LIGO SURF Final Report: Reducing Early Warning Retractions to Facilitate Multi-Messenger Astronomy

Anna Tosolini
(Dated: July 27, 2021)

ABSTRACT

Early warning gravitational wave detection pipelines are extremely important tools that could alert astronomers to a gravitational wave event before the event has occurred, facilitating a new frontier for multi-messenger astronomy. The aims of this project were to study the early warning pipeline, specifically GstLAL, and run the pipeline using data from Advanced LIGO and Advanced Virgo's third observing run colored to projected O4 sensitivities. This project involved varying the upper frequency bounds, which correspond to different early warning times, and using these different runs to test a new process of calculating the false alarm rate for events. The false alarm rate is a measure of how often we expect LIGO noise to produce a gravitational wave event. By decreasing the significance of noise events through this process, we are aiming to reduce the number of retractions for the next LIGO observing run.

INTRODUCTION

Gravitational waves, originally predicted by Albert Einstein almost a century ago, are ripples in the fabric of spacetime produced when massive objects such as black holes or neutron stars collide. These waves propagate through spacetime at the speed of light, interacting weakly with matter as they travel through the universe. The Laser Interferometer Gravitational-Wave Observatory, or LIGO, can detect these waves by observing very minute changes in the length of the LIGO detector's 4 kilometer arms. The interferometer process begins when a laser is shone through a beam-splitter, sending light through the arms to mirrors at the ends and back. (To visualize this more clearly, see figure 1.) If there is no change in the length of the arms, no signal is detected through the photodetector. However, a minuscule change in the distance of the arms can result in a signal. Indeed, when gravitational waves reach the Earth, they are so small that only a tiny fraction of the original amplitude is detectable; therefore, the change in distance of the arms is also very small.

One way researchers extract a signal from the data collected by LIGO is by using a technique called matched filtering, which correlates the pre-calculated signal with the observed data. Matched filtering techniques use thousands of templates, which are estimates of gravitational wave signals, and these templates are compared to real

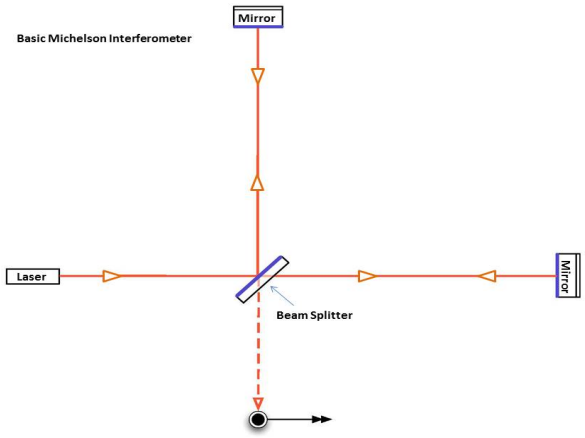

FIG. 1. This picture from the LIGO Caltech website displays the basic layout of the LIGO detectors.

data to identify candidate signals. Search pipelines are used to filter through thousands of gravitational waveform templates and pick out the best candidate signals. Pipelines can vary in the following processes, but this project will work specifically with GstLAL, which is a low-latency search pipeline that conducts matched-filtering using the time dependence of the templates. Ultimately, the pipeline computes a false alarm rate for each trigger, which calculates the probability that the trigger is due to a gravitational wave and not simply due to the background [1] [2].

This project involves working with the GstLAL early warning pipeline to increase the pipeline's sensitivity. These results will be worthwhile because improving the sensitivity of the search pipeline will facilitate electromagnetic observations, and observing binary neutron star and neutron star black hole events in low-latency is vital to LIGO research. This topic has already been researched in the past by the GstLAL team, and contains work of many previous papers written by my mentor and others. The results of this research will actively test information that could contribute to the next LIGO observing run [3].

To look into some early warning prospects for the next LIGO observing run, we can calculate and plot the number of detections LIGO expects per year expected for O4 versus the time before the merger (in seconds), with inputs such as chirp mass, signal-to-noise ratio, and various bounds for the frequencies. To calculate the number of detections per year, we take the binary neutron star

merger rate of $320\text{Gpc}^{-3}\text{yr}^{-1}$ and multiply by a volume function which takes the distance that LIGO can see as an input. This distance function involves interpolating the power spectral density data for specific frequencies, and then integrating through the data with various frequency bounds. This plot is attached below. As binary systems approach merger, their frequencies increase as a power law of $f^{11/3}$, so varying the bounds of these frequencies allows us to identify candidate binaries before merger. Since the frequency increase is monotonic, we know that the frequency will never decrease before merger, so we can simply look at the lower frequencies if we are attempting to find binaries before merger. Through this process, we have estimated the number of early detections LIGO can hope to achieve every year, as well as the fraction of those detections close enough for observation through the electromagnetic spectrum. With these graphs and equations, we can create a clearer picture of the binaries before they merge, which will facilitate greater understanding of processes we wish to learn more about [4] [5] [6].

