

Low significance of evidence for black hole echoes in gravitational wave data

Julian Westerweck,^{1,2,*} Alex Nielsen,^{1,2,†} Ofek Fischer-Birnholtz,^{1,2,‡}
Miriam Cabero,^{1,2} Collin Capano,^{1,2} Thomas Dent,^{1,2} Badri
Krishnan,^{1,2} Grant Meadors,^{1,3} and Alexander H. Nitz^{1,2}

¹*Max-Planck-Institut für Gravitationsphysik, D-30167 Hannover, Germany*

²*Leibniz Universität Hannover, D-30167 Hannover, Germany*

³*Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany*

Recent detections of merging black holes allow observational tests of the nature of these objects. In some proposed models, non-trivial structure at or near the black hole horizon could lead to echo signals in gravitational wave data. Recently, Abedi, Dykaar and Afshordi claimed tentative evidence for repeating damped echo signals following the gravitational-wave signals of the binary black hole merger events recorded in the first observational period of the Advanced LIGO interferometers. We reanalyse the same data, addressing some of the shortcomings of their method using more background data and a modified procedure. We find a reduced statistical significance for the claims of evidence for echoes, and calculate increased p-values for the null hypothesis of signal-free noise. We conclude that their analysis does not provide significant observational evidence for the existence of Planck-scale structure at black hole horizons.

I. INTRODUCTION

The detections of gravitational wave (GW) signals allow for new tests of the nature of black holes [1–8]. Black holes are characterised by their horizons. In vacuum general relativity these horizons are devoid of material structure. The possibility that additional structure may form at or near the horizon location has been widely discussed in the literature, motivated by a number of different models and theoretical considerations. The Advanced LIGO [9, 10] and VIRGO [11] detectors have detected gravitational wave signals from several binary black hole mergers [1–6]. These detections now make those ideas testable in the observational regime.

A generic set of models called Ultra Compact Objects (UCOs) [12–16] can mimic black holes in terms of their gravitational wave emission at early stages. These models are designed to match the properties of standard black holes at sufficiently large distances, but differ in the near-horizon regime. The gravitational wave signal from the inspiral of two UCOs is expected to be almost identical to that of standard black holes (for possible tidal modifications see [17]). However, the merger and ringdown signals may differ sufficiently to be detectable. Near-horizon material structures motivated by semi-classical and quantum gravity theories could lead to interactions in the near-horizon regime. This would lead to at least partial reflection of incoming waves, which in standard vacuum general relativity would be fully absorbed by the black hole.

In particular, recent works by Abedi, Dykaar and Afshordi (ADA) [18–20] have claimed to find evidence of near-horizon Planck-scale structure using data [21,

22] from the three Advanced LIGO events GW150914, LVT151012 and GW151226. In the simplified analysis of [18, 19], this near-horizon structure gives rise to so-called echoes [12, 13, 23, 24]. The data used by ADA is from the LIGO Open Science Center (LOSC) [21, 22], which contains a total of 4096 seconds of strain data from both Advanced LIGO detectors around each of the three events. Out of this data ADA used only 32 seconds centered around each event for their analyses. The authors claimed in [18] to find evidence for such echoes in data following the three events with combined significance of 2.9σ (p-value 3.7×10^{-3} ; with the one-sided significance convention used in [1–3, 25], this value corresponds to 2.7σ). If this claim were true, it would force a major re-evaluation of the standard picture of black holes in vacuum Einstein gravity.

Here we investigate concerns about the methods in [18] and ADA’s updated works [19, 20], and give a different significance estimate for the findings. Our initial caveats concerning [18] appeared as [26]. We do not examine the theoretical motivations for the existence of such near-horizon Planck-scale structure, nor the model templates for which ADA have chosen to search. Rather, we focus on the data analysis methods as reported and on the significance estimates assigned to them. We identify a number of shortcomings in the analysis and perform an improved analysis, which corrects for several of these problems. We estimate the significance of the findings in the gravitational wave data [21, 22] and conclude that there is as of yet no significant evidence for the existence of black hole echoes in these data.

II. ADA’S MODEL AND SEARCH PROCEDURE

The analysis of ADA [18–20] consists of three parts: a simple waveform model, a search procedure, and a significance estimation method. In this section we briefly review these.

* julian.westerweck@aei.mpg.de

† alex.nielsen@aei.mpg.de

‡ ofek.birnholtz@aei.mpg.de

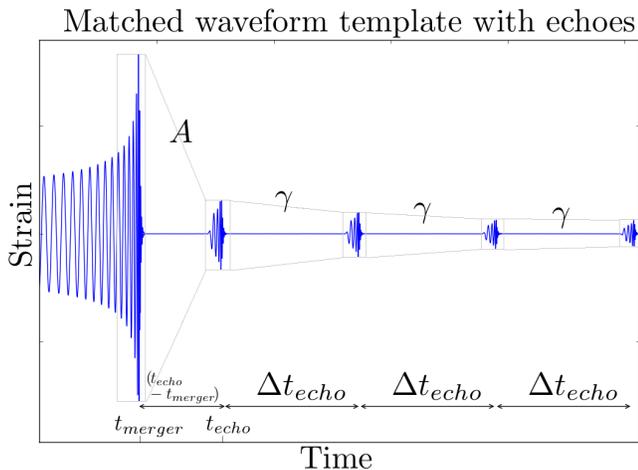


FIG. 1. A coalescence template extended to include echoes.

