

# Application of a Zero-latency Whitening Filter to Compact Binary Coalescence Gravitational-wave Searches

Leo Tsukada,<sup>1,2,\*</sup> Kipp Cannon,<sup>1,†</sup> Chad Hanna,<sup>3</sup> Drew Keppel,<sup>4</sup> Duncan Meacher,<sup>3</sup> and Cody Messick<sup>3</sup>

<sup>1</sup>*Research Center for the Early Universe (RESCEU), Graduate School of Science,  
The University of Tokyo, Tokyo 113-0033, Japan*

<sup>2</sup>*Department of Physics, Graduate School of Science,  
The University of Tokyo, Tokyo 113-0033, Japan*

<sup>3</sup>*The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

<sup>4</sup>*California Institute of Technology, Pasadena, California 91125, USA*

(Dated: August 14, 2017)

Joint electromagnetic and gravitational-wave (GW) observation is a major goal of both the GW astronomy and electromagnetic astronomy communities for the coming decade. One way to accomplish this goal is to direct follow-up of GW candidates. Prompt electromagnetic emission may fade quickly, therefore it is desirable to have GW detection happen as quickly as possible. A leading source of latency in GW detection is the whitening of the data. We examine the performance of a zero-latency whitening filter in a detection pipeline for compact binary coalescence (CBC) GW signals. We find that the filter reproduces signal-to-noise ratio (SNR) sufficiently consistent with the results of the original high-latency and phase-preserving filter for both noise and artificial GW signals (called “injections”). Additionally, we demonstrate that these two whitening filters show excellent agreement in  $\chi^2$  value, a discriminator for GW signals.

Keywords: gravitational waves, compact binary coalescence, whitening filter, low latency

## I. INTRODUCTION

The main target of ground-based gravitational-wave detectors are GW signals from CBC and this design concept has been put on firmer grounds than ever since the first detection of GWs from binary black holes, GW150914 [1]. In addition to the detection of GWs alone, association of GW signals with electromagnetic (EM) counterparts will provide scientific benefits for multiple fields of astronomy. One of the most promising EM counterpart candidates is a short gamma-ray burst (SGRB) [2].

SGRBs are intense and non-repeating flashes of  $\gamma$ -rays with a duration of  $\lesssim 2$  s. Their origin and mechanism have been investigated and, although no conclusive evidence has been discovered, it is widely accepted that SGRBs are emitted via a beamed relativistic jet from CBC events containing at least one neutron star (NS) [3, 4]. The association of a GW signal with its SGRB counterpart will establish the CBC progenitor model. Furthermore, it will also enable us to improve localization precision of the GW source, which leads to identification of the host galaxy and the progenitor’s local environment [2].

Despite the first successful GW detections by the advanced Laser Interferometer Gravitational-wave Observatory (aLIGO [5]), there are still a couple of challenges to tackle toward the realization of a follow-up search with EM waves. One is the poor localizability of GW sources. This challenge can be addressed by a worldwide detector network on Earth. The aLIGO consists

of two detectors at geographically separate sites: one in Hanford, WA, with 4 km arms; one in Livingston, LA, with 4 km arms. Furthermore, several other detectors are being constructed or upgraded. Virgo [6] in Italy, for instance, will soon participate in a joint observation run with aLIGO, and KAGRA [7] in Japan is now in the process of apparatus installation.

Another challenge, more relevant to this work, is the latency of a data-analysis pipeline, specifically the GstLAL pipeline for this work. With respect to latency, the pipeline consists of five components: data calibration, data distribution, whitening process (explained in Section. II B), trigger production, and alert distribution [8, 9]. The total latency is a little less than 70 s [10], the first three are the major bottlenecks taking  $\sim 30$  s in total.

Because the most sensitive EM telescopes are unable to survey the entire sky continuously, alerting them to the occurrence at a CBC in a timely manner increases the chance of a successful association. Theoretical work on several GRB models propose a vast range of the time lag between a GW emission and the onset of a following SGRB, from  $<10$  s [11] to  $10^3$  s to  $10^4$  s [12]. This motivates achieving alert latencies below 10 s.

