Targeted coherent search for gravitational waves from compact binary coalescences

I. W. Harry[*] and S. Fairhurst[†]

School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff, CF24 3AA, UK
(Received 22 December 2010; published 4 April 2011)

We introduce a method for conducting a targeted, coherent search for compact binary coalescences. The search is tailored to be used as a follow-up to electromagnetic transients such as gamma-ray bursts. We derive the coherent search statistic for Gaussian detector noise and discuss the benefits of a coherent, multidetector search over coincidence methods. To mitigate the effects of nonstationary data, we introduce a number of signal consistency tests, including the null signal-to-noise ratio, amplitude consistency, and several χ^2 tests. We demonstrate the search performance on Gaussian noise and on data from LIGO's fourth science run and verify that the signal consistency tests are capable of removing the majority of noise transients, giving the search an efficiency comparable to that achieved in Gaussian noise.

I. INTRODUCTION

There has been excellent progress toward gravitational-wave astronomy over recent years. The first generation of large scale gravitational-wave interferometers reached unprecedented sensitivities and have undertaken extended science runs. The U.S. Laser Interferometer Gravitational-wave Observatory (LIGO) [1], the French-Italian Virgo [2], and the German-British GEO600 [3] detectors now form a collaborative network of interferometers. The data from these detectors have been analyzed for gravitational waves from compact binary coalescence [4], stochastic background [5], unmodeled burst [6], and pulsar [7] sources. LIGO's sixth science run (S6) and Virgo's second and third science runs (VSR2 and VSR3) ended in October 2010 and yielded the most sensitive data yet taken; the analysis of this data is ongoing. In the meantime, the detectors are being upgraded to their advanced configurations [8–10], with the expectation of a ten-fold improvement in sensitivity. With these sensitivities, it is expected that gravitational waves will be observed regularly [11]. Furthermore, with a proposed advanced detector in Japan [12], a possible detector in Australia [13], and third-generation detectors on the horizon [14], future prospects are promising.

As the gravitational-wave community matures it is essential that a relationship is built between gravitational-wave (GW) and electromagnetic (EM) astronomers. The GW emission from a source is likely to provide complementary information to emission in various EM bands, and a joint observation is significantly more likely to answer outstanding astrophysical questions. Already this relationship is beginning to mature. A number of EM transients have already been followed up in GW data [15–17].

Additionally, infrastructure is also being put in place to allow for EM follow-up of GW observations [18].

Compact binary coalescences (CBC) are one of the most promising sources of gravitational waves, and also an ideal candidate for joint GW-EM astronomy. During the late stages of inspiral and merger, a compact binary emits a distinctive, chirping gravitational-wave signal. Furthermore, CBCs containing at least one neutron star (NS) are expected to emit electromagnetically. Specifically, binary neutron stars (BNS) and neutron star-black hole binaries (NSBH) mergers are the preferred progenitor model for the short gamma-ray burst (GRB) [19,20]. It is also possible that these mergers will be observable electromagnetically as orphan afterglows [19], optical [21] or radio transients [22]. Since GRBs are well localized both in time and on the sky by EM observations, the corresponding GW search can be simplified by reducing the volume of parameter space relative to an all-sky, all-time search. Targeted searches for CBC waveforms associated to short GRBs were performed using data from LIGO's fifth science run (S5) and Virgo's first science run (VSR1) [15,16].

Coherent search methods have been introduced by numerous authors [23–25]. In coherent searches, data from all operational detectors are combined in a coherent manner before searching for a signal. Additionally, coherent analyses naturally impose the restriction that gravitational waves have only two independent polarizations. However, primarily for computational reasons, many searches have instead made use of a coincidence requirement—namely that a signal with consistent parameters is observed in two or more detectors in the network. Indeed, all previous LIGO and Virgo CBC search result papers have used a coincidence search [15,16]. Coincidence searches can approach the sensitivity of a fully coherent analysis, but will generally not achieve the same sensitivity.

---
[*]ian.harry@astro.cf.ac.uk
[†]stephen.fairhurst@astro.cf.ac.uk

In this paper, we introduce a targeted coherent analysis for CBC signals, appropriate for searching for GW from EM transients such as GRBs. We begin by deriving the coherent analysis for a templated CBC search. This has been presented in the literature previously [23]. We present an alternative derivation based on the $\mathcal{F}$-statistic formalism [26] introduced for continuous-wave analyses and now widely used. This allows for a more straightforward derivation of the coherent detection statistic. It also allows for a simple comparison to the coincident search and a straightforward derivation of the null stream [27] which by definition contains no gravitational-wave signal.

The data output by gravitational-wave interferometers are neither stationary nor Gaussian, but are contaminated by noise transients of instrumental and environmental origin. This makes the task of analyzing the data a complex one, and matched filtering alone is not sufficient to distinguish signal from noise—the most significant events by signal-to-noise ratio (SNR) would typically be due to non-Gaussian transients or "glitches" in the data. A significant effort goes into understanding the cause of these glitches [28] and removing times of poor data quality from the analysis. While these efforts greatly reduce the number of glitches they cannot remove them entirely. Therefore the analysis must also employ methods to distinguish signal from noise transients. In previous CBC searches, signal consistency tests [29,30] have proved very effective at removing the non-Gaussian background. We provide an overview of the formalism for these χ^2 consistency tests, and then extend these tests to the coherent analysis introduced earlier. We demonstrate that these tests continue to be effective in separating signal from nonstationary noise in the coherent analysis. In addition, coherent analyses naturally lend themselves to multidetector consistency tests, such as the null stream. We describe a number of consistency tests for the coherent CBC search and again demonstrate their efficacy.

This paper presents the first implementation of a modeled coherent detection search for CBC signals, which has been run on real LIGO data. We present results from running this analysis on both simulated data and real data from Gaussian data and real detector data taken from LIGO's fourth science run (S4). We are able to show that the signal consistency tests we have implemented are sufficient to remove the majority of non-Gaussian transients and render the search almost as sensitive as if the data were Gaussian and stationary. This search is available to be used to search for CBCs associated to GRBs in S6 and VSR2 and VSR3.

The layout of this paper is as follows. In Sec. II, we describe the formulation of a targeted coherent triggered search for CBC signals. In Sec. III we discuss an implementation of the null stream formalism and other multidetector consistency tests. In Sec. IV we describe a number of χ^2 tests that can be applied in a coherent search to try to

separate and veto glitches. Finally, in Sec. V we outline an implementation of a targeted, coherent search for CBC and present results on both simulated and real data.

II. COHERENT MATCHED FILTERING

In this section, we describe the coherent matched-filtering search for a gravitational-wave signal from a coalescing binary in data from a network of detectors. We restrict attention to binaries where the component spin can be neglected. The description is primarily tailored toward searches where the sky location of the gravitational-wave event is known *a priori*, as is appropriate when performing a follow-up of an EM transient such as a GRB [16,19]. Finally, since all previously published CBC search results [4,31,32] have used a coincidence search between detectors, we compare the coherent analysis with the multidetector coincident analysis.

The coherent analysis for coalescing binary systems has been derived previously using a similar method in [23,33–35]. Our presentation makes use of the $\mathcal{F}$-statistic formalism, introduced in [26]. This was originally defined as a method for performing searches for continuous-wave searches and has been regularly used for this task (see for example [36]). It was noted in [37] that the $\mathcal{F}$-statistic and the multiple detector inspiral statistic derived in [33] are similar and the $\mathcal{F}$-statistic was adapted to searches for CBC signals in [38].

A. The binary coalescence waveform

The generic binary coalescence waveform depends upon as many as 17 parameters. However, we restrict attention to binaries on circular orbits with nonspinning components. This reduces the parameter space to nine dimensions: the two component masses M_1 and M_2; the sky location of the signal (θ, ϕ); the distance, D, to the signal; the coalescence time of the signal, t_o; the orientation of the binary, given by the inclination ι, the polarization angle ψ and the coalescence phase ϕ_o. We also assume that the sky location (θ, ϕ) of the signal is known, thereby reducing the number of unknown parameters to seven.

In the radiation frame, where the gravitational wave propagates in the $\mathbf{e}_z^R$-direction, the gravitational waveform is given by

$$\mathbf{h} = h_+ \mathbf{e}_+ + h_\times \mathbf{e}_\times, \qquad (2.1)$$

where

$$\mathbf{e}_+ = \mathbf{e}_x^R \otimes \mathbf{e}_x^R - \mathbf{e}_y^R \otimes \mathbf{e}_y^R, \qquad \mathbf{e}_\times = \mathbf{e}_x^R \otimes \mathbf{e}_y^R + \mathbf{e}_y^R \otimes \mathbf{e}_x^R, \qquad (2.2)$$

and the waveforms $h_{+,\times}$ depend upon seven parameters $(M_1, M_2, t_o, D, \iota, \psi, \phi_o)$. The three angles (ι, ψ, ϕ_o) give the relationship between the radiation frame and the source frame (in which $\mathbf{e}_z^S$ lies in the direction of the binary's angular momentum and $\mathbf{e}_x^S$ along the separation between

the binary components at t_o). Even for a known sky location, it is necessary to search a seven-dimensional parameter space of signals. Naively covering this space with a grid of templates would be prohibitively costly [39]. However, the analysis is greatly simplified by the observation that the last four parameters enter only as amplitude parameters which can be analytically maximized over at minimal cost.[1] Specifically, the two polarizations of the waveform can be expressed as

$$h_+(t) = \mathcal{A}^1 h_0(t) + \mathcal{A}^3 h_{\pi/2}(t)$$
$$h_\times(t) = \mathcal{A}^2 h_0(t) + \mathcal{A}^4 h_{\pi/2}(t). \tag{2.3}$$

The two phases of the waveform are written as h_0 and $h_{\pi/2}$. These depend upon the physical parameters of the system (in this case just the masses) as well as the coalescence time t_o.[2] $\mathcal{A}^\mu$ are constant amplitude terms and are given explicitly as [35,38]

$$\mathcal{A}^1 = A_+ \cos2\phi_o \cos2\psi - A_\times \sin2\phi_o \sin2\psi$$
$$\mathcal{A}^2 = A_+ \cos2\phi_o \sin2\psi + A_\times \sin2\phi_o \cos2\psi$$
$$\mathcal{A}^3 = -A_+ \sin2\phi_o \cos2\psi - A_\times \cos2\phi_o \sin2\psi \tag{2.4}$$
$$\mathcal{A}^4 = -A_+ \sin2\phi_o \sin2\psi + A_\times \cos2\phi_o \cos2\psi,$$

where

$$A_+ = \frac{D_o}{D} \frac{(1 + \cos^2\iota)}{2} \qquad A_\times = \frac{D_o}{D} \cos\iota, \tag{2.5}$$

and D_o is a fiducial distance which is used to scale the amplitudes $\mathcal{A}^\mu$ and waveforms $h_{0,\pi/2}$. Thus, the amplitudes $\mathcal{A}^\mu$ depend upon the distance to the source and the binary orientation as encoded in the three angles (ι, ψ, ϕ_0). For *any* set of values $\mathcal{A}^\mu$, the expressions (2.4) can be inverted to obtain the physical parameters, unique up to reflection symmetry of the system [38].

