



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Application of a zero-latency whitening filter to compact binary coalescence gravitational-wave searches

Leo Tsukada, Kipp Cannon, Chad Hanna, Drew Keppel, Duncan Meacher, and Cody Messick

Phys. Rev. D **97**, 103009 — Published 14 May 2018

DOI: [10.1103/PhysRevD.97.103009](https://doi.org/10.1103/PhysRevD.97.103009)

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

Leo Tsukada,^{1,2,*} Kipp Cannon,^{1,†} Chad Hanna,³ Drew Keppel,⁴ Duncan Meacher,³ and Cody Messick³

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

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

³*Department of Physics and Astronomy & Astrophysics
The Pennsylvania State University, University Park, Pennsylvania 16802, USA*

⁴*California Institute of Technology, Pasadena, California 91125, USA*

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 χ^2 value, a discriminator for GW signals.

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

I. INTRODUCTION

Detecting gravitational wave signals from coalescing compact objects such as neutron stars and black holes is a primary goal for ground-based gravitational-wave detectors. The advanced LIGO [1] and advanced Virgo gravitational wave detectors [2] have successfully detected gravitational waves from several black hole binaries [3–6] following the initial binary black hole discovery of GW150914 [7]. Most recently, LIGO and Virgo detected the first ever neutron star collision, known as GW170817 [8]. Quite remarkably, GW170817 was detected in coincidence with a short gamma-ray burst (SGRB), which has long been expected to be one of the most promising electromagnetic (EM) counterpart candidates to binary neutron star mergers [9]. Furthermore, subsequent electromagnetic emission spanning from radio up to X-ray was detected within hours to weeks after the GW emission [10].

The association of GW170817 with electromagnetic transients across many wavelengths and the discoveries that followed were initiated by the rapid identification of the gravitational wave signal by the GstLAL detection pipeline [11]. Until the detection of GW170817, the question of what data processing latency is required, scientifically, for such a GW detection system to be most effective at enabling multimessenger astrophysics has been open. Theoretical work on GRB models has proposed a vast range of the time delay between GW emission and the onset of the following SGRB, from <10 s [12] to 10^3 s

to 10^4 s [13]. Now it is known the time lag is approximately 2 s [8], which motivates achieving alert latencies below 2 s.

Within the low-latency analysis paradigm, the GW detection problem consists of three main stages: data calibration and distribution, candidate identification, and alert distribution [11, 14]. The GstLAL project (publicly available in [15]) began in 2009 in order to produce near zero-latency gravitational wave candidate identification [14]. The GstLAL candidate identification achieved latencies of approximately one minute in advanced LIGO’s first observing run. For example, the binary black hole known as GW151226 was detected within 70 seconds [3] by the GstLAL-based compact binary coalescence detection pipeline. Also, the GstLAL has inspired related works to pursue a low-latency search for CBC GW signals, such as J. Luan, *et al.* [16], which proposes an algorithm to reduce computational cost for matched filtering by introducing chains of Infinite Impulse Response filters.

This work addresses one of the open questions put forward in [14] for how to achieve theoretically zero latency candidate event identification, namely, data whitening, which was the largest source of candidate identification latency for the GstLAL pipeline in advanced LIGO’s first observing run. Data whitening is just one part GstLAL’s modular framework for gravitational wave candidate event identification which also includes matched filtering an orthogonal decomposition of gravitational wave templates [17], and using a multidimensional likelihood ratio based ranking statistic to distinguish noise from signal [18]. The overarching pipeline is described in [11] and we will not elaborate further on the pipeline here.

In this work, we describe how the data whitening filter

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

† kipp@resceu.s.u-tokyo.ac.jp

can be optimized in terms of latency. In particular, we implement a zero-latency algorithm with a linear finite-impulse-response (FIR) filter. We use the minimum-phase FIR filter approximation technique described by Damera-Venkata, *et al.* in [19]. Other algorithms for whitening filter approximation are available, for example [20–24]. In principle, the zero-latency algorithm can be used by any GW data-analysis pipeline. For the study presented here we used the GstLAL [11] detection pipeline and note that this whitening algorithm was used during the detection of GW170817. In order to verify the suitability of the zero latency whitening algorithm, we compare relevant parameters of simulated gravitational wave candidates, namely SNR and χ^2 , to the original frequency-domain whitening algorithm. Since, as will be shown below, we find the Damera-Venkata, *et al.*, algorithm to meet the needs of low-latency searches for compact-object mergers, we have not investigated the performance of other techniques at this time.

