

Prospects for searches for long-duration gravitational-waves without time slides

Michael Coughlin,^{1,*} Patrick Meyers,² Shivaraj Kandhasamy,³ Eric Thrane,⁴ and N. Christensen⁵

¹*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

²*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

³*Physics and Astronomy, University of Mississippi, University, Mississippi 38677-1848, USA*

⁴*School of Physics and Astronomy, Monash University, Clayton, Victoria 3800, Australia*

⁵*Physics and Astronomy, Carleton College, Northfield, Minnesota 55057, USA*

(Received 1 May 2015; published 17 August 2015)

The detection of unmodeled gravitational-wave transients by ground-based interferometric gravitational-wave detectors is an important goal for the advanced detector era. These searches are commonly cast as pattern recognition problems, where the goal is to identify statistically significant clusters indicating the presence of gravitational-wave transients in spectrograms of detector strain power when the precise signal morphology is unknown. In previous work, we have introduced a clustering algorithm referred to as seedless clustering, and shown that it is a powerful tool for detecting weak and long-lived (~ 10 – 1000 s) gravitational-wave transients. However, as the algorithm is currently conceived, in order to carry out a search on approximately a year of data, significant computational resources may be required for estimating background events. Currently, the use of the algorithm is limited by the computational resources required for performing background studies to assign significance to events identified by the algorithm. In this paper, we present an analytic method for estimating the background generated by the seedless clustering algorithm and compare the performance to both Monte Carlo Gaussian noise and time-shifted gravitational-wave data from a week of LIGO's 5th Science Run. We demonstrate qualitative agreement between the model and measured distributions and argue that the approximation will be useful to supplement conventional background estimation techniques for advanced detector searches for long-duration gravitational-wave transients.

DOI: [10.1103/PhysRevD.92.043007](https://doi.org/10.1103/PhysRevD.92.043007)

PACS numbers: 95.75.-z, 04.30.-w, 07.05.Bx

I. INTRODUCTION

Second-generation gravitational-wave detectors such as Advanced LIGO [1] and Advanced Virgo [2] will be coming online in the coming months and years. Some searches for gravitational-wave transients seek to detect gravitational-wave transients lasting ~ 10 – 1000 s. Compact binary coalescences of black holes (and/or neutron stars) are one example of long-lived gravitational-wave sources [3–5]. Uncertain models exist for more exotic sources of long-lived transients, including emission from rotational instabilities in protoneutron stars [6–9] and black-hole accretion disk instabilities [10–12]. When a matched filter search is not possible, searches for unmodeled long-lived transients [13–18] can be employed.

Searches for compact binary coalescences typically rely on performing time-slides of single detector *triggers*, generated by performing a matched filter (in the compact binary coalescence case) on single-detector time series data. Time-slides are the shift of either time-series data or triggers from one detector against those of another, and are designed to eliminate potential gravitational-wave content contaminating the background estimation. Burst searches

use clustering algorithms, as opposed to matched filtering techniques, on single detector time-frequency maps [19] or multidetector coherent maps [14,20,21]. Calculating the coherent statistic takes significant computational resources because the triggers are inherently multidetector and so the time-slides must be done on the time-series data itself, rather than the single-detector triggers. Because detector noise is generally non-Gaussian, it is difficult to know if an event in one detector is signal or noise. For this reason, multiple detectors are required to perform gravitational-wave searches. To estimate background, these searches time-shift the data of each detector with respect to the other (s), by some unphysical delay which is larger than the light travel time between the detectors. The coherent statistics are then computed between the time-shifted data in the same way as the original search algorithm, and in this way, false alarm rates can be estimated. Current searches use thousands of time-slides or more [4,22,23]. The main limitation on the number of time-slides that can be performed is limited computational resources, although the short-duration coherent burst pipelines have now been tested in the ten-thousand time-slides regime. Currently, they are moving toward the hundred-thousand time-slide regime, while the compact binary matched filtering pipelines have successfully generated 5σ background

*coughlin@physics.harvard.edu

distributions. 5σ means that the event has a SNR exceeding a significance threshold that corresponds to a 1 in 1,744,278 occurrence, after accounting for trials factors. For a single event, therefore, 1,744,278 time-slides must be performed to determine whether an event has a SNR exceeding a 5σ threshold. The difficulty in reaching these levels is due to computationally intensive calculations like the matched filter used in compact binary coalescence [24], calculation of the coherent SNR [25] used in burst searches, and potentially seedless clustering [15–18] in a long-duration transient search.

Was *et al.* demonstrated the limitation of using time-slides to perform background estimation in the single-detector trigger case [26,27]. Although coherent analyses do not use single-detector time-slides, background estimation for coherent searches rely on estimating the properties of the noise with finite measurements and therefore have error bars on their background estimation as well. They showed that the precision on the background estimation using time-slides of trigger streams is in fact limited and that the variance associated with their use saturates at some point. The computational limitations and the potential problems with time-slides motivate a search for potential alternative forms of background estimation in gravitational-wave searches. Gravitational-wave searches for isotropic stochastic gravitational wave backgrounds [28,29] and directional searches toward Sco X-1, the galactic center, and SN1987A [30] have assumed that the detection statistic is normally distributed with a known mean and variance that can be calculated from first principles when performing the searches. These searches sum up data from long stretches of time, and combined with the use of long time segments (60 s) and Gaussianity cuts, these statistics are Gaussian by the central limit theorem. This has the significant computational cost-saving benefit of not requiring time-slides to perform the search, although limited time-shift analyses are used as sanity checks and to ensure that particularly non-Gaussian frequency bins can be removed from the analysis.

