PHYSICAL REVIEW D **93,** 064004 (2016)

Fully-coherent all-sky search for gravitational-waves from compact binary coalescences

D. Macleod,[1,2] I. W. Harry,[1,3,4] and S. Fairhurst[1]

[1]*School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, United Kingdom*
[2]*Louisiana State University, Baton Rouge, Louisiana 70803, USA*
[3]*Department of Physics, Syracuse University, Syracuse, New York 13244, USA*
[4]*Max Planck Institut für Gravitationsphysik, Am Mühlenberg 1, D-14476 Potsdam-Golm, Germany*
(Received 28 September 2015; published 3 March 2016)

We introduce a fully coherent method for searching for gravitational wave signals generated by the merger of black hole and/or neutron star binaries. This extends the coherent analysis previously developed and used for targeted gravitational wave searches to an all-sky, all-time search. We apply the search to one month of data taken during the fifth science run of the LIGO detectors. We demonstrate an increase in sensitivity of 25% over the coincidence search, which is commensurate with expectations. Finally, we discuss prospects for implementing and running a coherent search for gravitational wave signals from binary coalescence in the advanced gravitational wave detector data.

DOI: 10.1103/PhysRevD.93.064004

I. INTRODUCTION

In recent years, the Laser Interferometer Gravitational-Wave Observatory (LIGO) [1] and Virgo [2] have operated as a network of ground-based gravitational wave (GW) detectors in an attempt to detect and study signals of astrophysical origin. The data have been searched for evidence of gravitational waves from compact binary coalescences (CBCs) [3–5], unmodeled GW bursts [4,6], non-axisymmetric spinning neutron stars [7], and a stochastic gravitational-wave background (SGWB) [8]. While these searches yielded no direct detections, great strides were made in both instrumental science and data analysis techniques, paving the way for highly-anticipated second-generation, or advanced, detectors [9,10]. The advanced LIGO detectors began operation in late 2015, with Virgo expected to join a year or two later, and all detectors evolving to their design sensitivity over the following years [11]. In addition, the KAGRA detector [12] in Japan and a third LIGO detector in India [13] are expected to join the global network.

In the coming years, the first direct observations of gravitational waves are expected [14] and binary mergers of neutron stars and/or black holes are the most promising astrophysical sources. It has long been argued that the most sensitive way to search for gravitational waves from a network of detectors is to use a coherent search [15–19] in which data from all detectors are combined in an optimal way prior to performing the search. While this method has been applied to searches for unmodeled burst sources [20], it has proven more difficult to use for the binary merger search. Instead, a coincidence search has been used [3,21], whereby the data from individual detectors are searched independently and the events recorded in the different detectors checked for time and mass consistency appropriate for a gravitational wave signal.

The benefit of the coincidence search is that it reduces the computational cost, at the expense of some loss in sensitivity. The dominant cost of the search is in performing a matched filter of the data against a bank of template waveforms [22]. The coincidence search performs this task once per detector per template. A search over the sky location of the signal is then trivially done by time-shifting the results from the different detectors accordingly [21]. Similarly, the noise background is estimated by applying larger time shifts to the data (significantly longer than the light travel time between detectors) to search for noise coincidences. A naive implementation of the coherent search would require an independent filtering of the data for each template and each sky point, with the search repeated for each time-shift used to estimate the noise background. Computationally, this is not feasible.

There are, however, good reasons to believe that the coherent search will be more sensitive than the coincident one [16], providing the motivation to overcome the computational challenges of the coherent analysis [17,23]. In a coincidence search, it is necessary to place a threshold on the signal-to-noise ratio (SNR) of events which will be stored by the analysis prior to identifying coincidences. This means that the power in the GW signal will only be accumulated in those detectors where there was an event above threshold. In comparison, the coherent analysis naturally incorporates the SNR from all operational detectors. In the coincidence analysis, each detector is searched independently, and there is no guarantee that the observed signals in each detector are compatible with a GW source with two polarizations. In the coherent analysis the data from all detectors is combined to extract the two physical GW polarizations. When there are more than two detectors, it is then possible to generate a null stream (or streams) [24,25] which will contain only noise. Removing

these additional noise contributions from the SNR will enhance the sensitivity of the search. The benefit of a coherent analysis becomes more significant as the number of detectors in the global network increases. Thus, with the realistic prospect of a five detector network operating in the next few years, there is increased motivation to overcome the challenges posed by a coherent CBC search.

In the past few years, a coherent search for gravitational waves associated with gamma-ray bursts (GRBs) has been developed [26,27] and used in analyses of LIGO and Virgo data [4,28]. The observed GRB signal is used to restrict the sky location and arrival time of the GW signal, which significantly reduces the computational cost of the analysis. Nonetheless, in developing a coherent, targeted search for binary merger signals, many of the issues involved in performing an all-sky, all-time analysis have been addressed. In particular, the analysis has been constructed so that each template is filtered independently through the data from each detector and the single detector (complex) SNR time series are used to search over the sky and also perform a small number of time shifts. Thus, the calculation of the single detector SNR remains the dominant computational cost. Furthermore, the algorithm was shown to improve search sensitivity by around 30% when compared to a coincidence-based search over the same data [26]. In this work, we extend the targeted, coherent search to an all-sky search, demonstrating the first fully coherent all-sky search for GWs from CBCs.

This paper is laid out as follows. In Sec. II, we begin by briefly reviewing the coherent search as implemented in [26], then we describe our new methods to extend that work to an all-sky coherent search. This requires searching over amplitude parameters, sky position, time, component masses and spins. In addition, we briefly recap the signal consistency tests that are used to mitigate the effects of nonstationary data. In Sec. III, we describe in detail the search as performed with the three-detector, two-site LIGO network as it existed in LIGO's fifth science run, giving the results of the search in Sec. IV. We also discuss the prospects for a coherent analysis of advanced detector data, and evaluate the likely sensitivity improvements. We end, in Sec. V, with a summary and discussion of future prospects.

II. A COHERENT SEARCH FOR COALESCING BINARY SYSTEMS

The gravitational waveforms emitted by a coalescing binary can be calculated by the post-Newtonian formalism when the two compact objects are well separated [29]. As they spiral closer, higher order terms in the post-Newtonian expansion become increasingly important and numerical relativity is used to calculate the waveform in the final stages of inspiral and through the merger and ringdown of the merged system [30,31]. Using this information, a number of phenomenological models have been constructed that accurately describe the gravitational waveform over a large region of the parameter space of binaries which do not precess (i.e. the components are either nonspinning or have spins aligned with the orbital angular momentum) [32,33]. Consequently, when performing a search for these systems, matched filtering techniques are generally used.

The binary coalescence waveform for binaries in circular orbit depends upon at least fifteen parameters, and possibly more if we include the equation of state for neutron stars. These parameters are the two masses, six components of the spin (encoding magnitude and orientation of the two spins), the location (distance, right ascension and declination) and orientation (inclination, polarization and phase) of the binary, and the time of coalescence. The search we describe below is restricted to binaries where the spin-induced precession of the orbit can be neglected, for which the component spins must be aligned with the binary's orbital angular momentum [32,33]. This restricts the number of parameters to eleven. In addition, we focus only on the dominant harmonic of the waveform and ignore higher modes [34–36].

In the rest of this section, we describe the method by which each of these 11 dimensions of the parameter space is covered in the search. Finally, we discuss additional features which have been developed to mitigate the presence of non-Gaussian artefacts in the detector data.

A. Maximizing over amplitude parameters: D,ι,ψ,ϕ

We begin by considering four parameters—the distance, binary inclination, polarization and coalescence phase—that affect only the observed amplitude and phase of the waveform in the various detectors. For the purposes of a coherent search, the amplitude of a GW signal from a nonprecessing binary inspiral can be decomposed into two polarizations [26],

$$h_+(t) = \mathcal{A}^1 h_0(t) + \mathcal{A}^3 h_{\pi/2}(t), \tag{1a}$$

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

Here, $h_0(t)$ and $h_{\pi/2}(t)$ describe the two phases of the waveform, which depend upon the masses and spins of the binary components and are usually assumed to be orthogonal. The amplitudes $\mathcal{A}^i$ are

$$\mathcal{A}^1 = \frac{D_0}{D}\frac{(1+\cos^2\iota)}{2}\cos 2\phi_0 \cos 2\psi - \frac{D_0}{D}\cos\iota \sin 2\phi_0 \sin 2\psi, \tag{2a}$$

$$\mathcal{A}^2 = \frac{D_0}{D}\frac{(1+\cos^2\iota)}{2}\cos 2\phi_0 \sin 2\psi + \frac{D_0}{D}\cos\iota \sin 2\phi_0 \cos 2\psi, \tag{2b}$$

$$\mathcal{A}^3 = -\frac{D_0}{D}\frac{(1+\cos^2\iota)}{2}\sin 2\phi_0 \cos 2\psi - \frac{D_0}{D}\cos\iota \cos 2\phi_0 \sin 2\psi, \quad (2c)$$

$$\mathcal{A}^4 = -\frac{D_0}{D}\frac{(1+\cos^2\iota)}{2}\sin 2\phi_0 \sin 2\psi + \frac{D_0}{D}\cos\iota \cos 2\phi_0 \cos 2\psi. \quad (2d)$$

These terms are dependent on four parameters of the source: its distance, D; the coalescence phase, ϕ_0; the polarization angle, ψ; and the inclination angle, ι. D_0 is a scaling distance, which is used in normalizing the waveforms $h_{0,\pi/2}$.

The GW signal seen by detector X is a combination of the two polarizations weighted by the detector antenna response, $F^X_{\{+,\times\}}$ [37],

$$h^X(t) = F^X_+ h_+(t^X) + F^X_\times h_\times(t^X), \quad (3)$$

where the time of arrival in detector X depends upon the sky location of the source relative to the detector and the time of arrival at a fiducial location, for example the Earth's center [26].