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 χ^2 test

One statistic used to estimate the detection significance of GW signals is the signal-to-noise ratio (SNR), ρ , computed using a matched filter [25]. For the calibrated strain data $s(t)$ and a template waveform $h(t)$, the output of the matched filter is

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

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

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

where

$$\sigma^2 \equiv (h(t), h(t)). \quad (3)$$

Additionally, strain data contain noise transients which do not obey a Gaussian distribution and may accidentally produce high ρ . 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, χ^2 , defined below, in order to distinguish the transients [11]. 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 χ^2 value is computed for each trigger using the time-dependent SNR $\rho(t)$, the peak SNR ρ_p at the timestamp of t_p , the noise-weighted auto-correlation function of a template $R(t)$.

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

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 μ is to normalize the χ value for a well-fit signal. The time window δt is a tunable parameter.

Both SNR and χ^2 values are used to derive a likelihood ratio necessary for ranking triggers [18].

B. Frequency-domain whitening filter

As can be seen in (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.

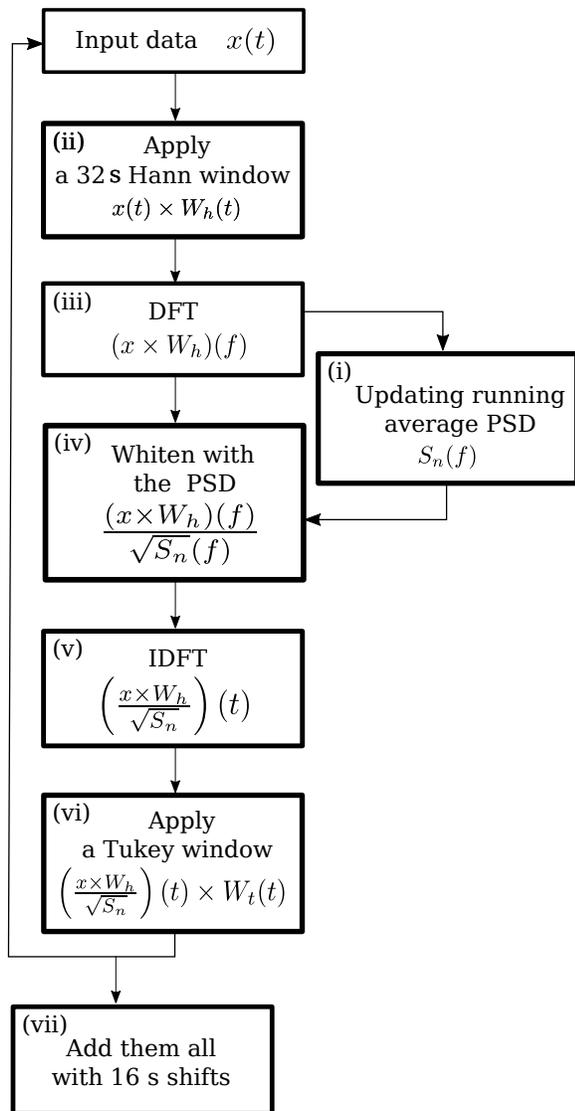


FIG. 1. A schematic diagram of the frequency-domain whitening algorithm. Each numbered process corresponds to the numbers mentioned in Section II B. $W_h(t)$ and $W_t(t)$ represent Hann and Tukey window functions, respectively.

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. For the given LIGO strain data, the square root of the inverse PSD is employed to construct the FIR of a linear-phase filter shown in Fig. 3. Therefore, it is possible to replace step (iv) 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. Fig. 2 is amplitude and phase responses of the filter. The phase response shows no phase shift during this whitening pro-

