



# CHORUS

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

## Detecting Gravitational-Wave Transients at $5\sigma$ : A Hierarchical Approach

Eric Thrane and Michael Coughlin

Phys. Rev. Lett. **115**, 181102 — Published 27 October 2015

DOI: [10.1103/PhysRevLett.115.181102](https://doi.org/10.1103/PhysRevLett.115.181102)

# Detecting gravitational-wave transients at five sigma: a hierarchical approach

Eric Thrane<sup>1, a</sup> and Michael Coughlin<sup>2</sup>

<sup>1</sup>*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

<sup>2</sup>*Department of Physics, Harvard University, Cambridge, MA 02138, USA*

As second-generation gravitational-wave detectors prepare to analyze data at unprecedented sensitivity, there is great interest in searches for unmodeled transients, commonly called bursts. Significant effort has yielded a variety of techniques to identify and characterize such transient signals, and many of these methods have been applied to produce astrophysical results using data from first-generation detectors. However, the computational cost of background estimation remains a challenging problem; it is difficult to claim a  $5\sigma$  detection with reasonable computational resources without paying for efficiency with reduced sensitivity. We demonstrate a hierarchical approach to gravitational-wave transient detection, focusing on long-lived signals, which can be used to detect transients with significance in excess of  $5\sigma$  using modest computational resources. In particular, we show how previously developed seedless clustering techniques can be applied to large datasets to identify high-significance candidates without having to trade sensitivity for speed.

*Introduction.* With second-generation gravitational-wave (GW) detectors coming online later this year, the first direct detection of GWs may be near. In order to establish the significance of a detection, it is common to report a false alarm probability (FAP), which quantifies the probability that a noise fluctuation could produce an event at least as loud as the observed candidate (as measured by some detection statistic). In some subfields, e.g., particle physics, “ $5\sigma$  significance,” corresponding to  $\text{FAP} \approx 2.9 \times 10^{-7}$ , is used as a detection threshold.

In order to estimate the FAP of a GW candidate, it is common to perform time-shifts in which the GW strain time series from one detector is shifted with respect to the series from a second detector by an amount greater than both the travel time between the detectors and the coherence time of the signals being targeted. Time-shifting preserves non-Gaussian and non-stationary features that characterize the zero-lag (no time-shift) noise, while simultaneously eliminating true GW signals. By performing  $N$  time-shifts, it is possible to generate a distribution of the detection statistic, which can be used to estimate FAP to a level of  $\geq 1/N$ . The  $5\sigma$  threshold corresponds to  $N \approx 3.5 \times 10^6$ . In many cases it is computationally impractical to carry out this many time-shifts, though, it has been accomplished in the “detection” of a LIGO blind injection with a matched filter search [1]. Despite the pervasive use of time-shifts, there are limitations [2].

For many transient GW searches, the significant computing costs incurred by background estimation arise from a desire to use a coherent detection statistic. Coherent algorithms utilize the complex-valued cross-power obtained by cross-correlating strain data from  $\geq 2$  detectors instead of or in addition to the incoherent auto-power observed in each detector separately; see, e.g., [3–6]. The extra phase information helps differentiate sig-

nal from background, improving the sensitivity of the search. However, the cost of background estimation for coherent searches is relatively large compared to a comparable incoherent search because the detection statistic must be recalculated for each time-shift, after the fresh application of a clustering algorithm. Some algorithms use single-detector auto-power exclusively, which allows for much more rapid background estimation [7, 8].

In this Letter, we describe a hierarchical approach to background estimation in the context of a search for long-lived, unmodeled GW transients using seedless clustering [9, 10]. First, we identify “events” using a computationally intensive, but incoherent, single-detector statistic. Second, we calculate a computationally fast, coherent detection statistic for each event identified with the single-detector statistic. The second, coherent detection statistic is used to evaluate significance. By splitting the calculation into an incoherent stage and a coherent stage, the computationally intensive calculations are carried out just once, allowing rapid background estimation without sacrificing the sensitivity gained by the use of coherence.