The gravitational waveform observed in a detector X is

$$h^X = h^{ij} D_{ij}^X, \tag{2.6}$$

where D_{ij}^X denotes the detector response tensor. For an interferometric detector, in the long wavelength limit, the response tensor is given by

$$\mathbf{D}^X = (\mathbf{e}_x^X \otimes \mathbf{e}_x^X - \mathbf{e}_y^X \otimes \mathbf{e}_y^X), \tag{2.7}$$

where the basis vectors $\mathbf{e}_x^X$ and $\mathbf{e}_y^X$ point in the directions of the arms of the detector. It is often convenient to reexpress the gravitational-wave signal observed in a given detector as

$$h^X(t) = F_+(\theta^X, \phi^X, \chi^X) h_+(t) + F_\times(\theta^X, \phi^X, \chi^X) h_\times(t), \tag{2.8}$$

where the detector response to the two polarizations of the gravitational wave is encoded in the functions

$$F_+(\theta, \phi, \chi) = -\frac{1}{2}(1 + \cos^2\theta) \cos2\phi \cos2\chi$$
$$- \cos\theta \sin2\phi \sin2\chi, \tag{2.9}$$

$$F_\times(\theta, \phi, \chi) = \frac{1}{2}(1 + \cos^2\theta) \cos2\phi \sin2\chi$$
$$- \cos\theta \sin2\phi \cos2\chi. \tag{2.10}$$

These response functions depend upon the three angles $(\theta^X, \phi^X, \chi^X)$ which relate the detector frame to the radiation frame: θ^X and ϕ^X give the sky location relative to the detector, while χ^X is the polarization angle between the detector and the radiation frames. We have, somewhat unconventionally, allowed for a polarization angle in transforming from source to radiation *and* radiation to detector coordinates. In what follows, we will often find it convenient to fix the angle χ^X by explicitly tying it to the detector (or equatorial) frame; for example, by maximizing the detector (or network) sensitivity to the + polarization. The angle ψ then describes the orientation of the source with respect to this preferred radiation frame.

Since we are considering CBC observed in ground-based detectors, the time that a potential signal would spend in the sensitivity band of any detector will be short (less than 60 s for the initial detectors). Thus the change in the source's sky location over the observation time may be neglected, and the angles $(\theta^X, \phi^X, \chi^X)$ can be treated as constants. When working with a network of detectors, it is often useful to work in the equatorial frame. The location of the source (θ, ϕ, χ) is measured relative to this frame, and coalescence time is measured at the Earth's center. In this case, the location and orientation of the detector X are specified by three angles, which we denote $\vec{\alpha}^X$, and the detector response will depend upon six angles $(\vec{\alpha}^X, \theta, \phi, \chi)$. Then, the observed signal in a given detector is[3]

$$h^X(t) = F_+(\vec{\alpha}^X, \theta, \phi, \chi) h_+(t^X) + F_\times(\vec{\alpha}^X, \theta, \phi, \chi) h_\times(t^X), \tag{2.11}$$

where t^X is the time of arrival of the signal at detector X,

$$t^X = t - dt(\vec{\alpha}^X, \theta, \phi, \chi), \tag{2.12}$$

and dt gives the difference in arrival time of the signal between the geocenter and detector, for the given sky position.

---

[1]This was observed for the inspiral signal in [33] and independently for continuous-wave signals in [26].
[2]This decomposition is actually valid for all binaries in which the plane of the orbit does not precess. Thus, binary coalescence waveforms in which the spins are aligned with the orbital angular momentum can also be expressed in this form. However, for generic spin configurations, the orbit will precess and this simple decomposition is no longer applicable

---

[3]We do not give the explicit formula for the response function dependent on the six angles $(\vec{\alpha}^X, \theta, \phi, \chi)$, as the expression is somewhat lengthy. It can be obtained by performing six successive rotations to the detector response tensor to transform from the detector frame, via the equatorial frame, to the radiation frame. The calculation is detailed in [40].

Combining the final expressions for the binary coalescence waveform (2.3) and the detector response (2.11), we can express the gravitational waveform observed in a given detector as

$$h^X(t) = \mathcal{A}^\mu(D, \psi, \phi_o, \iota)h_\mu^X(t), \quad (2.13)$$

where the $\mathcal{A}^\mu$ are defined in (2.4), h_μ^X are given by

$$
\begin{aligned}
h_1^X(t) = F_+^X h_0(t^X) \qquad & h_2^X(t) = F_\times^X h_0(t^X) \\
h_3^X(t) = F_+^X h_{\pi/2}(t^X) \qquad & h_4^X(t) = F_\times^X h_{\pi/2}(t^X),
\end{aligned}
\quad (2.14)
$$

and we use the standard summation convention over the repeated index μ.

B. Multidetector binary coalescence search

Matched-filtering theory [41] provides a method for determining whether the signal $h(t, \xi)$, parametrized by the time and other parameters ξ, is present in a noisy data stream. The data output by a detector are

$$s^X(t) = n^X(t) + h^X(t, \xi), \quad (2.15)$$

where $n^X(t)$ is the noise, taken to be Gaussian and stationary. The noise $n^X(t)$ of the detectors is characterized by the noise power spectral density (PSD) $S_h^X(f)$ as

$$\langle \tilde{n}^X(f)[\tilde{n}^X(f')]^\star \rangle = \delta(f - f')S_h^X(f). \quad (2.16)$$

With this, we define the single detector inner product between two time series $a(t)$ and $b(t)$ as

$$(a^X|b^X) = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}^X(f)[\tilde{b}^X(f)]^\star}{S_h^X(f)}. \quad (2.17)$$

Then, the likelihood ratio of there being a signal h present in the data is given by

$$\Lambda(h) = \frac{P(s|h)}{P(s|0)} = \frac{e^{-(s^X - h^X|s^X - h^X)/2}}{e^{-(s^X|s^X)/2}}. \quad (2.18)$$

For a known signal h, with no unknown parameters, in Gaussian noise, the likelihood ratio is optimal in the Neyman-Pearson sense. However, this statistic is not optimal in the presence of non-Gaussian noise, as is discussed in greater detail in Secs. III and IV. It is often more convenient to work with the log-likelihood,

$$\log \Lambda = (s|h) - \frac{1}{2}(h|h), \quad (2.19)$$

and we will do so.

The likelihood ratio for multiple detectors is a straightforward generalization of the single detector expression (2.18). Assuming that the noise in different detectors is independent, in the sense that

$$\langle \tilde{n}^X(f)[\tilde{n}^Y(f')]^\star \rangle = \delta^{XY}\delta(f - f')S_h^X(f), \quad (2.20)$$

the multidetector inner product is simply given by the sum of the single detector contributions[4]

$$(\mathbf{a}|\mathbf{b}) := \sum_X (a^X|b^X). \quad (2.21)$$

The multidetector log-likelihood is given by

$$\ln\Lambda = (\mathbf{s}|\mathbf{h}) - \frac{1}{2}(\mathbf{h}|\mathbf{h}). \quad (2.22)$$

Specializing to the case of binary coalescence, we can substitute the known waveform parametrization (2.13) into the general matched filter likelihood ratio (2.22). The multidetector likelihood ratio becomes

$$\ln\Lambda = \left[\mathcal{A}^\mu(\mathbf{s}|\mathbf{h}_\mu) - \frac{1}{2}\mathcal{A}^\mu \mathcal{M}_{\mu\nu}\mathcal{A}^\nu\right], \quad (2.23)$$

where the matrix $\mathcal{M}_{\mu\nu}$ is defined as

$$\mathcal{M}_{\mu\nu} := (\mathbf{h}_\mu|\mathbf{h}_\nu). \quad (2.24)$$

The derivative of (2.23) with respect to $\mathcal{A}^\mu$ provides the values of $\mathcal{A}^\mu$ which maximize the likelihood ratio as

$$\hat{A}^\mu = [\mathcal{M}^{\mu\nu}(\mathbf{s}|\mathbf{h}_\nu)], \quad (2.25)$$

where, following [42], we take $\mathcal{M}^{\mu\nu}$ to be the inverse of $\mathcal{M}_{\mu\nu}$. We then define the maximized "coherent SNR" via the maximum likelihood ratio as

$$\rho_{\text{coh}}^2 := 2\ln\Lambda|_{\max} = [(\mathbf{s}|\mathbf{h}_\mu)\mathcal{M}^{\mu\nu}(\mathbf{s}|\mathbf{h}_\nu)]. \quad (2.26)$$

An identical maximization is performed in the derivation of the $\mathcal{F}$-statistic, used in searches for gravitational waves from asymmetric neutron stars [26]. In that context, the quantity obtained in (2.26) is typically denoted $2\mathcal{F}$. However, to make a closer connection to previous CBC search methods, we denote this quantity ρ_{coh}^2.

It is not difficult to show that ρ_{coh}^2 follows a χ^2 distribution with 4 degrees of freedom in the absence of a signal, and a noncentral χ^2 distribution (again with 4 degrees of freedom) when a signal is present. See, for example, [42] for more details. Furthermore, ρ_{coh}^2 is now a function of only the waveform components h_μ and no longer the $\mathcal{A}^\mu$ parameters. Thus four of the original seven waveform parameters have been analytically maximized, leaving three to be searched over.

Calculating the maximized likelihood ratio, as well as estimating the parameters $\hat{\mathcal{A}}^\mu$ requires an inversion of the matrix $\mathcal{M}_{\mu\nu}$. CBC signals will spend a large number of cycles in the sensitive band of the detector and consequently the 0 and $\frac{\pi}{2}$ phases will be (close to) orthogonal. Since the frequency evolves slowly, the amplitudes of the two phases will be close to equal,[5] i.e.,

---

[4]Note that we will explicitly write out the summation over detectors X, and *do not* use implicit summation over these indices.

[5]Indeed, several CBC waveforms are generated directly in the frequency domain [43], making these equalities exact.

$$(h_0^X|h_{\pi/2}^X) \approx 0 \quad (h_{\pi/2}^X|h_{\pi/2}^X) \approx (h_0^X|h_0^X) =: (\sigma^X)^2. \quad (2.27)$$

Therefore, the matrix $\mathcal{M}$ simplifies to

$$\mathcal{M}_{\mu\nu} = \begin{pmatrix} A & C & 0 & 0 \\ C & B & 0 & 0 \\ 0 & 0 & A & C \\ 0 & 0 & C & B \end{pmatrix}, \quad (2.28)$$

where

$$A = \sum_X (\sigma^X F_+^X)^2$$
$$B = \sum_X (\sigma^X F_\times^X)^2 \quad (2.29)$$
$$C = \sum_X (\sigma^X F_+^X)(\sigma^X F_\times^X).$$

1. Dominant polarization

Since we have included a polarization angle in both the transformation between equatorial and radiation frame (χ) and between radiation and source frame (ψ), we have the freedom to specify one of these without placing any physical restriction on the signal. The coherent SNR is further simplified by introducing a dominant polarization (DP) frame which renders $\mathcal{M}_{\mu\nu}$ diagonal.

Under a rotation of the radiation frame by an angle χ^{DP}, the detector response functions transform as

$$F_+^X \rightarrow F_+^{DP,X} = F_+^X \cos 2\chi^{DP} + F_\times^X \sin 2\chi^{DP}$$
$$F_\times^X \rightarrow F_\times^{DP,X} = -F_+^X \sin 2\chi^{DP} + F_\times^X \cos 2\chi^{DP}. \quad (2.30)$$

The rotation through χ^{DP} will have an identical effect on all detectors. Thus, there exists a polarization angle χ^{DP} which satisfies

$$C^{DP} = \sum_X (\sigma^X F_+^{DP,X})(\sigma^X F_\times^{DP,X}) = 0. \quad (2.31)$$

This can be solved to give χ^{DP} as

$$\tan 4\chi^{DP} = \frac{2\sum_X (\sigma^X F_+^X)(\sigma^X F_\times^X)}{\sum_X [(\sigma^X F_+^X)^2 - (\sigma^X F_\times^X)^2]}. \quad (2.32)$$

This choice serves to diagonalize the matrix $\mathcal{M}$. To uniquely determine χ^{DP}, we impose an additional requirement that the network be more sensitive to the $+$ polarization than to the $\times$ polarization. The value of χ^{DP} is a function of the detector network, the source location and waveform; in particular, it depends upon $F_{+,\times}^X$ and σ^X. From now on, we assume that we are working in the dominant polarization frame and drop the DP superscript from our expressions.

The concept of the dominant polarization frame has been introduced previously in unmodeled burst searches [25,44,45]. While the idea is very similar, the actual implementation is somewhat different.

In the case that both A and B are nonzero, i.e., that the detector has some sensitivity to both polarizations, the coherent SNR can be written, in the dominant polarization, as

$$\rho_{coh}^2 = \frac{(\mathbf{s}|\mathbf{F}_+\mathbf{h_0})^2 + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_+\mathbf{h_0}|\mathbf{F}_+\mathbf{h_0})}$$
$$+ \frac{(\mathbf{s}|\mathbf{F}_\times\mathbf{h_0})^2 + (\mathbf{s}|\mathbf{F}_\times\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_\times\mathbf{h_0}|F_\times\mathbf{h_0})}. \quad (2.33)$$

The coherent SNR can then be seen to arise as the quadrature sum of the power in the two phases of the waveform (0 and $\frac{\pi}{2}$) in the two gravitational-wave polarizations ($+$ and $\times$).