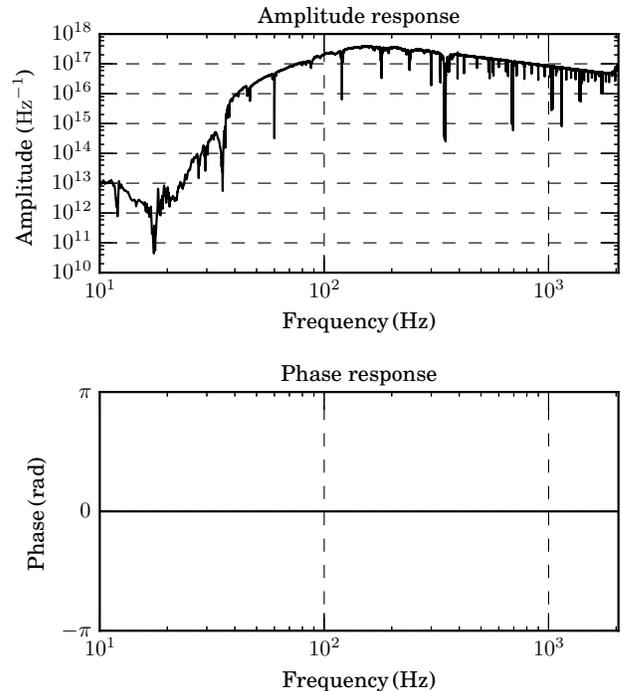


FIG. 2. The amplitude phase response of the linear-phase whitening filter. The Nyquist frequency is 2048 Hz.

cess, which guarantees preserving the phase of the data.

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. We adopt the technique of Damera-Venkata *et al.* [19] which derives a minimum-phase approximation of the desired filter by applying a 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. In both Fig. 3, 5, the time origin can be uniquely determined by the requirement that the timestamps at which artificial GWs (called “injections”) are recovered be preserved (described in Section. III C and Fig. 12, 13). Thus, Fig. 5 indicates that the FIR filter does not use any information of future input samples, which is the reason why it is called a “zero-latency whitening filter”.¹

¹ We ignore the computation time as the time require to compute each output sample is significantly less than the sample period

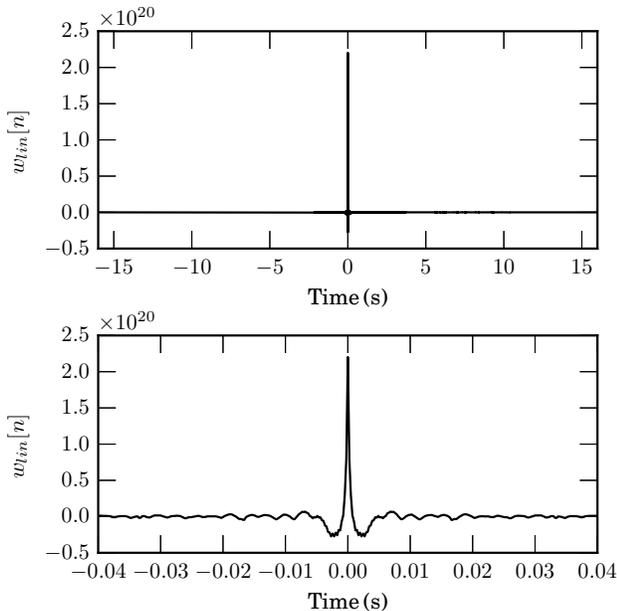


FIG. 3. The impulse response of the original whitening filter. This is symmetric about its peak. Negative times indicate the filter requires data from the future for its evaluation, meaning the output must be delayed with respect to the input. In this case, the latency is 16 s.

As described in Section II B, this whitening filter is equally applied to both of the templates and the strain data. This can be expressed by writing (1) as

$$z = 4 \int_0^\infty \left(\frac{\tilde{h}(f)}{\sqrt{S_n}} \right)^* \left(\frac{\tilde{s}(f)}{\sqrt{S_n}} \right) df. \quad (7)$$

The desired whitening transformation is a frequency-dependent scale factor ($\sqrt{S_n(f)}$ is real-valued, it does not alter phases). We approximate this transformation with a purely causal FIR filter. The difference between the actual transformation and the desired transformation can be described by introducing an error factor in the frequency domain representation. If $1/\sqrt{S_n(f)}$ is the transformation we wish to apply to $\tilde{h}(f)$ and $\tilde{s}(f)$, let $\delta\sqrt{S_n(f)}/\sqrt{S_n(f)}$ be the transformation performed by the causal FIR approximation, so

$$z = 4 \int_0^\infty \left(\frac{\tilde{h}(f)}{\sqrt{S_n}} \delta\sqrt{S_n} \right)^* \left(\frac{\tilde{s}(f)}{\sqrt{S_n}} \delta\sqrt{S_n} \right) df \quad (8)$$

$$= 4 \int_0^\infty \left(\frac{\tilde{h}(f)}{\sqrt{S_n}} \right)^* \left(\frac{\tilde{s}(f)}{\sqrt{S_n}} \right) |\delta\sqrt{S_n}|^2 df. \quad (9)$$

of the data stream.