We demonstrate this technique by estimating the background—past the  $5\sigma$  level—for one week of simulated Monte Carlo data and one week of Initial LIGO noise, recolored to resemble data from Advanced LIGO [11]. We calculate the sensitivity of this mock search for several toy model waveforms and find that it is not adversely affected by the incoherent stage. The remainder of this Letter is organized as follows. We review the basics of transient identification with seedless clustering and describe the details of the new hierarchical detection statistic, we describe a mock data analysis carried out on one week of Monte Carlo data and one week of recolored Initial LIGO noise, and we present results demonstrating the ability to estimate background at the  $5\sigma$  level.

*Method.* In previous work, we have described seedless clustering [9, 10, 12–14], a technique in which GW transients are identified by looking for clusters of excess

---

<sup>a</sup>Electronic address: [eric.thrane@monash.edu](mailto:eric.thrane@monash.edu)

coherence integrated along many different parametrized curves through frequency-time space. In particular, cubic Bézier curves [15] provide a useful family of curves suitable for the detection of bursting sources lasting tens of seconds to weeks, which slowly evolve in frequency, but which are approximately narrowband on short time scales [9, 10, 12]. Such long-lived signals [5], created, e.g., by rotational instabilities in nascent neutron stars [16–19] or by the fragmentation of an accretion disk [20–22], are potentially detectable by second-generation detectors [9, 10].

Previous applications of seedless clustering have been employed to look for clusters of excess coherence using a coherent statistic [10, 12]:

$$\mathbf{p}(t; f) = \frac{2\sqrt{2}}{\mathcal{N}} \tilde{s}_I^*(t; f) \tilde{s}_J(t; f) / \sqrt{P'_I(t; f) P'_J(t; f)}. \quad (1)$$

Here,  $\tilde{s}_I(t; f)$  is the Fourier transform of the strain time series in detector  $I$  for a data segment centered on time  $t$  with frequencies  $f$ .  $P'_I(t; f)$  is the auto-power spectrum calculated using (typically nine) neighboring segments while  $\mathcal{N}$  is a Fourier normalization constant. By integrating  $\mathbf{p}(t; f)$  along suitable curves in frequency-time space (and applying a phase factor to “point” in different sky directions), one can define a signal-to-noise ratio for the cluster [25]

$$\text{SNR}_{\text{tot}} \equiv \frac{1}{N^{1/2}} \sum_{\{t; f\} \in \Gamma} \text{Re} [e^{2\pi i f \Delta\tau} \mathbf{p}(t; f)]. \quad (2)$$

Here,  $\Gamma$  is a template describing the set of spectrogram pixels in the parametrized (Bézier) curve,  $N$  is the number of pixels in the curve, and  $\Delta\tau$  is the direction-dependent time delay between two detectors.

By calculating  $\text{SNR}_{\text{tot}}$  for many different templates  $\Gamma$  and for many different time delays  $\Delta\tau$ , it is possible to find very weak signals buried in noise [10]. The loudest event in each  $\mathbf{p}(t; f)$  spectrogram is deemed to have a signal-to-noise ratio:

$$\text{SNR}_{\text{tot}}^{\text{max}} = \max_{\Gamma} [\text{SNR}_{\text{tot}}] \quad (3)$$

$\text{SNR}_{\text{tot}}^{\text{max}}$  is a coherent detection statistic. The observed value of  $\text{SNR}_{\text{tot}}^{\text{max}}$  is compared to a distribution of  $\text{SNR}_{\text{tot}}^{\text{max}}$  generated using time-shifts in order to assign a false-alarm probability.

To reap the benefits of seedless clustering, it is desirable to employ a large number of templates:  $\mathcal{O}(10^7)$  for a single  $\approx 300$  s spectrogram. The calculation can be sped up by parallelization, but it becomes prohibitive to repeat  $> 10^6$  times on time-shifted data, the number of realizations required to carry out  $5\sigma$  background estimation. In order to eliminate this computational bottleneck, we introduce a single-detector, incoherent statistic

$$\mathbf{l}_I(t; f) = \frac{2\sqrt{2}}{\mathcal{N}} |\tilde{s}_I(t; f)|^2 / P'_I(t; f). \quad (4)$$

Note that  $\mathbf{l}_I(t; f)$  is equivalent to  $\mathbf{p}(t; f)$  for the case where the two detector indices are equal  $I = J$ . It is the ratio of the auto-power in detector  $I$  at time  $t$  to the auto-power in detector  $I$  at the times neighboring  $t$ .