Some searches for long-duration gravitational-wave transients use the same cross-correlation technique as stochastic searches [14], although other methods exist [20,21]. They utilize cross-power spectrograms, computed from the cross-correlation of two gravitational-wave detectors, and use pattern recognition algorithms to search for clusters of excess strain cross-power [14]. One algorithm used to search for long-duration gravitational waves is *seedless clustering*, which integrate along many different paths in spectrograms. This algorithm is sensitive to signals that can be well-approximated by parametrized curves, and the advantage of seedless clustering is most pronounced for long and weak signals [15–18]. We have previously shown how seedless clustering algorithms can be used to significantly enhance the sensitivity of searches for signals of this type [15]. Although seedless clustering algorithms are

embarrassingly parallel [31], and therefore computations can be performed on graphical processor units, seedless clustering searches are still limited by computation of the noise background.

Cannon *et al.* [32] recently proposed a method to estimate the false alarm probability of compact binary coalescences without time-slides. They are able to approximate compact binary events as a Poisson process in order to convert the calculated false alarm probability into a false alarm rate. This in particular allows for a statistical detection of a population of events, which could be collectively more significant than the single most significant event alone. The method proposed in our paper is similar in that we measure events based on the data and then use a statistical approximation to the distribution of the measured tracks to make approximations to the noise background. There are also a number of notable differences. Because long-duration transient gravitational waves are typically searched for using a coherent combination of detector data, the trigger distributions no longer obey Poisson statistics. Instead, we will exploit the fact that seedless clustering sums many approximately statistically independent pixels to use Gaussian statistics to estimate the background. In this paper, we demonstrate a semianalytical approximation to the seedless clustering output from cross-correlation spectrograms. One potential criticism of the analysis that follows is the fact that we compare the approximation with data from time-slide analyses out to $\approx 3\sigma$, not to the 5σ distributions we present at the end of the paper. It would be necessary to perform 5σ worth of time-slides to verify the approximation. This calculation is currently very difficult to do computationally, and of course, if we could perform 5σ worth of time-slides, we would not need an approximation in the first place. Moreover, as we perform the analysis using a relatively clean week-long stretch of data, different sets of data could result in different results. Therefore, we consider the analysis that follows as a first test for the feasibility of an approximate method. As argued above, we expect the background distributions to be better behaved in long-duration analyses than in short-duration searches, and therefore perhaps less susceptible to significant deviations from empirical distributions. In the future, we can use distributions generated by future analyses that perform more time-slides and over longer periods to compare against the approximation to test its utility. Therefore, although time-slides are likely required to create confidence in a detection due to the problem just described, we now summarize several reasons why it is useful to consider alternative significance-estimation strategies.

A. Algorithm verification

The semianalytic method provides a verification for the pipeline in multiple ways. In the case where data-quality work is being performed correctly, in general, the data

should be generally well approximated by Gaussian noise, outside of some data transients which pass the data quality cuts. Therefore, background estimation should approximately follow the distribution if it is assumed that the data is Gaussian. Also, this provides a sanity check that the algorithm performs as expected on the data. By performing a limited number of time-slides or performing a simulated analysis on Gaussian noise, it should be clear that the model for the algorithm is correct, which can provide confidence that the algorithm is performing as expected (or not).

B. Sensitivity to waveform models

There are a number of papers contained in the literature about the sensitivity of gravitational-wave detectors to long-duration gravitational waves [6,15–18]. In general, the sensitivity studies have been performed by running the analysis on 1,000 ft -maps to reach a False alarm probability (FAP) of 0.1%, and the sensitivity to various waveform models are computed relative to this number. For a year of data, assuming ft -maps of 250 s with 50% overlap, and a desirable FAP of $\approx 3\sigma$ or FAP = 0.27%, there will be more than 10^8 maps analyzed. Before any analysis, either a search for gravitational-waves or a waveform sensitivity study, is performed, it is desirable to be able to estimate the background quickly. This estimate informs expectations of potential results as well as how to setup the analysis. Using the method described in this paper, we can analytically compute what we expect a threshold based on this number of maps, without needing to perform an analysis with many time-slides.

C. Event follow-up and electromagnetic alerts

There are preparations for joint electromagnetic and gravitational-wave observations in the advanced detector era [33]. Low-latency gravitational-wave searches are aiming for run-times ≤ 1 min. For significance estimates on this time-scale, rapid background estimation techniques are required. The method described in this paper is able to give an approximate FAR for any event on this time-scale. In the case where there is eventually interest in joint electromagnetic and gravitational-wave observations for generic long-duration transients, this method may be useful for making that happen.

The remainder of the paper is organized as follows. In Sec. II, we describe the formalism of an all-sky transient search and seedless clustering. In Sec. III, we present the results of a Monte Carlo and time-shifted study comparing the semianalytical model to seedless clustering. In Sec. IV, we explore the errors, both systematic and statistical, with our method. In Sec. V, we discuss our conclusions and suggest directions for future research.