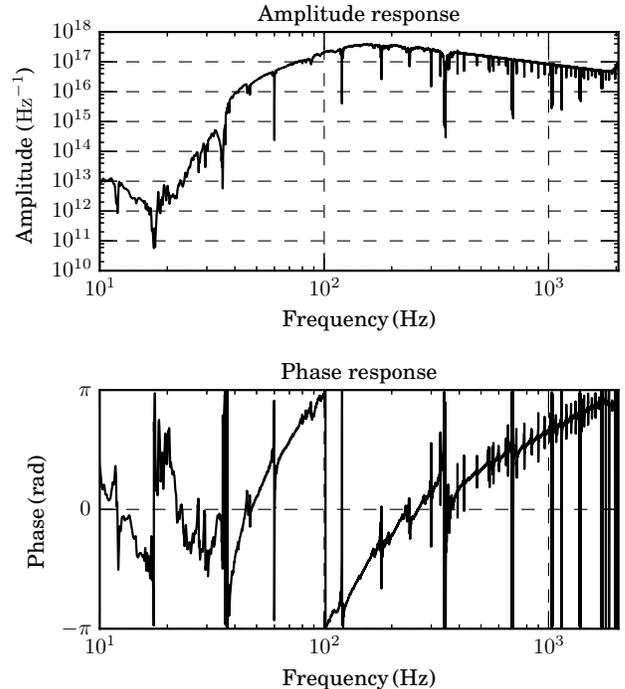


FIG. 4. The amplitude and phase response of the minimum-phase whitening filter. The Nyquist frequency is 2048 Hz.

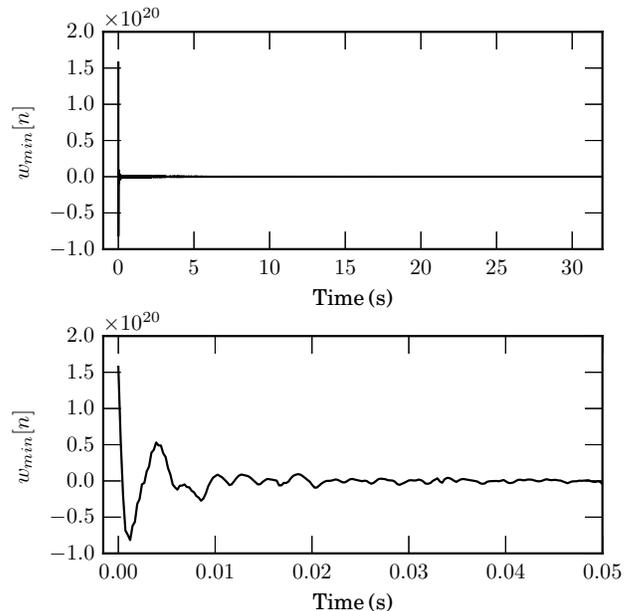


FIG. 5. The impulse response of the zero-latency whitening filter. The impulse response is zero for all negative times, indicating that this filter is purely causal.

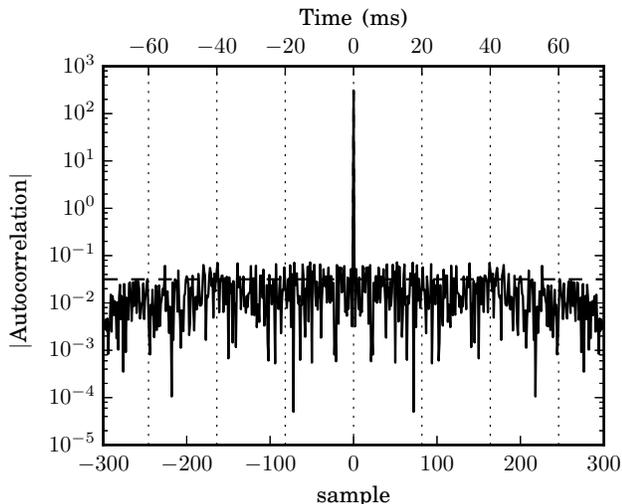


FIG. 6. Magnitude of auto-correlation of output stream from the zero-latency whitening filter. It should be noted that this does not show any peak at the frequency of ~ 100 Hz.