The matched-filter is described by an inner product between a template GW waveform of the above form, h, and the detector data s. In general, the inner product between two such time series, a^X and b^X, is given by

$$(a^X|b^X) = 4\mathrm{Re}\int_0^\infty \frac{\tilde{a}^X(f)\cdot[\tilde{b}^X(f)]^\star}{S^X_h(f)}, \quad (4)$$

where $S^X_h(f)$ is the noise power spectral density (PSD) in detector X. For a network of detectors, we define the multidetector inner product as the sum of the single detector inner products,

$$(\mathbf{a}|\mathbf{b}) \equiv \sum_{X=1}^{D}(a^X|b^X) \quad (5)$$

where D is the number of detectors in the network. The multidetector log-likelihood is then calculated as [26],

$$\ln\Lambda = (\mathbf{s}|\mathbf{h}) - \frac{1}{2}(\mathbf{h}|\mathbf{h}) = \mathcal{A}^\mu(\mathbf{s}|\mathbf{h}_\mu) - \frac{1}{2}\mathcal{A}^\mu\mathcal{M}_{\mu\nu}\mathcal{A}^\nu, \quad (6)$$

where $\mathbf{h}_\mu = (\mathbf{F}_+\mathbf{h_0}, \mathbf{F}_\times\mathbf{h_0}, \mathbf{F}_+\mathbf{h}_{\pi/2}, \mathbf{F}_\times\mathbf{h}_{\pi/2})$, and

$$\mathcal{M}_{\mu\nu} \equiv (\mathbf{h}_\mu|\mathbf{h}_\nu). \quad (7)$$

Maximizing the log-likelihood over the values of $\mathcal{A}^i$, the coherent SNR is defined as

$$\rho^2_{\mathrm{coh}} \equiv 2\ln\Lambda|_{\max} = (\mathbf{s}|\mathbf{h}_\mu)\mathcal{M}^{\mu\nu}(\mathbf{s}|\mathbf{h}_\nu), \quad (8)$$

where $\mathcal{M}^{\mu\nu}$ is the inverse of $\mathcal{M}_{\mu\nu}$.

We can rewrite Eq. (8) in a manner that makes it easier to compare to the coincident search. To do so, we introduce the complex SNR z^X in detector X as

$$z^X = (s^X|h^X_0) + i(s^X|h^X_{\pi/2}). \quad (9)$$

Then, we can write the coherent SNR as

$$\rho^2_{\mathrm{coh}} = \sum_{X,Y=1}^{D} z^X P_{XY} z^Y, \quad (10)$$

where P_{XY} is a projection of the SNR onto the 2-dimensional signal space:

$$P_{XY} = \left[\frac{(\sigma^X F^X_+)(\sigma^Y F^Y_+)}{\sum_Z(\sigma^Z F^Z_+)^2} + \frac{(\sigma^X F^X_\times)(\sigma^Y F^Y_\times)}{\sum_Z(\sigma^Z F^Z_\times)^2}\right], \quad (11)$$

and $\sigma^X = \sqrt{(h_0|h_0)_X}$ encodes the sensitivity of each detector.

Meanwhile the coincident SNR, obtained by filtering the same template waveform in all detectors, is given by

$$\rho^2_{\mathrm{coinc}} = \sum_{X,Y=1}^{D} z^X \delta_{XY} z^Y. \quad (12)$$

Thus, for a signal in the absence of noise (i.e. $P^{XY}z_Y = z^X$) the coincident and coherent SNRs are identical. For noise events, the coincident SNR includes all of the noise, while the coherent SNR incorporates only those contributions which are compatible with a coherent signal at all detectors.

B. Covering the sky

The coherent SNR derived above depends upon the sky location of the source in two ways: through its dependence on the detector sensitivities, encoded in $F_{+,\times}$, and through the differences in arrival time of the signal at the different detectors. Consequently, the value of the coherent SNR will change depending on the sky location of the source. There is no analytic way to maximize over the sky position and instead we must search over a discrete grid of sky points, much as we do for the binary masses and spins. The density of points required will depend upon the template's auto-correlation time [38]. For binary mergers, this depends upon the bandwidth of the signal [39] which is typically around 100 Hz and varies only slowly across the mass parameter space. In what follows, we neglect the mass dependence of the sky grid and instead place a grid which is sufficiently dense for all templates. While this results in a

(somewhat) overdense grid for higher mass templates, the effect is small.

The effect of timing on the coherent SNR is more significant than the change of the antenna response of the detectors [38]. In addition, the antenna responses change slowly in regions of the sky where the detectors have good sensitivity and most rapidly near the nulls of the detector. This further reduces the significance of the changes in $F_{+,\times}$ and consequently we ignore these effects and place points in the sky based solely on time delays. We do, however, make use of the appropriate detector responses when performing the search.

1. Sky tiling for a two-site network

Sky coverage is significantly easier for a two-site network than one with three or more sites. For such a network, there exists only a single time-delay baseline between the observatories meaning that we can use a one-dimensional sky grid. As we have remarked earlier, the coherent SNR calculated for a two *detector* network is no different from the coincidence SNR. However, for the first eighteen months of LIGO's Science Run 5 (S5) (November 2005—April 2007), the three LIGO detectors formed a two-site, three-detector network, with both LIGO Hanford Observatory (LHO) instruments taking part in the run alongside a single instrument at LIGO Livingston Observatory (LLO). In this case, there is a benefit to performing a coherent analysis. This is an ideal first test case as we can explore the effects of a coherent search, but with a reduced sky grid. The second-generation Advanced LIGO (aLIGO) instruments will form a two-detector, two-site network during the first observing run in 2015 [9,11]. However, there is no good reason that a coherent analysis would offer improved sensitivity to a coincidence one for this network.

For a two-site network, localization by triangulation will reconstruct only the difference in time of arrival between the sites, mapping to a ring on the celestial sphere. Furthermore, for the initial LIGO network, the maximized coherent SNR is completely independent of the values of $F_{+,\times}$. As a result, the most efficient tiling of the sky sphere for a two-site "all-sky" search is given by the one-dimensional space of physically allowed time-delays.

The allowed time delays are bounded by the light travel time, T, between sites. If we choose a temporal resolution δt, then the size of the grid is

$$N = \left\lfloor \frac{T}{\delta t} \right\rfloor. \tag{13}$$

The required temporal resolution depends upon the waveform's autocorrelation time. For a signal with a bandwidth around 100 Hz, template's autocorrelation falls off on time scales of ∼1 ms [39]. This gives an idea of the required spacing of points, and we have found empirically that $\delta t = 0.5$ ms is an appropriate value.[1] The light travel time between LHO and LLO is $T = 10$ ms so the two-site LIGO grid requires 40 sky points.

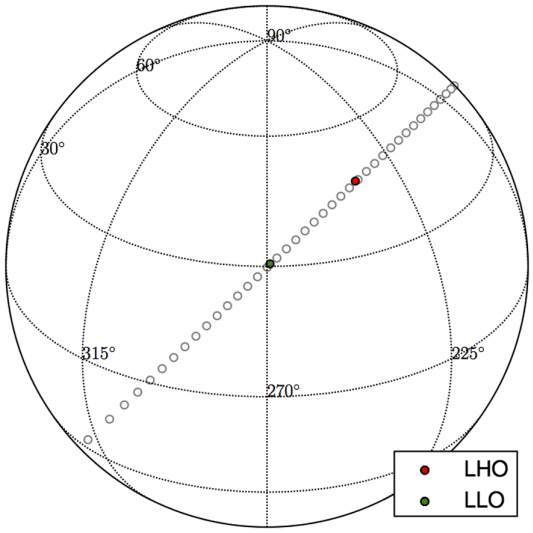

FIG. 1. A two-site all-sky grid for the LIGO detector network. The points span the allowed time-delays between sites, and are chosen such that they form a great circle passing overhead both.

To seed this grid, we lay points with time delay τ in the range $[-T, T)$, evenly spaced by δt. Each value of δt corresponds to a ring on the sky, but we wish to choose only a single point, at any location along that ring. To do this we find the intersection between each value of δt and the great circle containing both sites. The resulting set of points for the LHO, LLO two-site network is shown in Fig. 1. The grid has greatest density where the time-delay is smallest, with density dropping symmetrically in either direction, as the time delay grows.

2. Sky tiling for a three-site network

During the last six months of S5 [40] and throughout Science Run 6 (S6) [3], the LIGO and Virgo detectors operated a three-site network, allowing much more accurate time-delay triangulation, and better sky localization [41]. This is likely to be the same network running during the middle years of the advanced detector era, after both Advanced LIGO and Advanced Virgo are observing, but before other detectors are operating.

With three detectors, the network can triangulate any signal to a single point in the hemisphere above the plane of the network. A symmetry still exists in that plane, producing a second point in the other hemisphere. However, in many cases, the different detector responses allow us to distinguish between these points [42]. The unfortunate

[1]The GW data are downsampled to 4096 Hz when performing this analysis. Since it is more straightforward to shift by an integer number of samples, we actually use a time shift of $2/4096$ seconds.

consequence of better sky localization is the need for much larger sky grids for a full coherent analysis, increasing the computational cost of the search.

In order to map the sky for three sites, we follow the analytical models of [43]. Consider a network of D detectors, and define the time-delay vector

$$\boldsymbol{\tau} = \begin{pmatrix} \tau_2 \\ \vdots \\ \tau_D \end{pmatrix}, \tag{14}$$

where τ_n is the arrival time difference between detector 1 and detector n. Let T_m be the light travel time between detector 1 and detector m, and define α_{mn} as the angle separating the lines connecting detectors 1 and m, and detectors 1 and n. Then, we can construct a bounding ellipse for the physically-admissible time delays,

$$\boldsymbol{\tau}^T \mathbf{A}_D \boldsymbol{\tau} \leq B_D, \tag{15}$$

where, for the case of three detectors,

$$\mathbf{A}_3 = \begin{bmatrix} \frac{T_3^2}{T_2^2} & -\frac{T_3}{T_2}\cos(\alpha_{23}) \\ -\frac{T_3}{T_2}\cos(\alpha_{23}) & 1 \end{bmatrix} \tag{16a}$$

$$B_3 = T_3^2 \sin^2(\alpha_{23}). \tag{16b}$$

Here, $\cos\alpha_{23}$ effectively measures the correlation between the two time delays. When $\cos\alpha_{23} = \pm 1$, the three sites lie in a line, the time delay matrix $\mathbf{A}_3$ is singular and $B_3 = 0$. In this case, the time-delays are degenerate and the localization is no better than a two-site network. When $\alpha_{23} = \pi/2$, the time-delays are independent ($\cos\alpha_{23} = 0$) and the time-delay baselines form the major and minor axes of the bounding ellipse. For values of α_{23} between these two extremes, there is some correlation between the time delays observed in the two detectors, and the ellipse of permitted time delays will not be aligned with the baselines between detectors.