II. FORMALISM

We use the cross-correlation of two GW strain channels from spatially separated detectors to perform

searches for long-duration GW transients. We construct ft -maps of cross-power signal-to-noise ratio. We divide detector strain time series into segments and compute Fourier transforms of the segments to create the pixels, which we denote as $\tilde{s}_I(t; f)$, where we take strain data from detector I for the segment with a mid-time of t . Following [15–18], the segments are 50%-overlapping and Hann-windowed with duration of 1 s and a frequency resolution of 1 Hz.

The expression for the cross-power signal-to-noise ratio is as follows [14]:

$$\rho(t; f | \hat{\Omega}) = \frac{2\sqrt{2}}{\mathcal{N}} \operatorname{Re} \left[e^{2\pi i f \Delta \vec{x} \cdot \hat{\Omega} / c} \frac{\tilde{s}_I^*(t; f) \tilde{s}_J(t; f)}{\sqrt{P'_I(t; f) P'_J(t; f)}} \right], \quad (1)$$

where $\Delta \vec{x}$ is a vector describing the relative displacement of the two detectors, $\hat{\Omega}$ is the direction of the GW source, and c is the speed of light. The time delay between the two detectors, which is a direction-dependent phase factor, is in the $e^{2\pi i f \Delta \vec{x} \cdot \hat{\Omega} / c}$ term. $P'_I(t; f)$ and $P'_J(t; f)$ are the auto-power spectral densities for detectors I and J in the segments neighboring t . \mathcal{N} is a FFT normalization factor, $L \times Fs$, where L is the length of data in seconds and Fs is the sampling frequency. Additional details can be found in [14–18].

We write the total signal-to-noise ratio for a cluster of pixels as a sum over $\rho(t; f | \hat{\Omega})$:

$$\operatorname{SNR}_{\text{tot}}(\Gamma) \equiv \frac{1}{N^{1/2}} \sum_{\{t; f\} \in \Gamma} \rho(t; f | \hat{\Omega}), \quad (2)$$

where N is the number of pixels in Γ , which is chosen from a bank of parametrized frequency-time tracks, and each track is referred to as a *template*.

To modify the above algorithm to perform an all-sky search [16–18], we use complex signal-to-noise ratio:

$$\mathfrak{p}(t; f) = \frac{2\sqrt{2}}{\mathcal{N}} \left[\frac{\tilde{s}_I^*(t; f) \tilde{s}_J(t; f)}{\sqrt{P'_I(t; f) P'_J(t; f)}} \right]. \quad (3)$$

This statistic preserves the complex phase information, which encodes the direction of the source. As a proxy for the sky location, which is unknown, we add an additional variable $\Delta \tau$ which corresponds to the time delay between the detectors [16]. Therefore, we rewrite Eq. (2) as

$$\operatorname{SNR}_{\text{tot}}(\Gamma) \equiv \frac{1}{N^{1/2}} \sum_{\{t; f\} \in \Gamma} \operatorname{Re}[e^{2\pi i f \Delta \tau} \mathfrak{p}(t; f)], \quad (4)$$

and this sum is carried out for many randomly selected clusters Γ . We finally define $\operatorname{Max}[\operatorname{SNR}_{\text{tot}}]$ as the maximum of $\operatorname{SNR}_{\text{tot}}$ taken over all Γ .

A. Parametrizations

In any seedless clustering algorithm, Γ is chosen such that it is sensitive to the morphology of the gravitational waves being searched for. There are two types we will consider in this paper, although the method is generic enough to work for any parametrization.

1. Bézier curves

For generic narrow-band long transient gravitational waves [15,16], Γ is chosen randomly from the set of quadratic Bézier curves [34] subject to the constraint that the curve persists for a duration t_{\min} . Three time-frequency control points determine the template: $P_0(t_{\text{start}}, f_{\text{start}})$, $P_1(t_{\text{mid}}, f_{\text{mid}})$, and $P_2(t_{\text{end}}, f_{\text{end}})$, and the curve is parametrized by $\xi = [0, 1]$:

$$\begin{pmatrix} t(\xi) \\ f(\xi) \end{pmatrix} = (1 - \xi)^2 P_0 + 2(1 - \xi)\xi P_1 + \xi^2 P_2. \quad (5)$$

These arrays allow the sum in Eq. (2) to be computed for a large number of templates in parallel. For practical applications, the number of templates T is typically chosen to be $T = \mathcal{O}(10^4 - 10^8)$. To perform a computationally feasible all-sky analysis, $T = 2 \times 10^4$ templates are feasible, and we use this number in the analysis that follows.

2. Post-Newtonian templates for compact binary coalescences

Another parametrization for Γ currently in the literature creates templates based on a post-Newtonian model for chirp-like signals created by circular compact binary coalescences [17]. For searches for compact binary coalescences with seedless clustering, we can use a more specialized template bank consisting of parametrized chirps:

$$f(t) = \frac{1}{2\pi} \frac{c^3}{4GM_{\text{total}}} \sum_{k=0}^7 p_k \tau^{-(3+k)/8}, \quad (6)$$

where

$$\tau = \frac{\eta c^3 (t_c - t)}{5GM_{\text{total}}}, \quad (7)$$

where the expansion coefficients p_k can be found in [35], G is the gravitational constant and M_{total} is the total mass of the binary. These templates are parametrized by the coalescence time and the chirp mass. It was shown in [17] that while the waveform depends on the individual component masses, the main features of the signal can be well approximated by only the chirp mass, and we can approximate that the individual component masses are equal. This has similarities to the matched-filtering

template banks used in compact binary searches [3,4]. The key difference is that to first-order, chirp mass is the only term that contributes to the time-frequency evolution, and therefore the template bank is only one-dimensional. Combined with the fact that this method is only sensitive to the pixelated frequency evolution of the gravitational wave (instead of the phase), the templates used in this analysis are significantly coarser than that in traditional searches.