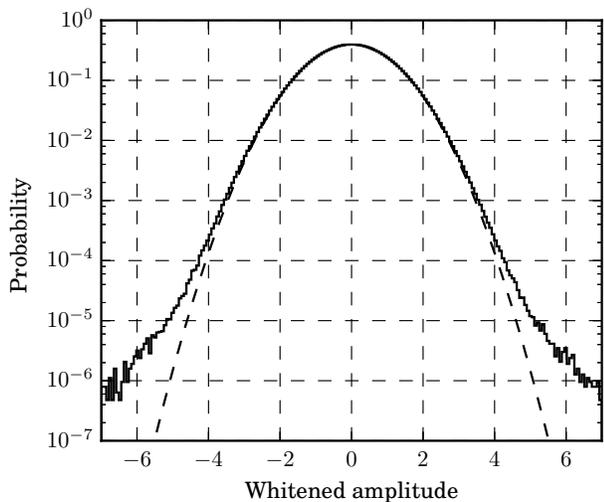


FIG. 7. Histogram of the output's amplitude. The departure of the observed counts from the expected counts outside $(-5, 5)$ is due to the presence of non-Gaussian “glitches” in the interferometer data.

We see the matched filter output is insensitive to $\arg \delta \sqrt{S_n}$, the phase errors arising from the causal FIR approximation (shown in Fig. 4), it is sensitive only to the amplitude errors. Fig. 6, 7, 8 show that the causal FIR filter approximation succeeds in whitening the data, producing zero mean, unit variance, stationary (nearly) white Gaussian noise, so the amplitude errors are not significant.

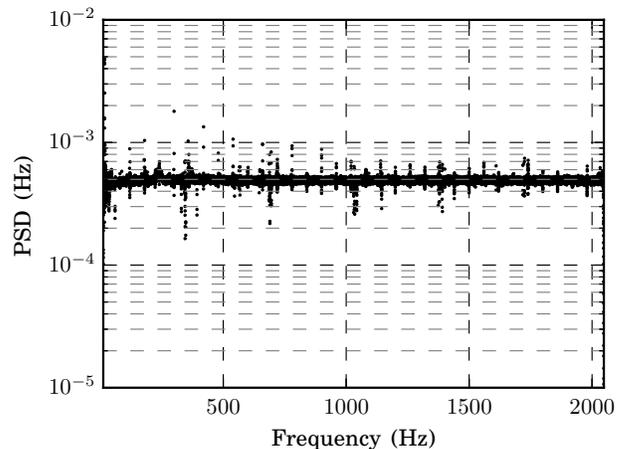


FIG. 8. The averaged power spectral density of output stream from the zero-latency whitening filter. The data below 12 Hz are dropped off to ignore the effect of a high pass filter applied before the whitening filter.

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 (10). 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.

$$s'(t) = \begin{cases} s_{\text{old}}(t) & (t < t_{\text{up}}) \\ s_{\text{old}}(t) \cos^2 \frac{\pi}{2} \frac{(t-t_{\text{up}})}{\Delta t_{\text{tr}}} + s_{\text{new}}(t) \sin^2 \frac{\pi}{2} \frac{(t-t_{\text{up}})}{\Delta t_{\text{tr}}} & (t_{\text{up}} \leq t < t_{\text{up}} + \Delta t_{\text{tr}}) \\ s_{\text{new}}(t) & (t_{\text{up}} + \Delta t_{\text{tr}} \leq t) \end{cases} \quad (10)$$

where $s'(t)$ is the resulting filter, $s_{\text{old}, \text{new}}$ 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 Δt_{tr} is the duration of the filter transition. Particularly, we set Δt_{tr} as 0.125s in this work so that the transition timescale lies outside the frequency band of interest, which starts at 10 Hz. Therefore, the detail of the transition does not affect the detectability of GW signals. 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 χ^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) [26]. The pipeline scans LIGO strain data for any GW signal candidates (called “triggers”) and computes SNR and χ^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 [27]. 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 χ^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. 6 and Fig. 7 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. 4). Also, the averaged power spectral density is shown in Fig. 8. The spectrum is flatten 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. 9, we plot SNR and χ^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 χ^2 of $10^{-1} \sim 10^3$. Fig. 9 presents good agreement of both SNR and χ^2 between the two whitening filters.