A grid of hexagonal time-delay tiles is laid in (τ_2, τ_3) coordinates, such that the distance between any two points matches the desired resolution. In addition, all points must lie within the ellipse defined by Eqs. (16a) and (16b). The resulting time-delay grid for the three-site LIGO-Virgo network is shown in Fig. 2.

The time-delay grid is then projected onto the celestial sphere by constructing a detector network coordinate system, as shown in Fig. 3, where the time-delay coordinates are related to network longitude, ϕ, and latitude θ, via

$$\phi = \pm\cos^{-1}\left(-\frac{T_2\tau_3 - T_3\tau_2\cos(\alpha_{23})}{T_3\sqrt{T_2^2 - \tau_2^2\sin(\alpha_{23})}}\right), \tag{17a}$$

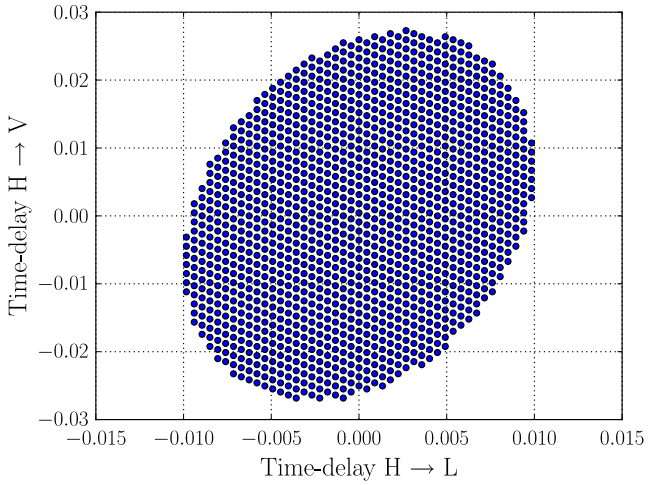

FIG. 2. Time-delay tiles for the LIGO-Virgo three-site network. All physically admissible points in this space are laid in a hexagonal grid, with a minimal 0.5 ms spacing between neighboring points.

$$\theta = \cos^{-1}\left(-\frac{\tau_2}{T_2}\right). \tag{17b}$$

The network coordinates (ϕ, θ) are then related to earth-fixed longitude and latitude with a simple rotation. This projection is done twice, once for each hemisphere above and below the plane of the detector network. In the end, we obtain a fixed grid in the Earth-centric frame.

Figure 4 shows an all-sky grid for the LIGO-Virgo network; 4a views the grid from nearly overhead the plane

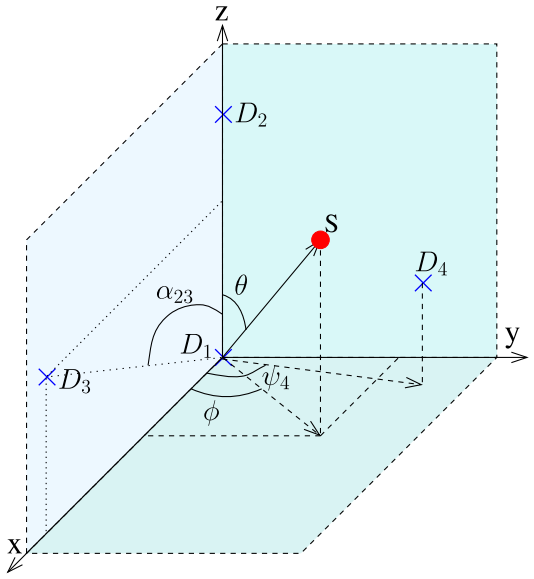

FIG. 3. The network coordinate system used in projecting points in time-delay space onto the sky [43]. A three-site network defines a right-handed coordinate system, with a potential fourth-detector breaking the symmetry x-z plane.

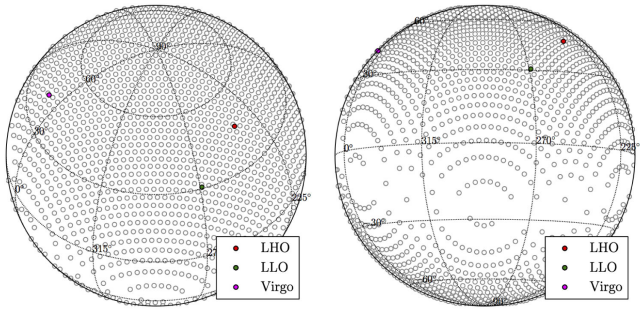

FIG. 4. The three-site all-sky coherent search sky grid for the LIGO-Virgo detector network. (a) and (b) view the same grid from different angles.

of the network, where the grid is densest, while 4b shows the relative sparsity of the grid parallel to the network plane. This grid, using a time-delay resolution of $\delta t = 0.5$ ms, contains over 2,700 points, representing a huge computational cost if applied naïvely to the coherent analysis.

C. Searching over coalescence time

The matched filter between a template signal with coalescence time t_c and the data can be written as:

$$(s|h_{t_c}) = 4\mathrm{Re}\int_0^\infty \frac{\tilde{s}(f)\cdot[\tilde{h}_{t_c=0}(f)]^\star e^{2\pi i f t_c}}{S_h(f)}. \qquad (18)$$

Therefore, it is possible to generate single detector SNR time series efficiently by using a Fourier transform [22]. In practice, this is done by dividing the data into segments of a fixed length and performing a fast Fourier transform (FFT) on each segment. Due to the finite duration of the templates (and also the inverse power spectrum), filter wraparound will lead to the corruption of the SNR time series at the beginning and end of each segment. This effect is mitigated by simply overlapping the FFT segments.

These single detector SNR time series are then used to calculate the coherent SNR as a function of time. In practice, we find that it is not necessary to calculate the coherent SNR for every sky point and every time sample. Instead, we require that the single detector SNRs are above a threshold prior to proceeding with the calculation of the coherent SNR. This greatly reduces the computational cost of the calculation so that the computation of coherent SNR, for a three-site sky grid, remains dominated by the calculation of the single detector SNR time series.

D. Searching over mass and spin

The amplitude and phase evolution of the waveform depends sensitively on the masses and spins of the binary components. To search over the mass and (aligned-)spin parameters, we make use of a discrete bank of template waveforms, which covers the parameter space of binaries sufficiently densely that the difference between any system and the nearest template is small enough that minimal signal power is lost [44,45].

For the coincidence searches performed on the initial detector data, a separate template bank was constructed for each detector based upon its sensitivity [21], where the density of templates depends upon the noise power spectrum of the data from the detector. For a coherent search, we must use the same bank for all detectors in the network. However, at different points in the sky, the detectors have different antenna responses and so contribute differently to the coherent analysis. Thus, in principle, the template bank should be dependent upon the sky location. However, as discussed above, for simplicity we do not do this. Indeed, a detailed investigation [38] showed that the effect is minimal. Instead, one would naturally use the harmonic mean of the detector PSDs to construct an average PSD for generating the template bank [26]. For the analysis presented here, where the PSDs of the LIGO instruments have very similar shape, we make a more straightforward choice and simply used a bank for one of the detectors as we found that this had little effect on the results.

E. Signal consistency tests and null SNR

In Gaussian data, the coincident or coherent SNR would serve as a detection statistic: the greater the value of the SNR the less likely to arise due to noise fluctuations alone. However, in real detector data, there are numerous nonstationarities in the data which can lead to high SNR events. Various techniques have been developed to mitigate the effect of these "glitches" and get the search as close to the Gaussian limit as possible [21,26,46,47].

1. Null SNR

A coherent gravitational wave search involves combining the data from the detectors in the network to produce data streams that are sensitive to the two polarizations of gravitational radiation. When there are more than two detectors in the network, it is possible to construct additional data streams which do not contain any gravitational wave contribution [24]. Using the framework above, we denote the null SNR as

$$\rho^2_{\mathrm{null}} = \rho^2_{\mathrm{coinc}} - \rho^2_{\mathrm{coh}}. \qquad (19)$$

In Gaussian noise, this would be χ^2-distributed with $2(d-2)$ degrees of freedom. Removing this noise contribution from the coherent SNR increases the sensitivity of the coherent search. When the data contain nonstationary transients, they will tend to be observed in a single detector and will not be consistent with a coherent signal. They will, therefore, have a large null SNR, and we can remove events with high null SNR from the search results.

2. Single detector SNR threshold

We also make use of thresholds on the single detector SNRs to reject events which are unlikely to be caused by real signals. By requiring that the SNR is above a given threshold in at least two detectors, we can eliminate the vast majority of events caused by nonstationary transients, as they appear in only a single detector. As discussed above, an additional benefit of the single detector thresholds is that they are very cheap to compute. Consequently, we apply them before calculating the coherent SNR in order to reduce the computational cost of the analysis.

3. χ^2 tests

Finally, we make use of signal consistency tests, typically called χ^2 vetoes [46,47]. The basic concept is to subtract the template that matches the observed signal and then check that what remains is consistent with Gaussian noise. This is done by filtering the residual data with test waveforms, T^i, that are orthogonal to the best fit template and calculating the sum of squares of SNRs in those templates. In Gaussian noise, the value will be χ^2-distributed with 4N degrees of freedom (where N is the number of test waveforms used). Any mismatch between the signal and the template will lead to imperfect cancellation of the signal and a contribution to the χ^2. This mis-match could be due to the discreteness of the template-bank, differences between the true waveforms and those used as templates, or errors due to the calibration of the detectors. To account for this, an SNR-dependent threshold is typically used when rejecting events with a large χ^2 value.

In this analysis, we make use of three different χ^2 tests, as implemented for the targeted coherent search [26]:

(1) *Frequency bins:* The test waveforms T^i are generated by chopping up the template $h(t)$ into $(N+1)$ sub-templates in the frequency domain, each of which contains an equal fraction of the power. From these, we can construct N orthonormal waveforms, all of which are orthogonal to $h(t)$.

