LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	**LIGO–T1500055** **VIR-0093A-15**

The LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics
(2015-2016 edition)

The LSC-Virgo Search Groups, the Data Analysis Software Working Group,
the Detector Characterization Working Group and the Computing Committee

WWW: http://www.ligo.org/ and http://www.virgo.infn.it

Processed with LaTeX on 2015/12/03

Contents

1 The LSC-Virgo White Paper on Data Analysis

Gravitational wave searches and astrophysics in the LIGO Scientific Collaboration (LSC) and Virgo collaboration are organized by astrophysical source classification into four working groups. The **Compact Binary Coalescence (CBC)** group searches for signals for merging neutron stars or black holes by filtering the data with waveform templates. The **Burst (Burst)** group searches for generic gravitational wave transients with minimal assumption on the source or signal morphology. The **Continuous Waves (CW)** group targets periodic signatures from rotating neutron stars. The **Stochastic (SGWB)** group looks for a gravitational wave background of cosmological or astrophysical origin. Joint teams across two or more working groups exist where the science suggests overlap between sources or methods. In addition, the **Detector Characterization (Detchar)** group collaborates with the detector commissioning teams and works to improve searches by identifying and mitigating noise sources that limit sensitivity to astrophysical signals.

The *LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics*, which is updated yearly, describes the astrophysical search plans of the LSC-Virgo working groups. This document is its executive summary. For each group, it provides a mission statement and scientific priorities in the Advanced Detector Era, as well as statements from Detector Characterization, Calibration and Hardware Injection teams.

We refer to the Advanced Detector Era (ADE) as the epoch of Advanced LIGO and Advanced Virgo science data acquisition, scheduled to start in the second half of 2015. Table 1 shows the planned schedule of science runs, as provided by the LSC-Virgo Joint Running Plan Committee, which includes representatives from the laboratories, the commissioning teams and search groups.

Epoch	Estimated Run Duration	Run Name	$E_{\rm GW} = 10^{-2} M_{\odot} c^2$ Burst Range (Mpc)		Binary Neutron Star Range (Mpc)		Number of Binary Neutron Star Detections
			LIGO	Virgo	LIGO	Virgo	
2015	3 months	O1	40 – 60	–	40 – 80	–	0.0004 – 3
2016–17	6 months	O2	60 – 75	20 – 40	80 – 120	20 – 60	0.006 – 20
2017–18	9 months	O3	75 – 90	40 – 50	120 – 170	60 – 85	0.04 – 100

Table 1: Plausible observing schedule and expected sensitivities with the advanced LIGO and Virgo detectors, which will be strongly dependent on their commissioning progress. The two LIGO detectors meet the minimum requirements for the first “O1" run now. [arXiv:1304.0670]

The LSC-Virgo scientific priorities for ADE observations are summarized in Table 2, by search group, in three categories:

- **Highest priority:** searches most likely to make detections or yield significant astrophysical results.
- **High priority:** promising extensions of the highest priority goals that explore larger regions of parameter space or can further the science potential of LIGO and Virgo.
- **Additional priority:** sources with low detection probability but high scientific payoff.

Computing needs and resource allocations are derived from the science priorities presented in this table. Scientific motivations, details on methods and the strategy for result validation are provided in the search plans that constitute the white paper.

We note that the LSC-Virgo Collaboration has adopted a *Multiple Pipeline Policy* [LIGO-M1500027], which calls for all astrophysical results to be validated with a different analysis, using independent methods and tools when possible. In some cases this may require the same data to be analyzed by more than one pipeline for the same science target.