B. Semianalytical approximation

We now describe a semianalytical approximation to the background of our seedless clustering algorithms. Seedless clustering, which computes the sum of pixels in a track, divided by the square root of the number of pixels in the track, lends itself to modeling due to its simplicity. By the central limit theorem, the sum of a sufficiently large number of independent random variables, each with a well-defined expected value and well-defined variance, will be approximately normally distributed. Therefore, we expect that the sum of many pixels will approach a normal distribution, given by

$$P(z)dz = \frac{1}{\sigma\sqrt{2\pi}} e^{-(z-\mu)^2/2\sigma^2} dz. \quad (8)$$

Because seedless clustering measures the maximum SNR_{tot} of many tracks, here we seek the extreme value distribution for SNR_{tot} . This is motivated by the desire for a distribution with which to compare those measured from an analysis using the algorithm. We can analytically compute a probability distribution for this maximum value as follows. Given a random sample of SNR_{tot} drawn from many maps, (X_1, \dots, X_N) , from a continuous distribution with a probability density function $f(x)$ and cumulative density function $F(x)$, the cumulative density function of the maximum of SNR_{tot} is then given by

$$\begin{aligned} \text{CDF}_{\text{Max}[\text{SNR}_{\text{tot}}]}(z) &= P(\max(X_i) < z) \\ &= P(X_1 < z, \dots, X_N < z) = P(X_1 < z) \dots P(X_N < z), \end{aligned} \quad (9)$$

where the third equality assumes that the random samples are independent. In some cases where the random samples are not independent but the correlated distributions are conservative, as in some of our analyses below, an upper-bound can instead be derived

$$\begin{aligned} \text{CDF}_{\text{Max}[\text{SNR}_{\text{tot}}]}(z) &= P(X_1 < z, \dots, X_N < z) \\ &\leq P(X_1 < z) \dots P(X_N < z). \end{aligned} \quad (10)$$

We will discuss the effect of this assumption in Sec. IV. In the case where the probability density functions are identical, the equation becomes

$$\text{CDF}_{\text{Max}[\text{SNR}_{\text{tot}}]}(z) = [P(X < z)]^N = [F_X(z)]^N. \quad (11)$$

We show below that we can use this equation, where the CDFs are given by Gaussian CDFs, to approximate the seedless clustering distributions. Even though in our case $F_X(z)$ is derived from a Gaussian distribution, Eq. (11) is true for any general distribution represented by $F_X(z)$. Hence in cases where the analytic expression for $F_X(z)$ is difficult to derive or approximate, one can use the observed distribution.

III. BACKGROUND STUDY

We can test the approximations by performing the analysis on Monte Carlo Gaussian noise and initial LIGO noise from the Hanford, WA (H1) and Livingston, LA (L1) detectors. We create complex signal-to-noise ratio spectrograms $\mathfrak{p}(t; f)$ using Eq. (3) and analyze each with the various seedless clustering algorithms. Following [15,16], we create 250 s maps in a band between 100–250 Hz with spectrogram resolution of $1 \text{ s} \times 1 \text{ Hz}$ using 50%-overlapping Hann windows. The results for each are as follows.

A. Bézier parametrization

We begin by analyzing the performance of the analytic model on the Bézier parametrization.

We run the seedless clustering algorithm over hundreds of these maps and save SNR_{tot} for each of the tracks in the map. The left of Fig. 1 shows a histogram of the resulting SNR_{tot} distribution for both the Monte Carlo and initial LIGO data.

We fit Eq. (8) to the resulting distributions. We find best fits of $\mu = 0.0007$ and $\sigma = 0.99$. The fact that the distribution has approximately a mean of zero and a standard deviation of one is expected based on the fact that ρ has a mean of 0 and we use the \sqrt{N} normalization in the SNR_{tot} calculation. We find that the agreement is reasonable out to the tails of the distribution. The right of Fig. 1 shows the standard deviation of the distribution as a function of track length. The standard deviation differs on the order of a few percent across the track lengths considered. For the sake of simplicity, we assume that the distribution is approximately independent of track length.