C. Injection-based test

Fig. 10 shows SNR and χ^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. 11. 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 with the highest SNR among all located within 1 s of its true end time. After the trigger extraction, SNR, χ^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. 10. 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. 12, 13 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 ~ 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 χ^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. 10), 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 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

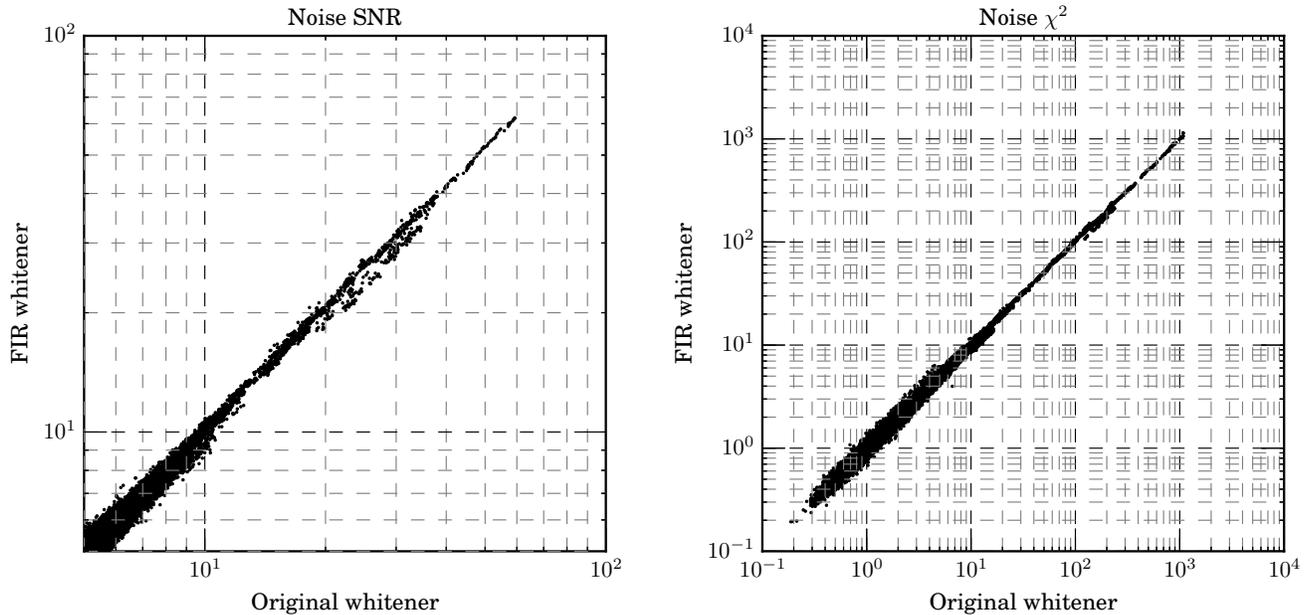


FIG. 9. Scatter plots of SNR and χ^2 computed by the frequency-domain and zero-latency FIR whiteners for noise triggers. The value of SNR and χ^2 range 5 – 60 and 0.2 – 10^3 respectively. We have associated a pair of triggers within the end-time window of 10^{-2} s and with identical component masses, m_1 and m_2

χ^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 Science 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.