This work will focus on the whitening process among the components mentioned above and describe how the whitening filter can be optimized in terms of latency. Thereafter, we investigate the influence of this optimization on the original detectability of GW signals. In particular, we implement the zero-latency algorithm with a Finite-impulse-response (FIR) filter and compare the values of SNR and  $\chi^2$  between the original (frequency-domain) and zero-latency algorithms. It should be noted that the zero-latency algorithm is applicable to general

---

\* tsukada@resceu.s.u-tokyo.ac.jp

† kipp@resceu.s.u-tokyo.ac.jp

CBC data-analysis pipelines.

This paper is organized as follows. In Section II, we give an overview of statistical method to analyze CBC GW signals and the comparison between the frequency-domain and zero-latency whitening algorithms. In Section III, we present the performance tests of the zero-latency whitening filter, including the comparison to the original whitening filter. Lastly, we conclude in Section IV.

## II. METHOD

### A. Matched filtering and $\chi^2$ test

One statistic used to estimate the detection significance of GW signals is the signal-to-noise ratio (SNR),  $\rho$ , computed using a matched filter [13]. For the calibrated strain data  $s(t)$  and a template waveform  $h(t)$ , the output of the matched filter is formulated as shown below.

$$z = (s(t), h(t)) \equiv 4 \int_0^\infty \frac{\tilde{h}^*(f)\tilde{s}(f)}{S_n(f)} df \quad (1)$$

where  $S_n(f)$  is the one-sided power spectral density of the detector strain noise.  $\rho$  is defined as the output of the matched filter in the case of a normalized GW signal.

$$\rho \equiv \frac{z}{\sigma} \quad (2)$$

$$\text{where } \sigma^2 \equiv 4 \int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)} df \quad (3)$$

Additionally, strain data contain noise transients which do not obey a Gaussian distribution and may accidentally produce high  $\rho$ . Such non-Gaussian transient (referred to as “glitches”) are distinct from real GW signals in that they do not have the morphology of the template  $h(t)$ . Making use of this distinction, we employ another statistic,  $\chi^2$ , defined below, in order to distinguish the transients [9]. The time-dependent SNR of data is compared with that expected from the a real signal using the auto-correlation function of the template at its time of peak amplitude,  $R(t)$ . A  $\chi^2$  value is computed for each trigger using the time-dependent SNR  $\rho(t)$ , the peak SNR  $\rho_p$  at the timestamp of  $t_p$ , the noise-weighted auto-correlation function of a template  $R(t)$ .

$$\chi^2 \equiv \frac{1}{\mu} \int_{t_p-\delta t}^{t_p+\delta t} |\rho(t) - \rho_p R(t)|^2 dt \quad (4)$$

$$\text{where } \rho(t) \equiv \frac{4}{\sigma} \int_0^\infty \frac{\tilde{h}^*(f)\tilde{s}(f)}{S_n(f)} e^{2\pi i f(t-t_p)} df \quad (5)$$

$$R(t) \equiv \frac{4}{\sigma^2} \int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)} e^{2\pi i f(t-t_p)} df \quad (6)$$

The factor  $\mu$  is to normalize the  $\chi$  value for a well-fit signal. The time window  $\delta t$  is a tunable parameter.

Both SNR and  $\chi^2$  values are used to derive a likelihood ratio necessary for ranking triggers [14].

### B. Frequency-domain whitening filter

As can be seen in Eq.(1), a matched filter is interpreted simply as an inner product between  $\tilde{s}(f)$  and  $\tilde{h}^*(f)$  with the weight of  $1/\sqrt{S_n(f)}$  for each. This weighting process is called “whitening”, named after the fact that the transformation ideally returns only white noise. Referring to Fig.1, the power spectrum density (PSD) of 32 s chunks of input data is measured for the subsequent whitening (i). Since the discrete Fourier transform (DFT) processes a chunk of a time series as if it were a periodic infinite series, we apply a Hann window function (ii) to suppress discontinuities on the boundaries of each period. These discontinuities will lead to click noise in the whitened data and may produce fake GW signals. Furthermore, in order to prevent any remaining discontinuity on the boundaries, some additional samples around the 32 s time series have been filled with zeros after applying the hann window. Then, the spectrum of the Hann-windowed block is computed (iii) and the block is whitened with the PSD (iv). After the inverse DFT (v), we further apply a Tukey window (vi). The purpose of this windowing is to suppress time-domain leakage which appears through the whitening and IDFT. The above procedure is repeatedly applied for every 32 s block with a 50% overlap and the PSD is updated every 16 s. In the end, all whitened data chunks, each of which is separately processed, are added with a consecutive 16 s shift in order to output a continuous time series of the whitened data (vii). The algorithm’s main drawback is latency. Since a 32 s block is processed all at once every 16 s, the latency depends on the sample’s location in the block, and can be anywhere from 16 s to 32 s.