We now simulate an all-sky search by performing 100 time-slides in a week of data. The data are processed with a glitch identification cut [36] as if it were a real analysis. In order to apply the algorithm from [36], we assume that the source is optimally oriented with an optimal sky position. To compare this to the analytic approximation in Eq. (11), we use the Gaussian fit shown in Fig. 1 to approximate the SNR_{tot} distribution. The steps required to turn the SNR_{tot} distribution into a p-value vs SNR distribution are as follows. To generate a SNR_{tot} value for a single simulated *map*, we generate N random numbers consistent with the Gaussian distribution of mean and variance as estimated above. We then take the maximum value of these values to compute the $\text{Max}[\text{SNR}_{\text{tot}}]$. To generate a p-value vs SNR distribution, $\text{Max}[\text{SNR}_{\text{tot}}]$ is generated for M instances of spectrograms, where $1/M$ is the smallest p-value required. $\text{Max}[\text{SNR}_{\text{tot}}]$ is placed in ascending order. The p-value is calculated as an array between $1/M$ and 1 with spacing given by $1/M$. For a Gaussian distribution where the mean




FIG. 1 (color online). The plot on the left is the background distribution for the seedless clustering algorithm cluster SNR defined in Eq. (4). Monte Carlo denotes Gaussian colored noise. Time-shift denotes real time-shifted data with vetoes to limit the effects of instrumental artifacts. The theoretical line corresponds to the Gaussian approximation to the distribution given by Eq. (8). The plot on the right is the standard deviation of SNR_{tot} as a function of track length. The standard deviation differs on the order of a few percent across the track lengths considered. In our analysis, we approximate the standard deviation of SNR_{tot} across track length as constant.

and standard deviation are the same across all trials, this process can also be performed analytically by simply computing Eq. (11) for the measured distribution.

We perform two search simulations using the Bézier parametrization. The first uses Bézier templates computed for a specific search direction. The second loops over time-delays for each template. By searching over 40 different time-delays, corresponding to 40 different sky rings, a computationally efficient all-sky search can be performed. This was demonstrated in [17] to be sufficient to recover signals in arbitrary directions. The top left of Fig. 2 demonstrates the analysis for the first simulation using 20,000 tracks, showing both empirical time-slides as well as the theoretical approximation method explained as before (both the 10th, 50th, and 90th percentiles). We find excellent agreement between the analytic model and empirical time-slides for the directed search. The distributions for the all-sky search, on the top right of Fig. 2, however, are not

generally within 1σ . We explore systematic errors related to this in the next section. Finally, we show the $\text{Max}[\text{SNR}_{\text{tot}}]$ required for a 5-sigma gravitational-wave detection using the Bézier parametrization in Fig. 3.

B. Compact binary coalescence parametrization

We now analyze the performance of the analytical model on the chirp templates. We create maps assuming Gaussian noise consistent with the design sensitivity of Advanced LIGO. Following [17], we create 660 s maps in a band between 10–150 Hz with a spectrogram resolution of $1 \text{ s} \times 1 \text{ Hz}$.

We perform a similar analysis to the above. We find best fits for the Gaussian distribution of $\mu = -0.004$ and $\sigma = 1.06$. The major difference between the Bézier and chirp-like template analysis is the degree of correlations between the drawn tracks. In the Bézier case, the tracks are




FIG. 2 (color online). Background distributions computed for the seedless clustering algorithm using both Bézier and chirp-like templates. The distributions are generated from time-shifted initial LIGO data. The top row corresponds to the Bézier templates and the bottom row chirp-like templates. The left column corresponds to a directed search (in a specific sky direction) and the right to an all-sky search performed looping over 40 time-delays for each template. The theoretical line corresponds to using Gaussian distributions with standard deviations presented in Fig. 3. The dotted lines correspond to 1σ error bars on the analytic approximation. These are derived from simulating the p-value vs. SNR distribution using many random seeds, in essence creating thousands of p-value vs SNR distributions consistent with the measured distributions, and computing the median and 1σ error bar for each p-value. The analytical background distributions for the directed searches are consistent with the measured background (within 1σ). The distributions for the all-sky searches, however, are not generally within 1σ . This is due to the assumption in the analytic-model that the loop over time-delays creates more independent trials, which is not the case and biases the result (please see the text for more details).