2. Synthetic $+$ and $\times$ detectors

In the dominant polarization the coherent SNR is comprised of separate $+$ and $\times$ components, with no cross terms. We can go one step further and interpret the coherent SNR as arising from two synthetic detectors, one sensitive to only the $+$ polarization and one sensitive to only the $\times$ polarization. These synthetic detectors are most easily formed by combining the "over-whitened" data streams o^X from the various detectors, where

$$o^X(f) = \frac{s^X(f)}{S_h^X(f)}. \quad (2.34)$$

The over-whitened synthetic data streams are simply

$$o_{+,\times}(f) = \sum_X F_{+,\times}^X o^X(f); \quad (2.35)$$

and the power spectra for these over-whitened data streams are

$$S_{+,\times} = \left(\sum_X \frac{(F_{+,\times}^X)^2}{S_h^X(f)}\right)^{-1}. \quad (2.36)$$

Using this, the unwhitened synthetic data streams are given as[6]

$$s_{+,\times}(f) = \sum_X \frac{F_{+,\times}^X s^X(f)}{S_h^X(f)} \left(\sum_Y \frac{(F_{+,\times}^Y)^2}{S_h^Y(f)}\right)^{-1}. \quad (2.37)$$

In terms of these synthetic detectors the detection statistic becomes

$$\rho_{coh}^2 = \frac{(s_+|h_0)_+^2 + (s_+|h_{\pi/2})_+^2}{(h_0|h_0)_+}$$
$$+ \frac{(s_\times|h_0)_\times^2 + (s_\times|h_{\pi/2})_\times^2}{(h_0|h_0)_\times}, \quad (2.38)$$

---

[6]There is some ambiguity in fixing the overall normalization of the synthetic detectors. We require that our synthetic detectors have the same sensitivity to the two polarizations as the original network did by requiring $(h_0^{+,\times}|h_0^{+,\times})_{+,\times} = \sum_X (F_{+,\times}^X \sigma^X)^2$.

where the subscripts $+$, $\times$ on the inner products denote the fact that the power spectrum of the synthetic detectors is used in their evaluation.

3. Network degeneracy

In many cases, a detector network is much more sensitive to one gravitational-wave polarization than the other. In the extreme limit (e.g., colocated and coaligned detectors such as those at the Hanford site) the network is entirely insensitive to the second polarization. In the dominant polarization frame, the network becomes degenerate as $B \rightarrow 0$ or equivalently

$$\sum_X (\sigma^X F_\times^X)^2 \rightarrow 0. \qquad (2.39)$$

Thus the network will only be degenerate if $F_\times^X = 0$ for all detectors X. If the network is degenerate then it is easy to see that the detection statistic will be degenerate as well. In this case it is logical to remove the $\times$ terms from the detection statistic reducing it to

$$\rho_{\text{coh}}^2 = \frac{(\mathbf{s}|\mathbf{F}_+\mathbf{h_0})^2 + (\mathbf{s}|\mathbf{F}_+\mathbf{h}_{\pi/2})^2}{(\mathbf{F}_+\mathbf{h_0}|\mathbf{F}_+\mathbf{h_0})}, \qquad (2.40)$$

which is χ^2 distributed with 2 degrees of freedom.

In this formalism the coherent SNR changes abruptly from (2.40) to (2.33). If there is any sensitivity, no matter how small, to the $\times$ polarization, there is an entirely different detection statistic. This arises due to maximization over the $\mathcal{A}^\mu$ parameters, allowing them to take any value. Thus, even though a network may have very little sensitivity to the $\times$ polarization, and consequently there be little chance of observing the waveform in the $\times$ polarization, this is not taken into account in the derivation. A possible modification is to place an astrophysical prior on the parameters (D, ι, ψ, ϕ_0) and propagate this to the distribution of the $\mathcal{A}^\mu$ [46]. This would provide a smooth transition from the degenerate to nondegenerate search.

C. Comparison with coincident search

The single detector search is a special case of the degenerate network (2.40) and can be written as

$$\rho_X^2 = \frac{(s^X|h_0)^2 + (s^X|h_{\pi/2})^2}{(\sigma^X)^2}. \qquad (2.41)$$

A coincidence search requires a signal to be observed in two or more detectors, without requiring consistency of the measured waveform amplitudes in the different detectors. In many cases, coincidence searches have made use of different template banks in the different detectors [4,16,31] and required coincidence between the recovered mass parameters [47]. A comparison with the coherent analysis discussed above is facilitated if we consider a coincident search where an identical template is used in all detectors,

as was done in an analysis of early LIGO data [48]. In this case, the multidetector coincident SNR is given by

$$\rho_{\text{coinc}}^2 = \sum_X \rho_X^2 = \sum_X \frac{(s^X|h_0)^2 + (s^X|h_{\pi/2})^2}{(\sigma^X)^2}. \qquad (2.42)$$

This is not immediately comparable to the coherent SNR given in (2.33). However, both can be recast into similar forms by writing the coincident SNR as

$$\rho_{\text{coinc}}^2 = \sum_{X,Y} \sum_{i=0,\pi/2} \left(s^X \left| \frac{h_i}{\sigma^X}\right)[\delta^{XY}]\left(s^Y \left| \frac{h_i}{\sigma^Y}\right), \qquad (2.43)\right.\right.$$

and the coherent SNR can be written as

$$\rho_{\text{coh}}^2 = \sum_{X,Y} \sum_{i=0,\pi/2} \left(s^X \left| \frac{h_i}{\sigma^X}\right)[f_+^X f_+^Y + f_\times^X f_\times^Y]\left(s^Y \left| \frac{h_i}{\sigma^Y}\right),\right.\right.$$
$$(2.44)$$

where we have defined the orthogonal unit vectors (in detector space) f_+^X, $f_\times^X$ as

$$f_{+,\times}^X = \frac{\sigma^X F_{+,\times}^X}{\sqrt{\sum_Y (\sigma^Y F_{+,\times}^Y)^2}}. \qquad (2.45)$$

The SNR of the coincident search (2.43) is simply the sum of all power consistent with the template waveform in each detector. The coherent SNR (2.44) makes use of the fact that gravitational waves have only two polarizations to restrict the accumulated SNR to the physical subspace spanned by f^+ and $f^\times$. For a signal, the power will lie entirely in this subspace, while noise in the detectors will contribute to all components of the coincident SNR. Thus, the coherent analysis obtains precisely the same contribution to SNR from the gravitational wave signal, but has a reduced noise background. Specifically, the coherent SNR acquires contributions from four noise degrees of freedom, while the coincident SNR has $2N$ noise degrees of freedom, where N indicates the number of active detectors. For a nondegenerate two-detector network, the coincident and coherent SNRs are equal as in this case $f_+^X f_+^Y + f_\times^X f_\times^Y = \delta^{XY}$.

In the case where a network is sensitive to only one polarization, the coherent SNR is constructed solely from the f_+^X direction and coherent SNR is χ^2 distributed with 2 degrees of freedom.

Finally, we note that restricting to the coherent SNR can help to separate transients from gravitational-wave signals as those transients which do not contribute power to the signal space will be ignored. However, many noise transients will contribute to the coherent SNR and more active methods of removing them are required. These methods are the focus of Secs. III and IV.

III. SIGNAL CONSISTENCY BETWEEN DETECTORS

As discussed in the introduction, due to the presence of non-Gaussian noise transients, it is essential to make use of signal consistency requirements within search algorithms to distinguish glitches from gravitational-wave signals.

Multidetector analyses have made good use of signal consistency between detectors (see e.g. [25]). A particularly powerful test is the use of a ''null stream'' [27] which, by construction, contains no gravitational-wave signal. Many noise transients will contribute power to the null stream and can therefore be eliminated as candidate events. In addition, requiring that the gravitational-wave signal is recovered consistently between detectors can eliminate other noise transients; in our case this is equivalent to imposing restrictions on the recovered values of the parameters $\hat{\mathcal{A}}^{\mu}$. These two methods will be considered in turn. For matched-filtering searches, requiring consistency between the observed signal and template waveform has also proven very powerful [29]. A full description of waveform consistency tests is presented in the next section.

A. Null stream consistency

The gravitational waveform consists of two polarizations. Thus for networks comprising three or more detectors it is possible to construct one or more null data streams which contain no gravitational-wave signal [27]. In the context of a coherent search for CBC signals, the null consistency tests arise quite naturally. In Sec. II C, we noted that the coherent SNR can be thought of as a projection of the coincident multidetector SNR onto a four-dimensional signal subspace. The remaining dimensions in the coincident search do not contain any gravitational-wave signal, but will be subject to both Gaussian and non-Gaussian noise. Thus, we can define the null SNR as

$$\rho_{\mathrm{N}}^2 = \rho_{\mathrm{coinc}}^2 - \rho_{\mathrm{coh}}^2 = \sum_{X,Y} \sum_{i=0,\pi/2} \left(s^X \left| \frac{h_i}{\sigma^X} \right.\right) [N^{XY}] \left(s^Y \left| \frac{h_i}{\sigma^Y}\right.\right),$$

(3.1)

where

$$N^{XY} = \delta^{XY} - f_+^X f_+^Y - f_\times^X f_\times^Y.$$ (3.2)

A gravitational-wave signal matching the template h will provide no contribution to the null SNR, so we expect that, for signals, this statistic will be χ^2 distributed with $(2N - 4)$ degrees of freedom. A noise transient that is incoherent across the data streams may give a large coherent SNR, but it is likely to also give a large null SNR. Thus requiring a *small* null SNR will prove effective at distinguishing incoherent noise transients from real gravitational-wave signals. Since the definition of the null SNR (3.1) makes use of the template waveform, gravitational waveforms which do not match the template h can contribute to the null SNR.

We can go one step further and introduce synthetic null detectors in analogy with the synthetic $+$ and $\times$ detectors. For concreteness, we describe the three-detector case, but this can be extended in a straightforward manner to larger networks. When working with a three-detector network, there is a single null direction

$$n^X = \sum_{Y,Z} \epsilon_{XYZ} f_+^Y f_\times^Z,$$ (3.3)

where ϵ_{XYZ} denotes the Levi-Cevita symbol, and the projection onto the null space is given by $N^{XY} = n^X n^Y$. Then, the over-whitened synthetic null detector is

$$o_N(f) = \sum_X \frac{n^X}{\sigma^X} o^X(f).$$ (3.4)

The power spectrum of the null stream is[7]

$$S_N(f) = \left(\sum_X \frac{(n^X)^2}{(\sigma^X)^2 S_h^X(f)}\right)^{-1}$$ (3.5)

and the unwhitened null stream is

$$s_N(f) = \left(\sum_X \frac{n^X s^X(f)}{\sigma^X S_h^X(f)}\right) \cdot S_N(f).$$ (3.6)

Finally, the null SNR can be written as

$$\rho_{\mathrm{N}}^2 = \frac{(s_N|h_0)_N^2 + (s_N|h_{\pi/2})_N^2}{(h_0|h_0)_N}.$$ (3.7)

The null SNR described above differs from the multidetector null stream formalism introduced in [27] and used by several other authors. A null stream is constructed to be a data stream which contains no contribution from the h_+ and $h_\times$ gravitational waveforms, regardless of the details of the waveform. To provide a concrete comparison between the null stream and null SNR, we again restrict attention to a three detector network. The null stream is explicitly constructed as

$$s_{\mathrm{Null}}(f) = \sum_{X,Y,Z} \epsilon_{XYZ} F_+^X F_\times^Y s^Z(f).$$ (3.8)

By comparing the null stream in (3.8) with the synthetic null detector (3.6), it is clear that these will generically differ. To get an insight into the differences, consider a network with two colocated detectors A and B, with power spectra $S_h^A(f)$ and $S_h^B(f)$, respectively, and a third detector C which is sensitive to the other polarization of gravitational waves. For this network, the null stream will be a combination of only the A and B detector data. The power spectrum of the null stream is

$$S_{\mathrm{Null}}(f) = S^A(f) + S^B(f),$$ (3.9)

while for the synthetic null detector it is

$$\frac{1}{S_{\mathrm{N}}(f)} = \frac{1}{(\sigma^A)^2 S^A(f)} + \frac{1}{(\sigma^B)^2 S^B(f)}.$$ (3.10)

---

[7]In this case, there is a normalization ambiguity. For the synthetic plus and cross streams, it was natural to require that the synthetic detectors have the same sensitivities as the original network. For the null stream this is not feasible as the network has zero sensitivity to a signal in the null stream, so we normalize such that $(h_0|h_0)_N = 1$.