### C. Time-domain whitening filter

#### *FIR-filter-based algorithm*

Here, we present an alternative FIR-filter-based algorithm to the frequency-domain whitening described above. Fig.2 shows the square root of the inverse PSD computed from the input LIGO strain data. This amplitude response has been employed to construct the FIR of a linear-phase filter shown in Fig.3. Therefore, it is possible to replace step (4) in Fig.1 with this FIR filter. It should be noted that due to its peak location, the FIR-filter-based algorithm still has a latency of 16 s.

#### *Zero-latency algorithm*

According to the discussion in the previous paragraph, the peak of the filter must be moved to the left for the latency reduction. It is not possible to change the filter’s latency without changing the whitening transformation. The result will be an approximation of the original filter.

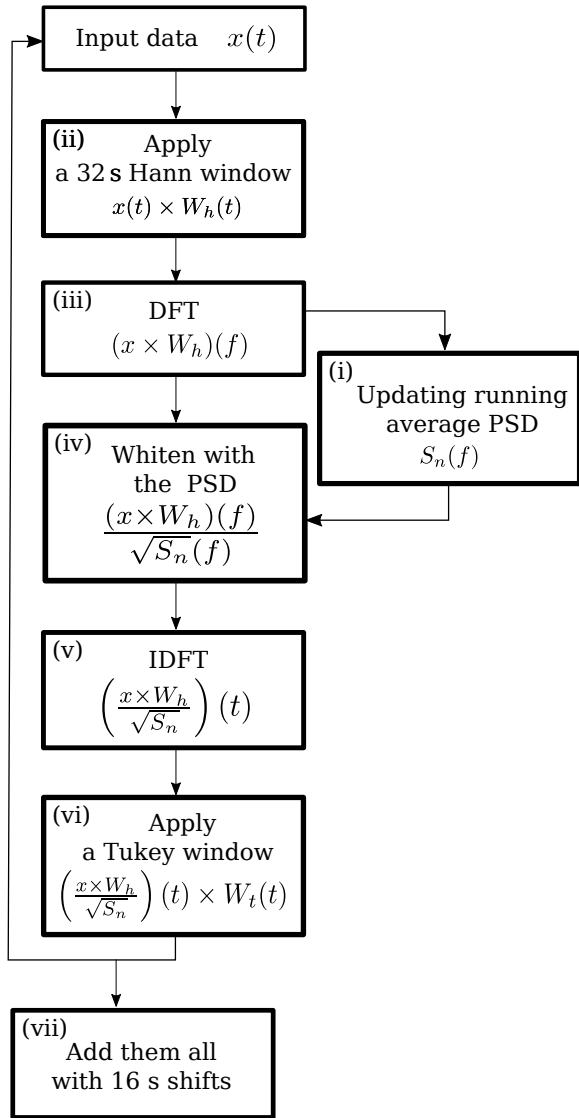


FIG. 1. A schematic diagram of the frequency-domain whitening algorithm. Each numbered process corresponds to the numbers mentioned in Section II B.  $W_h(t)$  and  $W_t(t)$  represent Hann and Tukey window functions, respectively.

We adopt the approximation technique of N.Damera-Venkata *et al.* [15] which approximately derives a phase response of a minimum-phase filter by applying the discrete Hilbert transform to the logarithm of a given magnitude response. Using more samples for the given magnitude response, one can more accurately approximate the magnitude response of the computed minimum-phase filter. The result is shown in Fig.5. Fig.5 indicates that the FIR has its peak at the exact beginning, which is the reason why it is called a “zero-latency whitening filter” in this paper. As described in Section II B, this whitening filter is equally applied to both of the templates and the strain data.




