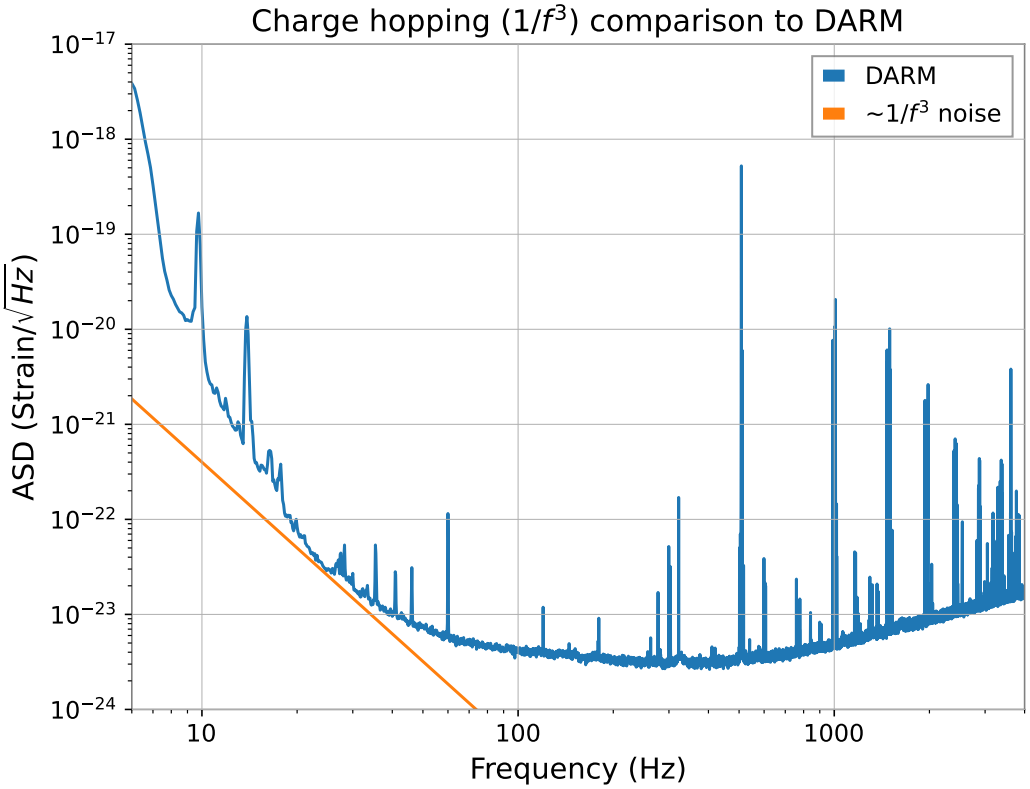

Figure 6: An arbitrary $1/f^3$ line plotted with LHO's ampltidue spectral density (ASD) to understand how we might see the charge hopping noise model proposed [3].

4.1 Correlations with Glitches

Cosmic rays can be considered to be passing through the test mass instantaneously relative to LIGO's sampling rate, so we consider glitches are that are short in duration, broadband in frequency, and with no known source. This includes blips, low-frequency blips, repeating blips, and tomtes.

In our first test, we consider the amplitude of PMT signals near glitches in time. Here we are using a list of glitches reported by GravitySpy from September through December of 2023. We take cosmic ray data within ± 1 second of a glitch, which we call our on-source data, and compare it to 20 seconds of background time far away from each glitch. We then use the Kolmogorov–Smirnov (KS) test to compare the two distributions. If we expect large cosmic ray showers to be the source of these glitches, then we could also expect to see an overabundance of high amplitude cosmic ray events in the on-source population.

Figure 7 shows an example for cosmic ray amplitude around Blips in one high-gain PMT (all cosmic ray amplitudes plots can be found in Appendix A). The following table lists the resulting KS test p-value for each glitch class:

Glitch Class	KS test p-value
Blips	0.281
Low-frequency blips	0.551
Repeating blips	0.534
Tomtes	0.211

No significant difference was found in these two distributions for any glitch type. Cosmic ray activity is measured to be the same around glitches as it is during times when there are no glitches in the data.

Next, we test if large cosmic ray showers are correlated in time with glitches. We take a month of PMT data and consider the top 10% of cosmic ray events from minute trend data. We limit our dataset to cosmic rays that are measured in multiple PMTs to study the largest showers (about 2000 in total). We then take the time difference between cosmic ray showers and the next nearest glitch in time to see if there is any temporal correlation. Figure 8 shows an example with blips (the same test with all glitches is shown in Appendix B). For comparison we did the same with random times throughout the month and found no significant difference from the random distribution.

4.2 Analyzing DARM

Next, we want to use our set of the highest energy cosmic ray showers to take a closer look at DARM. We search for any low level noise that may be present in DARM by combining information around all of the times of cosmic ray showers, looking across 1,000+ time series.

We started by looking at ASDs, both 10 seconds before cosmic ray showers and 10 seconds after, and average these ASDs over all cosmic ray showers. The results of are shown in Figure 9a. Figure 9b shows the ratio of these ASDs in each frequency bin. If cosmic rays