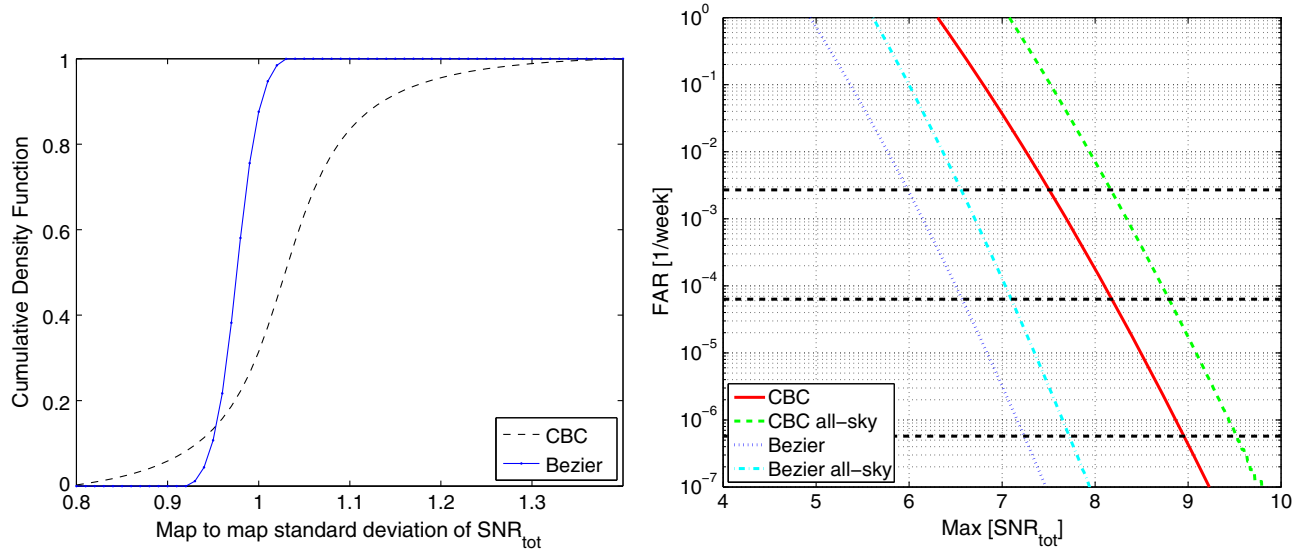


FIG. 3 (color online). The plot on the left is the cumulative density function of the standard deviation of the SNR_{tot} for the two parametrizations. Distribution is significantly broader for chirp-like templates due to track correlations. The plot on the right is the background distribution using the analytic approximation for the cases considered here. It shows the $\text{Max}[\text{SNR}_{\text{tot}}]$ required for a 3σ , 4σ , and 5σ gravitational-wave detection using seedless clustering in dotted horizontal lines. The four different lines corresponds to the four-different seedless analysis types considered in this paper (a directed and all-sky Bézier template search and a directed and all-sky compact binary search), which each have significantly different distributions of $\text{Max}[\text{SNR}_{\text{tot}}]$.

drawn randomly and the degree of correlation is simply determined by the overlap between the tracks. In the chirp-like template case, the degree of correlation is much higher, despite the significantly fewer templates used in the analysis. This correlation arises from the step in time and overlap in parameter space between chirp-like templates of similar chirp mass. This correlation is important because it changes the standard deviation of the SNR_{tot} of the tracks in individual maps and therefore the final distribution of $\text{Max}[\text{SNR}_{\text{tot}}]$. Figure 3 demonstrates the cumulative density function of the standard deviation of the SNR_{tot} for the two parametrizations. In the Bézier case, the standard deviation of SNR_{tot} is approximately a step function, which allows for the use of a single standard deviation to cover all cases. The distribution is significantly broader for chirp-like templates due to track correlations. It is for this reason that we modify the Bézier p-value algorithm by drawing from the measured distribution of standard deviations of the maps when drawing from the Gaussian distribution. The bottom of Fig. 2 demonstrates the algorithm using both empirical Monte Carlo noise as well as the theoretical approximation method explained as before (both the 10th, 50th, and 90th percentiles) for both a directed and all-sky search. Similar to the Bézier case, we find excellent agreement between the analytic model and empirical time-slides for the directed search, while the distributions for the all-sky searches, however, are not generally within 1σ . We explore systematic errors related to this in the next section. We show the $\text{Max}[\text{SNR}_{\text{tot}}]$ required for a 5-sigma gravitational-wave detection using the chirp-like parametrization in Fig. 3.

IV. UNCERTAINTIES: STATISTICAL AND SYSTEMATIC

We now explore the systematic and statistical errors in our measurement. The measurement of the statistical errors in this analysis is straightforward. Each simulation is computationally cheap, as it involves generation and manipulation of matrices of random numbers, and therefore can be performed over and over again to generate distributions. This was done in order to generate the 1σ distributions in Fig. 2, for example.

Of more interest, perhaps, is consideration of the systematic errors in the method. There are a number of reasons we might expect small disagreements between the theoretical model and the empirical results. A major source of systematic error is in the rotation of pixels in the all-sky searches, where time-delays are looped over. This involves a rotation in the complex plane of the individual pixels that make up the tracks. This creates difficulty for the analytic analysis. Because the analysis amounts to a rotation, the 40 time-delays do not correspond to 40 independent trials (which would simply multiply the number of tracks by 40). In the analysis above, we simply multiplied the number of tracks by 40, corresponding to the 40 time-delays in the analysis. In the left of Fig. 4, we explore this effect by simulating tracks without any rotation (the directional case) and with 40 rotations (the all-sky case), and compare this to the distribution of 40 random tracks. We show that using 40 random tracks is conservative relative to using the 40 time-delays case. This indicates that multiplying the number of trials by 40, as is done in the analysis, is conservative.