(2) *Template bank:* The test waveforms T^i are binary merger waveforms with different mass and spin parameters. In general, they will not be exactly orthogonal to $h(t)$, so we simply remove the part proportional to $h(t)$. Note, however, that the test waveforms will also not be orthonormal, and consequently the expected distribution is not χ^2-distributed with $4N$ degrees of freedom. We do not attempt to resolve this issue but instead apply an empirical threshold.

(3) *Autocorrelation:* The test waveforms T^i are time-shifted copies of the original waveform $h(t)$. As with the template bank test, these waveforms are neither orthonormal nor orthogonal to $h(t)$. We proceed as above.

III. SEARCH IMPLEMENTATION AND TESTING

In this section we demonstrate the first implementation of a fully-coherent, all-sky search for signals from binary neutron star inspirals. The full analysis calculates the coherent SNR for each template in a bank, along with a number of signal-consistency statistics that allow distinction between non-Gaussian noise artefacts and true signals [26,27].

A. Data selection

The coherent all-sky analysis was used to search one month of data from S5, during which the three LIGO instruments formed a two-site, three detector network [1]. As usual, the 4 km detectors at Hanford and Livingston will be denoted H1 and L1 respectively, while the 2 km Hanford detector is denoted H2. Only those data segments during which all three detectors were operating nominally are used, with the additional requirement that each segment was longer than 2176 seconds to allow for accurate measurement of the detector noise PSD; Fig. 5 shows the segments that were selected. The sensitivity of the detectors to neutron star mergers is characterized by the sky and source orientation averaged distance at which the signal from a 1.4–1.4$M_\odot$ binary would be observed with a SNR of 8 in a single detector. Figure 6 shows the sensitive range for each detector in the network during this period. The smaller H2 detector maintained a range between 6–7 Mpc throughout, while the larger instruments improved as the run progressed, with H1 peaking at 16 Mpc.

B. Generating a template bank

The coherent search was performed over the parameter space of neutron star binaries. We take a mass range of 1–3$M_\odot$ for the binary components, and neglect the effect of spins on the waveform [48]. A bank of template inspiral waveforms was laid using the methods of [26,45], with a single-detector bank for the L1 detector used as a simple approximation to a fully-coherent template bank. For the analysis, non-spinning, 3.5 post-Newtonian (PN)-order,

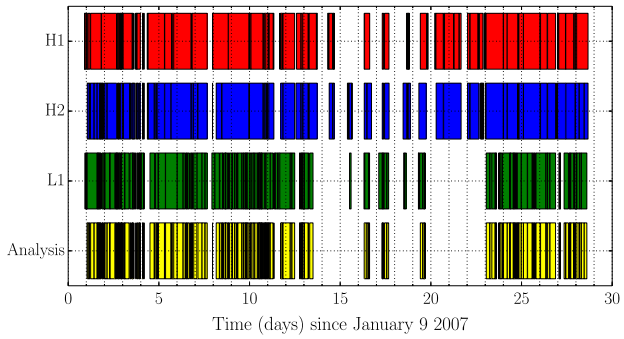

FIG. 5. Analysis segments for coherent all-sky search of one month of S5 data.

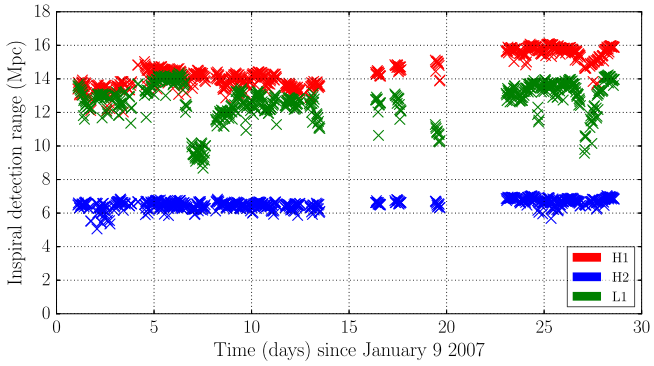

FIG. 6. Inspiral averaged sensitive range for the LIGO network during one month of S5.

binary neutron star (BNS) inspiral waveforms were placed with a maximum combined mass of $6.2M_{\odot}$, resulting in $\sim 2,200$ templates, as shown in Fig. 7. These waveforms were generated from the point at which the gravitational-wave frequency reaches 40 Hz, below which initial LIGO sensitivity degrades rapidly, and the longest template in the bank is therefore approximately 45 seconds long.

C. Performing the coherent analysis

For computational ease, each analysis segment is divided into chunks of 2176 seconds, overlapping by 64 seconds on each end; a template bank is generated for each of these chunks (due to changing detector sensitivity), with each chunk processed separately. The data from each detector are used to estimate its PSD, using the average of sixteen 50% overlapped 256-second sensitivity estimates.

For each of the 256-second blocks, the data from each detector are filtered against each template in turn, producing time series of single-detector SNR. The SNR time series at the beginning and end of each block is corrupted due to wraparound of the (up to) 45 s filters and effects of data conditioning [22], and only the middle 128-seconds is retained. When the single detector SNRs are above threshold, the sky grid is used to construct coherent combinations of these time series and calculate the coherent SNR. At any time sample where the coherent SNR for a given sky point is above threshold, the value is recorded and the signal consistency tests calculated.

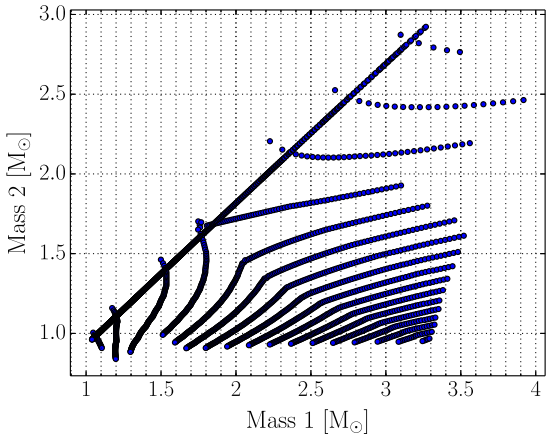

FIG. 7. The template bank used for the coherent all-sky search of one month of S5 data.

1. Background estimation with time-slides

The noise background is measured using time-shifted data. Since a fully-coherent search combines detector data at the filtering stage, each time-slide analysis requires recomputing the coherent matched-filter SNR. This represents a huge computational cost, in comparison to the coincidence-based analysis where time-shifts are performed on the single-detector events produced after filtering.

In this analysis, 10 time-shifts were constructed, each sliding data from the L1 detector forward by a multiple of 128-seconds. Such large slides are computationally simple and performed by combining data from the LHO instruments with those from L1 in a different 256-second block. The slides are performed on a ring (formed by the sixteen analysis blocks for a single chunk), whereby any L1 data slid off the end of the analysis chunk is reinserted at the start and filtered against the data from the LHO instruments in the first block. The computational cost for each of these time shifts is equal to the original analysis.

2. Simulations

A set of nonspinning BNS inspiral simulations were used to inform tuning of the signal-consistency cuts and the detection statistic, and measure search performance through simulation recovery efficiency. The signals were uniformly distributed in mass (with component masses between $1M_{\odot}$ and $3M_{\odot}$), sky location and orientation. Observed signals are expected to be distributed uniformly in volume. However, if simulations are distributed uniformly in volume, the vast majority of simulated signals will be at large distances and below the detection threshold of the pipeline. Instead, we generate simulations with distances uniformly distributed between 1–60 Mpc.

Throughout the following descriptions of signal-based and data quality cuts, we use time-shifted and simulated events to assess the sensitivity of the analysis, and to determine appropriate thresholds to separate signal from background.

D. Event down-selection

While the data from gravitational-wave detectors are often modeled as Gaussian, in practice this is rarely the case. The data are regularly contaminated with nonstationary, and non-Gaussian noise artefacts ("glitches") that will be detected, even in a coherent analysis, with high SNR. As a result, the rate of events identified as significant

by the matched-filter is too high to be either useful or practical—large noise glitches will trigger across the full template bank, producing multiple events from a single noise input.

Events are down-selected by identifying those most significant relative to the surrounding data. The full list of events are divided into 100 ms bins, with an event selected only if it is the loudest in its own bin, and louder than all events in a 100 ms window around itself. This selection method typically reduces the event rate by a factor of 100 or more, by identifying those event triggers that represent the peak of an excess power transient in the data (either noise or GW signal).

1. Signal-consistency cuts

Each of the signal-consistency tests outlined in Sec. II E are applied equally to the event triggers from the foreground data, each of the time-slide background trials, and the simulations. The SNR thresholds for the search are

(i) single-detector SNR ≥ 5 in one detector, and ≥ 4.5 in a second detector,
(ii) coherent SNR ≥ 7.

These limits are chosen empirically to avoid the number of event triggers being stored from becoming unmanageable, and as a way of rejecting noise artefacts, which are often seen with a loud SNR in only one detector. The first signal-consistency cut, on single-detector SNR, identifies those events with power in a single detector only, typically removing more than 90% of all events. Figure 8 shows the impact of this cut on the events from the S5 analysis, applied to both H1 and L1 data, including both the background from time slides and from the simulation set. Those background *trigger* (blue dots) on the diagonal in each figure represent events with power only in that detector, and fail the cut (black line) in the other detector. The simulated signals are below the diagonal because their coherent SNR is accumulated from a strong component in each detector.

The χ^2-based signal consistency tests are used to further reduce the impact of noise triggers. To do so, we first calculate the reweighted SNR [21], in a similar way as was done for the coincidence analysis, as:

$$\rho_{\chi^2} = \begin{cases} \rho, & \chi^2 \leq n_{\rm dof} \\ \dfrac{\rho}{\left[\left(1+\left(\frac{\chi^2}{n_{\rm dof}}\right)^3\right)/2\right]^{1/6}}, & \chi^2 > n_{\rm dof}, \end{cases} \tag{20}$$

where $n_{\rm dof}$ denotes the number of degrees of freedom for the χ^2 tests.