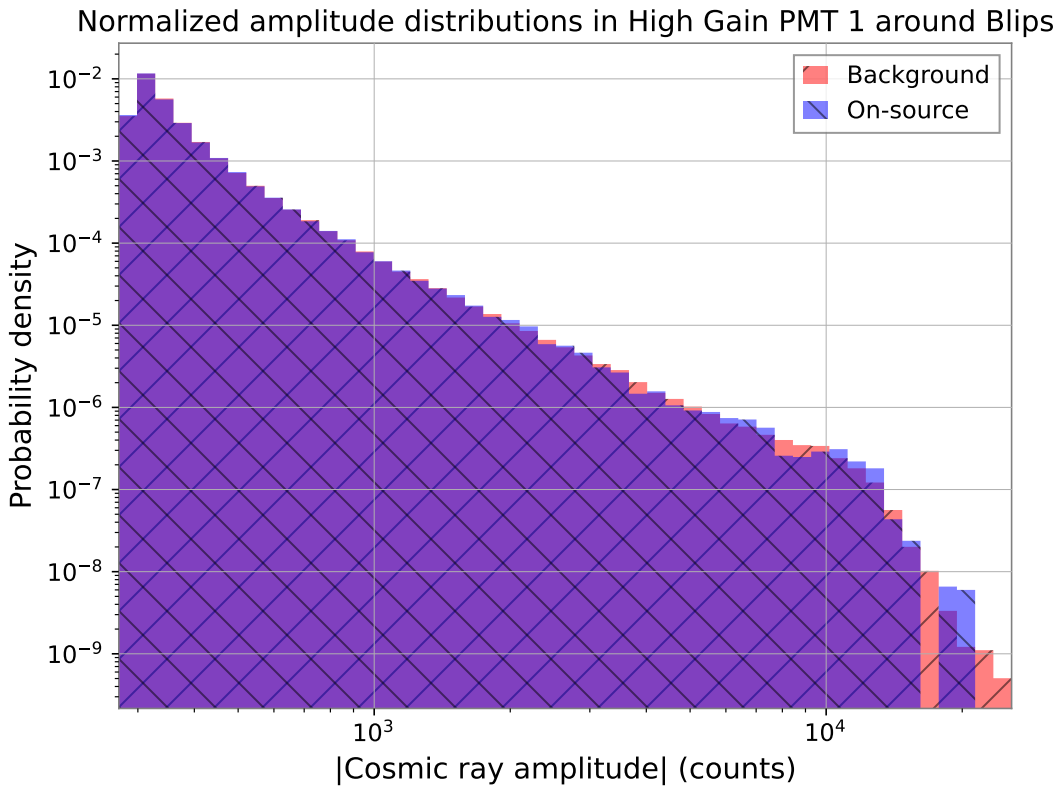

Figure 7: Histogram of cosmic ray amplitudes within 1 second of blip glitches (on-source) compared to times with no glitches (background).

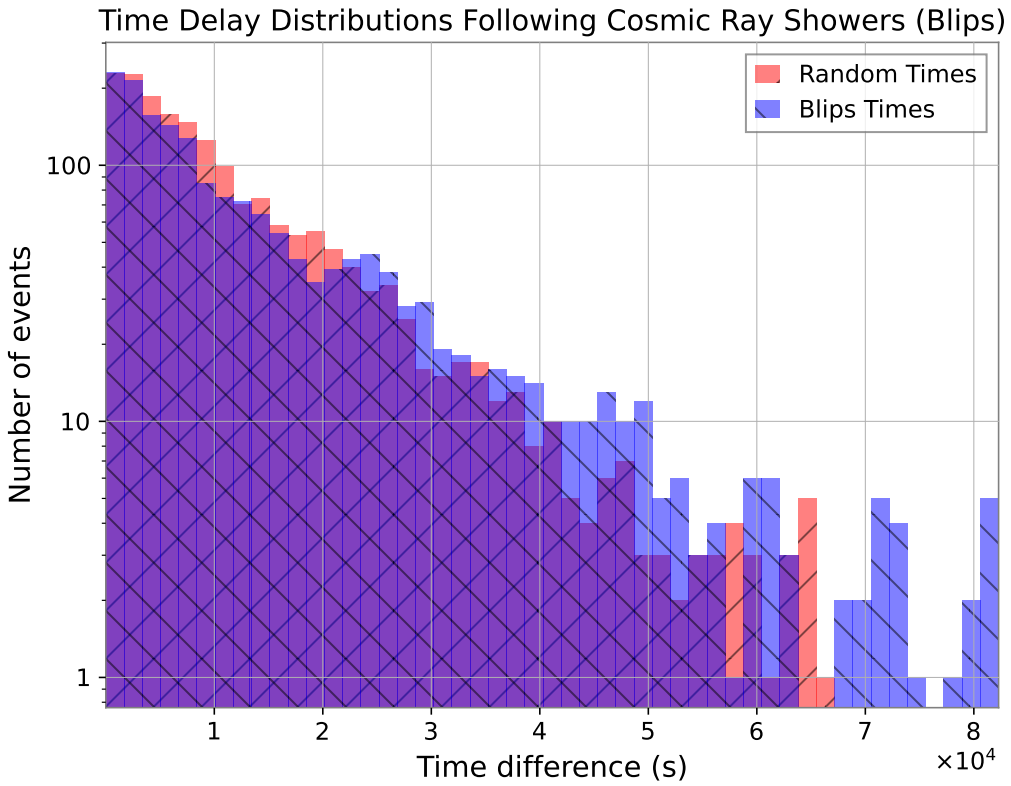

Figure 8: For every large cosmic ray shower, the time difference between the shower and the nearest glitch following the shower was calculated, and the distribution of these values is shown. The same test was done comparing the time difference between cosmic ray showers and random times throughout the month. Both of these follow the expected shape of a Poissonian process, indicating there's no correlation in time.