While the seedless clustering algorithms described above were developed in the context of a coherent search using  $\mathbf{p}(t; f)$ , it is straightforward to apply them to the new incoherent statistic  $\mathbf{l}_I(t; f)$  by simply defining a new single-detector signal-to-noise ratio:

$$\text{SNR}_{\text{tot}}^{(I)} \equiv \frac{1}{N^{1/2}} \sum_{\{t; f\} \in \Gamma} \mathbf{l}_I(t; f). \quad (5)$$

The phase factor from Eq. 2 vanishes because the time delay between two (now identical) detectors is zero. As before, we identify the most significant cluster in each spectrogram. It is necessary to do this separately for each detector:  $\text{SNR}_{\text{tot}}^{\text{max}(I)} = \max_{\Gamma} [\text{SNR}_{\text{tot}}^{(I)}]$ .

While  $\mathbf{l}_I(t; f)$  and  $\mathbf{p}_I(t; f)$  are similarly defined, it is important to note differences in their statistical behavior. If no signal is present, the distribution of  $\text{Re}[\mathbf{p}(t; f)]$  is peaked symmetrically about zero. However, the distribution of  $\mathbf{l}_I(t; f)$  is positive definite and asymmetric with a peak near one. Thus,  $\text{SNR}_{\text{tot}}^{\text{max}}$  [calculated from  $\mathbf{p}_I(t; f)$ ] and  $\text{SNR}_{\text{tot}}^{\text{max}(I)}$  [calculated from  $\mathbf{l}_I(t; f)$ ] have very different distributions. While  $\text{SNR}_{\text{tot}}^{\text{max}} = 10$  corresponds to a highly significant event,  $\text{SNR}_{\text{tot}}^{\text{max}(I)} = 29$  is typical for noise.

If we stopped here, and used  $\text{SNR}_{\text{tot}}^{\text{max}(I)}$  as a detection statistic, the method would rely on an incoherent statistic. However, we can do much better by using clusters identified with  $\text{SNR}_{\text{tot}}^{\text{max}(I)}$  to calculate a coherent statistic. The loudest cluster as ranked by the single-detector statistic is denoted  $\Gamma_I$ . The coherent signal-to-noise ratio for this cluster is:

$$\Lambda^{(I)} \equiv \frac{1}{N^{1/2}} \max_{\Delta\tau} \sum_{\{t; f\} \in \Gamma_I} \text{Re} [e^{2\pi i f \Delta\tau} \mathbf{p}(t; f)]. \quad (6)$$

The  $\max_{\Delta\tau}$  term indicates that the sum is carried out for 400 evenly spaced values of  $\tau$ , sufficient to match the diffraction-limited resolution.  $\Lambda^{(I)}$  represents the coherent signal-to-noise ratio in detectors  $I$  and  $J$  for the loudest cluster identified using only the auto-power from detector  $I$ . Last, we define the detection statistic as the maximum value of  $\Lambda^{(I)}$  among the two detectors:  $\Lambda = \max_I \Lambda^{(I)}$ .

To illustrate why this hierarchical design is useful, it is helpful to describe the procedure of background estimation as a series of numbered steps.

1. We break the coincident data into conveniently long spectrograms to be analyzed for clusters. In this Letter, we use 50%-overlapping,  $\approx 300$  s spectrograms.
2. For each spectrogram, we identify the loudest cluster  $\Gamma_I$  in each detector using the single-detector, incoherent statistic  $\text{SNR}_{\text{tot}}^{\text{max}(I)}$ .

3. If  $\text{SNR}_{\text{tot}}^{\max(I)}$  is less than some pre-determined threshold  $\text{SNR}_{\text{th}}^{(I)}$ , proceed no further. The cluster is not promising enough to spend time calculating the coherent statistic.
4. For each spectrogram—if there is a cluster passing the cut  $\text{SNR}_{\text{tot}}^{\max(I)} < \text{SNR}_{\text{th}}^{(I)}$ —we calculate the coherent detection statistic  $\Lambda$ .
5. Take the spectrogram data produced in step 1 and time-shift the clusters in one detector to create a new noise realization. Repeat steps 2–4 to generate a background distribution for  $\Lambda$ .