With a partially reflective surface outside the horizon, echo signals may be found as wave solutions in a cavity formed by the near-horizon membrane barrier and the angular momentum barrier (“photon sphere”) that exists further out [12, 13]. In the geometric wave picture, at each barrier the wave is partially reflected and partially transmitted. Each partially transmitted wave from the outer angular momentum barrier would be detected by distant observers as an echo. The delay time between subsequent echoes results from the travel time between the two barriers. This time may be different for the first echo due to non-linear effects during the merger, as may further parameters of the echo signal such as the dampening between successive echoes. For a description of the echoes as poles of the propagators see [15].

An example of such an echo template is shown in Fig.1 and several parameters define its features:

1. Δt_{echo} : The delay time between subsequent echoes, resulting from the travel time between the barriers. $\Delta t_{\text{echo,theory}}$ is the expected value found in [18, 19], based on the inferred final mass and spin parameters for each event [1–3]. In the search, the parameter Δt_{echo} is allowed to vary around the theoretical value $\Delta t_{\text{echo,theory}}$ to account for uncertainties.
2. t_{echo} : The time of the first echo. This is expected to be $t_{\text{merger}} + \Delta t_{\text{echo}}$, where t_{merger} is the time of the merger. It is allowed to deviate from this expectation in the search to account for non-linear effects close to the merger [18].
3. A : The amplitude of the first echo relative to the original signal amplitude.
4. γ : The relative amplitude between subsequent echoes.
5. t_0 : Only the last part of the original waveform is used to produce the echo waveform; this parameter

describes how far before t_{merger} the original waveform is tapered down to 0.

In addition, the phase is inverted between subsequent echoes. Likewise, the phase-difference between the original signal and the first echo is fixed to π , $\Delta\phi = \pi$. We use an abbreviated notation for the combination of parameters $x := (t_{\text{echo}} - t_{\text{merger}})/\Delta t_{\text{echo}}$, with an expected value for the first echo of $x = 1$.

The **ADA-search** procedure used in [18, 19] consists of the following steps:

1. Produce a pure echo template for given echo-parameters. The original event waveform is removed.
2. Produce a bank of these templates, with an evenly spaced grid in the parameters listed above.
3. Perform matched filtering with the echo templates. The original event is removed from the data prior to this.
4. Maximise SNR^2 over all parameters for each value of x .

The maximisation uses either each single event or combinations of events. The combination assumes some parameters to be different between events, namely A and Δt_{echo} . The parameters x , $t_0/\Delta t_{\text{echo,theory}}$ and γ are kept identical for each event.

The **ADA-estimation** uses the following method to estimate the significance of their findings [18]:

1. Find the maximum SNR^2 value in the range $x \in (0.99, 1.01)$ after the event.
2. Calculate and maximise SNR^2 over the time range $9 \leq \frac{t_{\text{echo}} - t_{\text{merger}}}{\Delta t_{\text{echo,theory}}} \leq 38$. The maximisation is slightly adapted for this step.
3. Divide this last range into 1450 segments, each of duration 2% of $\frac{t_{\text{echo}} - t_{\text{merger}}}{\Delta t_{\text{echo,theory}}}$.
4. A p-value is found as the number of segments with higher peak SNR^2 , divided by the total number of segments.

III. GENERAL REMARKS

A first immediate problem arises regarding how strong the relative signal should be for the three events. The two binary black hole events GW150914 and GW151226 were detected by the Advanced LIGO detectors with significance levels $> 5.3\sigma$ and signal-to-noise ratios of 23.7 and 13.0 respectively [3]. The other event, LVT151012, had a reported significance of only 1.7σ and a signal-to-noise ratio of 9.7 combined between the two Advanced LIGO detectors. However, in Table II of [19] we see that the

signal-to-noise ratio of the claimed echo signal is actually largest for LVT151012.

The higher SNR cannot be due to the different projected number of echoes between the events. The different Δt_{echo} leads to differing numbers of echoes in a given duration: the 32 seconds of data used would contain (~ 180) for LVT151012 and (~ 110) for GW150914. But although the number of echoes is larger for LVT151012, late echoes are strongly dampened. They decrease by a factor of 10 over ~ 22 echoes for the claimed $\gamma \sim 0.9$. Thus in order for the echoes of LVT151012 to have a higher SNR than the echoes of GW150914, their amplitude must be very high. In fact to account for the reported SNRs, the initial amplitude for the first echo of LVT151012 would have to be about 10% higher than that of GW150914 [26], while the original event's peak is about 2-3 times lower for LVT151012 in comparison to GW150914's. This would require their parameter A to be about 2-3 times larger for LVT151012 than for GW150914. This seems to be confirmed by the best fit search results in Table II of their updated work [19], where they give $A[\text{GW150914}] = 0.091$, $A[\text{LVT151012}] = 0.34$.

The inferred amplitude parameters suggest that a lot of gravitational wave energy was emitted in the echoes: a very rough calculation implies that the amount of energy emitted in the echoes was approximately 0.1 solar masses (for GW150914) and 0.2 solar masses (for LVT151012). This should be compared to the total estimated energy emitted by the original signal of 3 solar masses (for GW150914) and 1.5 solar masses (for LVT151012).

IV. VALIDATION OF THE MATCHED-FILTER ANALYSIS

We wrote a separate implementation of the **ADA-search** procedure, that we refer to as **ADA_{AEI}-search**. No changes were made to the algorithm as described before, while the implementation itself is independent. The SNR^2 -results obtained with our implementation are similar to those shown in [18].