Thus, if the power spectra of detectors A and B are identical, then the two null streams are also identical. In the extreme case that the sensitivity bands of the two detectors do not overlap at all, then there is no null stream ($S_{\text{Null}} \to \infty$). However, the null SNR need not vanish and is similar to a two-bin version of the χ^2 test described in Sec. IV D. Thus, it is possible to construct scenarios in which these two null stream formulations differ significantly.

For the most part, the power spectra of the ground-based detectors are rather comparable. So, in general there will not be a significant difference between these two forms. There are advantages to both methods. The null stream is designed to cancel all gravitational-wave signals from the data, thus making it more robust when the signal is not well known. However, by making use of the template signal, there are instances in which the null SNR provides a more powerful consistency test. Furthermore, it has a computational benefit in that it does not require the production of a null stream—all manipulations are performed on the single-detector SNR data streams which are subsequently separated into coherent and null components. In practice we have found very little difference in performance, and choose to compute the null SNR (3.1) for computational simplicity.

Finally, we note that both null stream formalisms will perform optimally only if the three detectors have similar sensitivities. In the case where one detector is significantly less sensitive than the others, the null stream will generally tend to the data of that least sensitive detector. Also, the null formalisms described here will only completely cancel a gravitational-wave signal provided that the calibration of the data streams is accurate; any error in calibration will lead to a signal surviving in the null stream.

B. Amplitude consistency

The four amplitude parameters $\mathcal{A}^\mu$, encoding the distance to and orientation of the binary system, can take any values. Indeed, any set of $\mathcal{A}^\mu$ corresponds to unique values of the distance, inclination angle, coalescence phase, and polarization angle, up to symmetries of the system. However, some of these values will be significantly more likely to occur astrophysically than others. For example, the number of binary coalescence events is expected to be approximately proportional to star formation rate [11] and consequently should be roughly uniform in volume. Thus, events are more likely to occur at a greater distance. Similarly, the gravitational-wave amplitude, at a fixed distance, is greater for face-on signals than edge-on ones, as is clear from (2.5). Therefore, we are more likely to detect face-on signals at a large distance than nearby, edge-on ones. Consequently, certain values of $\mathcal{A}^\mu$ are astrophysically more likely than others. This can be taken into account in a Bayesian manner by marginalizing over an appropriate distribution for the $\mathcal{A}^\mu$ rather than performing a maximization [46].

The distribution of noise events will follow its own characteristic distribution. For Gaussian noise, the expected distributions of the $\mathcal{A}^\mu$ can be explicitly calculated. Nonstationarities in the data will again produce a different distribution of amplitudes which cannot be analytically modeled in a straightforward manner. Specifically, the majority of transients is caused by a disturbance or glitch in a single detector with little or no signal in the other detectors. For networks with three or more detectors, this will typically be inconsistent with a coherent signal across the network, leading to a large value of the null SNR. In certain scenarios, most notably for two-detector networks, there will be a consistent set of values for the $\mathcal{A}^\mu$. However, these values carry the characteristic signature of a glitch. Specifically, the SNR contributions will typically be consistent with a nearby, close to edge-on system ($A_\times \approx 0$), with a very specific orientation to provide essentially no response in all but one detector. Thus, the glitch distribution of the $\mathcal{A}^\mu$ parameters will be significantly different from the distribution expected for gravitational-wave signals. In the remainder of this section, we explore the possibility of making use of the extracted $\mathcal{A}^\mu$ parameters to distinguish between glitches and signals. Unlike the null stream, amplitude consistency tests are available for two detector networks. They should be especially useful in the case of the two 4 km LIGO instruments, which have similar sensitivities to the majority of points on the sky.

We have argued that the majority of gravitational-wave signals will originate from (close to) face-on binaries while the majority of noise transients will mimic (close to) edge-on binaries. The recovered value of the inclination angle ι should then serve to separate signals from noise. To investigate this, we simulated a large number of simulated CBC signals and a large number of noise glitches, added Gaussian noise, and plotted the recovered inclination angle in Fig. 1. The glitches were generated as events with a large SNR in one detector coincident with Gaussian noise in a second detector. The signals were separated into two groups: the first with only face-on binaries ($|\cos\iota| = 1$) and the second a uniform distribution over the two sphere (uniform in $\cos\iota$ and ψ) of the binary orientation. In both cases, they were distributed uniformly in volume and orbital phase. We also consider two different network configurations, both containing two equally sensitive detectors. In the first case one detector is sensitive to $+$ and the other to $\times$ polarization; in the second case both detectors have strong and equal sensitivity to the $+$ polarization and weak but opposite sensitivity to the $\times$ polarization—rather typical for the Hanford-Livingston network. For both sets of signals and choices of network, there is a clear distinction between signal and glitch distribution. However, there is a clear downward bias on the recovered values of ι. This can be understood by looking at the expressions for A_+ and $A_\times$. For face-on binaries, these
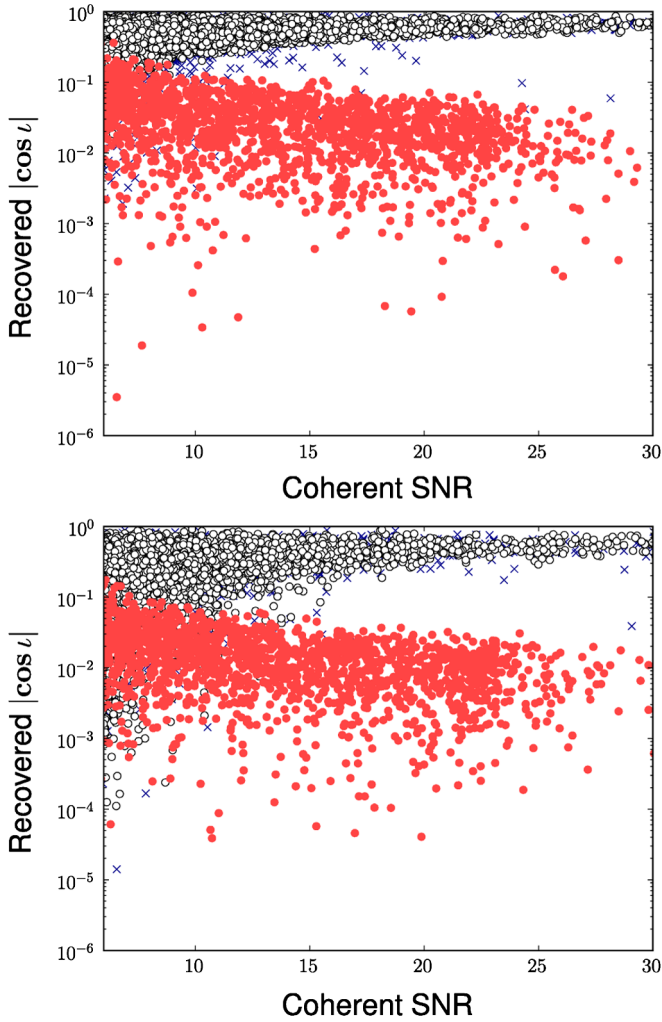
FIG. 1 (color online). The distribution of the recovered inclination angle plotted against coherent SNR for optimally oriented signals (open circles), uniformly distributed orientations (blue crosses), and simulated glitches (red solid circles). The top panel shows a network configuration which is equally sensitive to both gravitational-wave polarizations. The bottom panel shows a configuration which is 5 times more sensitive to the + polarization than to the ×.

will be equal but, in the presence of noise, $A_\times$ will be reconstructed to be somewhat smaller than A_+. A relative difference of only 5% leads to a recovered inclination of 45°, so even for loud signals there can be large discrepancy between the actual and recovered inclination angle.

Despite the difference in distribution between signal and noise, there is also a significant overlap of the populations at low SNRs. Consequently, any threshold imposed on the recovered inclination angle is liable to either reject a fraction of signals or pass a fraction of glitches. It is, however, quite possible that knowledge of these expected distributions could be folded into the detection statistic in a Bayesian manner.

We have found that using the observed SNR in the individual instruments to be a more effective discriminator

of signal and noise. To demonstrate the efficacy of such an approach, in Fig. 2 we plot the single-detector SNR as a function of the coherent SNR for the same population of glitches and the two classes of signals (face-on and uniformly distributed orientation) described above. The glitches fall into two groups depending upon which detector suffered the glitch. Since our model detectors are equally sensitive, then on average one expects each detector to accrue $1/\sqrt{2}$ of the coherent SNR. Even allowing for nonoptimally oriented signals and the addition of Gaussian


FIG. 2 (color online). The distribution of single-detector SNR plotted against coherent SNR for optimally oriented signals (open circles), uniformly oriented signals (blue crosses), and glitches (red solid circles). The top panel shows a network of two equally sensitive detectors, where one detector sees only the + polarization and the second detector sees only the × polarization. The bottom panel shows a similar network where both detectors have strong and equal sensitivity to the + polarization and weak but opposite sensitivity to the × polarization. This is typical for the Hanford-Livingston network. The diagonal solid black line shows the expected SNR for the optimally oriented signals. The horizontal dashed black line indicates a SNR of 4.

noise, the signals follow this expectation. Only a small number of signals are found with SNRs inconsistent with the expected values; these are ones that have very specific orientations. Overall, the signal and glitch populations are very well separated, at least until the coherent SNR becomes rather small.

The most effective strategy we have found is to require that all events have a SNR above 4 in the two most sensitive detectors in the network. The cut is illustrated in Fig. 2. This strategy removes the majority of glitch signals while having a negligible effect on the signal population at large SNR. For lower SNR the signals which are lost due to this cut would be unlikely to be detection candidates as Gaussian noise alone produces similar events.

IV. COHERENT χ^2 TESTS

Data from gravitational-wave detectors contain numerous nonstationarities due to both instrumental and environmental causes. These nonstationarities, or glitches, typically do not match well with the CBC waveform. However, they often contain enough power that, even though the match with the template is poor, a large SNR is observed. In the previous section, we have seen how the use of various coherent consistency tests can mitigate this problem. Additionally, a number of other signal consistency tests have been implemented [29,30] and used in searches for CBC signals [4,31,32]. These tests are all designed to eliminate glitches which have a different signal morphology than the template waveform. This is essentially done by testing whether the detector data orthogonal to the signal are well described as Gaussian and stationary—for a glitch, there will be residual power which does not match the template waveform. These tests are commonly known as "χ^2 tests" as they construct a statistic which is χ^2 distributed in the presence of Gaussian noise plus a signal matching the template waveform. If the data contain a glitch, the χ^2 statistic will generally have a large value, thereby allowing for differentiation of signal from nonstationary noise. In this section, we briefly review the general formulation of χ^2 tests before presenting a detailed description of three such tests which have been implemented for the coherent search described in Sec. II.