	Burst	CBC	CW	SGWB
Highest priority	All-sky search for generic GW transients, in low latency for EM followup and deep, offline for 4σ detection confidence	All-sky matched-filter search for binary neutron star (BNS) systems, deep and low latency	All-sky search for isolated neutron stars, both as a *quick-look* on owned resources and as a deep/broad search on Einstein@Home	Directional search for stochastic GW background
	Parameter estimation for the astrophysical interpretation of detected burst events	All-sky matched filter search for binary neutron-star and black-hole (NSBH) systems, deep and low latency	Targeted search for high value, known pulsars	Isotropic search for stochastic GW background
	Search for GW bursts triggered by outstanding GRB alerts	All-sky matched-filter, deep search for binary black-hole (BBH) systems	Directed searches for Cas-A	Constraints of a detected background of astrophysical origin with long transients
	Searches triggered by outstanding astrophysical events (a galactic supernova, neutron star transients, an exceptional high energy neutrino alert)	Parameter estimation of detected CBC events	Directed searches for X-ray binaries SCO-X1 and J1751-305	
	Search for cosmic string kinks and cusps	CBC searches triggered by all GRB alerts		
		Tests of General Relativity with CBC events		
High priority	Searches triggered by high energy neutrinos, extra-galactic supernovae, and GRB observations	All sky search for spinning binary neutron star systems (deep and low latency)	Targeted search for other known pulsars	Long transient follow up of CBC and burst candidates
	Burst search for intermediate mass and eccentric black hole binary systems	Matched filtered search for intermediate mass black hole binary systems	Directed searches for other isolated stars and X-ray binaries	
	All-sky search for long bursts of > 10 s duration			
Additional priority	GRB-triggered search for long-duration bursts and plateaus	Exploring effects of detector noise on parameter estimation	All sky search for isolated stars (alternative approaches)	
	Hypermassive neutron star followup		All-sky search for binaries	
	Burst searches triggered by radio transients and by SGR/SGR-QPO		Spotlight deep sky-patch search **	
	Burst tests of alternative gravity theories **		Search for Supernova post birth signals **	
			Search for continuous wave transients **	

Table 2: Science priorities of the LIGO-Virgo collaboration, for the four astrophysics search groups: Bursts, Compact Binary Coalescences (CBC), Continuous Waves (CW), and Stochastic Gravitational Wave Background (SGWB). The targets are grouped in three categories (highest priority, high priority, additional priority), based on their detection potential with Advanced Detectors. There is no additional ranking within each category in this table. Critical for accomplishing these science priorities are the detector characterization, calibration and injection activities described in this document.
** Future searches under development, not included in ongoing production computing requests.

1.1 Searches for Generic Transients, or Bursts

The mission of the Burst group is to detect gravitational wave transients, or *bursts*, and gain new information on populations and emission mechanisms of astrophysical objects, as well as to test theories of gravity. Central to the Burst group philosophy is the assumption of minimal information on the source, so that searches for gravitational wave bursts typically do not require a well-known or accurate waveform model and are robust against uncertainties in the gravitational wave signature. Burst searches thus have the potential to see events that other groups cannot. We refer to this as the “eyes wide open" approach.

For example: the complexity of Supernovae makes it difficult to reliably map the dynamics of a core-collapse into a gravitational wave signal. The merger of precessing intermediate-mass black holes ($\geq 100\,\mathrm{M}_{\odot}$) produces gravitational wave transients which appear as short, sub-second bursts in the data. Long gamma-ray bursts could be associated with a gravitational wave transient lasting more than 10 seconds. Since robust models are not available for many plausible sources, we need data analysis methods that are able to detect emission mechanisms that have not been envisioned yet.

The Burst group implements a variety of methods to identify instances of statistically significant excess power, localized in the time-frequency domain. To discriminate between gravitational waves and noise fluctuations, the analysis requires the signal to appear coherently in multiple detectors. The confidence of a candidate event is established by repeating the analysis on many instances of background, obtained by shifting the data from different detectors with non-physical delays. In a few special cases when an accurate signal model is available, such as for cosmic string cusps or neutron star ring-downs, a search can be done using matched filtering with a bank of templates.

Although burst search algorithms are designed to detect a wide range of signals, their tuning and interpretation benefit from considering how they perform for plausible astrophysical signals. Therefore, the group's science program involves an active collaboration with the theoretical astrophysics, source modeling and numerical relativity communities.

Many gravitational wave burst sources should also be observable in more traditional channels, from Earth-based astronomical data, through sensitive GRB/X-ray satellite detections, to neutrino signals. Knowledge of the time and/or sky position of the astrophysical event producing a gravitational wave burst can be used to increase the sensitivity of a triggered burst search compared to an untriggered, all-sky search, and the association with a known astrophysical event may be critical in establishing our confidence in a gravitational-wave burst detection. Most importantly, joint studies of complementary data enable scientific insight that cannot be accessed through gravitational waves or other messengers alone. Therefore, in addition to searches using only the gravitational wave data, a significant part of the Burst group's science program involves connecting with other observations and working closely with the astronomy and astrophysics communities.