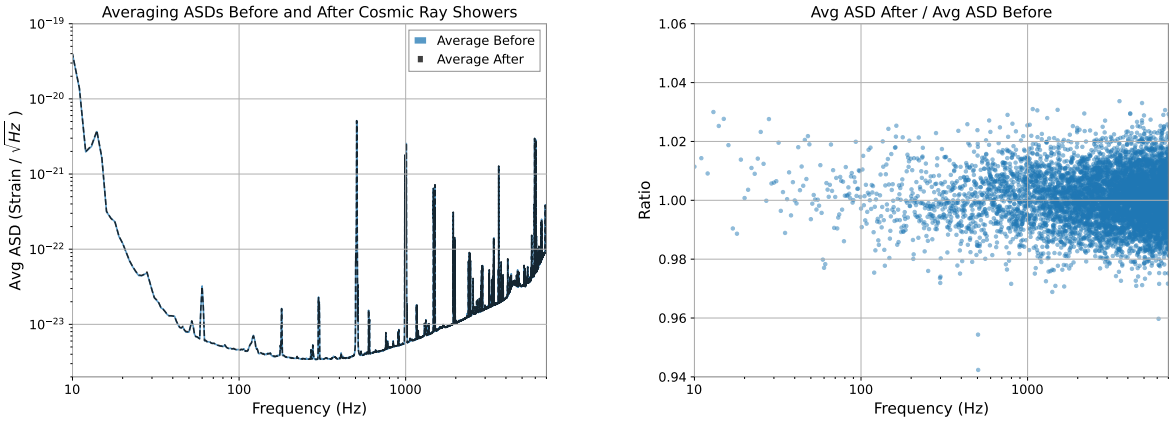

(a) (b)

Figure 9: For every large cosmic ray shower in the month, two ASDs were created using 10 seconds of DARM before and after the shower. Note that there are two lines present that are almost entirely overlapping. The average for for all cosmic ray showers is shown on the left, and the ratio of the ASD average after cosmic ratio showers compared to before for each frequency bin is shown on the right.

were creating noise a particular frequency (like in the heating vibrations case), we would expect the average ASD to be consistently higher than 1 around that frequency. Examining these two plots, we see no evidence of excess noise at any frequency. The averaged ASDs both before and after cosmic ray showers are consistent with the sensitivity curve we expect.

To look for broadband effects (like charge hopping), we also integrate the PSD before and after cosmic ray showers. We consider the most sensitive frequency bands to integrate over that avoid large peaks in the PSD. We then apply the KS test to the distributions of integrated PSD values before and after cosmic ray showers. An example distribution is plotted in Figure 10 (see Appendix C for all frequency bands tested) and the results for KS test p-valeus are listed in the table below:

Frequency Band	KS test p-value
10-30 Hz	0.453
20-55 Hz	0.734
65-115 Hz	0.240
125-450 Hz	0.057
550-950 Hz	0.543
1050-4000 Hz	0.574

There appears to be no difference integrated PSD values after cosmic ray showers when compared to values before.

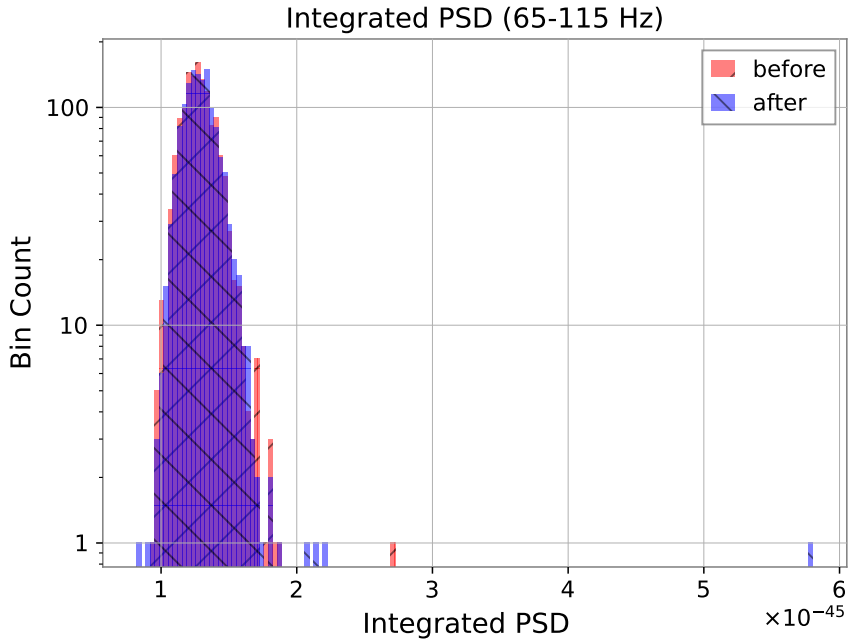

Figure 10: Distribution of PSDs integrated from 65 to 115 Hz both before and after cosmic ray showers.

4.2.1 Injected Signals Test

To test the methods previous mentioned, we revisted the same cosmic ray shower data with injected signals. In every DARM time series, we injected a damped sinusoid signal at the time of the cosmic ray shower with a frequency of ∼100 Hz. Thus, it is similar to the effect shown in Figure 5 but at a higher frequency. We increased the amplitude of the injection until an effect was observable in the ASD and PSD tests, shown in Figure 11. At around an injection amplitude of 10^{-23} strain, the KS test's reported p-value falls below 0.01. The results of the ASD and PSD tests with the injected signals at 10^{-23} are shown in Figure 12. The ASD test shows a peak at 100 Hz that is starting to rise over the normal distribution of noise. Similarly, when we integrated the PSD over the frequency band that contained our injections, the reported p-value was 0.00370. If cosmic rays created an impulse at this frequency, it would correspond to a shower of about 150 muons for every injection, which is a realistic number to expect in a typical cosmic ray shower. However the predicted effect of this mechanism from Section 3 is at a much lower frequency ($\sim$ 1 Hz).