Since steps 2–3 use a single-detector statistic, we obtain the same list of single-detector clusters every time—it does not matter if the data streams are shifted with respect to one another. This means that we can carry out steps 2–3 just once and reuse the results for subsequent time-slides. This is important because step 2 (cluster identification) is the computationally expensive step. We still need to recalculate the coherent statistic  $\Lambda$  for each time-slide, but this is a cheap calculation since we have reduced the number of templates from  $\mathcal{O}(10^7)$  per spectrogram to one (and only a fraction of spectrograms will contain a cluster passing the step 3 cut). The zero-lag data are analyzed identically with the same hierarchical process, ensuring that the background estimation can be used to identify detections.

*Mock data analysis.* We carry out two mock data analyses. For the first analysis, we analyze Monte Carlo noise. The simulated noise is Gaussian, stationary in time, and colored according to the design sensitivity of Advanced LIGO [11]. For the second analysis, we analyze time-shifted data from Initial LIGO, which has been recolored to resemble Advanced LIGO noise at design sensitivity. The recolored noise preserves non-Gaussian noise artifacts that are present in real data, but not in Monte Carlo simulations. Of course, the character of the non-Gaussian noise present in Advanced LIGO noise is likely to be different from that of Initial LIGO as the detectors are significantly different. The recolored noise results should therefore be taken as a plausible approximation of Advanced LIGO noise. In both cases, we use a two-detector network consisting of the LIGO Hanford and Livingston observatories.

In order to estimate the background for one week of data, we analyze 16.7 days of data, corresponding to  $10^4$ , 288 s-long, 50%-overlapping spectrograms. The spectrograms span a frequency range of 100–1800 Hz and are composed of  $1 \text{ Hz} \times 1 \text{ s}$  pixels, overlapping 50% in time. This choice of spectrographic resolution is suitable for many long-lived transient waveforms with slowly varying frequency [9, 10, 18, 23].

In order to impose the constraints of a realistic search, we assume that the seedless clustering parameters are tuned just once. We employ cubic Bézier curves with a minimum duration of 40 s. For each spectrogram, we employ  $10^7$  templates to find the loudest auto-power cluster

waveform	duration (s)	$f_{\min}$ – $f_{\max}$ (Hz)	median strain
ADI 1	39	130–170	$4 \times 10^{-24}$
ADI 2	230	110–260	$2 \times 10^{-23}$
FA 2	200	790–1080	$2 \times 10^{-22}$

TABLE I: Summary of the waveforms used in this study; see [9, 10] for additional details. Accretion disk instability (ADI) waveforms are down-chirping while fallback accretion (FA) waveforms are up-chirping. The ADI waveforms have been normalized so that the isotropic equivalent energy is  $0.1M_{\odot}$ . The median strain is for face-on systems at 100 Mpc.

for each detector  $\Gamma_I$ . The difference in frequency between the first and third Bézier points determine the extent to which the signal may vary in time. We limit this variation so that the frequency of the third control point is within  $\pm 50\%$  of the frequency of the first control point.

We consider three waveforms, which we have previously used to benchmark past studies using seedless clustering [9, 10]. In particular, we employ two down-chirping accretion disk instability (ADI) signals from [24] and one up-chirping fallback accretion (FA) signal from [18]. Following [9, 10], the accretion disk instability signals are scaled so that the isotropic equivalent energy is  $0.1M_{\odot}$ . The spectrographic properties of the signals are described in Tab. I. We follow the naming conventions adopted in [9, 10].

We do not include the results for a fourth waveform FA 1 from [9, 10], which persists for 20 s and evolves from 790–1080 Hz. The implementation described above performed poorly detecting this waveform because it is relatively short and it evolves rapidly in frequency. It is possible that we could effectively detect signals of this type using a different tuning from the one we describe above, but we seek to apply a uniform set of parameters, which are optimized for longer, less-quickly-evolving signals.