As a first check, we verify that the **ADA_{AEI}-search** procedure can distinguish between pure noise and simulated echo-signals. For this, a known signal is injected into simulated noise. We simulate Gaussian noise with a Power Spectral Density (PSD) similar to that found for the detector data around each event (calculated from the LIGO data). The **ADA_{AEI}-search** is then applied to simulated data of both pure noise and also the same noise with added injections of different amplitudes. In this test, we only use echo waveforms with parameters similar to the best-fit results of [18, 19]. Fig. 2 shows the dependence of the SNR^2 peak on the injection amplitude A_{inj} . The effectiveness of the method in finding a signal depends on A_{inj} . This test was performed for different realisations of the simulated noise. The minimum A_{inj} required to find a peak rising above the noise

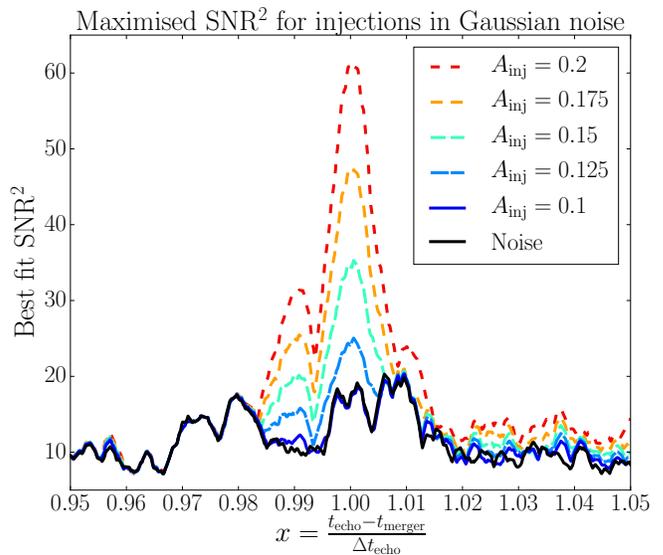


FIG. 2. The matched filtering technique is able to recover signals with a variety of amplitudes. As shown here, the SNR depends on the amplitude of the signal. The amplitude found by ADA ($A = 0.1$) is close to the level that is found in pure Gaussian noise. An amplitude twice as large as this would be clearly identifiable in the data.

background also depends on the noise instantiation. We find that $A_{\text{inj}} \sim 0.1$ can yield a visible peak. This is the best-fit value of A reported for GW150914 in [18]. In one out of five trials, a higher amplitude was necessary to distinguish the signal from noise, as shown in Fig. 2, where noise and smallest injection have almost identical SNR^2 results.

V. PRIOR RANGES AND TEMPLATE SPACING

Values for each echo parameter are determined from within a prior range. Each template in the bank is produced for a specific value for each parameter. The matched filtering method then finds a higher SNR for data similar to the template, and thus each template can recover signals with a range of similar parameters. The three parameters γ , t_0 and t_{echo} are determined by maximisation, with γ and t_0 kept fixed between the different events. In this, the parameters recovered are defined as the values corresponding to the template in the bank which yields the highest SNR. Thus the maximisation is performed over all templates in the bank and thus all values in the parameter grid used to create the bank. The boundaries of the parameter grid are determined by a prior range, where the ranges chosen by ADA are displayed in Table I of [18].

The values resulting from this maximisation for γ and t_0 are found to lie very close to the boundary of this prior range, 0.9 and -0.1 respectively [26]. This suggests that

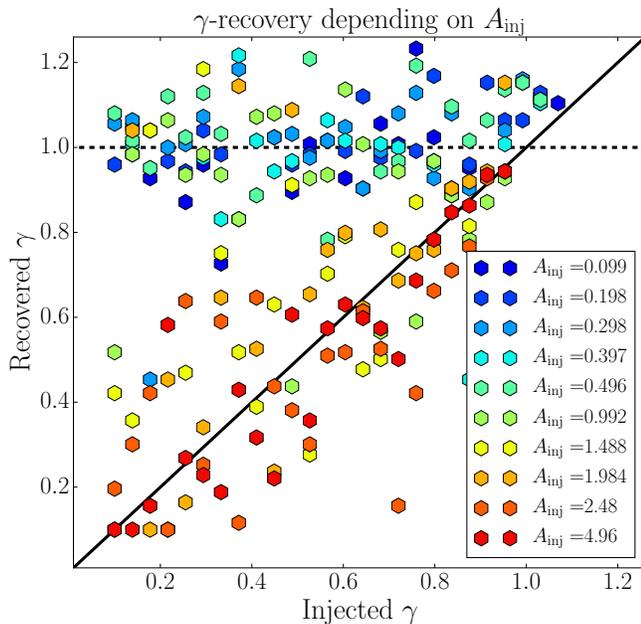


FIG. 3. Injected and recovered values for γ , correct recovery would be seen as a diagonal. The preference for $\gamma = 1$ at lower injection amplitudes is clearly seen.

there may be support for values of these parameters that lie outside of this range. If these values reflect the priors rather than the data, then they cannot be reliably considered as evidence for a detection claim. Furthermore, a value greater than unity for γ means that each successive echo has an amplitude greater than the previous echo. Such a result would require the echo signal to be extracting energy from the black hole spacetime.

We tested whether the preference for these parameter values is an artefact of the method, again using known signals injected into simulated noise. The noise is Gaussian with a PSD estimated from the 4096 seconds of LIGO data around GW150914. The injected signals are pure echo signals based on the LIGO GW150914-template for various echo parameters. The range of γ is widened to $\gamma \in (0.1, 2.0)$ both in the prior of the search and the injections. The relative amplitude of the injections $A = A_{inj}$ ranges from ADA's recovered value 0.01 to about 50 times this amplitude. We then compare the best-fit value of γ from the search with the value of the injection. This comparison is shown in Fig. 3.