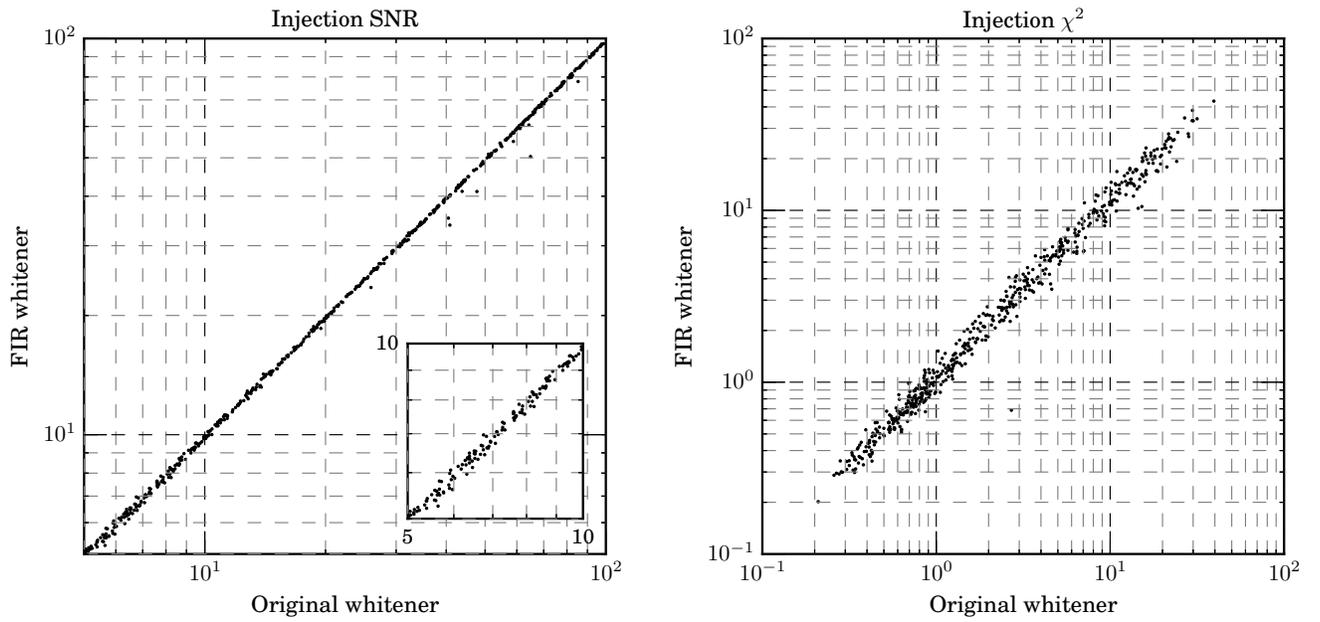


FIG. 10. Scatter plots of SNR and χ^2 computed by the frequency-domain and zero-latency FIR whiteners for injection triggers. First, 1289 injections with the SNR of $5 - 10^4$ were generated and truncated so that the maximum SNR would be 100. This is because triggers with higher SNR do not help to examine the critical discernibility of GW signals. Consequently, the values of SNR and χ^2 range $5 - 100$ and $0.3 - 400$ respectively. The smaller plot located on the left panel shows the data with SNR < 10

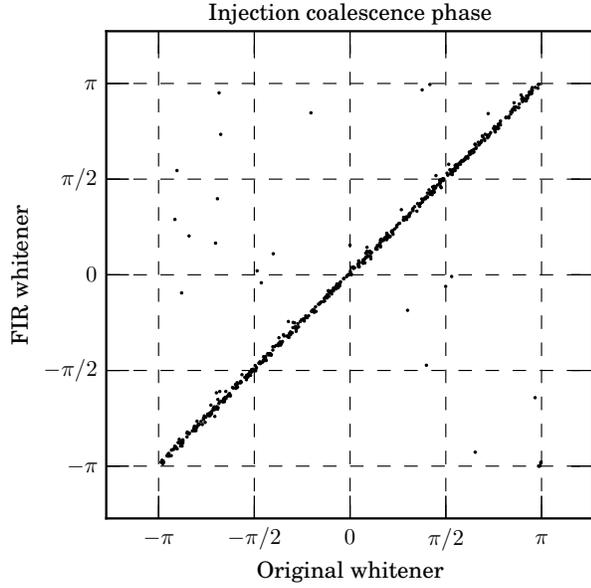


FIG. 11. The scatter plots of coalescence phase computed by the frequency-domain and zero-latency FIR whiteners for injection triggers. These triggers, whose SNR fall into the range from 5 to 100.

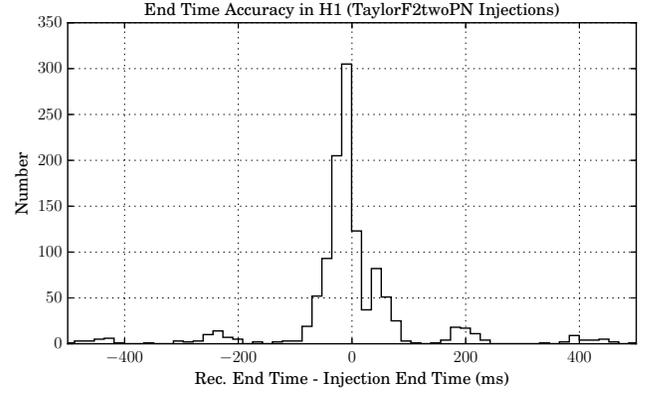


FIG. 13. Histogram of the discrepancy between the true and estimated end time for each injection in case of the zero-latency whitening filter. Each bin has the number of injection triggers whose end-time difference lies in the bin's range.