Above, we listed five steps necessary for carrying out the hierarchical search. The blue curve in Fig. 1a shows the FAP associated with the single-detector statistic  $\text{SNR}_{\text{tot}}^{\max(I)}$  obtained after steps 1–2 with recolored noise. FAP is defined for a single 288 s spectrogram so  $\text{FAP} = 10^{-2}$  corresponds to a false alarm rate of  $\text{FAR} = 1/8 \text{ hr}$ . The vertical lines in Fig. 1a show the median values of the single-detector statistic  $\text{SNR}_{\text{tot}}^{\max(I)}$  for data that includes an injected ADI 2 signal injected face-on at an optimal sky location. The different colors indicate different source distances.

For step 3, it is necessary to define a threshold  $\text{SNR}_{\text{th}}^{(I)}$  for the single-detector statistic. This threshold is applied in order to control the computational costs. By throwing out below-threshold candidates, we reduce the number of surviving clusters to include in subsequent steps. The choice of  $\text{SNR}_{\text{th}}^{(I)}$  is a balancing act. Increasing the threshold reduces the computational burden, but it can also reduce sensitivity by inadvertently discarding true signals. For the Monte Carlo analysis described here, we

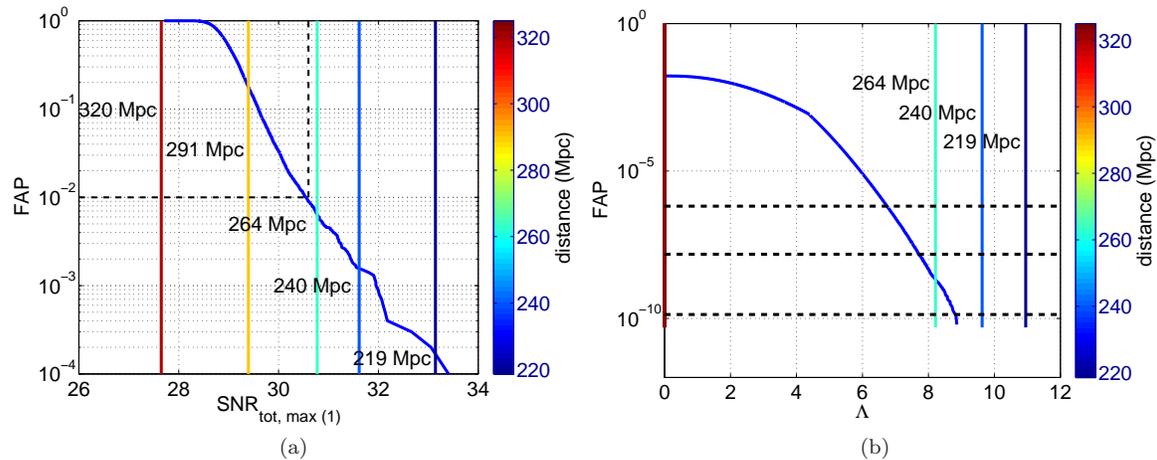


FIG. 1: Left (a): FAP versus the single-detector statistic  $\text{SNR}_{\text{tot}}^{\max(I)}$ . The blue curve indicates the distribution of recolored noise for Advanced LIGO at design sensitivity. The vertical lines show the median values for data including an accretion disk instability signal (ADI 2; face-on, optimal orientation) with different colors corresponding to different distances. The dashed black lines show which data are removed via the single-detector cut. Right (b): FAP versus the coherent statistic  $\Lambda$ . The blue curve indicates the distribution of recolored noise for Advanced LIGO (Hanford and Livingston) at design sensitivity. The vertical lines show the median values for data including an accretion disk instability signal (ADI 2; face-on, optimal orientation) with different colors corresponding to different distances. The dashed black lines indicate the threshold required for  $3\sigma$ ,  $4\sigma$ , and  $5\sigma$  detection in one week of data.

find that  $\text{SNR}_{\text{th}}^{(I)} = 29.4$  eliminates 99% of background triggers (thereby reducing the computational burden of steps 4–5 by a hundred fold) without severely harming sensitivity. The step 3 cut is indicated by dashed black line on Fig. 1a. Note that (face-on, optimally oriented) signals with distance  $\lesssim 260$  Mpc are likely to survive the cut. For recolored noise, the threshold is selected to be  $\text{SNR}_{\text{th}}^{(I)} = 30.6$ .