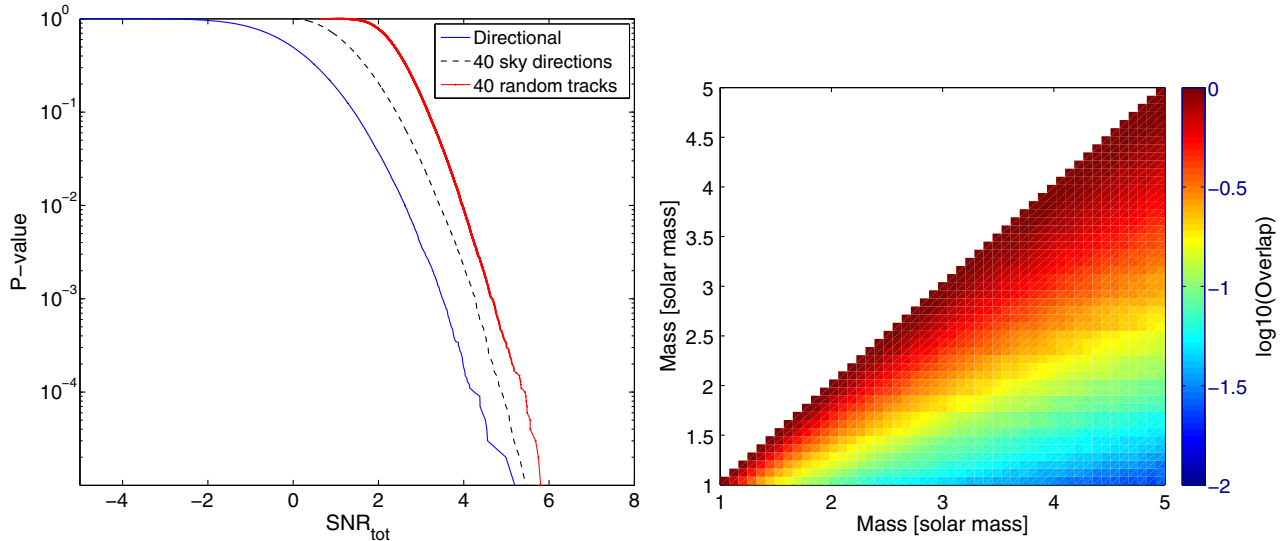


FIG. 4 (color online). The plot on the left is the p-value vs SNR_{tot} distribution for three different cases (for a single map): a directional search, a search looping over 40 sky-directions, and a search using 40 different tracks. As is expected, a search searching over 40 sky-directions results in higher background values for SNR_{tot} than the directional case. Similarly, a search using 40 different tracks is conservative relative to the 40 sky-direction case, indicating that it is appropriate to consider such a method conservative. The plot on the right is the overlap between the templates used in the compact binary search algorithm. In general, there is significant overlap between templates, which biases the computation of the theoretical SNR_{tot} distribution.

In this case, using Eq. (10), which places an upper-bound on the distribution is more accurate than Eq. (9), which assumes that the trials are independent.

One possibility to do even better is to actually measure the covariances between the rotated pixels. This situation is similar to [30]. In this work, the authors place limits on gravitational-wave strain from different portions of the sky. This was difficult because the distribution of maximum SNR for a sky map contains nonzero covariances that exist between different pixels (or patches) on the sky. They simulate the covariance between pixels numerically, by simulating many realizations that have expected covariances described by the Fisher matrix. In this case, we could numerically compute a covariance matrix, which we can diagonalize to create a basis of noncovariant variables. Then, one would generate random realizations of these noncovariant variables and use the covariance matrix to convert them into a set of randomly generated covariant variables. One difficulty is that the distribution of the noncovariant variables might not be the same as the covariant variables. With this method, we could determine the set of covariant variables which describe the distribution of SNR_{tot} .

Another potential systematic error arises from the use of Eq. (9), in particular the assumption that the trials are independent. One way in which this manifests is that real detectors have noise transients and nonstationary noise, which violate some of the approximations used here. Severe non-Gaussianity which eludes both the pixel thresholds and the Gaussianity cuts applied in the analysis would show up not only potentially as a loud background trigger, but would increase the correlation between tracks (as any

track that passes through those pixels would have an increased SNR_{tot}). Generating purely random numbers to approximate SNR_{tot} is an approximation. This is because the tracks are analyzed on the same map, and therefore overlapping tracks will have correlated SNR_{tot} values. This has the effect in the analysis of changing the effective σ from map to map. One potential conservative solution would be to measure the σ in each map and then generate $\text{Max}[\text{SNR}_{\text{tot}}]$ distributions using that value. Another implicit assumption is that the pixels in the tracks are uncorrelated. This assumes that the noise is Gaussian and stationary and ignores the correlation between pixels in the maps. However, the cross-power statistic uses PSD's from adjacent pixels (in the time direction) to estimate σ [from Eq. (2)]. This is to avoid a bias in pixel SNR for an isolated loud pixel. This means there is a correlation between adjacent pixels. In the right of Fig. 4, we explore this effect by computing the overlaps between the templates used in the compact binary search. Due to the need for maximal coverage of the parameter space, there is significant overlap between templates. The use of Eq. (9) biases the analysis in this case. Computations such as this could be used to modify Eq. (9) to account for the lack of independence between templates by determining the number of *effective trials* from the data.

V. CONCLUSIONS AND FUTURE WORK

In previous work, we showed how a seedless clustering algorithm could significantly improve the sensitivity of searches for long-lived, unmodeled gravitational-wave

transients [15–18]. Here, we show how the simplicity of the search statistic allows for the development of a semi-analytic approximation to the background generated by the algorithm and compared the performance using a week of LIGO S5 time-shifted data. We described algorithmic subtleties not addressed by this model and quantify the errors between the model and the measured distribution. We argued that it will be useful for pipeline characterization, as well as potentially for low-latency FAP reporting for gravitational-wave searches.

In the future, we will move beyond the simple models presented here to more complicated models. Some examples could be using non-Gaussian distributions, such as the student-t distribution, to better approximate the tails of the distribution, which is where we expect the strongest disagreement [37]. Other ideas include using the Edgeworth expansion to put bounds on the deviation from Gaussianity. As the tracks in individual maps are correlated (due to the

fact that some will overlap and use the same pixels), we could also consider generating correlated random values when generating our distributions for SNR_{tot} .