The **ADA-search** method is biased towards finding γ to be close to 1. Ideally, the recovered parameter value would be closest in the grid to the value of the injection. In Fig. 3, this would mean lying as close as possible to the plotted diagonal. In this figure, we see that the recovered values are closer to the injected ones for higher injection amplitudes. Thus for very high echo amplitudes, the recovery method could in principle be effective. For lower injection amplitudes, we see a preference for recovered values of γ close to 1, independently of γ for the injection. Thus finding $\gamma \sim 1$ as the best-fit value in the

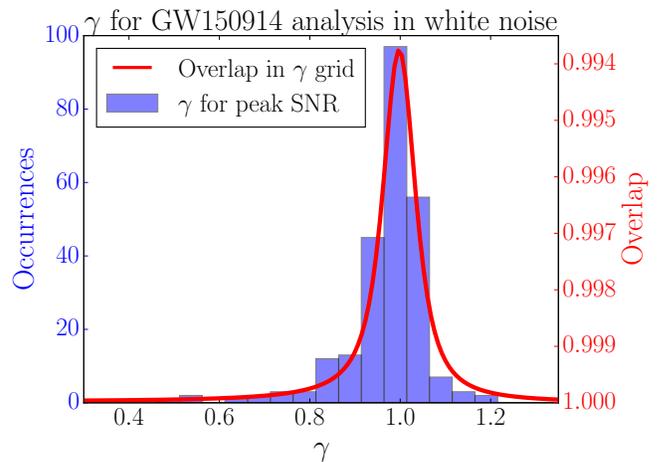


FIG. 4. Demonstration of the preference for $\gamma = 1$ in white Gaussian noise and the correspondence of the recovered distribution with the spacing in the template grid.

search does not necessarily mean that this is indeed the correct value for an existing signal. The method is biased such that we should expect to find these values for γ in almost all cases. Especially, this is true for echo amplitudes as found by ADA, and even significantly higher signal amplitudes. We interpret the recovery of $\gamma \sim 1$ as a generic property of the method. Thus finding such a value should not be considered evidence for a present signal.

The bias is introduced through the spacing between templates in the bank, as can be found through an analysis on white Gaussian noise and calculating the overlap between the templates. An analysis in white Gaussian noise using the same parameter range as ADA also shows a strong preference for $\gamma = 0.9$. Further extending the range to $\gamma \in (0.1, 2.0)$ displays preference for $\gamma = 1$ in white noise. The distribution of recovered γ values is shown in Fig. 4. The reason for this is revealed by calculating the overlap between neighbouring templates in the parameter grid for different γ , while keeping the other parameters fixed. As can be seen in Fig. 4, the match between neighbouring templates follows the same distribution as the preference for $\gamma = 1$. Templates with γ close to 1 lie further apart in the noise weighted match metric than other templates. Each template near γ close to 1 covers a larger region of the signal space than other templates and thus, more noise realisations are best matched by the (morphology-wise) more scarcely placed templates close to $\gamma = 1$.

Extending the template bank to a wider range in γ and t_0 and performing the same analysis as before leads to a modified SNR structure in x , where additional and higher peaks appear further away from the predicted value for the echoes in GW150914. As we will see below, applying a wider parameter range also for the background estimation results in a further factor ~ 3.5 for the p-value for the combination (1,2,3) (using the 32-second dataset es-

timation). The modified p-values for the wider priors of different combinations are found in Table I.

VI. EXTENDING THE BACKGROUND ESTIMATION

To calculate a significance for the match found in the templated search, an assessment of the background should be performed. Since an analytical noise model is not known, real data away from the possible signal is used. This data is assumed to be pure noise of a form similar to that at the event. Then one observes how often a similar or larger SNR value is obtained in this data. ADA chose to do this in a short period of time of approximately 16 seconds of data after each event. To obtain sufficient background statistics this period of time was used intensively: they consider 16 second stretches of data as independent when shifted by only 0.1 seconds.

This background estimation is problematic [26] for two reasons: potential contamination of background samples by an existing echo signal, and lack of independence between the background samples. The estimation uses a range of t_{echo} values that is only $\mathcal{O}(10)$ echo periods away from the merger. If there is indeed an echo signal in the data then this region will not be entirely free of the signal being searched for. At the beginning of the region the amplitude of the echoes would only have dropped by a factor $0.9^9 \sim 0.4$. One therefore expects a contaminated background estimation. Even in the absence of echoes, a random feature mistaken for echoes in one segment may well extend to neighbouring segments, and they should not be treated as independent (see discussion of template auto-correlations below for the problem of insufficient independence of samples). Each of the data sets released at [21] consist of 4096 seconds of data. Both GW150914 and LVT151012 are located 2048 seconds into this data, equivalent to $\mathcal{O}(10^3)$ echo periods, thus for large stretches of the data such contamination would be negligibly small.

We have performed a different background estimation as an independent test, which uses the full period of 4096 seconds of LOSC data available around each event. A comparison of the different choices of data used for background estimation is shown in Fig. (5). The obtained p-value and background estimate are only meaningful if the data in the background is comparable to that at the time of the event. A plot of the noise variations over the full 4096 seconds of data released for GW150914 is shown in Figs. 6 and 7. The variations are seen to be small and we conclude that our background estimate is indeed characteristic of the noise at the time of the events.

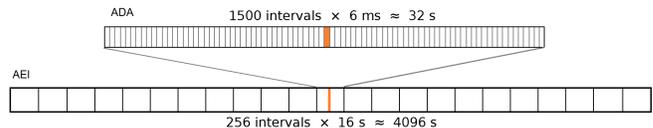


FIG. 5. Comparison of the data segments used to estimate the background between ADA [18, 19] and this work (AEI). Compared to ADA, we extend the amount of data used for background estimation to the full 4096 seconds for each event available from the LOSC [27].