Below is the distance equation that I used when determining the volume function for my program, which led to the plot below, seen in Fig.2.

$$D = \frac{1}{8}\left(\frac{5\pi}{24c^3}\right)^{1/2}(GM)^{5/6}\pi^{-7/6}\sqrt{4\int_{flow}^{fhigh}\frac{f^{-7/3}}{Sn(f)}df}$$

This equation is a measure of the range of LIGO's viewing capabilities. This distance function is used as the radius in the volume function which determines the number of detections LIGO will observe per year. We then multiply the volume function by the BNS merger rate we used earlier, $320\text{Gpc}^{-3}\text{yr}^{-1}$, to determine the detections per year. In this equation, we are assuming that the signal-to-noise ratio is 8, and that both masses are 1.4 solar masses. To find the equation for the red line on the plot, we find the maximum distance by creating a fraction,

$$D = \frac{200Mpc}{MaximumDistance}$$

which measures only the events that could be observed through electromagnetic observers, since 200 Mpc is the upper bound for viewing mergers in the electromagnetic spectrum. The maximum distance involves integrating the previous equation from 10Hz to 1024Hz. Then we use this distance in the volume function to determine the number of multi-messenger astronomy events per year that we could observe; we estimate from this graph that there will be one of these events in O4 [6].

METHODS

LIGO's third observing run in 2017 marked a turning point for gravitational wave science: this was the first

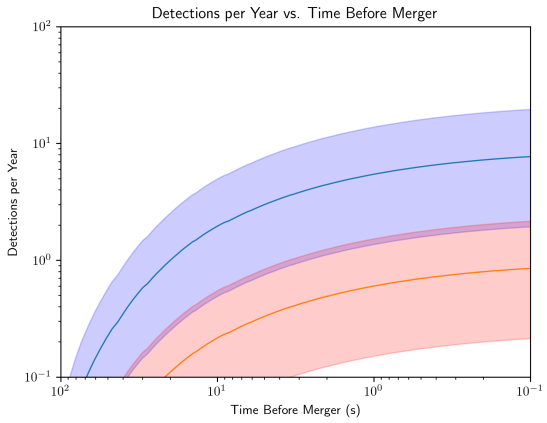

FIG. 2. This is a plot of the detections per year that we anticipate LIGO will receive during O4, the next observing run, plotted against the time before merger in seconds. The orange line corresponds to the events we will detect within 200 Mpc, which is an optimal range for observing binaries through the electromagnetic spectrum. The blue line corresponds to all events LIGO will detect. This plot shows us that we will detect around 8-9 events per year, and one detection per year that is 10-20 seconds early. Additionally, this plot tells us that we will detect one multi-messenger astronomy event in O4.

gravitational wave event detected from a binary neutron star system. Such systems are important in facilitating multi-messenger astronomy because phenomena are observed with gravitational waves as well as through the electromagnetic spectrum, and such observations carry insight about short gamma-ray bursts, r-process nucleosynthesis, the final state of the remnant, and many more astrophysical mysteries. In 2017, there was a 4.5 hour window between the gravitational wave's arrival to Earth and distribution of sky localization of the event. Reducing this time window will further our knowledge of the early stages of binary neutron star mergers, and is the motivation of early warning gravitational-wave detection pipelines, such as GstLAL [1].

As discussed previously, early warning detection pipelines, specifically GstLAL, sift through thousands of candidate signals through the process of matched filtering, which matches raw LIGO data to different waveforms. Candidate events are then assigned false alarm rates through comparison to the background, and through these false alarm rates the pipeline can quickly determine whether the signal we have detected is a gravitational wave event or not. Once the pipeline has found a potential gravitational wave signal, LIGO sends alerts to astronomers around the globe with an estimation of the sky localization of the event, and this combination of LIGO detections and astronomical observations facilitates multi-messenger astronomy. However, sometimes LIGO alerts are retracted because further re-

search has determined that the potential signal was simply noise. These retractions are a large inconvenience to astronomers and this project will attempt to reduce retractions by studying a new way to calculate the false alarm rates.

The goal of this project is to understand how the pipeline works and eventually learn enough to run the pipeline using data from LIGO's most recent observing run, which is whitened and recolored to the sensitivities of the next observing run. The equation for data whitening is as follows,

$$Y = WX, \mathrm{W}^T W = \Sigma^{-1}$$

[7] where W is the whitening matrix. Whitened and recolored data is used here to preserve the glitches from the original segment of data; Gaussian data does not include these glitches and is less useful in approximating LIGO data. The aims of this research are to reduce the number of retractions in order to facilitate multi-messenger astronomy. The results of this project will give us insight into how many false alarm events we could detect in the next observing run by testing different ways of calculating the false alarm rates in order to increase pipeline efficiency [8].