We then apply the following thresholds:

(iii) frequency-bin χ^2 reweighted SNR > 7,
(iv) template-bank χ^2 reweighted SNR ≥ 7,
(v) autocorrelation χ^2 reweighted SNR ≥ 7.
(vi) an SNR dependent threshold on the null SNR:

$$\rho_{\rm null} \leq \begin{cases} 6, & \rho \leq 20, \\ 6 + \left(\frac{\rho-20}{5}\right) & \rho > 20, \end{cases} \tag{21}$$

Figures 9a–9c show the impact of the three χ^2 consistency cuts, each evaluated after the single-detector SNR cut has been applied. These statistics clearly differentiate between the recovered simulations and the noise background, removing those events inconsistent with a true signal. Figure 9d shows the impact of the null SNR cut, similarly evaluated after the single-detector SNR cut has been applied. The null SNR cuts a relatively small number of noise events that are incoherent between detectors.

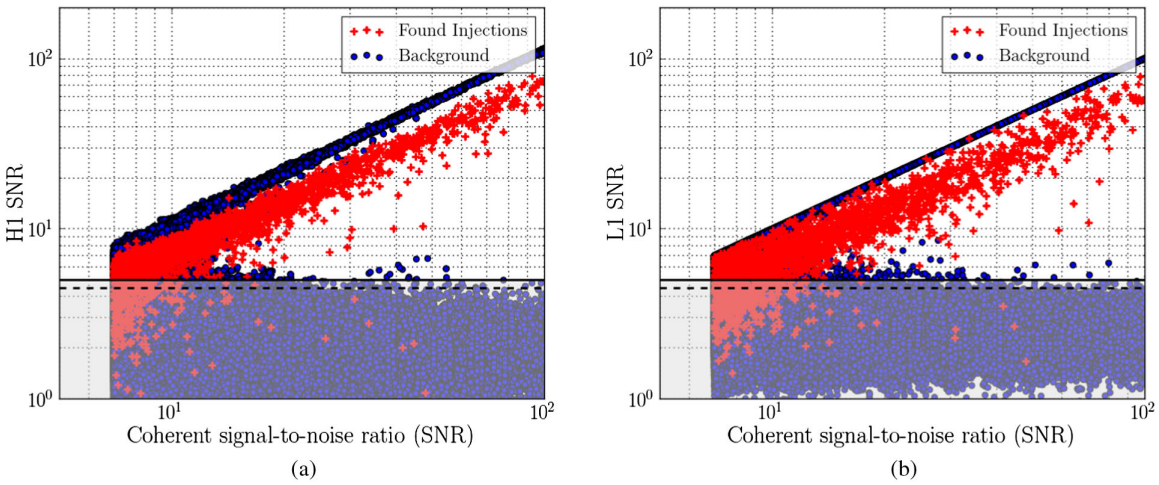

FIG. 8. The impact of the single-detector SNR cut on events from an all-sky coherent search. The blue dots are those from the noise background, while the red pluses are those from simulated BNS signals. The shaded region represents the single-detector cut as applied. All events with power only in a single-detector are vetoed as likely noise artefacts.

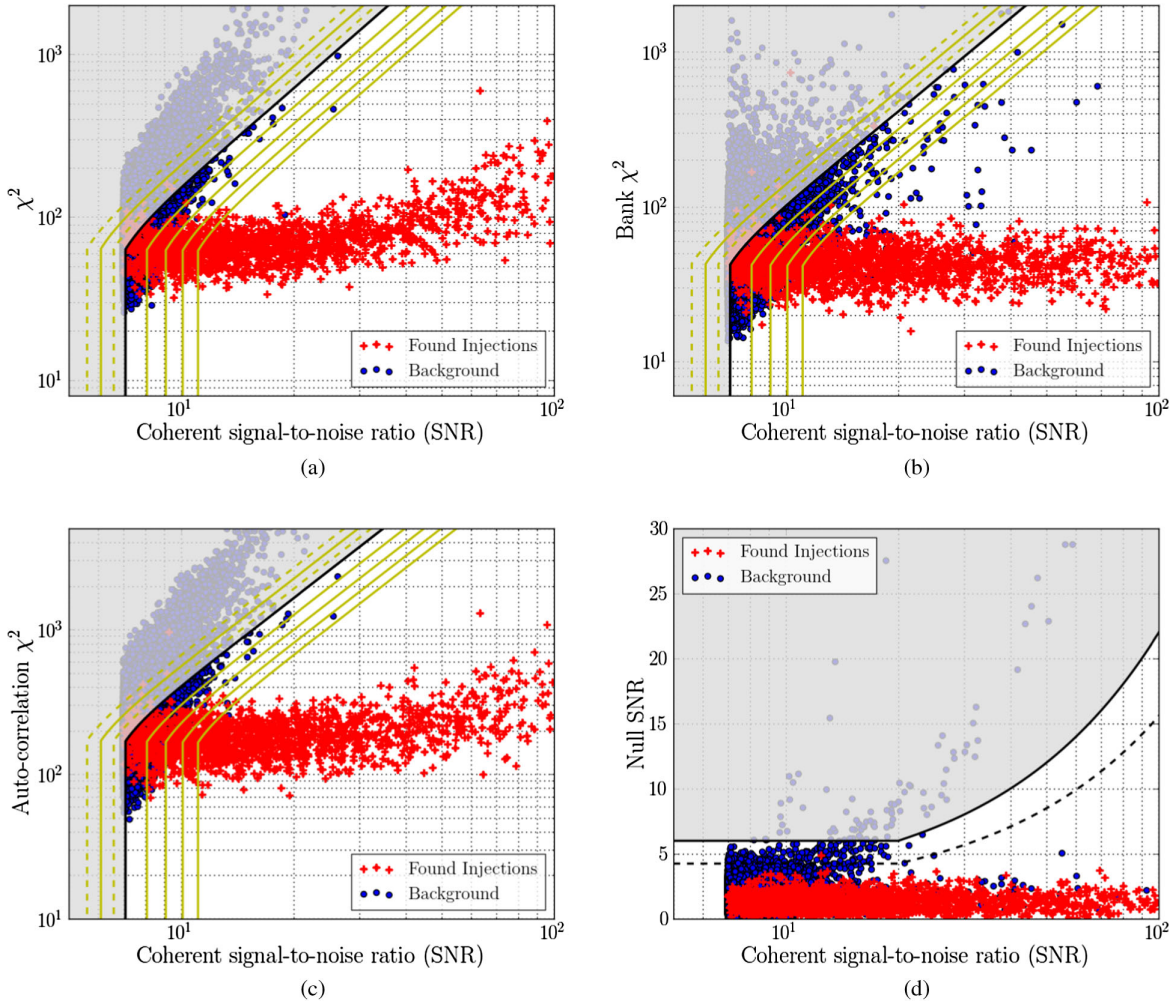

FIG. 9. The impact of signal-consistency cuts on the background of an all-sky coherent search. The blue dots are background events from time-slid analysis, and the red pluses are events from BNS simulations, and the shaded region covers those events failing the signal-consistency test. In figures (a) to (c), the contours represent constant values of the re-weighted SNR (dashed lines for half-integers, solid for integers). The detection statistic is constructed from the χ^2 re-weighted SNR, and the null SNR, figure (d), for which events above the dashed line are down-ranked using Eq. (23).

The final detection statistic is constructed using two signal-consistency statistics to down-rank likely noise events. First, we use the frequency-bin χ^2 test to reweight the SNR, using Eq. (20). The curves in Fig. 9a are lines of constant ρ_{χ^2} and show how high SNR noise events are downweighted while simulated signals have a reweighted SNR similar to their coherent SNR. Second, we use the null SNR to further reweight SNR into the detection statistic. We introduce an SNR dependent threshold, ξ_{null}:

$$\xi_{\text{null}} = \begin{cases} 4.25, & \rho \leq 20 \\ 4.25 + \left(\frac{\rho - 20}{5}\right), & \rho > 20 \end{cases} \tag{22}$$

and downweight any triggers for which the null SNR is greater than this value. The dashed line in Fig. 9d shows the threshold above which events are down-weighted based upon their null SNR. This downweighting reduces the significance of many background events, but affects only a handful of simulated signals.

The final detection statistic is

$$\rho_{\text{det}} = \begin{cases} \rho_{\chi^2}, & \rho_{\text{null}} \leq \xi_{\text{null}} \\ \frac{\rho_{\chi^2}}{\rho_{\text{null}} - (\xi_{\text{null}} - 1)}, & \rho_{\text{null}} > \xi_{\text{null}}. \end{cases} \tag{23}$$

2. Data quality cuts

Instrumental and environmental disturbances can lead to periods of poor data quality in the detectors. The signal consistency tests described above are used to mitigate the impact of these poor data. In addition, data quality (DQ) vetoes are also used to identify noise artefacts in the data, using instrumental and environmental correlations. We use the same data quality definitions as used in the coincidence search [49] (for a more detailed description of how they are

utilized, see [21,50]). Three different categories of data quality are generated. Data that are too poor to use at all, and would corrupt the PSD estimate are labelled *category 1*. When the impact is less severe, it is preferable to include the data in the analysis, and then remove any triggers at times of poor data quality. If we were to remove these data prior to filtering, we would lose an additional 64s on either side as discussed in Sec. III C. *Category 2* data quality vetoes identify times of instrumental problems with known correlations to the GW channel, while *category 3* vetoes identify times of poor data, often identified by statistical correlations with the GW channel.

These DQ vetoes are applied to data from each of the foreground, background and simulations such that if an event is vetoed in any one instrument then it is removed from the search [51]. Figure 10 shows the impact of category 2 and 3 DQ vetoes on the background events (from time slides), after the application of the signal-based vetoes. The category 2 vetoes are successful in removing the very loudest events, with the loudest event reduced from 13.9 to 11.8. In applying the category 2 vetoes, we discard around 1% of the available data. Category 3 vetoes are effective at removing the remaining tail of loud events in the coherent search, with the loudest event reduced to a detection statistic value of 9.3. However, to achieve this, we lose 25% of the available data since a DQ veto in any one of the three detectors leads to the data being discarded from the coherent analysis. One may question whether there is benefit in applying the vetoes if they lead to the removal of such a large amount of data. We have removed 25% of the data and succeeded in reducing the SNR of the loudest background events by a similar amount. This equates to improving the distance reach of the search by the same amount. Since sources are expected to be uniformly distributed in time and *volume*, imposing DQ cuts has improved the sensitivity of the search by around 50%.

The amount of time removed due to the data quality vetoes is still too high to be acceptable in a real search. However, for much of this time, the data quality is poor in only one detector. Therefore, it is possible to "recover" this time by performing a two-detector analysis on this data. Indeed, this is the procedure that has been followed in the coincidence search [21,49,50]. As we have noted previously, there is no reason for a two detector analysis to be performed coherently, so it would be natural to run the coincidence search over these times.