FIG. 2. The amplitude response of the whitening filter. Due to the sample rate of 4096 Hz, the Nyquist frequency is 2048 Hz. The spectrum is computed simply from square-root of the inverse PSD.

#### Smooth PSD tracking

In order to replace the whole algorithm in Fig.1, we have implemented an alternative method of the windowing process along with the whitening transformation. Specifically, we have allowed the PSD transition to occur continuously. Here, we have created a function which returns a linear combination of the newest and next newest filters during their transition as described by Eq.(7). The coefficient of the newer filter smoothly shifts from zero to one, sample by sample, according to a sinusoidal function. The zero-latency algorithm applies this function recursively any time a new whitening filter becomes available.

$$s'(t) = \begin{cases} s_{old}(t) & (t < t_{up}) \\ \cos^2 \frac{2\pi(t-t_{up})}{\Delta t_{tr}} s_{old}(t) + \sin^2 \frac{2\pi(t-t_{up})}{\Delta t_{tr}} s_{new}(t) & (t_{up} \leq t < t_{up} + \Delta t_{tr}) \\ s_{new}(t) & (t_{up} + \Delta t_{tr} \leq t) \end{cases} \quad (7)$$




FIG. 3. The FIR of the original whitening filter. This is symmetric about its peak, a feature of a linear-phase filter. The peak in the middle leads to the latency of 16 s.




FIG. 5. The impulse response of the zero-latency whitening filter. This FIR filter is purely causal.




FIG. 4. Magnitude of auto-correlation of output stream from the zero-latency whitening filter. Technically, the absolute value is taken. It should be noted that this does not show any peak at the frequency of  $\sim 100$  Hz.

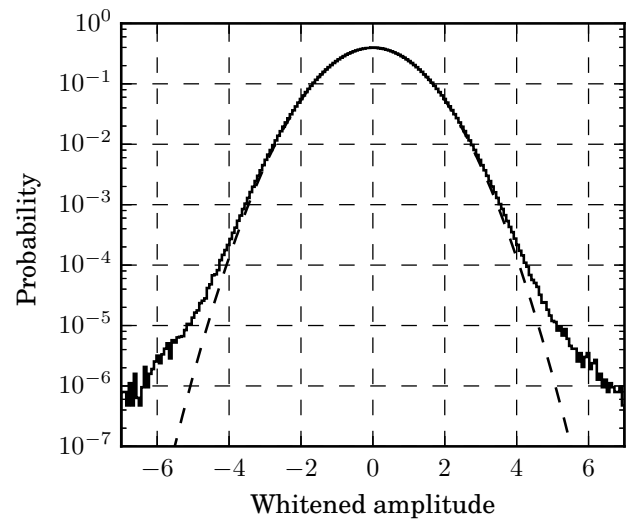


FIG. 6. The logarithmic histogram of the output's amplitude. The departure of the observed counts from the expected counts outside  $(-5, 5)$  is due to the presence of non-Gaussian “glitches” in the interferometer data.




FIG. 7. The averaged power spectral density of output stream from the zero-latency whitening filter. The data below 12 Hz are dropped off to ignore the effect of a high pass filter applied before the whitening filter.

where  $s'(t)$  is the resulting filter,  $s_{old,new}$  is the FIR of an older and newer filter respectively,  $t$  is the current time stamp,  $t_{up}$  is the time stamp when the PSD is updated and  $\Delta t_{tr}$  is the duration of the filter transition. Particularly, we set  $\Delta t_{tr}$  as 0.125 s in this work so that the transition timescale lies outside the frequency band of interest, which starts at 10 Hz. Note that this method is not unique in the application of a whitening filter but can be used for other time-dependent filtering.

### III. TESTS

Unlike the frequency-domain algorithm described in Section II B, the zero-latency whitening filter does not conserve the phase of input data, which potentially harms the GW detectability. Therefore, it is necessary to demonstrate how significantly this change affects the resulting SNR and  $\chi^2$ .