A. A general framework for χ^2 tests

Consider the data from a gravitational-wave detector at a time t which has produced a large SNR when filtered against a template $h(t)$. Generically, the data $s(t)$ can be decomposed as

$$s(t) = n(t) + Ah(t) + Bg(t), \qquad (4.1)$$

where $n(t)$ represents a Gaussian noise component, $h(t)$ is the template waveform, $g(t)$ is an additional non-Gaussian noise contribution to the data stream, and A and B are

amplitude factors. The glitch contribution $g(t)$ is taken to be the power orthogonal to $h(t)$ and both $g(t)$ and $h(t)$ are normalized, so that

$$(g|g) = 1, \qquad (h|h) = 1, \qquad (g|h) = 0. \qquad (4.2)$$

In order to construct a χ^2 test, we must introduce an additional set of waveforms T^i. These waveforms are required to be orthonormal and orthogonal to h,

$$(h|T^i) = 0, \qquad (T^i|T^j) = \delta^{ij}. \qquad (4.3)$$

Furthermore, for the χ^2 test to be effective, the T^i must have a good overlap with the glitch waveform $g(t)$.

The χ^2 discriminator is constructed as

$$\chi^2 = \sum_{i=1}^{N} (T^i|s)^2. \qquad (4.4)$$

When the data comprise only signal plus Gaussian noise, i.e. $B = 0$ in Eq. (4.1),

$$\chi^2 = \sum_{i=1}^{N} (T^i|n)^2 \qquad (4.5)$$

and the statistic is the sum of squares of independent Gaussian variables with zero mean and unit variance. Thus the test is χ^2 distributed with N degrees of freedom, with a mean and variance of

$$\langle \chi^2 \rangle = N, \qquad \mathrm{Var}(\chi^2) = 2N. \qquad (4.6)$$

This is true for *any* set of waveforms T^i given the above assumptions.

In the case where the data are not an exact match to the signal, we take both A and B nonzero, i.e. any signal or glitch can be decomposed into a part $Ah(t)$ proportional to the template under consideration plus a second orthogonal contribution $Bg(t)$. Clearly, for different glitches, the waveform $g(t)$ as well as the amplitude factor B will be different. In this case the χ^2 test takes the form

$$\chi^2 = \sum_{i=1}^{N} [(T^i|n)^2 + 2B(T^i|n)(T^i|g) + B^2(T^i|g)^2]. \quad (4.7)$$

This has a mean

$$\langle \chi^2 \rangle = N + B^2 \sum_{i}^{N} (T^i|g)^2 \qquad (4.8)$$

and a variance

$$\mathrm{Var}(\chi^2) = 2N + 4B^2 \sum_{i} (T^i|g)^2. \qquad (4.9)$$

The χ^2 test is distributed as a noncentral χ^2 distribution with N degrees of freedom and a noncentrality parameter [29]

$$\lambda = B^2 \sum_{i=1}^{N} (T^i|g)^2. \qquad (4.10)$$

The challenge in constructing a χ^2 test is to select the basis waveforms T^i such they have large overlaps with the observed glitches in the data. If this is done successfully, then any glitch producing a large SNR will also give a large value of χ^2, inconsistent with a signal in Gaussian noise.

In many cases, there is some uncertainty in the template waveform. For example, the post-Newtonian (PN) expansion used in generating CBC waveforms is truncated at a finite (typically 3 or 3.5 PN [49]) order and there will be differences between this analytically calculated waveform and the one provided by nature. There are similar uncertainties in waveforms obtained from numerical relativity simulations [50]. Additionally, to search the full parameter space of coalescing binaries, a discrete template bank is used which allows for some mismatch between the templates and any potential signal within the parameter space [39]. Normally the template bank is created so that the mismatch is no larger than 3% at any point in the parameter space. Finally, there are uncertainties in instrumental calibration [51] which will affect the match between signal and template.

We model these effects by parametrizing the signal as

$$H(t) = A\left(\sqrt{1 - \epsilon^2}h(t) + \epsilon m(t)\right), \qquad (4.11)$$

where $m(t)$ is the component of H that is orthogonal to h [$(m|h) = 0$] and ϵ encodes the mismatch between signal and template in the sense that

$$1 - \frac{(H|h)}{\sqrt{(H|H)(h|h)}} = 1 - \sqrt{1 - \epsilon^2} \approx \epsilon. \qquad (4.12)$$

In most cases, it is reasonable to assume a mismatch of less than 5%. The obvious counterexample is when searching for highly spinning systems using nonspinning waveforms; see e.g. [52,53].

Since (4.11) is a special case of (4.1) it follows directly that the mean and variance of the χ^2 test in the presence of a mismatched signal are

$$\langle \chi^2 \rangle = N + A^2 \epsilon^2 \sum_{i=1}^{N} (T^i|m)^2$$
$$\qquad\qquad\qquad\qquad (4.13)$$
$$\mathrm{Var}(\chi^2) = 2N + 4A^2 \epsilon^2 \sum_{i=1}^{N} (T^i|m)^2.$$

As the SNR of the signal is proportional to A, the expected χ^2 value for a mismatched signal increases with the strength of the signal. However, for mismatched signals $\chi^2 \propto \epsilon^2 A^2$ while for glitches $\chi^2 \propto B^2$ and provided $\epsilon A \ll B$ the two can be separated. See [29] for a more detailed discussion.

When introducing the χ^2 test, we assumed that the T^i were orthonormal and orthogonal to the template waveform h. In practice, this can be difficult to guarantee. The signal consistency tests discussed in the remainder of this section are constructed from gravitational waveforms.

If one picks a set of gravitational waveforms, t^i, there is no guarantee that they will be either orthonormal or orthogonal to h. We can, at least, construct waveforms which are orthogonal to h by introducing

$$T^i = \frac{t^i - (t^i|h)h}{\sqrt{1 - (t^i|h)^2}}. \qquad (4.14)$$

While this ensures $(h|T^i) = 0$ it does not guarantee orthonormality of the T^i, $(T^i|T^j) = \delta^{ij}$. Thus this method will not produce a χ^2 distribution, and will instead form a generalized χ^2 distribution. The mean of the distribution remains N but the variance is increased,

$$\mathrm{Var}(\chi^2) = 2N + 2\sum_{i \neq j}(T^i|T^j)^2. \qquad (4.15)$$

It has been found, however, that this does not present a significant obstacle to using these tests, especially as the thresholds are tuned empirically [54].

Multidetector χ^2 tests

In Sec. II, we derived a coherent multidetector search for coalescing binaries. The search involves filtering four waveform components h_μ against the multidetector data stream. Our initial discussion of χ^2 tests was limited to the description of a single phase template waveform h and test waveforms T^i. The extension to a two-phase waveform has been described previously [29] and here we extend that to a four-component waveform across multiple detectors, as is appropriate for this search. We begin by noting that the four waveform components h_μ are orthogonal in the dominant polarization basis. They are, however, not generally normalized, as

$$(\mathbf{h}_\mu|\mathbf{h}_\nu) = \mathcal{M}_{\mu\nu} = \mathrm{diag}(A, B, A, B), \qquad (4.16)$$

where A and B are defined in (2.29). Thus, we first normalize so that

$$(\hat{\mathbf{h}}_\mu|\hat{\mathbf{h}}_\nu) = \delta_{\mu\nu}. \qquad (4.17)$$

To construct a network χ^2 test, we require a set of (four-component) normalized, test waveforms $\hat{t}^i_\mu$. The components

$$T^i_\mu = \frac{\hat{t}^i_\mu - \sum_\nu (\hat{t}^i_\mu|\hat{\mathbf{h}}_\nu)\hat{h}_\nu}{\sqrt{1 - \sum_\sigma (\hat{t}^i_\sigma|\hat{\mathbf{h}}_\sigma)^2}}, \qquad (4.18)$$

constructed to be orthogonal to h_μ are used in the χ^2 test. Thus, the coherent, multidetector χ^2 test is

$$\chi^2 = \sum_{\mu=1}^{4} \sum_{i=1}^{N} (\mathbf{T}^i_\mu|\mathbf{s})^2. \qquad (4.19)$$

Provided the test waveforms are orthonormal, in the sense that

$$(\mathbf{T}^i_\mu|\mathbf{T}^j_\nu) = \delta^{ij}\delta_{\mu\nu}, \qquad (4.20)$$

the distribution for a signal matching h_μ plus Gaussian noise will be χ^2 distributed with $4N$ degrees of freedom. As for the single phase filter, we cannot always guarantee (4.20) is satisfied, although it is relatively simple to ensure the four components of a given template *are* orthogonal. This means that the statistic will not, in general, be χ^2 distributed: The mean remains $4N$ but the variance increases to

$$\text{Var}\,(\chi^2) = 8N + 2\sum_{i,j=1}^{N}\sum_{\mu,\nu=1}^{4}[(T_\mu^i|T_\nu^j)^2 - \delta^{ij}\delta_{\mu\nu}]. \quad (4.21)$$

B. The coherent bank χ^2 test

The bank χ^2 test was designed to test the consistency of the observed SNR across different templates in the bank at the time of a candidate signal. It was first described in [30] for the case of a single detector. A glitch will typically cause a high SNR in many templates across the bank, while a real signal will give a well-prescribed distribution of SNR across the template bank.

The bank χ^2 makes use of other CBC templates as the waveforms t^i to construct the χ^2 test. These N templates are taken from different points across the mass space. In implementing the bank χ^2, we choose a *fixed* set of template waveforms t^i which remain the same for every template h in the search template bank. The bank χ^2 statistic is then constructed following (4.18) and (4.19). The test is most effective when the set of T_μ^i is close to orthogonal [30] so we select templates which are well distributed


FIG. 3. The distribution of the bank χ^2 test for a single template h, with a bank of size 10. The plot shows the distribution of the bank veto calculated for every time sample in 128 s of simulated Gaussian data (with no signal present). In the case that the ten bank templates are orthogonal, the expected distribution is χ^2 with 40 degrees of freedom (shown as the solid black line). As can be seen, the actual distribution follows the expected one closely.

across the mass space, ensuring the overlaps $(T_\mu^i|T_\nu^j)$ are small for $i \neq j$. Figure 3 shows the distribution of the bank χ^2 for a single template filtered against Gaussian noise. The set of fixed bank waveforms consisted of ten waveforms distributed over the full mass parameter space. Using these waveforms, the deviation from a χ^2 distribution is negligible.

For the bank χ^2 to be effective, glitches in the data must have a good overlap with a reasonable fraction of the templates t^i. While, in general, it is difficult to predict the composition of glitches in the data, it seems reasonable to assume that glitches which produce a large SNR for the template h will also have a good overlap with other waveforms in the template space. Thus, the set of templates which is spread across the parameter space is suitable.

C. The coherent autocorrelation χ^2 test

Filtering a gravitational-wave template against data containing a matching signal produces a peak in the SNR at the time of the signal. Furthermore, there is a characteristic shape of this peak which depends upon the template waveform and also the noise power spectrum of the data. An example of this autocorrelation for a BNS template is shown in Fig. 4. A noise transient in the data will produce a peak in the SNR but it will typically lack the characteristic shape produced by a genuine CBC signal.

The "auto" χ^2 test was designed to test the consistency of the SNR peak [30]. It is a similar test to the bank χ^2, but where the bank χ^2 investigates consistency in SNR across the mass space, the auto χ^2 tests for consistency of the SNR time series. The set of templates t^i are chosen to be the original template h with time shifts δt^i applied. The values of δt^i are all unique and chosen to be of the same time scale as the autocorrelation of the template waveform


FIG. 4 (color online). The single detector autocorrelation of a gravitational-wave inspiral signal from a 1.4, 1.4 solar mass BNS. Both phases of the waveform are shown.

FIG. 5.   The distribution of the auto χ^2 test for a single template h, generated with 40 time-shifted templates, with shifts between 0.001 and 0.04 s. The plot shows the distribution of the auto veto calculated for every time sample in 128 s of simulated Gaussian data (with no signal present). In the case that the 40 time-shifted templates are orthogonal, the expected distribution is χ^2 with 160 degrees of freedom (shown in black). As can be seen, the actual distribution differs significantly from this due to the nonorthogonality of the t^i waveforms.

(typically 0.1 s or less) and the duration of nonstationarities in the data, which is similar.

In Fig. 5, we show the distribution of the auto χ^2 for a single template waveform filtered in Gaussian data. For this result, forty waveforms t^i were used, equally spaced with a 1 ms spacing, and all with coalescence times prior to that of h. Thus, the auto χ^2 is testing the consistency of the SNR time series for 0.04 s prior to the SNR peak. The overlap $(t^i|t^j)$ depends only upon the difference $\delta t^i - \delta t^j$ and Fig. 4 shows clearly that a significant fraction of the overlaps are far from zero. Consequently, the auto χ^2 test has a distribution with a large deviation from a χ^2 distribution with $4N$ degrees of freedom.