5 Discussion and Conclusions

In this study, we characterized the signals we observe from Hanford's cosmic ray detector and looked for any correlation with noise in LIGO Hanford. We confirmed that our electronics are capable of detecting individual cosmic rays, identifying cosmic ray showers, and

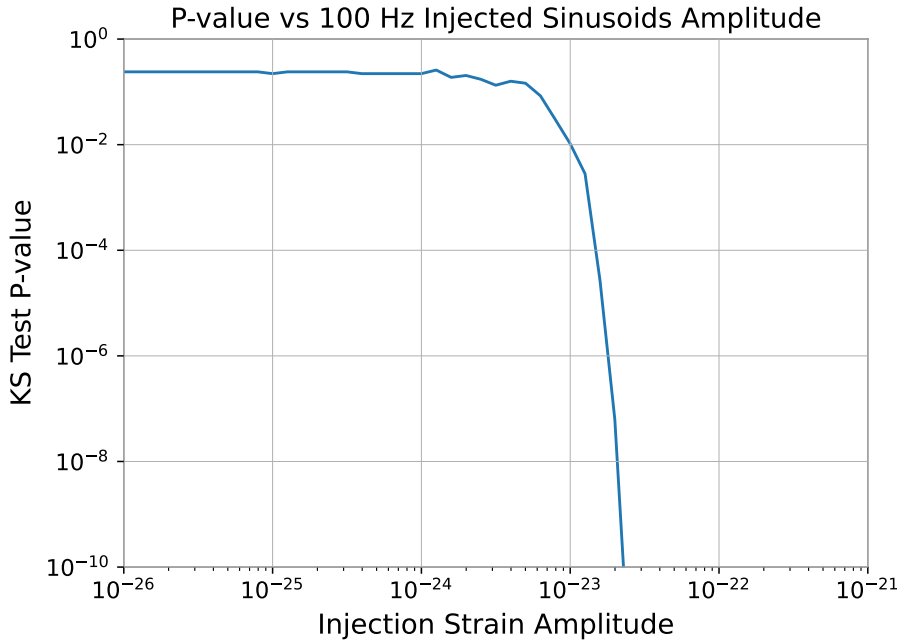

Figure 11: The KS test's reported p-value when applied to the distributions of integrated PSDs from 65 to 115 Hz with damped sinusoid injections at 100 Hz. At around 10^{-23} strain and above, the injections have a significant enough effect that we could say at a 99% confidence level that the two samples are drawn from different distributions.

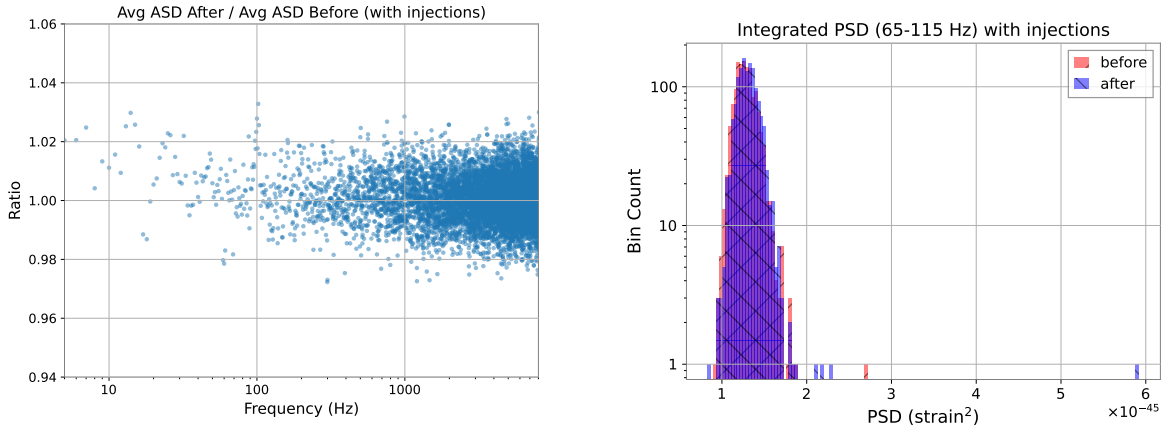

(a) (b)

Figure 12: (a) After damped sinusoids (at a strain amplitude of 10^{-23}) were injected, the ASDs were averaged across all cosmic ray showers, and the same ratio of the average ASD in Figure 9 was taken. A clear peak is seen at 100 Hz where the signals were injected. (b) The same test as Figure 10 where the PSD was integrated before and after cosmic ray showers over 65 to 115 Hz. In this case, the KS test reports a p-value of 0.00370.

distinguishing the intensities of these showers.

By comparing distributions of cosmic ray amplitude around blips and tomtes to times away from these glitches, we conclude there is no difference in cosmic ray activity around these classes of glitches. Similarly, we combined cosmic ray information from multiple PMTs to study the largest cosmic ray showers our detector observes. These have no temporal connection with the glitches.

We also studied DARM following large cosmic ray showers. By averaging our ASD both before and after the largest cosmic ray showers, we saw no difference stand out at any frequency. Similarly, integrating the PSD before and after cosmic ray showers in the most sensitive frequency bands of our detector produced no difference in these distributions either. At our current sensitivity, cosmic rays don't appear to be creating any noise in LIGO.

Given our predictions in section 3, we would therefore find it unlikely that cosmic ray showers would produce noise in Cosmic Explorer. However, given the results of our injection test, a realistic number of muons in a cosmic ray shower could create noise in future detectors via similar yet unpredicted mechanisms to the ones considered in this study. We are considering different modes cosmic rays could excite in the test mass pendulum system for future work. We would also recommend monitoring cosmic ray activity for future detectors.

References

[1] Braginsky, V. B., O. G. Ryazhskaya, and S. P. Vyatchanin. "Notes about noise in gravitational wave antennas created by cosmic rays." Physics Letters A 350.1-2 (2006): 1-4.

[2] Yamamoto, Kazuhiro, et al. "Effect of energy deposited by cosmic-ray particles on interferometric gravitational wave detectors." Physical Review D—Particles, Fields, Gravitation, and Cosmology 78.2 (2008): 022004.

[3] Rai Weiss, LIGO Technical Note, T960137 (1996).

[4] G. Harry, LIGO Technical Note, T080019 (2008).

[5] S. Klimenko, LIGO Technical Note, T1500449 (2015).

[6] S. M. Aston, et al. Update on Quadruple Suspension Design for Advanced LIGO, LIGO DCC P1200056 (2012).