We can qualify the overall impact of the combined signal-based and data-quality vetoes, and the effectiveness of the chosen detection statistic, by comparing the high coherent SNRs seen in Fig. 8 to the final distribution in Fig. 10.

IV. SEARCH PERFORMANCE

The performance of the search is measured using the results of the simulation run, after all signal consistency tests and data quality cuts have been applied. All simulations for which no event was recorded are classed as missed. Those simulations with an associated trigger with a larger value of the detection statistic than all of the background events are classed as recovered, while those events with an associated trigger that is not louder than all background events are "marginally" recovered.

Figure 11 shows the distribution of injections and their recovery as a function of the injected decisive distance and chirp mass. We use these two parameters as they best encode the expected sensitivity of the search. The overall amplitude of the GW signal scales with the chirp mass,

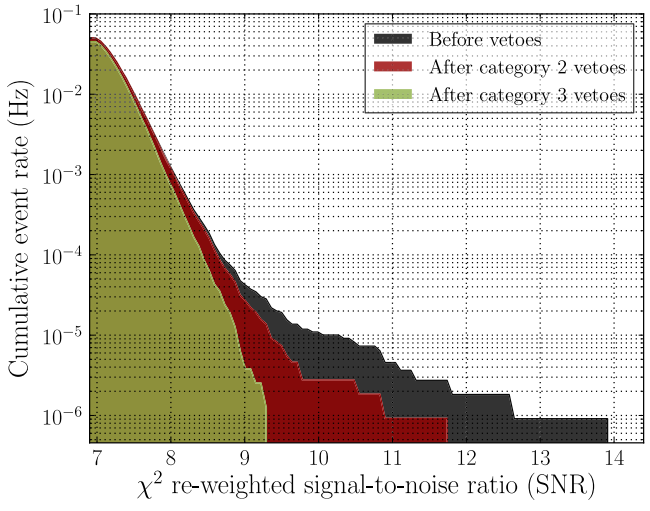

FIG. 10. The impact of data quality vetoes on the background of an all-sky coherent search. As seen, the loudest event is reduced from a reweighted SNR of 13.9–9.3.

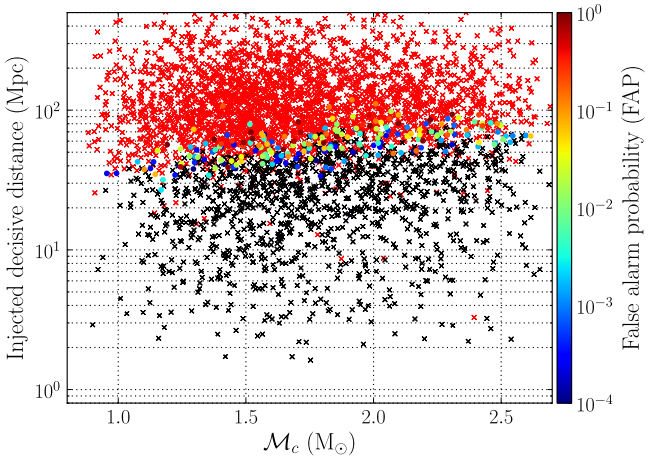

FIG. 11. Recovery of simulated BNS signals during the S5 coherent all-sky analysis. Successfully recovered signals as marked with black crosses, missed signals with red crosses, and marginally-recovered signals with dots colored by false alarm probability. The "decisive" distance is the second-largest effective distance for the network, as detailed in the text.

$\mathcal{M} = m_1^{3/5} m_2^{3/5} (m_1 + m_2)^{-1/5}$, of the system. The sensitivity of the search is typically limited by the SNR of the second most sensitive detector, and this is encoded in the decisive distance. The effective distance of a source is the distance at which an optimally-oriented and located signal would have the produced the same SNR as the given simulation [52]. The decisive distance is the second largest effective distance for the detectors in the network. Given a requirement of power in at least two detectors, the ability to detect an event will depend upon its decisive distance.

At lower masses, the majority of simulations injected below 30 Mpc (∼13 Mpc angle-averaged range) are successfully recovered, consistent with the network sensitive distance during S5 (Fig. 6),[2] with recovery improving as mass increases. With the background highly cleaned by the myriad cuts and vetoes, resulting in a low-significance loudest event, very few simulations are marginally recovered, with the transition rapidly made to completely missed signals at higher distances.

A. Comparison with the coincidence-based pipeline

We can compare the sensitivity of the coherent analysis with that achieved by the coincidence-based pipeline used to search data from the initial LIGO and Virgo detectors [3,40,49]. We have run the coincidence analysis with the same configuration as used in the LIGO-Virgo search results, to allow for a fair comparison between the coherent analysis and the coincidence-based `ihope` analysis, which is described in detail in [21]. Briefly, the search uses independent template banks in the operational detectors to identify triggers above a given SNR threshold in a single detector. Coincidences are formed from triggers which have comparable mass parameters and whose arrival time at the two detectors are consistent with an astrophysical origin [53]. For coincident triggers, a frequency-bin χ^2 test is performed and a reweighted SNR calculated for each detector that recorded the trigger. The final detection statistic is the quadrature sum of the reweighted SNRs in the individual detectors. Data quality cuts are applied in the same way as for the coherent analysis described above. Background is estimated through time-shifting triggers and repeating the coincidence and χ^2 stages of the pipeline, while simulated signals are used to test the sensitivity of the search. We note that, as independent template banks are used in this coincidence pipeline, the coincident SNR is not the same as in Eq. (12) because the template filtered in each detector will be different. "Exact-match" coincident searches, using the same template bank in each detector, and demanding that coincidences are formed only from triggers occurring from the same template in multiple detectors, are being developed for use in the advanced detector era [54,55]. Comparisons between "exact-match" coincident searches and searches using independent template banks in each detector are discussed in [18,19,55].

To compare the sensitivity of the coincidence-based and coherent analysis, an identical set of simulated signals were added to the data analysed by the two searches. These simulated signals are used to calculate the percentage of signals detected as a function of distance and, by integrating over distance, we obtain the cumulative volume to which the searches are sensitive. In order to perform the comparisons, the same data segments were analysed using the coincidence-based `ihope` pipeline [21]. The same data quality veto method was used, whereby those events vetoed in a single detector are removed from the search. Finally, the coincidence search was performed using identical template parameters and simulation parameters, allowing a direct comparison of search efficiency with the new coherent pipeline.

1. Computational cost

The coherent and coincident algorithms are both limited by the speed of the FFT—the computational core of the matched-filter. The coherent search has been implemented to ensure that the computation of the coherent SNR time series, even when considering a large number of sky points, is dominated by the FFT computations needed to obtain the single detector SNR time series. Therefore, in Gaussian noise, where coherent SNR is the optimal detection statistic, the computational cost of our coherent search and a coincidence search would be roughly equivalent. However, to counteract non-Gaussianities in the data we compute a set of χ^2 statistics, which themselves include FFT computations. As the number of sky-points increases, the number of times the χ^2 statistics must be computed also increases, and for large sky grids this can be the dominant computational cost. The coherent search has been implemented to ensure that the CPU-intensive calculation of the χ^2 statistics is minimized by first applying single detector thresholds and by applying the cheaper signal consistency tests—null SNR, auto χ^2 and bank χ^2—before computing the expensive frequency-bin χ^2. Nevertheless, for the 2-site coherent search presented here we found that the algorithm was roughly a factor of 2 more expensive than its coincidence-based predecessor due to additional frequency-bin χ^2 calculations. Recent work has demonstrated that a non-FFT based implementation of the frequency-bin χ^2 can greatly reduce the computational cost of that operation in the coincidence search [23]. We plan to investigate whether a similar implementation can provide a similar improvement when applied to the coherent search.

The implementation of background estimation via time shifts used in this analysis is computationally costly. Since we permute the order of the data segments, it is necessary to recompute the single detector SNR time series for each time shift. Thus, each time-slide is computationally equivalent to the zero-lag foreground, resulting in a further factor of ten

[2]The angle-averaged range shown in Fig. 6 is calculated for a BNS with mass $m_1 = m_2 = 1.4\mathrm{M}_\odot$, for which $\mathcal{M}_c \simeq 1.2$.

increase in computational cost for this search. Performing time shifts in the coincidence analysis has only a small impact in the total computational cost [21]. So, with ten time shifts, the coherent search is a factor of twenty more costly than the coincidence search. Of course, ten background trials is nowhere near sufficient to estimate event significance to detection level, where we might require a false rate of one per hundred or thousand years [3]. It will be a challenge to achieve this in a coherent analysis. Recent developments [27] have led to the implementation of shorter time shifts which can be performed without the need to recalculate the single detector SNR. This allows for ∼30 unique background trials at little additional cost. Using this method, a search including 1000 trials, using the current implementation, would result in around a ×60 computational cost in moving from coincidence-based to coherent.

2. Signal recovery

Figure 12 compares the efficiency of simulation recovery between the two analyses. The top and center panels compare the injection recovery and cumulative volume respectively as functions of distance between the two searches; the bottom panel shows the relative volume improvement of the coherent search. At low distances, the coherent search recovers slightly fewer simulated signals than the coincident analysis, but the differences are not statistically significant. However, the near 25% improvement in sensitive volume over the full simulation campaign highlights the advantage of this coherent pipeline compared to the previously published algorithm.

All comparisons are performed at a fixed false rate of one event per analysis time, corresponding to an SNR of around 9.3. Ideally, we would present the sensitivity comparison as a function of the false alarm rate, between one per month (as shown in Fig. 12) and one per hundreds or thousands of years. However, as discussed above, the computational cost of the coherent search, as implemented, makes this impractical and we present the results at a single false alarm rate.

We can compare the increased sensitivity of the coherent search with expectations. The coherent search employs a lower single detector SNR threshold than the coincident search. Specifically, the S5 search used a threshold of 5.5 in each detector, while we have required an SNR above 5 in one detector and 4.5 in a second. Furthermore, in the coherent analysis, the SNR of the third detector will contribute to the coherent SNR, regardless of its amplitude. In the coincident search, the SNR of the third detector will only contribute if it is above threshold. For the S5 data, where H2 was roughly half as sensitive as H1 and L1, this means that a large fraction of signals will be below threshold in H2.