Here, we implement each of the zero-latency and frequency-domain whitening filters in the CBC gstLAL pipeline, which is compiled in LSC Algorithm Library (LAL) [16]. The pipeline scans LIGO strain data for any GW signal candidates (called “triggers”) and computes SNR and  $\chi^2$  for each trigger. In this work, we employ strain data from H1 with a duration of 45.056 s (from 08:25:23 to 20:56:19 UTC on 2005/11/27) during the fifth science run, called S5 [17]. A template bank is used spanning: component masses  $3M_{\odot} \leq m_1, m_2 \leq 6M_{\odot}$ ; total mass  $6M_{\odot} \leq M_{total} \leq 12M_{\odot}$ ; a minimal match of 97%; sampling frequency of 2048 Hz; non-spinning waveform to second post-Newtonian (PN) order. Along with statistical tests described below, we conduct two kinds of tests, namely, a noise-based and an injection-based test. In the noise-based test, the pipeline computes SNR and  $\chi^2$  from the strain data with no GW signal. Therefore,

all triggers in this test arise from detector noise which accidentally produce higher SNR than the threshold. On the other hand, the injection-based test requires artificial GW signals, and so we add injections to the same strain data. An injection is generated once every 31.4 s in the data, so the number of injections amounts to 1435. These two tests examine the agreement between the two whitening filters for noise and signals.

#### A. Statistical tests

Fig.4 and Fig.6 show the auto-correlation and amplitude histogram created from an output stream of the zero-latency whitening filter. For comparison, expected curves are shown as dashed lines in the both figures, each of which indicates a delta-function with some variance and a Gaussian distribution respectively. Both plots show good agreement between the output and pure white noise. In particular, there is no apparent peak of the auto-correlation at around 100 Hz, at which the amplitude response shows its peak (See Fig.2). Also, the averaged power spectral density is shown in Fig. 7. The spectrum is flattened throughout the shown frequency domain. Therefore, we conclude that the zero-latency algorithm sufficiently functions as a whitening filter.

#### B. Noise-based test

In Fig.8, we plot SNR and  $\chi^2$  computed for each noise trigger by the zero-latency whitening filter versus the frequency-domain one. Here, we associate each counterpart by spotting a pair of triggers within the end-time window of  $10^{-2}$  s and with identical component masses ( $m_1$  and  $m_2$ ). The two whiteners produced triggers with an SNR of  $5 \sim 60$  and a  $\chi^2$  of  $10^{-1} \sim 10^3$ . Fig.8 presents good agreement of both SNR and  $\chi^2$  between the two whitening filters.

#### C. Injection-based test

Fig.9 shows SNR and  $\chi^2$  computed with the zero-latency whitening filter versus the frequency-domain one in the presence of injections. Also, coalescence phase for every injection is shown in Fig.10. The coalescence phase is a phase of injection waveform at the coalescence time and determined by the ratio between cosine and sine components of a chosen template. In this test, we first simulate waveforms, based on a collection of parameters chosen randomly from a given probability distribution:  $m_1$  and  $m_2$  from a Gaussian distribution with the mean of  $4.5M_{\odot}$  and the standard deviation of  $0.5M_{\odot}$ ; cosine of the inclination chosen from a uniform distribution; non-spinning waveform to second post-Newtonian (PN) order with the cut-off frequency of 30 Hz. Next, the pipeline searches for and extracts injection by spotting the one