Once we have obtained this segment of data, it is submitted into the GstLAL pipeline for analysis. For this project, it is vital to submit several different runs of this data with various upper frequency bounds; for early warning detection, we are mainly focused on the lower frequencies so the runs are of 29, 32, and 1024 Hz. After the runs have completed, the output files for the marginalized likelihood of each run is combined and this combined output can be compared to the individual data, with a specific focus on comparing the false alarm rates. To calculate the false alarm rate, the likelihood ratio and the probability density functions are compared for both the signal and the noise for each frequency. Specifically, this project is testing if marginalizing over multiple frequency bands and combining into one probability density function is more effective in calculating false alarm rates than comparing individual runs to their respective probability density functions.

RESULTS

Once the analyses have completed for the different upper frequency bounds, a comparison of the background plots for each run can be seen in figures 3 and 4, which are the two most drastic runs: 1024Hz (the full frequency spectrum) and 29Hz (an upper limit corresponding to early warning detection). From figures 3 and 4, it is observed that there is much more noise in the 29Hz run than there is in the 1024Hz run, a factor which hinders early warning detection. The green dot on the plots (corresponding to an SNR=8 and a chi2=1) is enveloped in

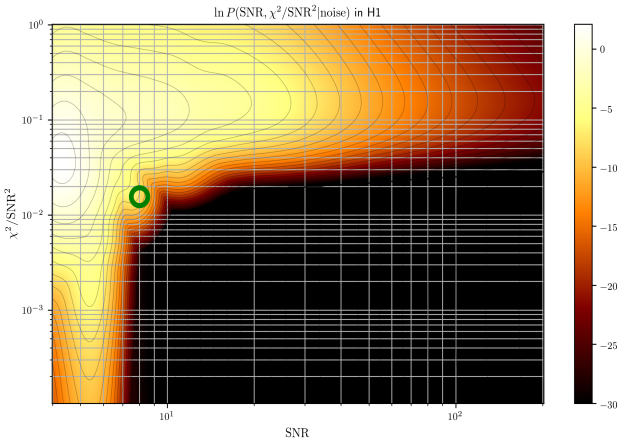

FIG. 3. Here is a background plot of a run through the GstLAL pipeline with a segment of O3a data whitened and recolored to O4 sensitivities, with an upper frequency bound of 29Hz. The orange area corresponds to the noise and background that we expect for O4, and we can see that the green dot (corresponding to an SNR=8 and χ^2=1) is obscured in the background, making early-warning detection very difficult.

noise (or the orange area of the plot) in the early-warning case, whereas for the full frequency spectrum it is just on the edge of noise and more easily detectable.

With these completed runs, the output files which contain the false alarm rates for every frequency are combined, and the likelihood ratios are marginalized for all of the frequencies to study how the false alarm rates change. This was completed for 3 frequencies: 29Hz, 32Hz, and 1024Hz. This study is testing whether three individual runs these frequencies or one run which marginalizes over these frequencies is safer in calculating the false alarm rates. From the histograms which compare the number of non-injection (or noise) false alarm rates for the individual runs against the combined runs, it is clear that the marginalization process makes the noise less significant, with one exception where the noise false alarm rate increased. These results can be seen in the comparison between figures 5 and 6. Though this up-ranked event is an outlier, it is still important to study to ensure this process decreases the significance of all noise. Additionally, it is noted that the injection false alarm rates are unchanged, so this process is safe for signals, as seen in figure 7.

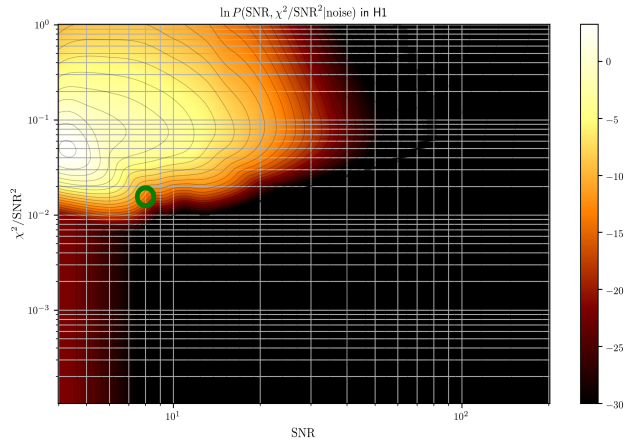

FIG. 4. This plot is similar to Fig 2, but this run took into account the entirety of the frequency spectrum LIGO can detect, with an upper bound of 1024Hz. We can see that the background is much smaller for this plot than the previous one, and that the green dot is located right on the outskirts of the noise, making detection more plausible.