Having eliminated 99% of the clusters, we proceed to calculate  $\Lambda$ . In Fig. 1b, we plot FAP as a function of  $\Lambda$  for recolored noise. The different colored vertical lines indicate the median value of  $\Lambda$  for injected signals (face-on, optimal orientation). If an injected signal does not pass the step 3 auto-power cut in at least one detector, it is assigned a value of  $\Lambda = 0$ . The blue curve starts at  $\text{FAP} \approx 10^{-2}$  because 99% of the clusters are removed by the auto-power cut. The horizontal dashed lines indicate the required significance for  $3\sigma$ ,  $4\sigma$ , and  $5\sigma$  detections in one week of data. Recall that FAP is measured in reference to a 288 s spectrogram. One week corresponds to 4800 (50%-overlapping) spectrograms and so  $5\sigma$  corresponds to  $\text{FAP} = 7 \times 10^{-11}$ .

From Fig. 1b, it is apparent that the  $\lesssim 260$  Mpc that we avoided cutting in step 3 (Fig. 1a) are detectable with the coherent  $\Lambda$  statistic at the  $>4\sigma$  level. From Figs. 1a and b, we see why the hierarchical approach is superior to the single-detector statistic. The coherent statistic yields a FAP which is many orders of magnitude smaller. Next, we consider how the sensitivity is affected by the use of an incoherent statistic as an intermediate step. We compare  $\Lambda$  to  $\text{SNR}_{\text{tot}}^{\max}$  for mock signals that are loud enough to detect with the hierarchical scheme at  $5\sigma$  and

find that they are comparable to within 37%. Somewhat surprisingly, the hierarchical scheme is actually slightly more sensitive in this regime because it searches the sky in a systematic way whereas the  $\text{SNR}_{\text{tot}}^{\max}$  statistic pairs each track with a random sky position guess [10].

The initial clustering procedure (corresponding to steps 1–3) is carried out on Kepler GK104s Graphics Processor Units (GPUs) at a cost of 1.2 continuously running GPUs. We find that 18 Intel Xeon E5-2670 Central Processing Unit (CPU) cores are required to match the performance of one GPU. Steps 4–5, in which we time-shift the data in order to calculate the coherent statistic  $\Lambda$ , are performed using single CPU cores at a cost of 5.8 continuously running cores in order to achieve the background estimation necessary to identify  $5\sigma$  detection candidates. The calculation in steps 4–5 does not, at present, benefit dramatically from GPU acceleration. To put this in perspective, a fully coherent search, calculating  $\text{SNR}_{\text{tot}}$  for each time-slide with seedless clustering would require  $2.2 \times 10^6$  continuously running GPUs.

In Tab. II, we summarize the results of a sensitivity study in which we estimate  $5\sigma$  detection distance: the distance at which we can detect a waveform with a FAP corresponding to  $5\sigma$  with a false dismissal probability (FDP) of 50%. We consider the cases of face-on, optimally oriented sources (useful for comparison with previous work [9, 10]) and also randomly oriented sources with random sky locations. The first column is the waveform, the second is the type of noise (MC = Monte Carlo or RC = recolored noise), and the second is the  $5\sigma$  detection distance. The RC detection distances are sometimes smaller than the MC distances because non-stationary noise arti-

facts present in real data tend to decrease the sensitivity compared to idealized Gaussian noise.

Waveform	Noise	$5\sigma$ distance (Mpc)
ADI 1 optimal	MC	250
	RC	250
ADI 1 random	MC	150
	RC	127
ADI 2 optimal	MC	232
	RC	232
ADI 2 random	MC	139
	RC	130
PT 2 optimal	MC	30
	RC	25
PT 2 random	MC	18
	RC	14

TABLE II: Five sigma detection distances (FDP = 50%) for the three different test waveforms summarized in Table I in Gaussian Monte Carlo noise (MC) and Initial LIGO noise, recolored to match Advanced LIGO at design sensitivity (RC). We assume Hanford and Livingston detectors operating at design sensitivity. “Optimal” means that the source is face-on and optimally oriented to maximize detectability. “Random” means that the orientation and sky location of the source are chosen from random distributions.