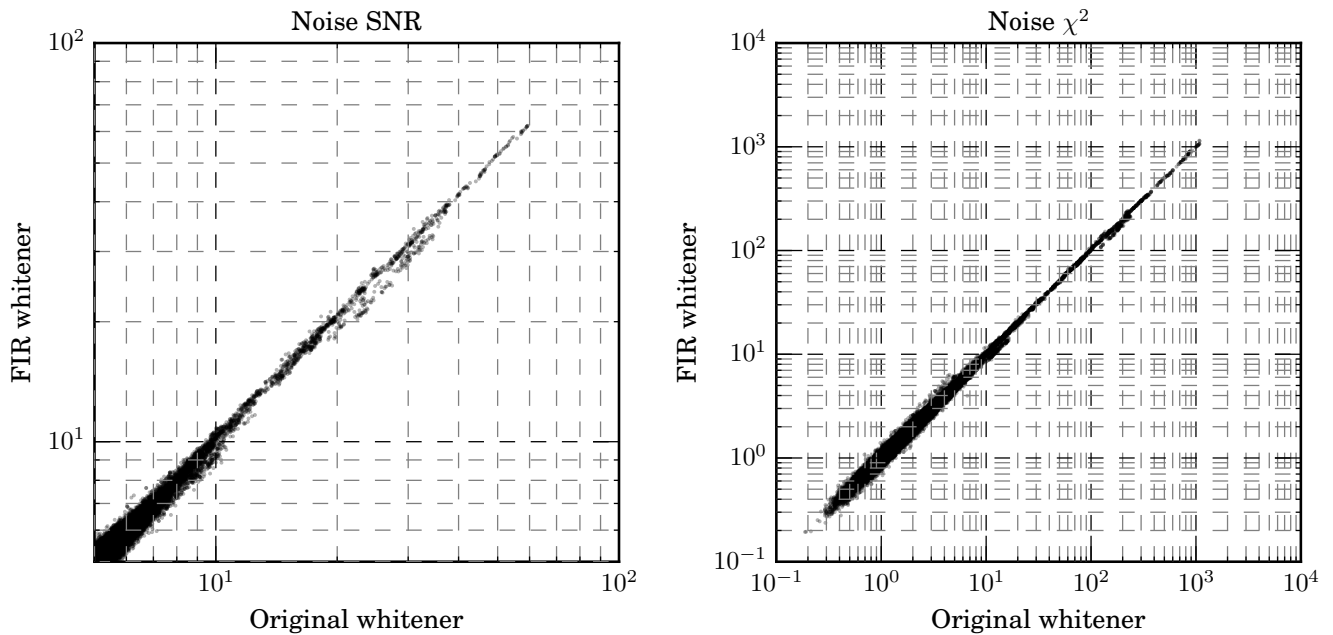


FIG. 8. Scatter plots of SNR and  $\chi^2$  computed by the frequency-domain and zero-latency FIR whiteners for noise triggers. The value of SNR and  $\chi^2$  range 5 – 60 and 0.2 –  $10^3$  respectively. We have associated a pair of triggers within the end-time window of  $10^{-2}$  s and with identical component masses,  $m_1$  and  $m_2$