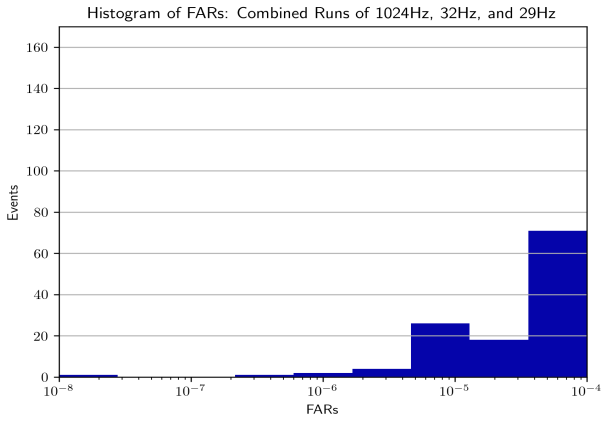

FIG. 6. This histogram displays the non-injection false alarm rates for the combined run. Comparing this to the previous histogram of the false alarm rates from the 1024 Hz run, it is noted that most of the false alarm rates decreased in significance. There is one event that was up-ranked, and this will need to be studied more.

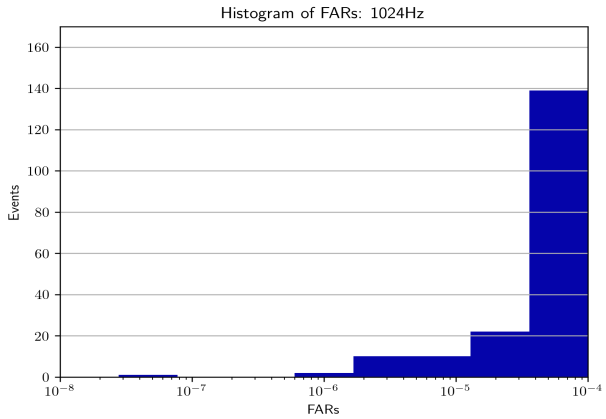

FIG. 5. This is a histogram of the non-injection false alarm rates for the 1024Hz run. It shows these noise events and their expected FAR, which can be compared to the histogram of the combined runs [will add that once I have it!]to analyze how this process has changed the false alarm rates for noise.

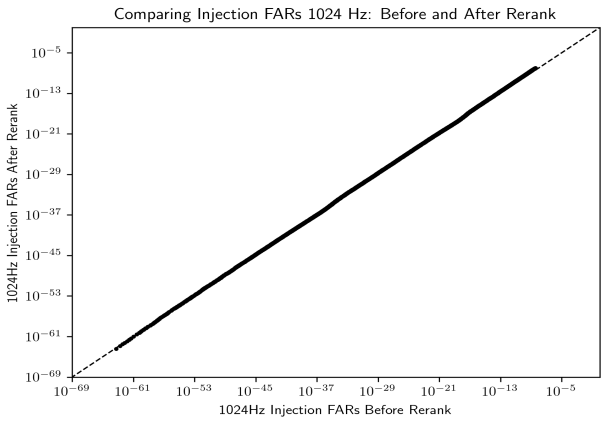

FIG. 7. This plot compares the injection false alarm rates for the 1024 Hz run before and after the re-ranking process. Since the points follow the diagonal line, there is no change in the false alarm rates of the injections through this process. This infers that the process is safe for injections, which means this process will preserve the signal false-alarm rates.

ACKNOWLEDGMENTS

This research would not be possible without the guidance of Ryan Magee, as well as support from the LIGO Laboratory, the California Institute of Technology, and the National Science Foundation.

REFERENCES

[1] S. Sachdev *et al.*, Astrophys. J. Lett. **905**, L25 (2020), arXiv:2008.04288 [astro-ph.HE].

[2] C. Messick *et al.*, Phys. Rev. D **95**, 042001 (2017), arXiv:1604.04324 [astro-ph.IM].

[3] R. Magee *et al.*, Astrophys. J. Lett. **910**, L21 (2021), arXiv:2102.04555 [astro-ph.HE].

[4] R. Abbott *et al.* (LIGO Scientific, Virgo), Astrophys. J. Lett. **913**, L7 (2021), arXiv:2010.14533 [astro-ph.HE].

[5] R. Magee, A.-S. Deutsch, P. McClincy, C. Hanna, C. Horst, D. Meacher, C. Messick, S. Shandera, and M. Wade, Phys. Rev. D **98**, 103024 (2018), arXiv:1808.04772 [astro-ph.IM].

[6] J. Abadie *et al.* (LIGO Scientific, VIRGO), (2010), arXiv:1003.2481 [gr-qc].

[7] Wikipedia Contributors, "Whitening transformation — Wikipedia, the free encyclopedia," (2021), [Online; accessed 21-September-2021].

[8] K. Cannon *et al.*, Astrophys. J. **748**, 136 (2012), arXiv:1107.2665 [astro-ph.IM].