D. The coherent χ^2 test

The "standard" χ^2 test originally proposed in [29] has been used as a discriminator in many gravitational-wave searches for CBCs. Given the template waveforms and the detector sensitivity, it is possible to predict the accumulation of SNR as a function of frequency. By calculating the observed SNR contribution from a number of frequency bins, and comparing to the predicted value, one can construct a χ^2 consistency test.

Formally, given a template h which produced a candidate signal with a SNR of ρ, calculate N nonoverlapping frequency windows such that the expected SNR is ρ/N in each. Then, calculate the actual SNR ρ^i in each of these frequency bins and compare with the expected value by calculating


FIG. 6.   The distribution of the χ^2 test for a single template h, split into 16 nonoverlapping frequency bins. The plot shows the distribution of the χ^2 test calculated for every time sample in 128 s of simulated Gaussian data (with no signal present). The observed distribution of values shows (shown in gray) is an excellent match with the expected χ^2 distribution with 60 degrees of freedom (shown in black).

$$\chi^2 = N \sum_{i}^{N} (\rho^i - \rho/N)^2. \qquad (4.22)$$

For a gravitational-wave signal matching the template h plus Gaussian noise, this statistic will be χ^2 distributed with $N - 1$ degrees of freedom. Written in the form (4.22) it appears different from the general case we discussed earlier. In [29] it was shown that it can be reexpressed in the form (4.4).

This χ^2 test can be extended to coherent, multidetector searches. Indeed, in [55], the construction was applied to a coherent search for continuous gravitational waves. Here, we present the extension to a coherent CBC search. First, define

$$\rho_\mu^i = \frac{(\mathbf{s}|\mathbf{h}_\mu^i)}{\sqrt{(\mathbf{h}_\mu|\mathbf{h}_\mu)}} \qquad (4.23)$$

to be the SNR contribution in the ith frequency bin to the SNR.[8] The coherent χ^2 statistic is then constructed as

$$\chi^2 = N \sum_{i=1}^{N} \sum_{\mu=1}^{4} (\rho_\mu^i - \rho_\mu/N)^2. \qquad (4.24)$$

---

[8]Strictly speaking the frequency bins for the F_+ and $F_\times$ components will be different because, as we have noted in Eq. (2.36), the PSDs for the synthetic + and × detectors are not equal. However, usually the difference between the two is small enough that it can be safely ignored to avoid computing twice the number of filters. Alternatively, in [29] a method was presented for calculating the standard χ^2 test using unequal frequency bins; that method could easily be incorporated into a coherent search.

As all the components are orthogonal it is easy to see that this statistic will be exactly χ^2 distributed with $4(N-1)$ degrees of freedom. One can interpret this as the sum of the single detector χ^2 values for the h_0 and $h_{\pi/2}$ waveforms in the synthetic $+$ and $\times$ detectors. Figure 6 shows the distribution of the standard χ^2, using sixteen frequency bins. The distribution matches the expected χ^2 with 60 degrees of freedom.

An alternative approach to applying the χ^2 test to a coherent search was proposed in [23]. This approach involves calculating the χ^2 values for each of the active detectors and using these values to veto glitches.

V. IMPLEMENTATION AND PERFORMANCE OF A COHERENT SEARCH

In this section, we describe an implementation of the targeted, coherent search for gravitational waves from CBCs. The search calculates the coherent SNR described in Sec. II and the various strategies for discriminating between signal and noise in non-Gaussian, nonstationary data introduced in Secs. III and IV. We demonstrate the efficacy of the search by performing test analyses of simulated data and data from LIGO's S4 run.

A. Implementation of a coherent triggered search for CBCs

Here, we describe the main steps by which the algorithms described in Secs. II, III, and IV have been implemented. The analysis is available in the LIGO Scientific Collaboration Applications Library (LAL) suite [56], and makes use of a large number of tools and methods previously implemented in that library.

1. Analysis setup

A targeted, coincident search for gravitational waves from CBCs associated to GRBs has been implemented, and used in a search of S5 and VSR1 data [16]. The coherent search pipeline uses many of the same definitions, and much of the same architecture as the coincidence search pipeline to determine the analysis details. Specifically, "on-source" time is $[-5, +1)$ seconds around the reported time of the GRB; this is when a gravitational-wave signal would be expected [20,57] and is the time over which we perform our search. The noise background is estimated using 1944 s of "off-source" data split into 324 trials of 6 s length each. These are used to calculated the significance of any event occurring in the on-source. As in Ref. [16] we impose a 48 s "buffer zone" between the on-source and off-source regions, which is not used in the analysis. To obtain an accurate estimate of the detectors' power spectra, we only analyze data from a detector if it has taken at least 2190 s of continuous data around the time of the GRB. Modulo this restriction, the coherent analysis is designed to make use of data from all detectors that were operating at the time of the GRB.

2. Template bank generation

The problem of placing a nonspinning template bank in (m_1, m_2) space for a single detector has been extensively studied [39,58–60]. However, less thought has been given to the problem of placing an appropriate bank for a coherent analysis, targeted or otherwise. For the targeted, coherent SNR statistic, described in Sec. II, the maximization procedure ensures that we need only place templates in (m_1, m_2) space. In the single detector search, the parameter space metric is independent of the amplitude and phase parameters that are maximized over. However, for the coherent analysis, this is generically *not* the case, and the metric depends upon the $\mathcal{A}^\mu$ parameters [42].

At present, we have not implemented an optimal template bank for a coherent analysis. Instead, simply make use of an overdense template bank generated for one of the detectors in the network. In the results presented later, we have used a bank generated with the initial LIGO design spectrum, with a maximum total mass of $40M_\odot$ and a minimum component mass of $1M_\odot$, which are the same values as used in previous searches for GRBs [16]. While this method of template placement enables us to demonstrate the efficiency of the coherent search, it is clearly not the optimal solution. A simple improvement would involve placing a template bank appropriate for the (maximally sensitive) synthetic $+$ detector defined in Eq. (2.36). For the relatively common situation where the network is significantly more sensitive to the $+$ polarization than the $\times$, this will be close to the optimal solution. However, for network and source configurations which provide good sensitivity to both polarizations, a more detailed investigation along the lines of [42] is warranted.

3. Coherent SNR and null streams

The data are read in and conditioned using the methods and algorithms developed for the S4 search for post-merger ringdowns from CBCs [61,62]. The data are then downsampled to a frequency of 4096 Hz and split into overlapping 256 s segments for analysis. The noise PSDs are calculated using the same method as in [61].

Each template in the bank is filtered against the data from each detector to generate the single detector filters $(s^X | h^X_{0, \pi/2})$ and sensitivities, σ^X [defined in Eq. (2.27)]. The algorithms used are taken from the LAL FindChirp library [63], specifically those written to perform a search for spinning waveforms [53] using the physical template family (PTF) waveforms [64].[9] The waveform templates are generated using the TaylorT4 post-Newtonian approximant [43]. The single detector filter outputs are shifted in time to account for the relative delays from the given GRB sky location. They are then combined to form the coherent and

---

[9]This choice stems from the desire to extend this search to incorporate a single spin. This is particularly appropriate for NSBH binaries where the spin of the NS can be safely neglected.

the null SNRs as described by (2.44) and (3.1). A "trigger" is recorded at any time the coherent SNR is greater than 6, and no louder event occurred in any template in the bank within 0.1 s.

4. Calculating the χ^2 tests

The analysis calculates signal based vetoes in the same manner that it does the coherent SNR: The necessary single detector filters are constructed and then these are combined together to create the χ^2 tests as described in Sec. IV.

Calculating the standard χ^2 test is computationally expensive. Therefore this veto is only calculated for a segment if there is at least one event with SNR above threshold and values of bank χ^2, auto χ^2, and null SNR that do not immediately lead to it being dismissed as a glitch.

5. Event significance

After the analysis, we have a set of triggers with associated masses and coalescence times. Each trigger is characterized by the coherent SNR as well as a number of other quantities: the null SNR, single detector SNRs, and values of the χ^2 tests. In order to identify candidate gravitational-wave events, these triggers are compared to those obtained from an identical analysis performed over the off-source times, and over the off-source data with simulated gravitational-wave signals added. A simulation is deemed to have been recovered by the analysis if there is any trigger within 0.1 s of the signal time; no attempt is made to guarantee a good match between simulated and recovered parameters. Although this means that signals may be found due to a nearby glitch, in practice we find that the effect is minimal, particularly when considering detection candidates, which are louder than all background.

By comparing the observed triggers from the off-source and simulations, we construct a detection statistic designed to best separate signal from noise. At present, this is done by performing a number of simple cuts in various parameters, as described in detail in the remainder of the section. In the future, we plan to investigate the use of multidimensional classifiers to improve the efficacy of the analysis. The detection statistic is used to calculate a false alarm probability for a given on-source observation, by comparing with the off-source results. We use 300 background trials, allowing for a false alarm probability as low as 3×10^{-3} to be assigned to an event. This suffices to identify GRBs for which an interesting gravitational-wave trigger has been observed. Realistically a false alarm probability closer to 10^{-4} or 10^{-5} would be required for a detection candidate. This could be achieved by performing additional background trials on time-shifted data from the detectors. This has not yet been implemented.

B. Analysis of simulated data

The analysis was first run on simulated data for the initial LIGO network, comprising a 4 km detector at the


FIG. 7 (color online). The SNR of triggers in the off-source region plotted against time for an analysis of simulated Gaussian noise in the initial LIGO (H1, H2, L1) network. The gap in the middle of the plot contains the on-source region and the buffer; this is not used in estimating noise background rates. The loudest trigger occurs with a SNR of 7.24.

Livingston site (L1) and 4 km and 2 km detectors at the Hanford site (H1 and H2, respectively). The coherent analysis pipeline analyzed a 2048 s stretch of stationary, Gaussian data as if a GRB had occurred in the middle of the data stretch.

Figure 7 shows the triggers produced by the pipeline in the off-source time. The loudest event in the approximately 2000 s of off-source data has a SNR of 7.24. In Gaussian noise, the signal consistency tests have no power as they are designed specifically to reduce the effect of non-Guassian transients in the data.

C. Analysis of real data

A test analysis was also performed on real data taken from S4. We analyzed an arbitrary stretch of 2048 s of data for which all three of the LIGO detectors were operating and ran the analysis as if a GRB had occurred during this time. The simulated sky location of the GRB was (184.6°, 42.34°) in right ascension and declination, respectively. For this chosen time and sky location the sensitivities of the H1 and L1 detectors were roughly equal and the H2 detector was half as sensitive as the other two.

The data from the detectors are neither Gaussian nor stationary. Thus, the goal of the analysis is to reduce the background from nonstationary data using the signal consistency tests described in Secs. III and IV. We proceed by investigating the various signal consistency tests one by one before combining these cuts into a detection statistic.

The final goal is to obtain a search sensitivity as close as possible to that obtained in Gaussian data. The sensitivity is assessed by evaluating the efficiency of observing simulated signals in the data. We make use of a set of simulated

BNS signals (component masses limited to be between 1 and $3M_{\odot}$) all oriented face-on to the detectors. We evaluate the efficiency of the search as the fraction of simulated signals observed with a detection statistic greater than any event in the off-source data.

1. Coherent SNR

Figure 8 shows the coherent SNR of triggers produced during the analysis of the S4 data. It clearly demonstrates that the data are not well characterized by Gaussian noise alone. A number of loud transients are present in the data and show up as short duration peaks of large SNR. The largest of these has a SNR of almost 40. If events were ranked by coherent SNR, a signal would have to be very loud to stand out above this non-Gaussian background. In addition to the loud peaks there are also a large number of smaller non-Gaussian peaks that occur rather frequently in the data. All of these affect the sensitivity of the search.