[7] D. Ugolini, et al. Charging Issues in LIGO, LIGO DCC P070068 (2007).

[8] M. Evans, et al. Cosmic Explorer Horizon Study, arXiv:2109.09882 (2021)

[9] R. Frey, The PEM Cosmic Ray Detectors, LIGO DCC T010088 (2001).

[10] R. Frey, LHO cosmic rays in S5 – Update, LIGO DCC G070526 (2007).

A Studying Cosmic Ray Amplitude Around Glitch Times

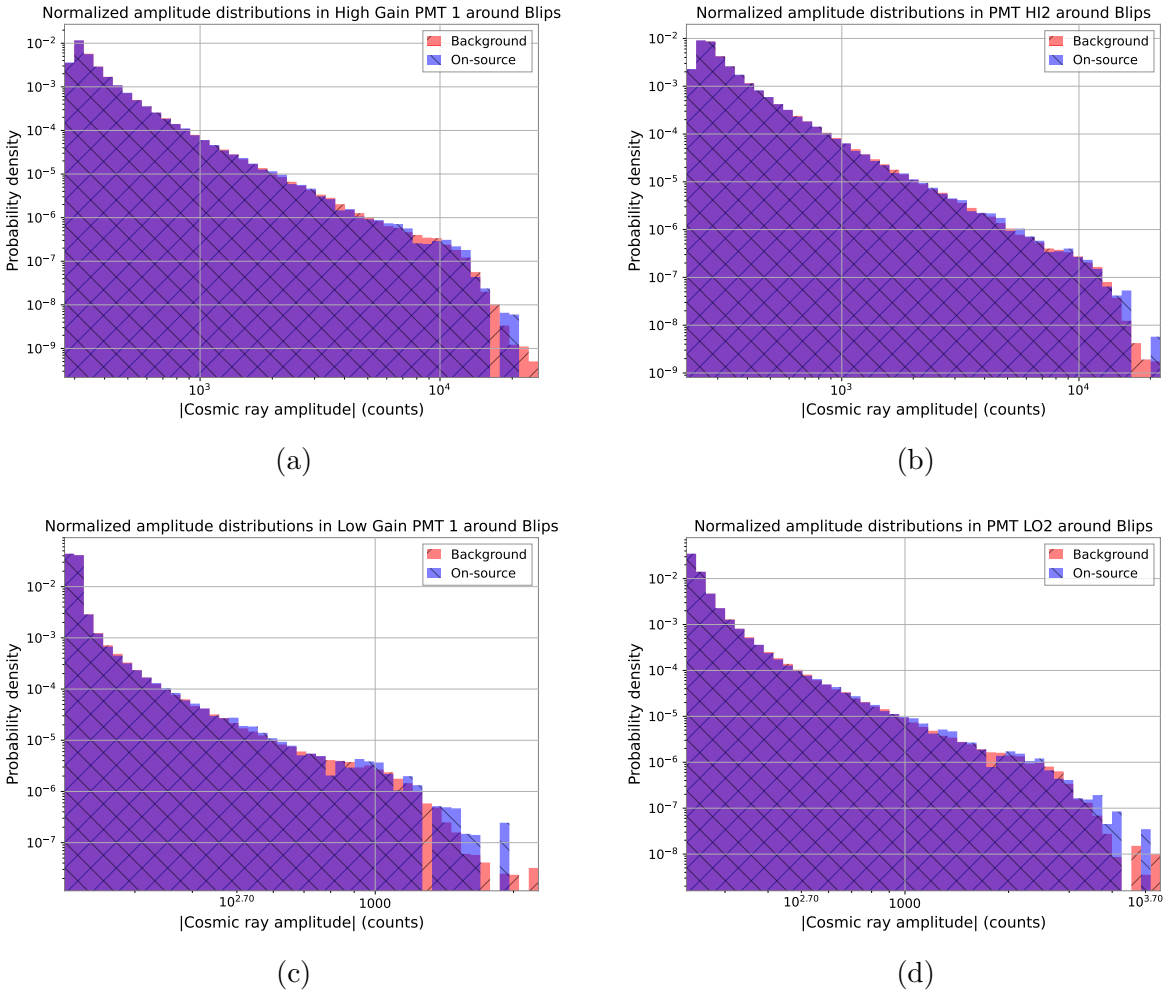

Figure 13: Blips Amplitude test

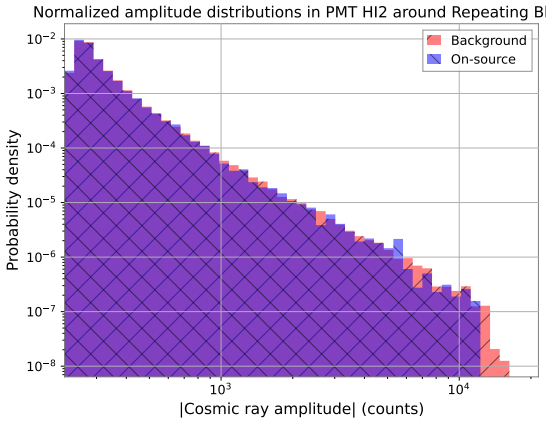

(a)

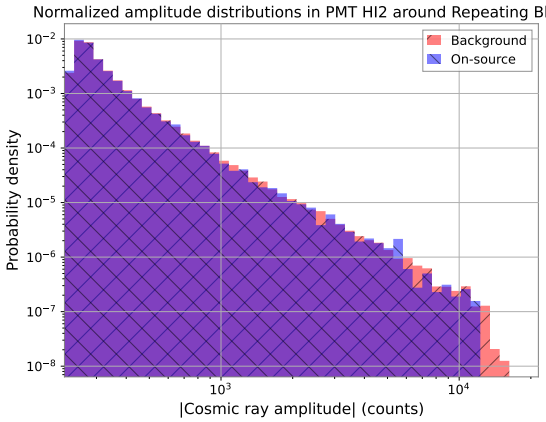

(b)