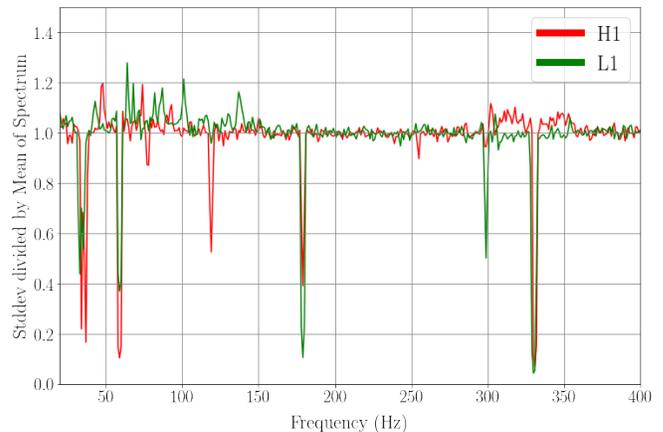


FIG. 7. Rayleigh plot of noise variation, showing the ratio of the standard deviation to the mean for frequency bins of the PSD estimated using 16 second segments of the 4096 second data stretch for GW150914.

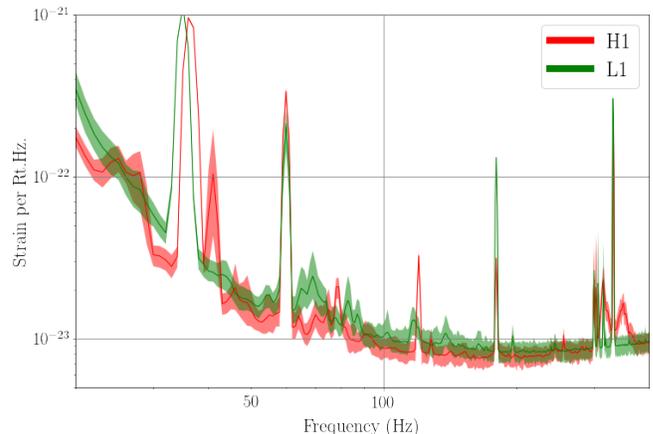


FIG. 6. Variation of the noise spectrum during the 4096s around GW150914 calculated using 16 second segments, showing the 1% and 99% percentile variations. In this sense, the data is sufficiently stationary for background estimation to be reliable during the full 4096 second data stretch.

In our case, the 4096 seconds of data for each event are divided into 128 independent, 32-second long segments. For each, the echoes analysis is performed as it was on the 32-second segment containing the event. The resulting peak SNR in $x \in (0.99, 1.01)$ is found for each

segment. Simply counting the number of segments containing a higher peak SNR in this interval yields an estimate for the p-value. For the combined first three events, GW150914, LVT151012 and GW151226, our resulting p-value of 0.032 is about a factor 3 larger than the value of 0.011 found in [18], where less data and overlapping intervals were used.

An estimate of the p-value significance in this way is susceptible to the small number statistics. The original LIGO templates, before introducing echoes, contain only about 16 seconds long waveforms followed by a flat zero template. Echoes were introduced only into this 16-second flat region when producing the echo templates. After removing the original event, we are left with a template of about 16 seconds of no signal, followed by the produced echoes. So we can double the number of background samples without losing independence between samples, by dividing the available data into 256 independent segments of 16 seconds length.

The exact number of data segments varies slightly for each of the events. This is due to the positions of the original signals, and a loud transient noise (inconsistent with the echo morphology) in one segment of the LVT151012 data, which was discarded. The transient noise was found by noticing a very high SNR outlier for one data segment. The short noise can be seen in the whitened time-domain data and the search procedure always aligns one of the echoes with the noise feature, yielding the high SNR. The effect of not discarding the noise dataset always is an increase of p-value due to the very high SNR, but the effect on the resulting p-values is minimal ($\sim 1/256$). Only the estimation with 16 seconds long segments is influenced by this noise. The total number of estimates when combining events thus is 125 to 126 for 32-second segments and 250 to 251 for 16-second segments.

The results of this alternative approach for the significance estimation, both using 32 and 16 seconds long segments, are shown in Table I. Different combinations of the available events are considered, denoted chronologically as (GW150914, LVT151012, GW151226) \rightarrow (1, 2, 3). For the combination (1, 2, 3), a p-value of 0.011 was found in [18]. With four points out of 125 trials giving maximised combined SNR values larger than immediately after each GW event, for our method, the p-value for the particular SNR value found is $4/125 \sim 3\%$. Using twice the number of samples of 16 seconds length each, we find $5/250 \sim 2\%$ for the p-value. The combined background is shown in Fig. 8 which shows the peak value of SNR^2 found in each segment for both the real detector data and Gaussian noise. Here we note that there is no obvious trend in the peak SNR over time. Thus also by this measure, there is no indication of the noise being unstable and preventing its use for a background estimation.

A second concern about the background estimation used in [18–20] arises from the very small shift in time between samples that are considered independent. In this, the quasi-periodic nature of the echo signal has to be considered, leading to potentially long templates with

Event	[19]	original 16s (32s)	widened priors 16s (32s)
GW150914	0.11	0.199 (0.238)	0.705 (0.365)
LVT151012	-	0.056 (0.063)	0.124
GW151226	-	0.414 (0.476)	0.837
(1,3)	-	0.159	0.801
(1,2,3)	0.011	0.020 (0.032)	0.18 (0.144)

TABLE I. p-values obtained by using 4096 seconds of LIGO data divided into segments of 16 or 32 seconds length. Results are given for the priors in t_0 and γ chosen in the original analysis and for widened priors, for the 3 events individually and for the combination of events. LVT151012, of the weakest-SNR, is of exceptionally high significance.