ACKNOWLEDGMENTS

We thank Bob Dobrow on helpful discussions about Order Statistics. M. C. is supported by National Science Foundation Graduate Research Fellowship Program, under NSF Grant No. DGE 1144152. E. T. was a member of the LIGO Laboratory, supported by funding from United States National Science Foundation. N. C.'s work is supported by NSF Grant No. PHY-1204371. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation and operates under cooperative agreement PHY-0757058. This paper has been assigned document number LIGO-P1500053.

-
- [1] J Aasi *et al.*, *Classical Quantum Gravity* **32**, 115012 (2015).
 - [2] F Acernese *et al.*, *Classical Quantum Gravity* **32**, 024001 (2015).
 - [3] J. Aasi *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. D* **87**, 022002 (2013).
 - [4] J. Abadie *et al.*, *Phys. Rev. D* **85**, 082002 (2012).
 - [5] J Abadie *et al.*, *Classical Quantum Gravity* **27**, 173001 (2010).
 - [6] A. L. Piro and E. Thrane, *Astrophys. J.* **761**, 63 (2012).
 - [7] A. L. Piro and C. D. Ott, *Astrophys. J.* **736**, 108 (2011).
 - [8] A. L. Piro and E. Pfahl, *Astrophys. J.* **658**, 1173 (2007).
 - [9] A. Corsi and P. Mészáros, *Astrophys. J.* **702**, 1171 (2009).
 - [10] K. Kiuchi, M. Shibata, P. J. Montero, and J. A. Font, *Phys. Rev. Lett.* **106**, 251102 (2011).
 - [11] M. H. P. M. van Putten, *Phys. Rev. Lett.* **87**, 091101 (2001).
 - [12] M. H. P. M. van Putten, *Astrophys. J. Lett.* **684**, L91 (2008).
 - [13] J. Aasi *et al.*, *Phys. Rev. D* **88**, 122004 (2013).
 - [14] E. Thrane, S. Kandhasamy, C. D. Ott *et al.*, *Phys. Rev. D* **83**, 083004 (2011).
 - [15] E. Thrane and M. Coughlin, *Phys. Rev. D* **88**, 083010 (2013).
 - [16] E. Thrane and M. Coughlin, *Phys. Rev. D* **89**, 063012 (2014).
 - [17] M. Coughlin, E. Thrane, and N. Christensen, *Phys. Rev. D* **90**, 083005 (2014).
 - [18] M. Coughlin, P. Meyers, E. Thrane, J. Luo, and N. Christensen, *Phys. Rev. D* **91**, 063004 (2015).
 - [19] W. G. Anderson, P. R. Brady, J. D. E. Creighton, and É. É. Flanagan, *Phys. Rev. D* **63**, 042003 (2001).
 - [20] P. Sutton *et al.*, *New J. Phys.* **12**, 053034 (2010).
 - [21] S. Klimentenko, I. Yakushin, A. Mercer, and G. Mitselmakher, *Classical Quantum Gravity* **25**, 114029 (2008).
 - [22] J. Abadie *et al.* (The LIGO Scientific Collaboration and The Virgo Collaboration), *Phys. Rev. D* **85**, 122007 (2012).
 - [23] J. Abadie *et al.*, *Astrophys. J.* **760**, 12 (2012).
 - [24] B. Allen, *Phys. Rev. D* **71**, 062001 (2005).
 - [25] L. Wen and B. F. Schutz, *Classical Quantum Gravity* **22**, S1321 (2005).
 - [26] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, N. Leroy, F. Robinet, and M. Vavoulidis, *Classical Quantum Gravity* **27**, 015005 (2010).
 - [27] M. Was, M.-A. Bizouard, V. Brisson, F. Cavalier, M. Davier, P. Hello, N. Leroy, F. Robinet, and M. Vavoulidis, *Classical Quantum Gravity* **27**, 194014 (2010).
 - [28] B. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Nature (London)* **460**, 990 (2009).
 - [29] J. Aasi *et al.* (LIGO and Virgo Collaboration), *Phys. Rev. Lett.* **113**, 231101 (2014).
 - [30] J. Abadie *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **107**, 271102 (2011).
 - [31] I. Foster, *Designing and Building Parallel Programs: Concepts and Tools for Parallel Software Engineering* (Addison Wesley, Reading, MA, 1995).
 - [32] K. Cannon, C. Hanna, and D. Keppel, *Phys. Rev. D* **88**, 024025 (2013).
 - [33] L. P. Singer, L. R. Price, B. Farr, A. L. Urban, C. Pankow, S. Vitale, J. Veitch, W. M. Farr, C. Hanna, K. Cannon *et al.*, *Astrophys. J.* **795**, 105 (2014).
 - [34] G. Farin, *Curves and Surfaces for CAGD, Fourth Edition: A Practical Guide* (Academic Press, New York, 1996).
 - [35] L. Blanchet, *Living Rev. Relativity* **9**, 4 (2006).
 - [36] T. Prestegard, E. Thrane, N. L. Christensen, M. W. Coughlin, B. Hubbert, S. Kandhasamy, E. MacAyeal, and V. Mandic, *Classical Quantum Gravity* **29**, 095018 (2012).
 - [37] C. Röver, *Phys. Rev. D* **84**, 122004 (2011).