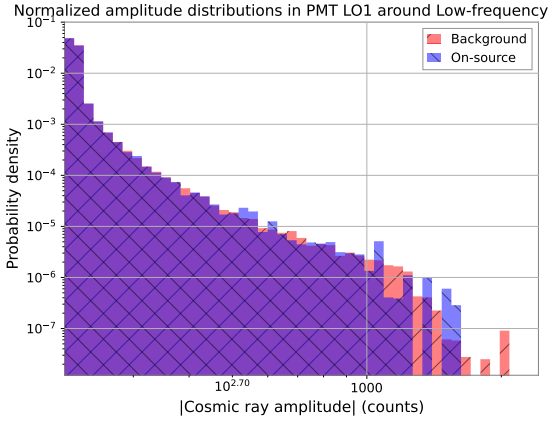

(c)

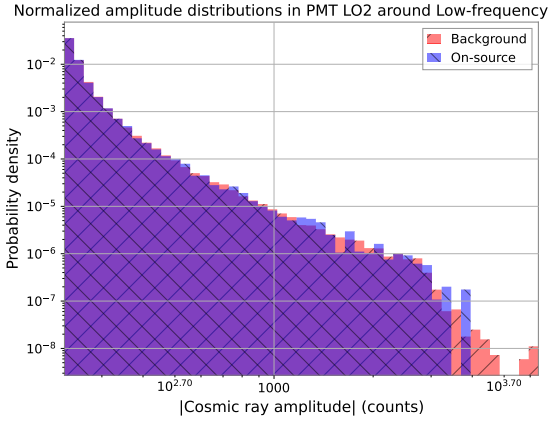

(d)

Figure 14: Low frequency Blips Amplitude test

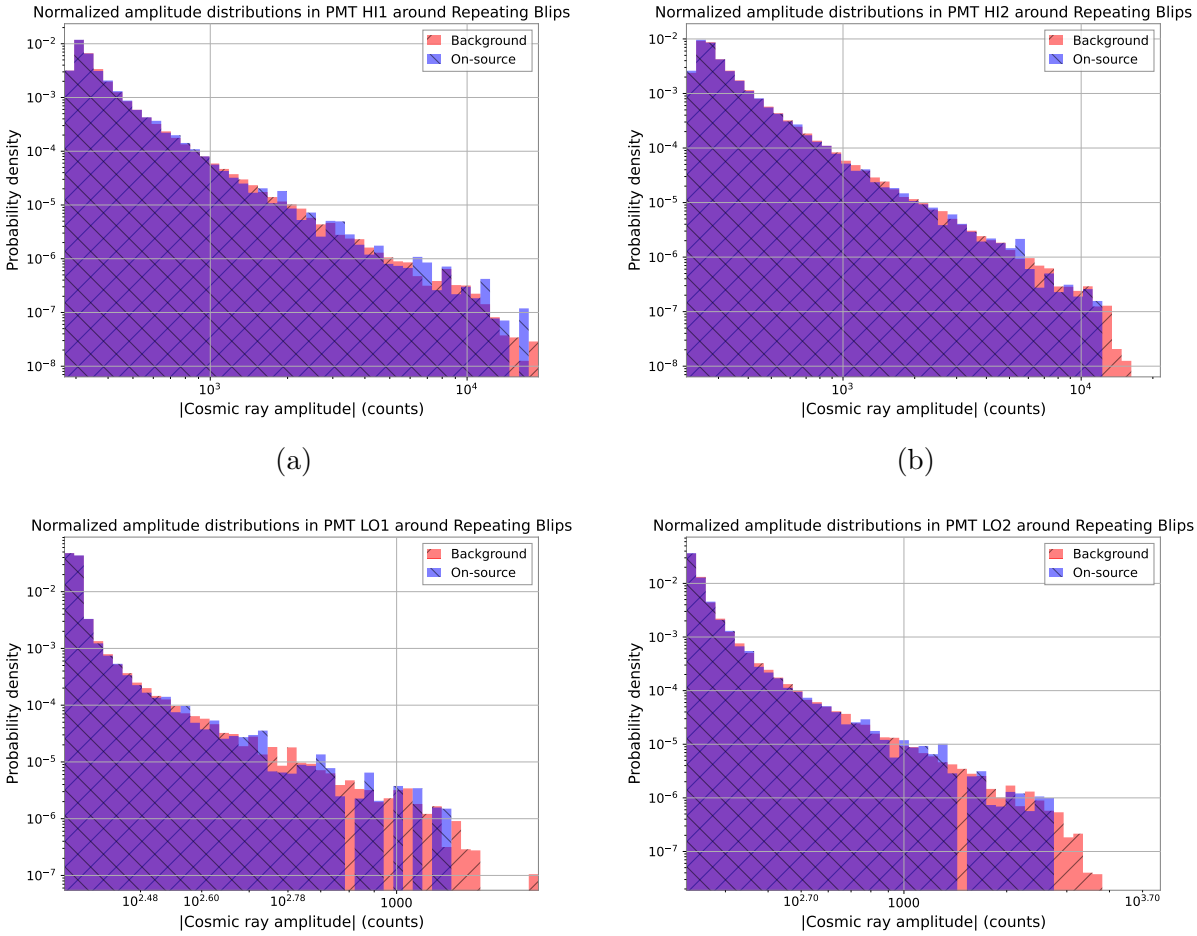

Figure 15: Repeating Blips Amplitude test

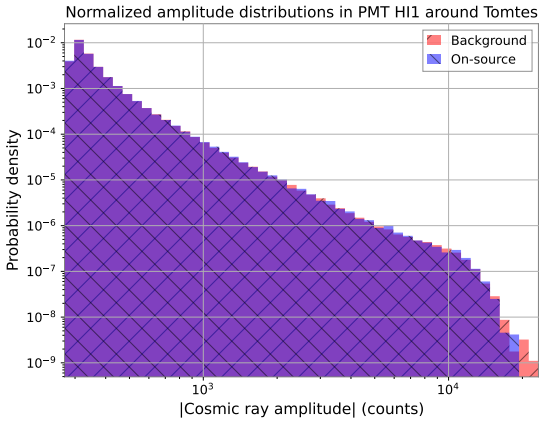

(a)