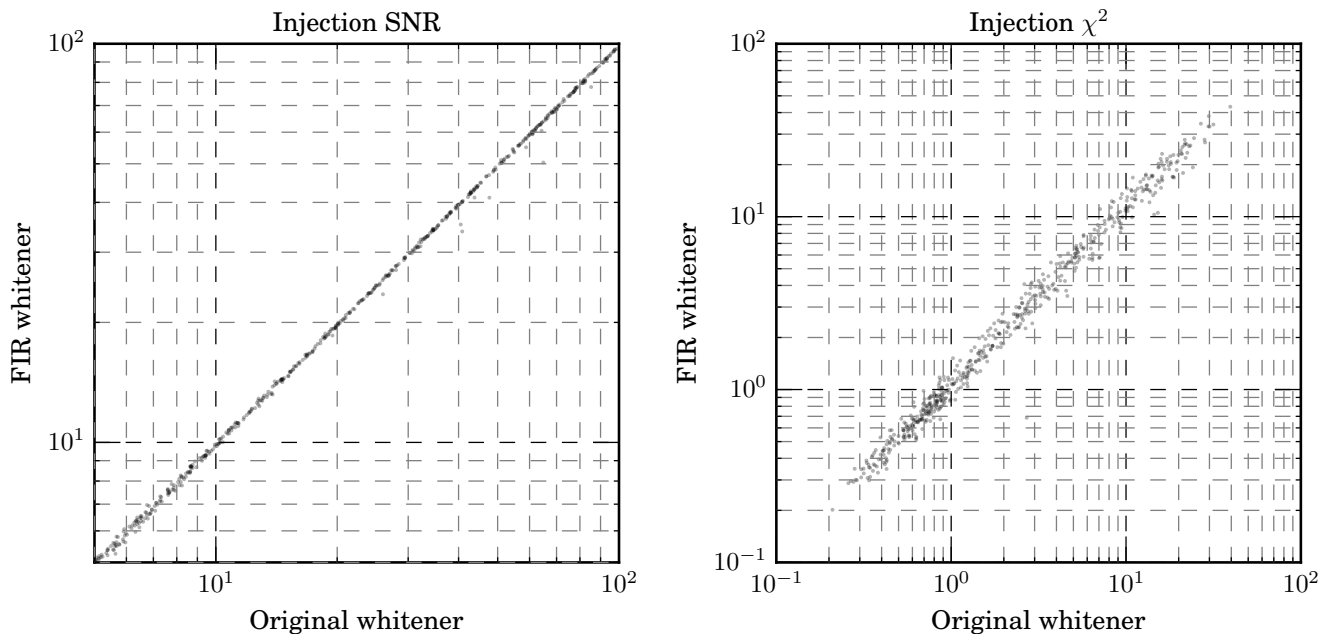


FIG. 9. Scatter plots of SNR and  $\chi^2$  computed by the frequency-domain and zero-latency FIR whiteners for injection triggers. First, 1289 injections with the SNR of 5 –  $10^4$  were generated and truncated so that the maximum SNR would be 100. This is because triggers with higher SNR do not help to examine the critical discernibility of GW signals. Consequently, the values of SNR and  $\chi^2$  range 5 – 100 and 0.3 – 400 respectively.

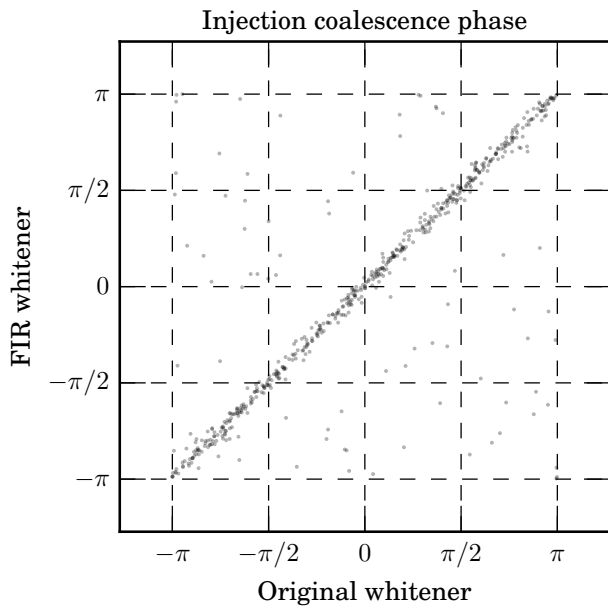


FIG. 10. The scatter plots of coalescence phase computed by the frequency-domain and zero-latency FIR whiteners for injection triggers. These triggers, whose SNR fall into the range from 5 to 100.

with the highest SNR among all located within 1 s of its true end time. After the trigger extraction, SNR,  $\chi^2$  and coalescence phase of every injection trigger is recorded. In the end, only those with an SNR less than 100 are left to fit into the scatter plot, Fig.9. The above procedure is conducted for both whitening filters and we identify each counterpart by the end time of each injection.

We have also conducted a consistency test for an injection's end time in the case of the two whitening filters. Fig.11, 12 show histograms of the discrepancy between the true and estimated end time of each of the filters. In the both figures, the central peak has a tail width of  $\sim 100$  ms, which is consistent with the typical tail width of the auto-correlation function of injected waveforms, suggesting that the pipeline properly generates and extract the injections from the triggers.

As a result of the injection-based test, we find the SNR and  $\chi^2$  computed with the zero-latency whitening filter to agree with those of the frequency-domain one. Although some triggers indicate that the new whitening filter slightly underestimates an SNR compared to the original one (See Fig.9), it will not harm the GW detectability since this case lies in the higher SNR regime.

#### IV. CONCLUSION

We have applied and implemented an algorithm that optimizes latency for the whitening filter in the CBC




FIG. 11. Histogram of the discrepancy between the true and estimated end time for each injection in case of the frequency-domain whitening filter. Each bin has the number of injection triggers whose end-time difference lies in the bin's range.