2. Null SNR

Figure 9 shows the performance of the null stream for both simulated signals and background noise. The ability of the null SNR to distinguish signal from noise is relatively poor in this example. The mock GRB analysis uses data from the two Hanford detectors and the detector at Livingston. As the two Hanford detectors are aligned, the null stream is derived from a combination of these detectors; the Livingston detector does not contribute. The loudest glitches during the time of this analysis originated in L1, and therefore do not contribute to the null SNR.


FIG. 8 (color online). The SNR of triggers in the off-source region plotted against time for an analysis of a mock S4 GRB. The axes on the plot are chosen to be identical to those for Fig. 7 to make the plots easier to compare. The S4 data have a large number of non-Gaussian features. The largest of these peaks extends to a coherent SNR of 40, although non-Gaussian structure is visible at SNRs as low as 7.


FIG. 9 (color online). The distribution of the null SNR plotted against coherent SNR. The solid line at null SNR of 5.25 is the line above which triggers are vetoed. The dashed line at 3.5 is the line above which triggers are downweighted (see Sec. V C 5).

However, quieter glitches in the Hanford detectors at a SNR around 10 do produce a large null SNR. Any trigger with a null SNR greater than 5.25 is eliminated from the analysis. In this example analysis, this removes a small fraction of the background and none of the simulated signals.

3. Single detector SNR

The most straightforward, and most effective, amplitude consistency test we have found is the requirement of a SNR greater than 4 in the two most sensitive detectors, in this analysis, the L1 and H1 detectors. Figure 10 demonstrates that this is a particularly effective strategy for removing noise glitches. Triggers arising due to glitches in the L1 detector have large coherent SNR but a negligible contribution from H1 and are consequently discarded. The single detector SNR threshold is very effective at removing background triggers—particularly those associated to glitches in L1. There is some loss of simulated signals, but these generally have a small enough SNR as to be indistinguishable from the background, even in Gaussian noise.

4. χ^2 tests

In Sec. IV we introduced three χ^2 tests designed to separate signals from noise glitches in the data. Figure 11 shows the distribution of the bank χ^2 for every time sample for a single template. This is directly comparable to Fig. 3 which shows the same for Gaussian data. The deviation from the predicted χ^2 distribution is due to the non-Gaussianity of the data.

The distribution of bank and auto χ^2 for both simulated signals and noise triggers is shown in Fig. 12. Both of these

FIG. 11. The distribution of the bank χ^2 test for a single template h, with a bank of size 10. The plot shows the distribution of the bank veto calculated for every time sample in 128 s of data. The observed distribution is inconsistent with the expected result in Gaussian noise (the black curve).
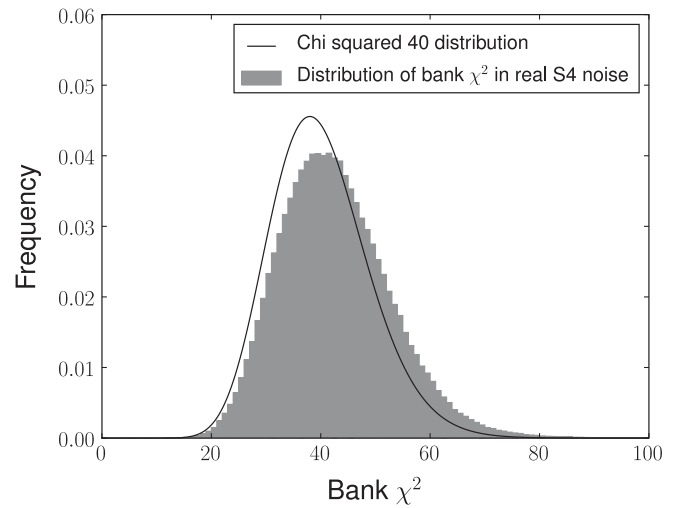
FIG. 10 (color online). The distribution of single detector SNR for the more sensitive H1 and L1 detectors, plotted against coherent SNR. The top panel shows the H1 SNR, the bottom panel shows the L1 SNR. The horizontal line indicates SNR = 4. Below this line triggers will be vetoed. The inclined dark gray line indicates the expected SNR of these face-on simulated signals.

tests are effective at separating the simulated signals from noise transients. In order to quantify this, we make use of the *newSNR* formalism that was developed for the latest coincident searches for CBCs [65,66]. For a signal matching the template waveform, the expected value of χ^2 is one per degree of freedom, while for noise transients this will be larger. The idea is to downweight the significance of noise triggers with large χ^2 values relative to signals. This is achieved by introducing the "*newSNR*":

$$\rho_{\text{new}} = \begin{cases} \rho, & \chi^2 \leq n_{\text{dof}} \\ \dfrac{\rho}{[(1 + (\frac{\chi^2}{n_{\text{dof}}})^{4/3})/2]^{1/4}}, & \chi^2 > n_{\text{dof}} \end{cases}, \qquad (5.1)$$

where n_{dof} is the number of degrees of freedom of the χ^2 test. For signals, *newSNR* will be similar to the SNR, while

noise transients with a large χ^2 value are significantly downweighted.

The *newSNR* can be calculated with all of the χ^2 tests introduced in Sec. IV. Any trigger with *either* an auto or bank *newSNR* less than 6 is discarded. The curves in Fig. 12 show this *newSNR* threshold for the two χ^2 tests. Finally, we turn to the standard χ^2 test. As this is rather costly to compute, we only do so for triggers which have passed all of the previously described thresholds (on coherent, null, and single detector SNR, bank and auto *newSNR*). Figure 13 shows the distribution of the standard χ^2 test for simulated signals and noise. The preceding tests have succeeded in removing the vast majority of non-Gaussian triggers from the data. A threshold of 6 on *newSNR* serves to eliminate a few more.

We have found that the standard χ^2 is the most effective of our χ^2 tests at separating signal from background, so we also make use of it in the final ranking of events. Figure 13 shows contours of constant *newSNR* which will be used in the final ranking.

5. Detection statistic

In the preceding discussion, we have imposed a number of cuts on the initial candidate events produced by the analysis pipeline. Let us briefly recap those cuts.

(1) Coherent SNR: generate a trigger at any time for which $\rho > 6$. Only keep the loudest trigger in each 0.1 s.
(2) Null SNR: discard any triggers with $\rho_N > 5.25$.
(3) Single detector SNR: discard any triggers for which $\rho_{H1} < 4$ or $\rho_{L1} < 4$.
(4) χ^2 tests: discard any triggers for which $\rho_{\text{new}} < 6$ for the bank or auto χ^2.
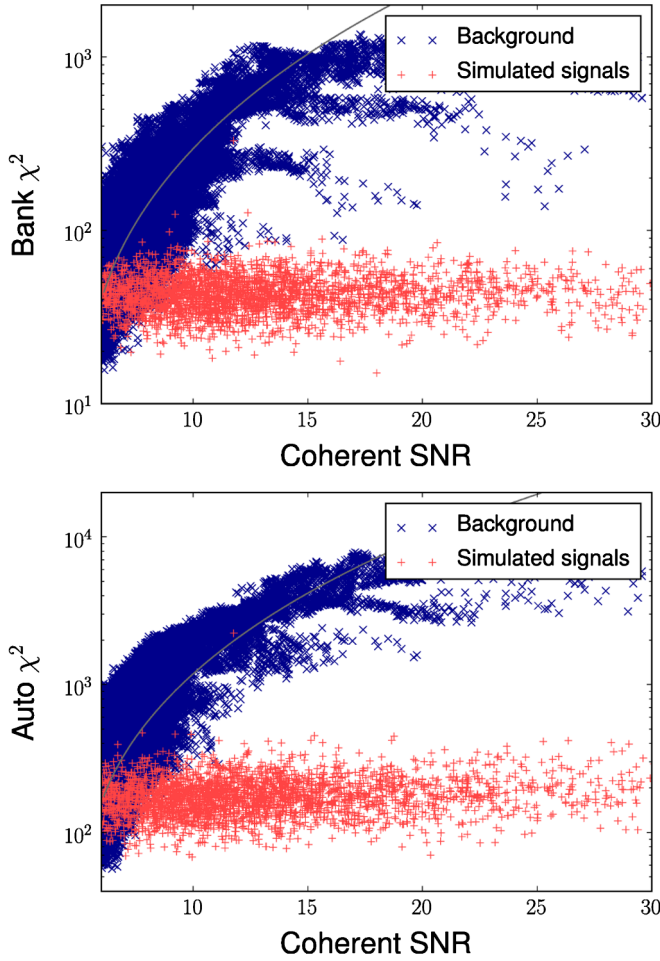
FIG. 12 (color online). The distribution of bank (top panel) and auto (bottom panel) χ^2 test plotted against SNR. There is a clear separation between simulated signal and background at coherent SNRs above 10. The solid line shows the line of *newSNR* = 6. Triggers with *newSNR* < 6 (above and to the left of the line) are discarded.

Finally, we rank the remaining triggers based on the *newSNR* calculated using the standard χ^2 and the null SNR:

(5) Detection statistic: rank remaining triggers using a detection statistic ρ_{det} given by

$$\rho_{det} = \begin{cases} \rho_{new}, & \rho_N \leq 3.5 \\ \frac{\rho_{new}}{\rho_N - 2.5}, & 3.5 < \rho_N < 5.25 \end{cases} \quad (5.2)$$

The length of time a CBC spends in the sensitive band of the detector varies greatly with the mass, and it has been found that the shorter, high-mass templates are more susceptible to occurring with large SNR at the time of glitches [67]. Also, the various signal consistency tests are less effective for these short templates. Therefore, we follow Ref. [67] and split the template bank into three regions


FIG. 13 (color online). The distribution of standard χ^2 test plotted against SNR. The majority of the background has already been removed by utilizing other signal consistency tests. The dashed lines show contours of *newSNR*, with the value of *newSNR* increasing toward the bottom right of the plot. The solid line shows the line of *newSNR* = 6; triggers above this line are discarded.

based on the chirp mass of the template, and calculate the false alarm probability for each mass bin separately:

(6) False alarm probability: for each trigger, calculate the false alarm probability by counting the fraction of off-source trials with a louder trigger in the same mass bin.

It is this false alarm probability which allows us to assess the significance of any events in the on-source data.

6. Performance of search

To quantitatively assess the sensitivity of the analysis we evaluate the efficiency of recovering simulated signals at a fixed false alarm probability. Any simulated signal which is found louder than the loudest off-source event is considered as found by the analysis—this ensures that the false alarm probability is less than 1 in 324 (the number of off-source trials). We use a population of optimally oriented BNS signals from the location of the fake GRB. Since all BNS injections are recovered with triggers in the low-mass bin (with chirp mass less than $3.5M_\odot$), we compare to off-source triggers in that bin.

Figure 14 shows the efficiency with which simulated signals are recovered as a function of distance for a variety of search methods. We begin by considering the sensitivity of a search in Gaussian noise; this will provide a benchmark against which to compare the searches in real data. The largest SNR recorded in the analysis of simulated, Gaussian data was 7.24 (see Fig. 7). Thus, the efficiency of a search in Gaussian data is given by the fraction of simulations found with a SNR greater than this. We have seen that the real detector data are *not* Gaussian, so only if the signal consistency tests are perfectly able to separate

FIG. 14 (color online). Efficiency of injection recovery is shown against distance at a false alarm probability of less than 1 in 324. This is calculated from a set of simulated signals in real data. The efficiency is shown for four different cases: (i) signals found above a SNR of 7.24, the loudest background trigger in Gaussian noise (black solid line); (ii) signals found above a SNR of 12.88, the loudest background trigger in real noise (red dashed line); (iii) signals found with a value of the detection statistic [given by Eq. (5.2)] above 7.41, the loudest background event (blue dot-dashed line); and (iv) signals found louder than all background in an H1-L1 search (green dotted line).

signals from noise transients can we hope to achieve a similar sensitivity in real data.