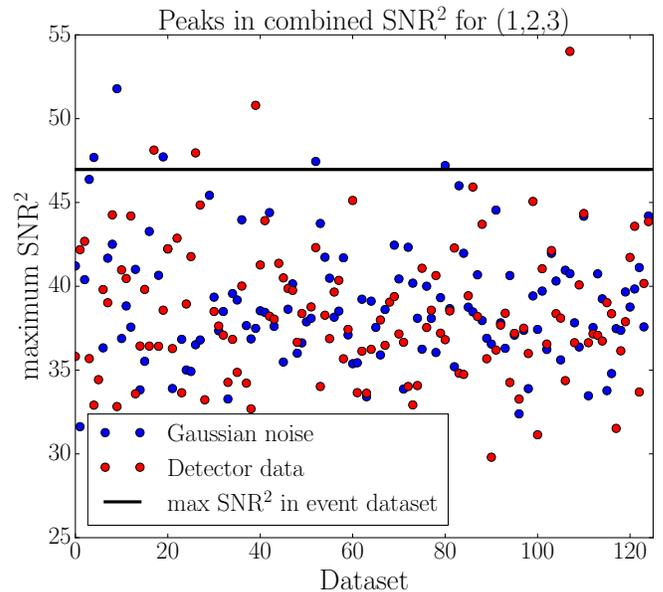


FIG. 8. The distribution of SNR peaks in the data is similar to Gaussian noise. The p-value is calculated from the number of points lying above the black line which indicates the combined SNR value found immediately after each GW event. There is no obvious trend in the peak SNR over time.

equally quasi-periodic autocorrelation in time. The autocorrelation of the echo templates is shown in Fig. 9. This affects the significance estimation as performed in [18, 19]. This method is problematic because the template is significantly longer than the shift in time between background sample intervals. If the autocorrelation between templates used in different background samples does not vanish, the results from these samples cannot be considered truly independent: a feature of the data at one point in time then influences the SNR found for several background samples. The total number of samples to be considered independent is thus reduced by this method.

To estimate the influence of this effect on the number of independent samples, we calculate the overlap between the template and a time-shifted version thereof, as shown in Fig. 9. The echo signal model leads to a series of peaks

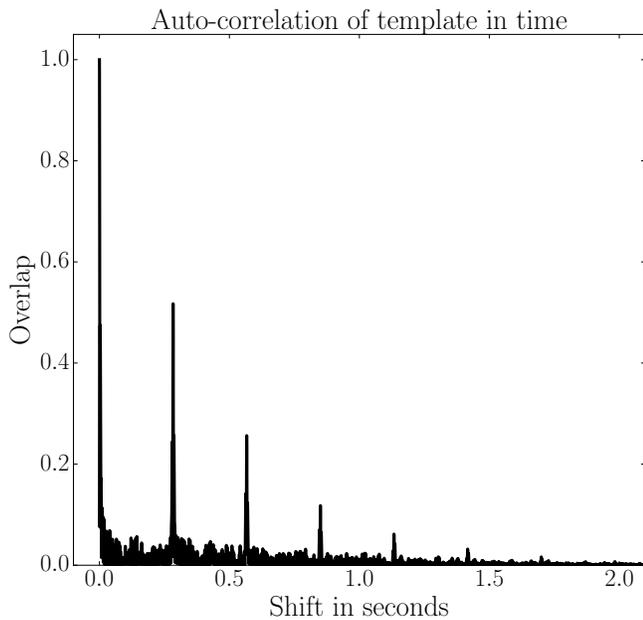


FIG. 9. Autocorrelation of the pure echo template depending on timeshift. The peaks demonstrate that there is significant correlation between background samples for small shifts.

in the overlap of original and time-shifted wavefunction, depending on the parameters γ , Δt_{echo} and t_0 . We choose a value of overlap to be considered sufficiently independent, e.g. 1%. To achieve this value, there are two ways of placement for the template with respect to the original position in time. For $\gamma < 1$, i.e. a dampened echo signal, applying a shift in time by a sufficiently large multiple of Δt_{echo} leads to reduction of the correlation. Using the GW150914 template and $\gamma = 0.5$ shows that at least $7\Delta t_{\text{echo}}$ are necessary. In Fig. 9, this corresponds to the very small peak close to 2.0 seconds of time-shift.

Alternatively, the templates can be interlaced, such that the echoes of one template are placed within the time between echoes of the other. This corresponds to the small overlap values between the peaks in Fig. 9. Here we again use the GW150914 template and the most favourable values in the prior range, i.e. the shortest echoes ($t_0 = -0.1\Delta t_{\text{echo, theory}}$) and the longest delay ($\Delta t_{\text{echo}} = 0.30166\text{s}$). Then about 6 echoes can be placed between the peaks of the original template. For these parameters we now consider shifts in time of the echo template up to $29\Delta t_{\text{echo}}$, which is the maximum shift used in ADA’s significance estimation. We find about 4 independent samples through sufficient timeshift and a factor 7 through interlacing, giving ~ 28 independent samples.