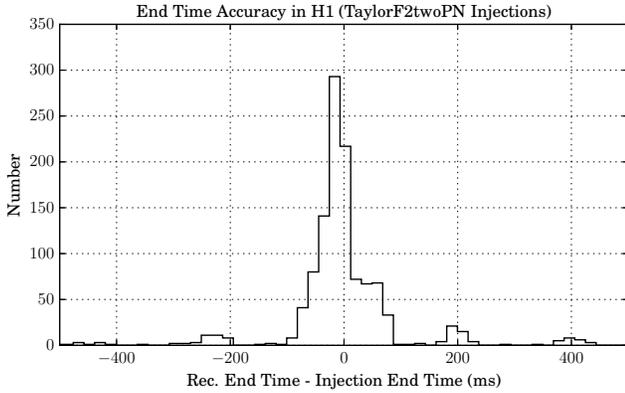


FIG. 12. Histogram of the discrepancy between the true and estimated end time for each injection in case of the frequency-domain whitening filter. Each bin has the number of injection triggers whose end-time difference lies in the bin's range.

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Class. Quant. Grav.* **32**, 074001 (2015).
- [2] F. Acernese *et al.* (Virgo Collaboration), *Class. Quant. Grav.* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].
- [3] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [4] *Physical Review Letters* **118** (2017), 10.1103/PhysRevLett.118.221101.
- [5] B. P. Abbott *et al.*, *The Astrophysical Journal Letters* **851**, L35 (2017).
- [6] *Physical Review Letters* **119** (2017), 10.1103/PhysRevLett.119.141101.
- [7] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [8] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **119**, 161101 (2017).
- [9] B. D. Metzger and E. Berger, *Astrophys. J.* **746**, 48 (2012), arXiv:1108.6056 [astro-ph.HE].
- [10] B. Abbott *et al.*, *Astrophysical Journal Letters* **848**, L12 (2017).
- [11] C. Messick *et al.*, *Phys. Rev. D* **95**, 042001 (2017).
- [12] X. Li, Y.-M. Hu, Y.-Z. Fan, and D.-M. Wei, *Astrophys. J.* **827**, 75 (2016), arXiv:1601.00180 [astro-ph.HE].
- [13] L. Rezzolla and P. Kumar, *Astrophys. J.* **802**, 95 (2015), arXiv:1410.8560 [astro-ph.HE].
- [14] 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).
- [15] LSC, (<https://wiki.ligo.org/DASWG/GstLAL>).
- [16] S. Hooper *et al.*, *Physical Review D* **86** (2012), 10.1103/PhysRevD.86.024012.
- [17] K. Cannon, *Physical Review D* **85** (2012), 10.1103/PhysRevD.85.081504.
- [18] K. Cannon, C. Hanna, and J. Peoples, arXiv:1504.04632.
- [19] N. Damera-Venkata, B. L. Evans, and S. R. McCaslin, *IEEE Trans. Signal Processing* **48**, 1491 (2000).
- [20] S. C. Pei and H. S. Lin, *IEEE Transactions on Circuits and Systems II: Express Briefs* **53**, 1113 (2006).
- [21] S. C. Pei, H. S. Lin, and P. H. Wang, *IEEE Transactions on Circuits and Systems I: Regular Papers* **57**, 925 (2010).
- [22] S. Kidambi and A. Antoniou, *IEEE Transactions on Circuits and Systems II: Express Briefs* **64**, 472 (2017).
- [23] B. C. Garai, P. Das, and A. K. Mishra, *Signal Processing* **91**, 1812 (2011).
- [24] M. Abo-Zahhad and Q. Al-Zoubi, *Digital Signal Processing* **16**, 211 (2006).
- [25] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2005), arXiv:gr-qc/0509116.
- [26] LSC, (<https://wiki.ligo.org/DASWG/LALSuite>).
- [27] B. P. Abbott *et al.* (LIGO Scientific Collaboration), *Rep. Prog. Phys.* **72**, 076901 (2009).