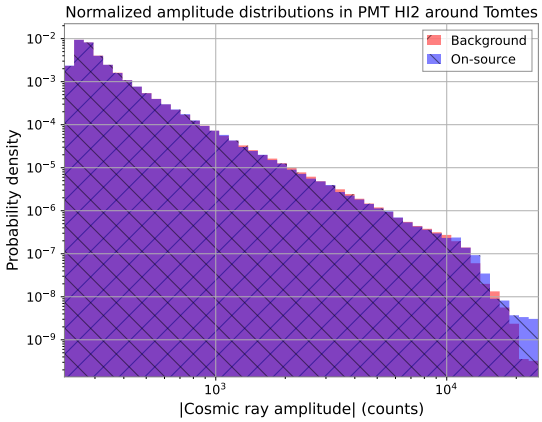

(b)

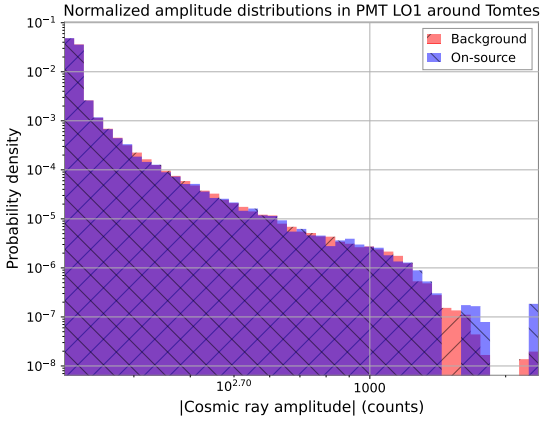

(c)

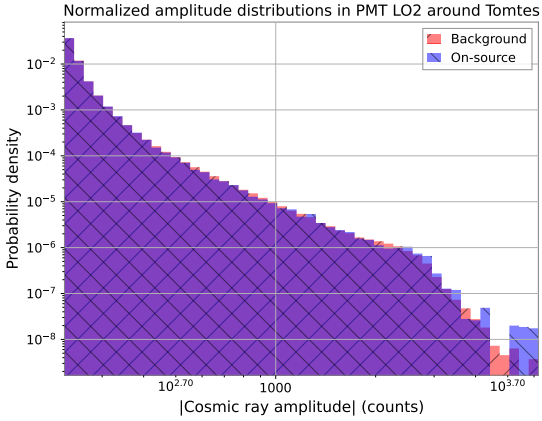

(d)