In order to put Tab. II in context, we compare the face-on, optimal sky location,  $5\sigma$  detection distances presented here—estimated for one week of data—to the  $3.3\sigma$  detection distances from [10]—estimated for a sin-

gle 288 s spectrogram with a smaller 150 Hz-wide band. Despite the fact that the results here include an additional  $4.6 \times 10^8$  trial factor, the hierarchical  $5\sigma$  detection distances for Monte Carlo are only 0.45 – 0.82 times less than the  $3\sigma$  detection distances from [10] obtained using seedless clustering. For recolored noise the ratios are 0.50 – 0.96. The event rate and loudness of long-lived transients are unknown, but the sensitivity we demonstrate here is sufficient to potentially detect long-lived signals with second-generation detectors [18, 22].

For illustrative purposes, we focus here on the detection of  $>5\sigma$  events with minimal resources. As a result, the aggressive auto-power cut eliminates some fraction of  $<5\sigma$  signals that could be detected with a different tuning. However, the algorithm can be tuned differently to balance cost with sensitivity for different FAPs. For future work, it is worthwhile considering if this strategy can be usefully employed in other situations where the character of the noise is less well-behaved. Moreover, the general strategy we outline here—splitting a search into a computationally expensive “incoherent” step followed by a computationally cheap “coherent” step in order to facilitate rapid background estimation—may find use in the broader community. For example, a similar scheme could prove useful in determining the significance of potentially faint electromagnetic signatures found (by  $\geq 2$  telescopes) in coincidence with gravitational-wave detections. We thank Anthony Piro for sharing the fallback accretion waveforms used in this analysis. We thank Peter Shawhan and Jonah Kanner for helpful comments.

- 
- [1] J. Abadie et al., Phys. Rev. D **85**, 082002 (2012).  
[2] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, N. Leroy, F. Robinet, and M. Vavoulidis, **27**, 015005 (2010).  
[3] Sutton P. et al., New Journal of Physics **12**, 053034 (2010).  
[4] Klimentenko S., Yakushin I., Mercer A., and Mitselmakher G., **25**, 114029 (2008).  
[5] E. Thrane, S. Kandhasamy, C. D. Ott, et al., Phys. Rev. D **83**, 083004 (2011).  
[6] J. Abadie et al., Phys. Rev. D **85**, 122007 (2012).  
[7] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and É. É. Flanagan, Phys. Rev. D **63**, 042003 (2001).  
[8] J. Abadie et al., Phys. Rev. D **81**, 102001 (2010).  
[9] E. Thrane and M. Coughlin, Phys. Rev. D **88**, 083010 (2013).  
[10] E. Thrane and M. Coughlin, Phys. Rev. D **89**, 063012 (2014).  
[11] J. Aasi et al., **32**, 074001 (2015).  
[12] E. Thrane, V. Mandic, and N. Christensen, Phys. Rev. D **91**, 104021 (2015).  
[13] M. Coughlin, E. Thrane, and N. Christensen, Phys. Rev. D **90**, 083005 (2014).  
[14] M. Coughlin, P. Meyers, E. Thrane, J. Luo, and N. Christensen, Phys. Rev. D **91**, 063004 (2015).  
[15] G. Farin, *Curves and Surfaces for CAGD, Fourth Edition: A Practical Guide* (Academic Press, 1996).  
[16] A. L. Piro and E. Pfahl, Astrophys. J. **658**, 1173 (2007).  
[17] A. L. Piro and C. D. Ott, Astrophys. J. **736**, 108 (2011).  
[18] A. L. Piro and E. Thrane, Astrophys. J. **761**, 63 (2012).  
[19] A. Corsi and P. Mészáros, Astrophys. J. **702**, 1171 (2009).  
[20] K. Kiuchi, M. Shibata, P. J. Montero, and J. A. Font, Phys. Rev. Lett. **106**, 251102 (2011).  
[21] M. H. P. M. van Putten, Phys. Rev. Lett. **87**, 091101 (2001).  
[22] M. H. P. M. van Putten, Astrophys. J. Lett. **684**, L91 (2008).  
[23] J. Aasi et al., Phys. Rev. D **88**, 122004 (2013).  
[24] L. Santamaría and C. D. Ott, *Gravitational wave emission from accretion disk instabilities - analytic models* (2011), <https://dcc.ligo.org/LIGO-T1100093-v2/public>.  
[25] We correct the normalization factor given in Eq. 5 of [10]: the coefficient is  $1/N^{1/2}$  as opposed to  $1/N$ .