To estimate the impact of these different thresholds, we generate a large number of simulated signals, uniformly

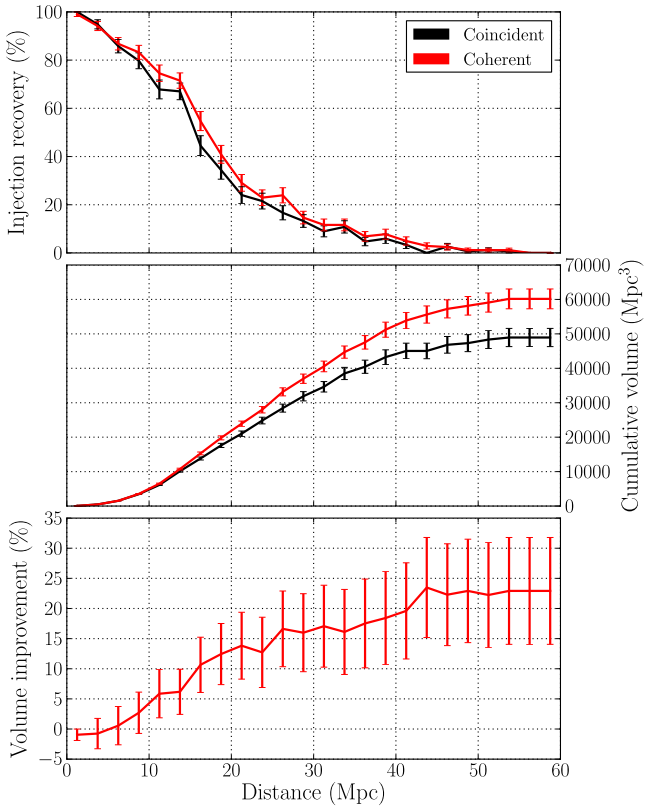

FIG. 12. Comparison of search performance between the coincident and coherent all-sky searches for one month of S5 data. The top panel shows the injection recovery (efficiency) as a function of distance, while the second panel shows the cumulative volume, comparing between the coincident (black) and coherent (red) searches. The bottom panel shows the relative volume improvement of the coherent search. The coherent search outperforms the coincident with nearly 25% greater volume sensitivity.

distributed in volume, with uniform binary orientation. For each, we calculate the expected SNR in each detector (ignoring noise contributions) and count the number which would be "detected" by the coincident and coherent searches. For the coherent search, we require the SNR to be greater than 5 in one detector, and greater than 4.5 in a second, with the combined SNR of 9.3 or more. In the coincident search, the SNR in each detector must be above 5.5 before it contributes in the combined SNR, which must be greater than 9.3. We do not account for the discreteness of the template bank, which will lead to a loss of SNR due to a mismatch between the signal and template waveforms. Similarly, we neglect the loss in SNR due to the discreteness of the sky grid used in the coherent search.

The majority of events that pass the coherent search threshold have a SNR below 5.5 in H2: 25% are observed above threshold in H2, compared to over 90% in both H1 and L1. This means that only 25% of the coherent sources are recovered as three detector coincidences. The remainder are observed in only two detectors (typically H1 and L1,

although there is a very small fraction that are seen only in H1 and H2). For these events, the SNR from the third detector is not included in the coincident network SNR, and consequently many of them will not be found above threshold. Of the simulated events observable by the coherent search, only 80% are observed by the coincident search. This is in excellent agreement with the results obtained by a full analysis on real data, where we find the coherent analysis to be 25% more sensitive.

B. Future prospects

We have demonstrated the benefits of performing a fully coherent search and shown that it leads to a 25% increase in sensitivity for the H1-H2-L1 network that operated during the initial detector era. However, since the H2 detector does not form part of advanced LIGO, this network will not operate in the future. Indeed, it seems unlikely that there will be two colocated detectors until the Einstein Telescope is operational [56]. Prior to that, the coherent analysis presented here will need to be extended to a three (or more) site analysis before it will be useful. Before undertaking this effort, it is worthwhile to investigate the likely benefits.

A planned evolution for the advanced LIGO and Virgo detectors is laid out at [11]. The first science run, in late 2015, involved only the advanced LIGO detectors. Following that, the advanced Virgo detector will join the network with a sensitivity around a half of the advanced LIGO detectors during early runs, rising to two thirds when the detectors are operating at design sensitivity at the end of the decade. A third advanced LIGO detector, located in India [13], is expected to be operational around 2022 with a similar sensitivity to the other LIGO detectors. The Japanese KAGRA [12] detector is being built and is expected to operate with a range similar to that of the advanced LIGO detectors. Therefore, the advanced detector network is expected to evolve from a two detector network in the first run to a five detector network by early next decade. Since the coherent analysis is not expected to benefit a two-detector search, we consider only networks of three or more detectors. Concretely, we evaluate the benefit of the coherent analysis for the LIGO-Virgo (HLV), LIGO-KAGRA-Virgo (HKLV), LIGO (with India)-Virgo (HILV) and LIGO (with India)-KAGRA-Virgo (HIKLV) network.

Table I provides the relative sensitivities of the coincident and coherent searches. As before, we generate a large number of simulated signals and identify those which would be observed by the searches. Again, we ignore any loss in SNR due to the discreteness of templates in the mass space or points in the sky grid. The sensitive volume available to each network is calculated assuming the only threshold is on the coherent network SNR. We evaluate the search sensitivity at a network threshold of 12, corresponding to a very low false rate, as might be expected for first detections, and 10, which might be more realistic during the routine detection era. For the coincident and coherent searches, we apply additional thresholds on the single detector SNRs and we calculate the fraction of sources which would be detected by the search. In the coherent search, we require an SNR above 5 in at least two detectors for the event to be detected, but the SNR from *all* detectors contributes to the network SNR. It is clear from Table I that this requirement has minimal impact on the search sensitivity. It is only for the five detector network that the single detector thresholds will reduce the detection rate by greater than 5%, and these events would be recovered if we were able to lower the single detector threshold to four.

TABLE I. The relative performance of the coherent and coincident searches for various future detector networks operating at their design sensitivity. For each network, we consider a detection threshold of 10 and 12. For the coherent search, we require two detectors to observe the signal above the SNR threshold of 5, but include the SNR from *all* detectors. For the coincident search, we only include the SNR contribution from detectors where the signal would be above threshold. The percentages give the sensitivity relative to a search which imposes only a threshold on the network SNR.

Network	Network threshold	Sensitive volume (10^6 Mpc^3)	Coherent search SNR 5	Coincident Search SNR 4	Coincident Search SNR 5.5
HLV	12	40	>99%	97%	90%
	10	65	>99%	94%	82%
HILV	12	55	>99%	94%	79%
	10	95	97%	87%	67%
HKLV	12	60	>99%	94%	80%
	10	103	96%	87%	67%
HIKLV	12	80	>99%	90%	70%
	10	135	93%	80%	57%

For the coincident search, we impose a single detector threshold of 5.5 in each detector and only those detectors with a signal above threshold contribute to the network SNR. This has a significant impact on the search sensitivity, with 10%/20%/30% of sources lost in a three/four/five detector search at SNR 12, increasing to 20%/30%/40% if the threshold can be lowered to 10. As an alternative to implementing a fully coherent analysis, one could simply lower the single detector threshold in the coincidence search. This would necessitate storing significantly more single detector triggers prior to performing the coincidence step. If the single detector threshold can be lowered to 4, then the majority of sources are recovered. Only with the five detector network, with observations being made at SNR 10, do we lose 20% of possible sources. We note that requiring a larger SNR threshold, say 5 or 5.5, in two detectors may help to reduce noise background and will have minimal effect on signals.

There are two additional effects that must be taken into account when doing a careful comparison of searches: the computational cost of the searches, and their noise

background. The sensitivity of the coincidence search will be reduced due to the fact that the noise contributions from all detectors are incorporated into the network SNR. Thus, in Gaussian noise, the coherent search background will be χ^2 distributed with four degrees of freedom, while for the coincidence search it will be χ^2 distributed with $2D$ degrees of freedom.[3] The impact of this is investigated in detail in [57], where noise background for the coincidence search is shown to be several orders of magnitude higher that the coherent search, at a fixed SNR. Thus, a comparison of the searches at *fixed false alarm rate* requires a higher threshold on the coincidence search, further reducing the sensitivity. However, as discussed above, the computational cost of the coherent analysis will be higher than the coincidence search. The computational cost of a search can be reduced by laying templates more sparsely in the mass-spin parameter space. This will lead to a loss in sensitivity as a signal is likely to have a poorer match with the closest template waveform. When computing resources are limited, comparison of the searches at *fixed computational cost* would favor the coincidence search. However, without a full implementation of the coherent analysis, it is not possible to perform the comparison. We have argued that the current implementation of the coherent analysis, with 1000 background trials, is around 60 times that of the coincidence search. Thus, strategies to reduce the computational cost of the coherent analysis are required before we can make such a comparison.

As an alternative to implementing the full coherent analysis, we could instead calculate the coherent SNR for coincident events observed in three or more detectors. The time delays between the detectors give a unique sky location (up to a reflection symmetry in the three detector case), and this determines the detector responses, $F^X_{+,\times}$. Then, given the complex SNR, z^X, from each detector, we can calculate the coherent SNR from Eq. (11), and also the null SNR from Eq. (19). This will reduce the noise background of the coincident search and provide the null SNR as an additional signal consistency test, and should yield many of the benefits of the coherent search, but with significantly lower overhead. We note, however, that the maxima of the single detector SNRs need not correspond to the maximum coherent SNR. Additionally, for four or more detectors, the measured time delays may not be consistent with a physical sky location. Thus, it may be necessary to keep a short stretch of the SNR time series around each trigger to reconstruct the coherent SNR [58] and, at this stage, it may be easier to simply implement the coherent search. Nonetheless, it is certainly worth investigating this approach, as it could give a significant boost to the sensitivity of coincidence searches at minimal additional computational cost. We note that this would not have been possible with the `ihope` pipeline [21] as it used different template banks for each detector. However, the analysis pipelines developed for the advanced detector era [54,55] do make use of a common template bank for all detectors.