Figure 16: Tomtes Amplitude test

B Searching For Temporal Correlation Between Cosmic Rays and Glitches

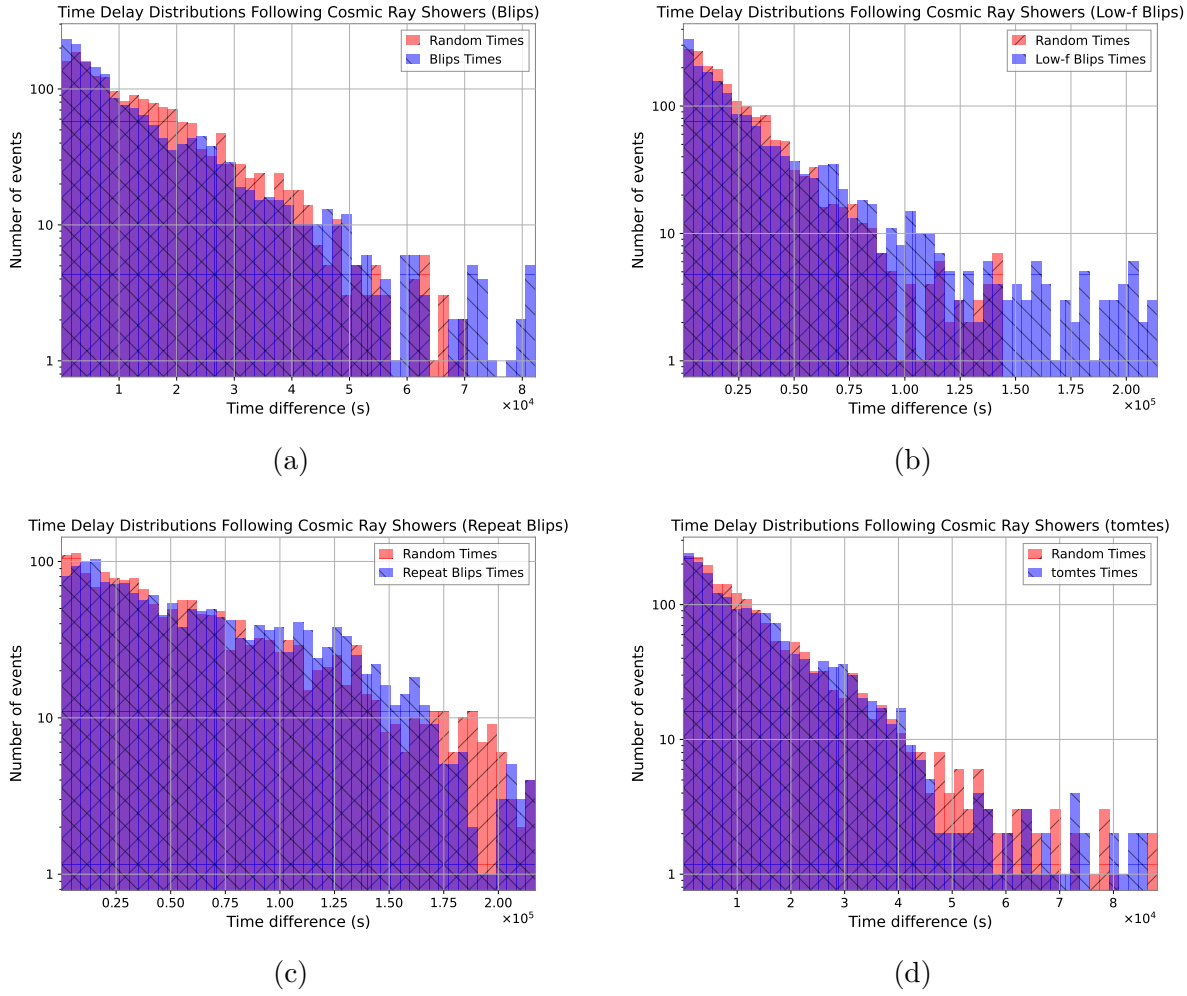

Figure 17: Time difference test between each cosmic ray shower and the next nearest glitch in time

C Integrated PSDs Near Large Cosmic Ray Showers

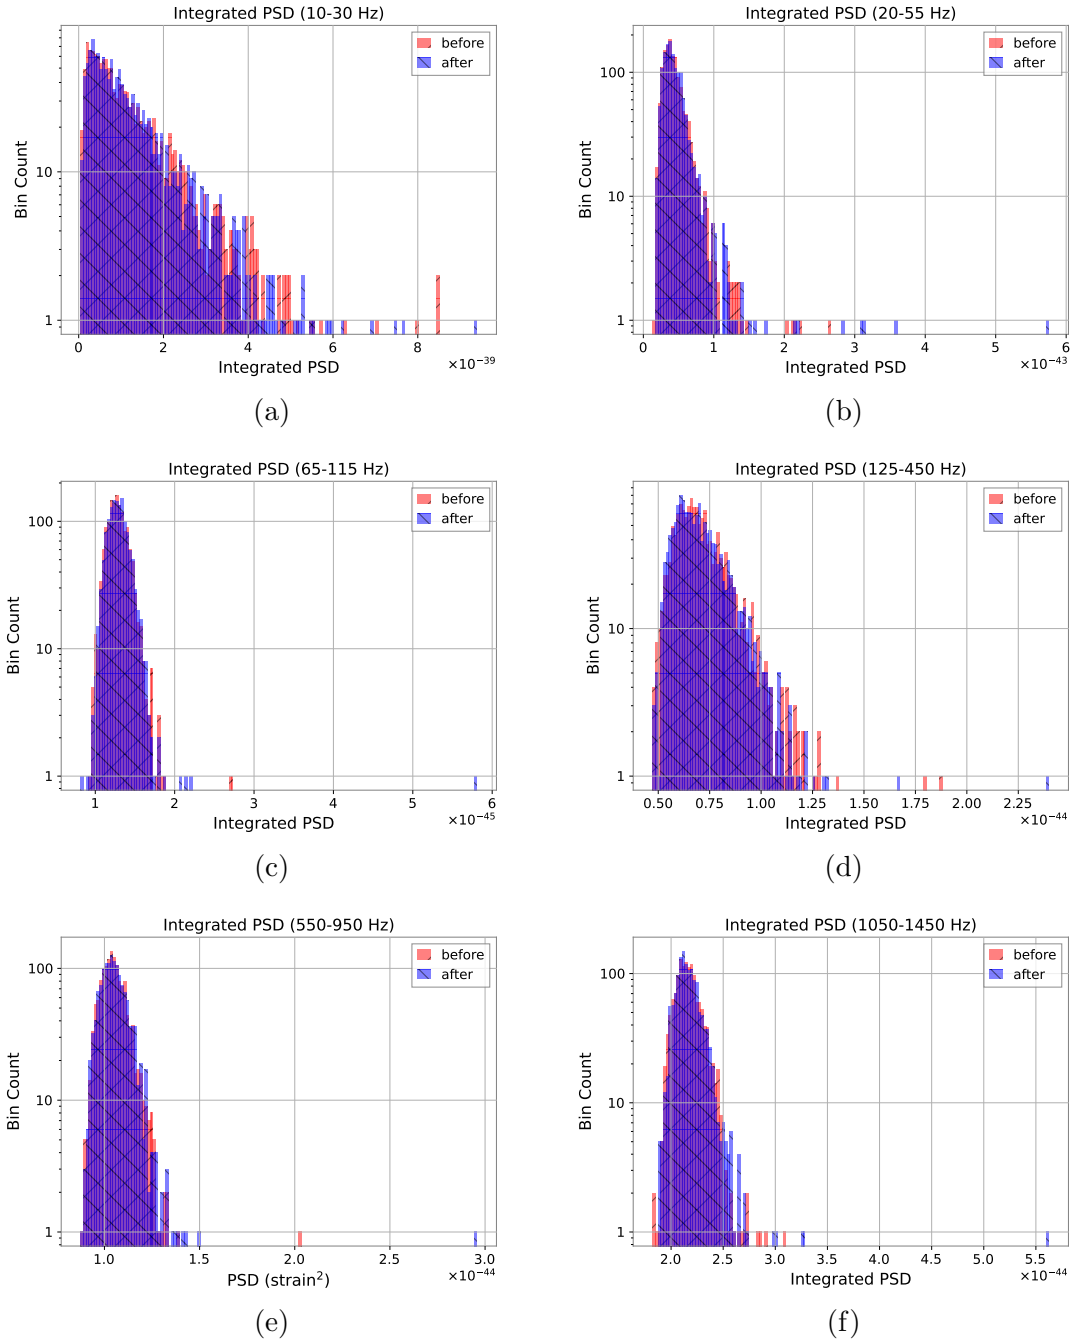

Figure 18: Integrated PSDs following large cosmic ray showers