As the maximisation is performed over a range of parameters, exactly determining the total number of independent samples would need further consideration. The parameters chosen for this estimate, however, are favourable, as smaller dampening or smaller time delay further lower the total number. For the maximisation

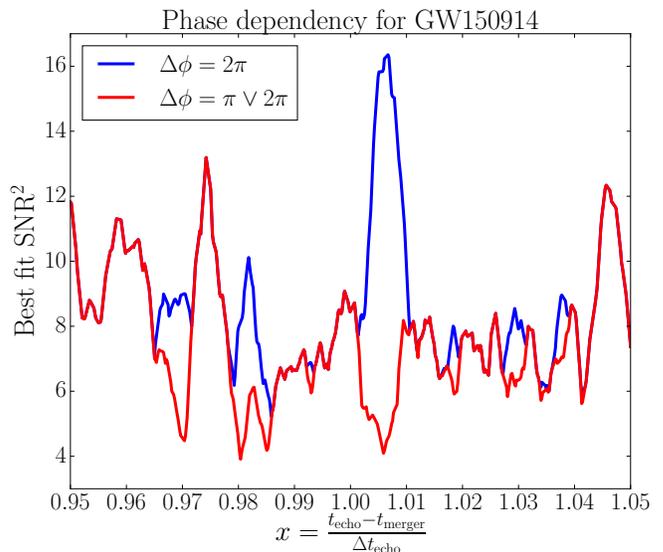


FIG. 10. The maximised SNR for GW150914 is shown with different phase change between event and first echo $\Delta\phi$. The original analysis allows for $\Delta\phi = \pi \vee 2\pi$ for each x . It finds a prominent peak near $x = 1$. However the model requires $\Delta\phi = \pi$. When we enforce this value in the analysis, the peak vanishes.

combining the different events, Δt_{echo} may vary independently between events, obstructing a clear estimate on the number of samples; the same considerations, however, still apply. These considerations suggest a small total number of independent background samples on the order of a few tens of samples.

The method we employed to estimate the background precludes this effect by only applying the matched filtering procedure to separate sets of data. The template thus is always placed in only one of the background samples and the resulting SNR cannot be influenced by data features in the remaining samples.

VII. REMARK ON ECHO TEMPLATES

The **ADA-search** procedure does not distinguish between inversion and non-inversion of the first echo’s phase. The waveform templates used here are based on the simple model described in [18, 19]. Within this model, the phase-change of the gravitational wave between the original signal and the first echo is described as a simple phase-flip. However, as only the square of the SNR from the matched filter analysis is considered for the maximisation, the result is insensitive to this phase inversion. Repeating the analysis for GW150914 and enforcing the phase inversion as required by this model, we find that the prominent peak in SNR at $x = 1$ vanishes. This is shown in Fig. 10. More sophisticated models would be needed to determine whether this phase flip is truly required or not. Still, it is worth noting that the peak

that forms the basis of evidence for echoes in [18, 19], does not contain this phase inversion as required by the simple model, but actually the opposite phase.

VIII. CONCLUSIONS

A full analysis of the data at a level necessary to confidently detect echo signals is outside the scope of this work. However, we have analysed the data using a simple templated search similar to [18]. Using an extended background in the full 4096 seconds of data released publicly by the LIGO team we find a p-value for the null noise only hypothesis of 0.02, higher than that reported using the restricted background of 0.011 in [18]. We have demonstrated the importance of LVT151012, the weakest LIGO candidate event, in obtaining this p-value. Excluding LVT151012 from the combined analysis yields an even larger p-value of 0.199, fully consistent with noise. We have also identified a number of weaknesses in the analysis method of [18] including the role of the prior boundaries and the density of templates. In particular we have examined the role of the γ parameter and found that the clustering of γ values near $\gamma \sim 1$ is entirely consistent with noise. The signal manifold is such that in pure Gaussian noise, one would expect many more triggers with values of $\gamma \sim 1$. This perhaps would not be

expected for true echo signals, although a full model of echo signals would be needed to make a quantitative prediction for this.

In conclusion, we find that the tentative evidence as presented in [18–20] is lacking in several key aspects. Hence their current methodology cannot provide observational evidence for or against the existence of near-horizon Planck-scale structure in black holes.

We hope some of the concerns explored here may be useful to vet other searches for echoes, such as [28], and help in the development of methods which would place black hole near-horizon physics within the realm of gravitational wave observations.

ACKNOWLEDGMENTS

The authors thank Andrew Lundgren and Vitor Cardoso and the authors of [18–20] for useful discussions. Some of the discussions particularly enjoyed the hospitality of meetings at Nikhef and at the Perimeter Institute. This research has made use of data, software and/or web tools obtained from the LIGO Open Science Center (<https://losc.ligo.org>), a service of LIGO Laboratory and the LIGO Scientific Collaboration. LIGO is funded by the U.S. National Science Foundation.

-
- [1] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.* **116** (2016) no.6, 061102 doi:10.1103/PhysRevLett.116.061102 [arXiv:1602.03837 [gr-qc]].
- [2] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence,” *Phys. Rev. Lett.* **116** (2016) no.24, 241103 doi:10.1103/PhysRevLett.116.241103 [arXiv:1606.04855 [gr-qc]].
- [3] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “Binary Black Hole Mergers in the first Advanced LIGO Observing Run,” *Phys. Rev. X* **6** (2016) no.4, 041015 doi:10.1103/PhysRevX.6.041015 [arXiv:1606.04856 [gr-qc]].
- [4] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2,” *Phys. Rev. Lett.* **118**, no. 22, 221101 (2017) doi:10.1103/PhysRevLett.118.221101 [arXiv:1706.01812 [gr-qc]].
- [5] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence,” *Astrophys. J.* **851**, no. 2, L35 (2017) doi:10.3847/2041-8213/aa9f0c [arXiv:1711.05578 [astro-ph.HE]].
- [6] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence,” *Phys. Rev. Lett.* **119**, no. 14, 141101 (2017) doi:10.1103/PhysRevLett.119.141101 [arXiv:1709.09660 [gr-qc]].
- [7] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “Properties of the Binary Black Hole Merger GW150914,” *Phys. Rev. Lett.* **116** (2016) no.24, 241102 doi:10.1103/PhysRevLett.116.241102 [arXiv:1602.03840 [gr-qc]].
- [8] B. P. Abbott *et al.* [LIGO Scientific and Virgo Collaborations], “Tests of general relativity with GW150914,” *Phys. Rev. Lett.* **116** (2016) no.22, 221101 doi:10.1103/PhysRevLett.116.221101 [arXiv:1602.03841 [gr-qc]].
- [9] G. M. Harry [LIGO Scientific Collaboration], “Advanced LIGO: The next generation of gravitational wave detectors,” *Class. Quant. Grav.* **27**, 084006 (2010). doi:10.1088/0264-9381/27/8/084006
- [10] J. Aasi *et al.* [LIGO Scientific Collaboration], “Advanced LIGO,” *Class. Quant. Grav.* **32**, 074001 (2015) doi:10.1088/0264-9381/32/7/074001 [arXiv:1411.4547 [gr-qc]].
- [11] F. Acernese *et al.* [VIRGO Collaboration], “Advanced Virgo: a second-generation interferometric gravitational wave detector,” *Class. Quant. Grav.* **32**, no. 2, 024001 (2015) doi:10.1088/0264-9381/32/2/024001 [arXiv:1408.3978 [gr-qc]].
- [12] V. Cardoso, E. Franzin and P. Pani, “Is the gravitational-wave ringdown a probe of the event horizon?,” *Phys. Rev. Lett.* **116** (2016) no.17, 171101 Erratum: [Phys. Rev. Lett. **117** (2016) no.8, 089902]