If we were to use *only* coherent SNR to rank events, and ignore all signal consistency tests, the sensitivity of the search would be about a factor of 2 worse than in Gaussian data. This is expected as the loudest low-mass trigger had a coherent SNR of 12.88, almost double the value observed in Gaussian noise. Once all of the signal consistency tests are taken into account, the loudest off-source event has a detection statistic of 7.41. Any simulated signal which produces a trigger which passes all signal consistency cuts and has $\rho_{\rm det} > 7.41$ is considered found. The use of signal consistency tests and a new detection statistic greatly improves the sensitivity of the search. The distance at which 50% efficiency is achieved is only about 10% less than in Gaussian noise. The 10% loss in sensitivity can be attributed to a slightly louder off-source event (7.41 rather than 7.24) and a *small* number of the simulated signals being vetoed by the signal consistency cuts.

Finally, we would like to illustrate the benefits of a coherent search over a coincidence search. A coincidence search filters each detector independently and records single detector triggers before searching for coincidence between the triggers in different detectors. For the initial LIGO network, signals close to detection threshold would be unlikely to been seen in the less sensitive H2 detector. Specifically, for a signal with a coherent SNR of 7.5, the expected SNR in H2 would be around 2.5 which would

be insufficient to generate a trigger. Therefore the sensitivity of a coincidence search would be limited by events observable in the two detector H1-L1 network. For a two detector search, the coherent and coincident SNRs are equal, and the null stream test is not available. Consequently, the performance of two-detector coherent and coincident searches should be comparable. The efficiency of this two-detector search is also shown in Fig. 14; it is about 10% less sensitive than the three-detector coherent search. This demonstrates that the coherent analysis, which incorporates the H2 data, can increase the distance reach of the search by 10%. For a network of three approximately equally sensitive detectors, we would expect an even greater sensitivity improvement from employing coherent techniques.

VI. DISCUSSION

We have presented a formulation of a targeted coherent search for compact binary coalescences. For Gaussian noise, the coherent SNR would be ideal for separating signals from the noise background. However, since data from gravitational-wave interferometers are neither Gaussian nor stationary, we have also discussed a number of methods of separating the nonstationary noise background from the signal population. These tests include various χ^2 tests, which were originally designed for use in single detectors. We have extended them to the network analysis and demonstrated their continued efficacy. Additionally, the coherent analysis naturally allows for tests which are not readily available in the coincidence case. The most significant of these is the null SNR which can be used to reject events which are not consistent with two gravitational-wave polarizations. We also explored consistency tests between the recovered amplitudes of the gravitational wave and found that a simple SNR threshold on the two most sensitive detectors gave excellent results. There are various other glitch rejection techniques, which have been recently discussed in the literature [68–70], and it might be possible to utilize such methods to improve the coherent search described here.

The analysis described in this paper has been implemented and in the final section we showed results of a test run. This made use of the S4 data from the LIGO detectors. Although the data were far from Gaussian, after the application of all of the signal consistency tests, the results were remarkably close to what would be expected in Gaussian noise. This analysis is available to be used in searches for GW inspiral signals associated with GRBs in more recent LIGO and Virgo data, such as S6 and VSR2 and VSR3.

There are a number of ways in which this analysis could be enhanced to broaden its use and increase its sensitivity. First, a number of GRBs, particularly those observed by Fermi [71] and IPN [72], are not localized sufficiently accurately that the error box can be treated as a point on the sky. Thus, it would be nice to extend this analysis to

cover a region of the sky. This would require looping over the relevant sky points, incorporating the correct detector sensitivities $F_{+,\times}$ and time delays. In principle, this would not greatly slow down the analysis as the majority of time is taken in performing the single-detector filters and these would *not* need to be recalculated. As well as looking at a patch on the sky, the analysis could be extended to cover the whole sky, as appropriate for an untriggered search. This brings in a host of new complications which have been met and dealt with by other coherent search methods [24,25]. In order to obtain a good estimate of the background for an all-sky, untriggered search, we would need to implement background estimation and time shifting the data would likely be the best way to do this.

The $\mathcal{F}$-statistic technique described in Sec. II B is formulated by maximizing over the extrinsic parameters of the system. From a Bayesian perspective this would imply that we have placed a uniform prior on the distributions of the $\mathcal{A}^{\mu}$. If instead we were to place an astrophysical prior on these amplitudes based on the expected distributions of (D, ι, ψ, ϕ_0) we would expect to increase the efficiency of the search [46]. Alternatively, since GRBs are thought to be rather tightly beamed, it is reasonable to take them as being face-on, or close to. In this case, the gravitational waves are circularly polarized and there is, in effect, only a single polarization. This opens the possibility of limiting the signal space to just this one polarization and adding an extra ''null'' test.

The progenitors of short GRBs are thought to be BNS or NSBH. The search we have described is ideal for the BNS case as the spins of the neutron stars are unlikely to have a significant effect on the waveform. However, when one of the components of the binary is a black hole, the spin could be large. Furthermore, the mass ratio is likely to be relatively large. In this case, the spin of the black hole can have a significant effect on the observed waveform [64]. Consequently, we would like to extend this search to incorporate spin effects. The infrastructure described in this paper can already accept spinning waveforms, but the implementation of signal based vetoes proves somewhat more complex. Work is underway on this [73].

[1] B. Abbott *et al.* (LIGO Scientific Collaboration), Rep. Prog. Phys. **72**, 076901 (2009).

[2] F. Acernese *et al.* (Virgo Scientific Collaboration), Classical Quantum Gravity **23**, S635 (2006).

[3] B. Willke (LIGO Scientific Collaboration), Classical Quantum Gravity **24**, S389 (2007).

[4] J. Abadie *et al.* (LIGO and Virgo Scientific Collaboration), Phys. Rev. D **82**, 102001 (2010).

[5] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration), Nature (London) **460**, 990 (2009).

[6] J. Abadie *et al.* (LIGO and Virgo Scientific Collaboration), Phys. Rev. D **81**, 102001 (2010).

[7] B. Abott *et al.* (LIGO and Virgo Scientific Collaboration), Astrophys. J. **713**, 671 (2010).

[8] G. M. Harry (LIGO Scientific Collaboration), Classical Quantum Gravity **27**, 084006 (2010).

[9] Advanced LIGO, http://www.advancedligo.mit.edu/.

[10] Advanced Virgo, http://wwwcascina.virgo.infn.it/advirgo/.

[11] J. Abadie *et al.* (LIGO and Virgo Scientific Collaboration), Classical Quantum Gravity **27**, 173001 (2010).

[12] LCGT: Large Scale Cryogenic Gravitational Wave Telescope, http://gw.icrr.u-tokyo.ac.jp/lcgt/.

[13] AIGO: Australian International Gravitational Observatory, http://www.aigo.org.au/.

[14] S. Hild, S. Chelkowski, and A. Freise, arXiv:0810.0604.

[15] B. Abbott *et al.* (LIGO Scientific Collaboration), Astrophys. J. **681**, 1419 (2008).

[16] J. Abadie *et al.* (LIGO and Virgo Scientific Collaboration), Astrophys. J. **715**, 1453 (2010).

[17] B. P. Abbott *et al.* (LIGO and Virgo Scientific Collaboration), Astrophys. J. **715**, 1438 (2010).

[18] J. Kanner *et al.*, Classical Quantum Gravity **25**, 184034 (2008).

[19] E. Nakar, Phys. Rep. **442**, 166 (2007).

[20] M. Shibata and K. Taniguchi, Phys. Rev. D **77**, 084015 (2008).

[21] B. D. Metzger *et al.*, arXiv:1001.5029 [Mon. Not. R. Astron. Soc. (to be published)].

[22] V. Predoi *et al.*, Classical Quantum Gravity **27**, 084018 (2010).

[23] A. Pai, S. Dhurandhar, and S. Bose, Phys. Rev. D **64**, 042004 (2001).

[24] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, Classical Quantum Gravity **25**, 114029 (2008).

[25] P. J. Sutton *et al.*, New J. Phys. **12**, 053034 (2010).

[26] P. Jaranowski, A. Krolak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).

[27] Y. Guersel and M. Tinto, Phys. Rev. D **40**, 3884 (1989).

[28] L. Blackburn *et al.*, Classical Quantum Gravity **25**, 184004 (2008).

[29] B. Allen, Phys. Rev. D **71**, 062001 (2005).

[30] C. Hanna, Ph.D. thesis, Louisiana State University, 2008.

[31] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **80**, 047101 (2009).

[32] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **79**, 122001 (2009).

[33] S. Bose, A. Pai, and S. V. Dhurandhar, Int. J. Mod. Phys. D **9**, 325 (2000).

[34] S. Bose, S. V. Dhurandhar, and A. Pai, Pramana J. Phys. **53**, 1125 (1999).

[35] S. Bose, T. Dayanga, S. Ghosh, and D. Talukder (unpublished).

[36] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **80**, 042003 (2009).

[37] C. Cutler and B. F. Schutz, Phys. Rev. D **72**, 063006 (2005).

[38] N. J. Cornish and E. K. Porter, Classical Quantum Gravity **24**, 5729 (2007).

[39] B. J. Owen and B. S. Sathyaprakash, Phys. Rev. D **60**, 022002 (1999).

[40] J. C. Livas, Ph.D. thesis, Massachussetts Institute of Technology, 1987.

[41] L. A. Wainstein and V. D. Zubakov, *Extraction of Signals from Noise* (Prentice-Hall, Englewood Cliffs, NJ, 1962).

[42] R. Prix, Phys. Rev. D **75**, 023004 (2007).

[43] T. Damour, B. R. Iyer, and B. S. Sathyaprakash, Phys. Rev. D **63**, 044023 (2001).

[44] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, Phys. Rev. D **72**, 122002 (2005).

[45] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, J. Phys. Conf. Ser. **32**, 12 (2006).

[46] R. Prix and B. Krishnan, Classical Quantum Gravity **26**, 204013 (2009).

[47] C. A. K. Robinson, B. S. Sathyaprakash, and A. S. Sengupta, Phys. Rev. D **78**, 062002 (2008).

[48] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **72**, 082001 (2005).

[49] L. Blanchet, Living Rev. Relativity **5**, 3 (2002), http://relativity.livingreviews.org/Articles/lrr-2002-3/.

[50] M. Hannam *et al.*, Phys. Rev. D **79**, 084025 (2009).

[51] J. Abadie *et al.* (LIGO Scientific Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **624**, 223 (2010).

[52] C. Van Den Broeck *et al.*, Phys. Rev. D **80**, 024009 (2009).

[53] D. Fazi, Ph.D. thesis, Università di Bologna, 2009.

[54] B. Abbott *et al.* (LIGO Scientific Collaboration), LIGO Report No. LIGO-T070109-01, 2007.

[55] Y. Itoh, M. A. Papa, B. Krishnan, and X. Siemens, Classical Quantum Gravity **21**, S1667 (2004).

[56] LSC Algorithm Library Suite, https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.

[57] M. B. Davies, A. J. Levan, and A. R. King, Mon. Not. R. Astron. Soc. **356**, 54 (2005).

[58] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).

[59] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash, Classical Quantum Gravity **23**, 5477 (2006).

[60] T. Cokelaer, Phys. Rev. D **76**, 102004 (2007).

[61] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **80**, 062001 (2009).

[62] L. M. Goggin, arXiv:0908.2085.

[63] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, arXiv:gr-qc/0509116.

[64] Y. Pan, A. Buonanno, Y. Chen, and M. Vallisneri, Phys. Rev. D **69**, 104017 (2004); **74**, 029905(E) (2006).

[65] J. Abadie *et al.* (unpublished).

[66] J. Abadie *et al.* (unpublished).

[67] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **79**, 122001 (2009).

[68] J. Clark, I. S. Heng, M. Pitkin, and G. Woan, Phys. Rev. D **76**, 043003 (2007).

[69] J. Veitch and A. Vecchio, Phys. Rev. D **81**, 062003 (2010).

[70] T. B. Littenberg and N. J. Cornish, Phys. Rev. D **82**, 103007 (2010).

[71] Fermi Gamma-ray Space Telescope, http://www.nasa.gov/mission_pages/GLAST/main/index.html.

[72] The Interplanetary Network Progress Report, http://ipnpr.jpl.nasa.gov/index.cfm.

[73] I. Harry and S. Fairhurst, arXiv:1101.1459.