[3]Since the searches impose single detector thresholds, the noise distribution will not be exactly χ^2 distributed, even in Gaussian noise. Nonetheless, the background from the coincidence search will be elevated, relative to the coherent search.

V. DISCUSSION

We have demonstrated the first implementation of a fully-coherent all-sky search for gravitational waves from the inspiral of two compact objects. This search extends the previously published targeted search for GW signals associated with short GRB [26,27] to the untargeted all-sky, all-time analysis. This fully-coherent, two-site search was seen to improve sensitive volume by nearly 25% over a coincidence-based search of the same data.

We have argued that a similar improvement is to be expected for three detector networks in the advanced detector era. The benefits for four and five detector networks are expected to be even greater. Nonetheless, the computational cost per template is significantly higher for the coherent analysis than the coincidence search. Additionally, estimating the noise background through a time-shifted analysis of the data further increases the cost of the coherent analysis. It will be difficult to obtain a background to "detection level" of one per hundred or thousand years using this method, and alternatives [59,60] may be needed. We have also discussed methods by which the sensitivity of the coincidence search could be enhanced, most notably by lowering the single detector thresholds (particularly on the least sensitive detectors), implementing a coherent follow up to all coincident events and incorporating the null SNR. These possibilities deserve detailed investigation, in order to determine the best way to implement a coherent analysis. In addition, the search presented here made use of much of the `ihope` infrastructure [21] used for the analysis of initial LIGO and Virgo data. In the meantime, there has been significant effort to modernize and optimize the coincident analysis [23,54]. Any coherent analysis of the advanced detector data will have to build upon this new analysis infrastructure.

The first advanced detector runs will feature only the two LIGO detectors and, for this network, a coherent analysis is equivalent to the coincident analysis. Nonetheless, the Virgo, KAGRA and LIGO India detectors will soon join the global network and, at this time, a coherent analysis has the potential to significantly increase the rate at which gravitational wave signals from binary mergers are observed.

ACKNOWLEDGMENTS

The authors would like to thank Patrick Sutton, Duncan Brown, B. Sathyaprakash, Valeriu Predoi, Sukanta Bose,

Chad Hanna, Patrick Brady, Frank Ohme, Erin Macdonald, James Clark, Andrew Williamson for discussions during the development of the coherent analysis presented here. In this work, D. M. M. was supported by the Science and Technology Facilities Council, UK, and by NSF Award No. PHY-1104371. S. F. would like to acknowledge the support of the Royal Society and STFC grant ST/L000962/1. I. H. acknowledges support from NSF Awards No. PHY-0854812, No. PHY-0847611, No. PHY-1205835 and thanks the Max Planck Gesellschaft for support.

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration), Rep. Prog. Phys. **72**, 076901 (2009).
[2] F. Acernese, P. Amico, M. Alshourbagy, F. Antonucci, S. Aoudia *et al.*, Classical Quantum Gravity **24**, S381 (2007).
[3] J. Abadie *et al.* (LIGO Collaboration, Virgo Collaboration), Phys. Rev. D **85**, 082002 (2012).
[4] J. Abadie *et al.* (LIGO Scientific Collaboration), Astrophys. J. **760**, 12 (2012).
[5] J. Aasi *et al.*, Phys. Rev. D **87**, 022002 (2013).
[6] J. Abadie *et al.* (VIRGO, LIGO Scientific), Phys. Rev. D **85**, 122007 (2012).
[7] J. Aasi *et al.* (VIRGO, LIGO Scientific), Phys. Rev. D **87**, 042001 (2013).
[8] B. P. Abbott *et al.* (LIGO Scientific Collaboration, VIRGO Collaboration), Nature (London) **460**, 990 (2009).
[9] J. Aasi *et al.* (The LIGO Scientific Collaboration), Class. Quantum Grav. **32**, 074001 (2015).
[10] F. Acernese *et al.* (Virgo), J. Phys. Conf. Ser. **610**, 012014 (2015).
[11] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Living Rev. Relativity **19**, 1 (2016).
[12] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto (The KAGRA Collaboration), Phys. Rev. D **88**, 043007 (2013).
[13] B. Iyer, T. Souradeep, C. Unnikrishnan, S. Dhurandhar, S. Raja, A. Kumar, and A. S. Sengupta, LIGO-India, Technical Report, 2011.
[14] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Classical Quantum Gravity **27**, 173001 (2010).
[15] S. Bose, A. Pai, and S. V. Dhurandhar, Int. J. Mod. Phys. D **09**, 325 (2000).
[16] L. S. Finn, Phys. Rev. D **63**, 102001 (2001).
[17] A. Pai, S. Dhurandhar, and S. Bose, Phys. Rev. D **64**, 042004 (2001).
[18] H. Mukhopadhyay, N. Sago, H. Tagoshi, H. Takahashi, and N. Kanda, Phys. Rev. D **74**, 083005 (2006).
[19] H. Mukhopadhyay, H. Tagoshi, S. Dhurandhar, and N. Kanda, Phys. Rev. D **80**, 123019 (2009).
[20] B. P. Abbott *et al.* (LIGO Scientific), Phys. Rev. D **80**, 102001 (2009).
[21] S. Babak, R. Biswas, P. R. Brady, D. Brown, K. Cannon *et al.*, Phys. Rev. D **87**, 024033 (2013).
[22] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Phys. Rev. D **85**, 122006 (2012).
[23] T. Dal Canton, A. H. Nitz, A. P. Lundgren, A. B. Nielsen, D. A. Brown *et al.*, Phys. Rev. D **90**, 082004 (2014).
[24] Y. Guersel and M. Tinto, Phys. Rev. D **40**, 3884 (1989).
[25] L. Wen and B. F. Schutz, Classical Quantum Gravity **22**, S1321 (2005).
[26] I. W. Harry and S. Fairhurst, Phys. Rev. D **83**, 084002 (2011).
[27] A. R. Williamson, C. Biwer, S. Fairhurst, I. W. Harry, E. Macdonald, D. Macleod, and V. Predoi, Phys. Rev. D **90**, 122004 (2014).
[28] J. Aasi *et al.* (The LIGO Scientific Collaboration, the Virgo Collaboration), Phys. Rev. Lett. **113**, 011102 (2014).
[29] L. Blanchet, Living Rev. Relativity **17**, 2 (2014).
[30] M. Hannam, Classical Quantum Gravity **26**, 114001 (2009).
[31] J. M. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter, Rev. Mod. Phys. **82**, 3069 (2010).
[32] L. Santamaria, F. Ohme, P. Ajith, B. Bruegmann, N. Dorband, M. Hannam, S. Husa, P. Moesta, D. Pollney, C. Reisswig, E. L. Robinson, J. Seiler, and B. Krishnan, Phys. Rev. D **82**, 064016 (2010).
[33] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroue, H. P. Pfeiffer, M. A. Scheel, B. Szilagyi, N. W. Taylor, and A. Zenginoglu, Phys. Rev. D **89**, 061502 (2014).
[34] C. Van Den Broeck and A. S. Sengupta, Classical Quantum Gravity **24**, 155 (2007).
[35] C. Capano, Y. Pan, and A. Buonanno, Phys. Rev. D **89**, 102003 (2014).
[36] V. Varma, P. Ajith, S. Husa, J. C. Bustillo, M. Hannam, and M. Pürrer, Phys. Rev. D **90**, 124004 (2014).
[37] P. Jaranowski, A. Krolak, and B. F. Schutz, Phys. Rev. D **58**, 063001 (1998).
[38] D. Keppel, arXiv:1307.4158.
[39] S. Fairhurst, New J. Phys. **11**, 123006 (2009).
[40] J. Abadie *et al.* (LIGO Scientific Collaboration, Virgo Collaboration), Phys. Rev. D **82**, 102001 (2010).
[41] F. Cavalier, M. Barsuglia, M.-A. Bizouard, V. Brisson, A.-C. Clapson, M. Davier, P. Hello, S. Kreckelbergh, N. Leroy, and M. Varvella, Phys. Rev. D **74**, 082004 (2006).
[42] J. Veitch, I. Mandel, B. Aylott, B. Farr, V. Raymond, C. Rodriguez, M. van der Sluys, V. Kalogera, and A. Vecchio, Phys. Rev. D **85**, 104045 (2012).
[43] O. Rabaste, E. Chassande-Mottin, and A. Pai, arXiv:0905.4832.
[44] B. J. Owen and B. S. Sathyaprakash, Phys. Rev. D **60**, 022002 (1999).
[45] T. Cokelaer, Phys. Rev. D **76**, 102004 (2007).
[46] B. Allen, Phys. Rev. D **71**, 062001 (2005).
[47] C. Hanna, Ph.D. thesis, LSU, Baton Rouge, LA, 2008.

[48] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz, Phys. Rev. D **86**, 084017 (2012).
[49] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **80**, 047101 (2009).
[50] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **79**, 122001 (2009).
[51] J. Slutsky, L. Blackburn, D. A. Brown, L. Cadonati, J. Cain, M. Cavaglia, S. Chatterji, N. Christensen, M. Coughlin, S. Desai, G. González, T. Isogai, E. Katsavounidis, B. Rankins, T. Reed, K. Riles, P. Shawhan, J. R. Smith, N. Zotov, and J. Zweizig, Classical Quantum Gravity **27**, 165023 (2010).
[52] P. R. Brady and S. Fairhurst, Classical Quantum Gravity **25**, 105002 (2008).
[53] C. A. K. Robinson, B. S. Sathyaprakash, and A. S. Sengupta, Phys. Rev. D **78**, 062002 (2008).
[54] K. Cannon *et al.*, Astrophys. J. **748**, 136 (2012).
[55] S. A. Usman *et al.*, arXiv:1508.02357.
[56] M. Punturo *et al.*, Classical Quantum Gravity **27**, 194002 (2010).
[57] T. Dal Canton, Ph.D. thesis, Leibniz University Hannover, Hannover, Germany, 2015.
[58] S. Bose, T. Dayanga, S. Ghosh, and D. Talukder, Classical Quantum Gravity **28**, 134009 (2011).
[59] K. Cannon, C. Hanna, and D. Keppel, Phys. Rev. D **88**, 024025 (2013).
[60] K. Cannon, C. Hanna, and J. Peoples, arXiv:1504.04632.