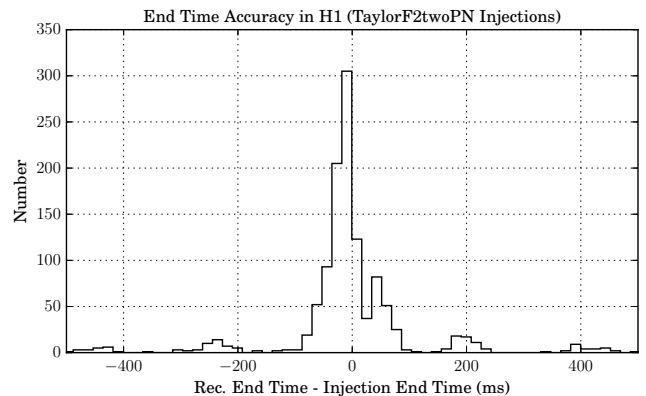


FIG. 12. Histogram of the discrepancy between the true and estimated end time for each injection in case of the zero-latency whitening filter. Each bin has the number of injection triggers whose end-time difference lies in the bin's range.

data-analysis pipeline. Through statistical tests between the frequency-domain and zero-latency whitening filters, we have found that the two statistical values, SNR and  $\chi^2$ , are in sufficient agreement for both noise and injection triggers. As a result, we have achieved a 16 s latency reduction in the whitening process. It should be noted that this work has yielded the first confirmation that a zero-latency whitening filter can be employed in the data-analysis pipeline for CBC GW searches.

#### ACKNOWLEDGMENTS

The authors are grateful to the LIGO lab for an offer of the S5 strain data. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Sci-

ence Foundation. This research was supported by the National Science Foundation through PHY-1454389 and

ACI-1642391. Funding for this project was provided by the Charles E. Kaufman Foundation of The Pittsburgh Foundation.

- 
- [1] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, Physical review letters **116**, 061102 (2016).
- [2] B. D. Metzger and E. Berger, Astrophys. J. **746**, 48 (2012), arXiv:1108.6056 [astro-ph.HE].
- [3] D. Eichler, M. Livio, T. Piran, and D. N. Schramm, Nature (London) **340**, 126 (1989).
- [4] E. Berger, 10.1146/annurev-astro-081913-035926, 1311.2603v1.
- [5] B. P. Abbott *et al.*, Classical and Quantum Gravity **32**, 074001 (2015).
- [6] F. Acernese *et al.* (VIRGO), Class. Quant. Grav. **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [7] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (The KAGRA Collaboration), Phys. Rev. D **88**, 043007 (2013).
- [8] K. Cannon, R. Cariou, A. Chapman, M. Crispin-Ortuzar, N. Fotopoulos, M. Frei, C. Hanna, E. Kara, D. Keppel, L. Liao, S. Privitera, A. Searle, L. Singer, and A. Weinstein, The Astrophysical Journal **748**, 136 (2012).
- [9] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. E. Creighton, R. Everett, C. Hanna, D. Keppel, R. N. Lang, T. G. F. Li, D. Meacher, A. Nielsen, C. Pankow, S. Privitera, H. Qi, S. Sachdev, L. Sadeghian, L. Singer, E. G. Thomas, L. Wade, M. Wade, A. Weinstein, and K. Wiesner, Phys. Rev. D **95**, 042001 (2017).
- [10] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **116**, 241103 (2016).
- [11] X. Li, Y.-M. Hu, Y.-Z. Fan, and D.-M. Wei, Astrophys. J. **827**, 75 (2016), arXiv:1601.00180 [astro-ph.HE].
- [12] L. Rezzolla and P. Kumar, Astrophys. J. **802**, 95 (2015), arXiv:1410.8560 [astro-ph.HE].
- [13] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, (2005), 10.1103/PhysRevD.85.122006, gr-qc/0509116v2.
- [14] K. Cannon, C. Hanna, and J. Peoples, 1504.04632v1.
- [15] N. Damera-Venkata, B. L. Evans, and S. R. McCaslin, IEEE Transactions on Signal Processing **48**, 1491 (2000).
- [16] LSC, URL <https://wiki.ligo.org/DASWG/LALSuite>.
- [17] B. P. Abbott *et al.*, Reports on Progress in Physics **72**, 076901 (2009).