- doi:10.1103/PhysRevLett.117.089902, 10.1103/PhysRevLett.116.171101 [arXiv:1602.07309 [gr-qc]].
- [13] V. Cardoso, S. Hopper, C. F. B. Macedo, C. Palenzuela and P. Pani, “Gravitational-wave signatures of exotic compact objects and of quantum corrections at the horizon scale,” *Phys. Rev. D* **94** (2016) no.8, 084031 doi:10.1103/PhysRevD.94.084031 [arXiv:1608.08637 [gr-qc]].
- [14] V. Cardoso and P. Pani, “The observational evidence for horizons: from echoes to precision gravitational-wave physics,” [arXiv:1707.03021 [gr-qc]].
- [15] Z. Mark, A. Zimmerman, S. M. Du and Y. Chen, “A recipe for echoes from exotic compact objects,” [arXiv:1706.06155 [gr-qc]].
- [16] S. H. Völkel and K. D. Kokkotas, “Ultra Compact Stars: Reconstructing the Perturbation Potential,” [arXiv:1704.07517 [gr-qc]].
- [17] V. Cardoso, E. Franzin, A. Maselli, P. Pani and G. Raposo, “Testing strong-field gravity with tidal Love numbers,” *Phys. Rev. D* **95**, no. 8, 084014 (2017) Addendum: [*Phys. Rev. D* **95**, no. 8, 089901 (2017)] doi:10.1103/PhysRevD.95.089901, 10.1103/PhysRevD.95.084014 [arXiv:1701.01116 [gr-qc]].
- [18] J. Abedi, H. Dykaar and N. Afshordi, “Echoes from the Abyss: Evidence for Planck-scale structure at black hole horizons,” [arXiv:1612.00266v1 [gr-qc]], 1 Dec 2016.
- [19] J. Abedi, H. Dykaar and N. Afshordi, “Echoes from the Abyss: Tentative evidence for Planck-scale structure at black hole horizons,” *Phys. Rev. D* **96**, no. 8, 082004 (2017) doi:10.1103/PhysRevD.96.082004 [arXiv:1612.00266v2 [gr-qc]], 3 Oct 2017.
- [20] J. Abedi, H. Dykaar and N. Afshordi, “Echoes from the Abyss: The Holiday Edition!,” [arXiv:1701.03485 [gr-qc]].
- [21] LIGO Scientific Collaboration, “LIGO Open Science Center - Data Releases for Observed Transients”, 2017, doi:10.7935/K5MW2F23, 10.7935/K5CC0XMZ, 10.7935/K5H41PBP, 10.7935/K53X84K2 <http://losc.ligo.org/events>
- [22] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein and B. Stephens, “The LIGO Open Science Center,” *J. Phys. Conf. Ser.* **610**, no. 1, 012021 (2015) doi:10.1088/1742-6596/610/1/012021 [arXiv:1410.4839 [gr-qc]].
- [23] T. Nakamura, H. Nakano and T. Tanaka, “Detecting quasinormal modes of binary black hole mergers with second-generation gravitational-wave detectors,” *Phys. Rev. D* **93** (2016) no. 4, 044048 doi:10.1103/PhysRevD.93.044048 [arXiv:1601.00356 [astro-ph.HE]].
- [24] B. Holdom and J. Ren, “Not quite a black hole,” [arXiv:1612.04889 [gr-qc]].
- [25] V. Connaughton *et al.*, “Fermi GBM Observations of LIGO Gravitational Wave event GW150914,” *Astrophys. J.* **826** (2016) no.1, L6 doi:10.3847/2041-8205/826/1/L6 [arXiv:1602.03920 [astro-ph.HE]], Appendix B, ‘Significance of two-parameter coincidence’.
- [26] G. Ashton, O. Birnholtz, M. Cabero, C. Capano, T. Dent, B. Krishnan, G. D. Meadors, A. B. Nielsen, A. Nitz and J. Westerweck, “Comments on: ”Echoes from the abyss: Evidence for Planck-scale structure at black hole horizons”,” [arXiv:1612.05625 [gr-qc]].
- [27] LIGO Scientific Collaboration, “LOSC Event tutorial v1.4”, 2016, http://losc.ligo.org/s/events/GW150914/LOSC_Event_tutorial_GW150914.html
- [28] R. S. Conklin, B. Holdom and J. Ren, “Gravitational wave echoes through new windows”, [arXiv:1712.06517 [gr-qc]].