1. **Highest priority**

   The Burst group is focused on an *eyes wide open* approach to detecting gravitational wave transients. To maximize its discovery potential, the Burst group employs a strategy of multiple searches, overlapping in parameter space to allow for cross-validation of search outputs. Highest priority goals for the analysis of advanced detector data include:

   - a statement on the transient gravitational wave sky, with population studies if we have several detections, a rare-event detection significance if we have one candidate or an upper limit on the rate of gravitational wave bursts if there is no detection;
   - deployment of multiple analyses for cross validation of the all-sky search results, including verifying the significance of any observed events, across a wide parameter space. This is especially

important for events that are not matched to a specific source model;
- the astrophysical interpretation of any detected signals, leveraging signal characterization and parameter estimation;
- a prompt analysis, trigger production and sky localization, to enable the electromagnetic follow-up of gravitational wave transients;
- prompt reports on astrophysically significant events, such as nearby gamma ray bursts, soft gamma repeater hyperflares, galactic supernovae as well as exceptional bursts of low (MeV) or high (GeV–PeV) energy neutrinos;
- a dedicated search for gravitational wave bursts originating from cosmic strings.

2. **High priority**

   The Burst group will extend the parameter space of the all-sky search to include longer duration transients ($\geq 10\,\mathrm{s}$) which may originate from various astrophysical sources such as long gamma-ray bursts. Long-duration burst searches share similar complexities with their short-duration counterparts. Since the long-duration search is not as mature as the short-duration one, multiple analyses will be deployed to cross-validate the results.

   The Burst group will also pursue, with the burst analysis approach, some classes of compact binary coalescence sources that are not well covered by the current waveform template banks. These include intermediate mass binary black holes, binary black holes with eccentric orbits and intermediate mass ratio inspirals.

   Finally, the Burst group will pursue multi-messenger searches for gravitational wave bursts in conjunction with signatures such as generic gamma ray bursts, fast radio transients, low- and high-energy neutrino observations, and electromagnetic observations of nearby core-collapse supernovae. The Burst group will use information on the astrophysical event to reduce the parameter space over which searches must be performed, leading to a reduction in the false alarm rate and, consequently, an improvement in search sensitivities.

3. **Additional Priority**

   Additional priorities include the search for gravitational waves in association with neutron star transients (eg. pulsar glitches, type I X-ray bursts and soft gamma ray repeater flares) and testing alternative theories of gravity with gravitational wave bursts.

Several of these science targets – intermediate mass black hole binaries, GRBs, electromagnetic followup – overlap with the CBC group, and joint teams are working together across the two groups on these targets.

1.2 Searches for Signals from Compact Binary Coalescences

The inspiral and merger of a binary containing stellar-mass compact objects (i.e., neutron stars and black holes) generates gravitational waves which sweep upward in frequency and amplitude through the sensitive band of ground-based gravitational-wave detectors. The highly relativistic speeds and strongly-curved space-times of compact object mergers generate gravitational waves that encode the dynamics of strong-field gravity. With extreme densities of matter completely inaccessible to terrestrial experiments, mergers involving neutron stars hold the key to understanding the equation of state of nuclear matter. Compact object mergers may also explain the origin and distribution of rare heavy elements and reveal the engine powering gamma-ray bursts. Measuring the masses and spins of a population of compact objects in the Universe can help explain how stellar collapse forms neutron stars and black holes.

At design sensitivity, Advanced LIGO will be able to detect binary neutron star (BNS) mergers to an angle-averaged range of ~ 180 Mpc, neutron star–black hole (NSBH) binaries to ~ 450 Mpc, and stellar-mass binary black holes (BBHs) at luminosity distances over 900 Mpc. LIGO and Virgo conduct their searches jointly with a three-detector network that can be used to localize sources on the sky through methods akin to triangulation. A wide variety of electromagnetic counterparts are expected to accompany the gravitational waves from compact object mergers, ranging from radio, through optical to x-rays and gamma-rays. The joint observation of a source by LIGO, Virgo, high-energy satellites, optical, and radio observatories will be a watershed event in astrophysics.

The scientific program of the LSC/Virgo Compact Binary Coalescence (CBC) Group aims to identify gravitational wave signals from compact binary sources in the detector data, measure the waveform parameters, and use detected signals to study the nature of gravity and the astrophysics of nature's most compact objects. This requires accurate modeling of gravitational wave sources to maximize detection rates, and to accurately measure parameters. The CBC group has an active collaboration with the theoretical astrophysics, source modeling and numerical relativity communities to address these challenges.

1. **Highest priority**

   The detection of gravitational waves from compact binary coalescence using the LIGO and Virgo detectors is the main goal of the CBC group. The highest-priority sources are systems containing neutron stars and/or stellar-mass black holes.

   Although a coincident electromagnetic (EM) counterpart is not required to detect gravitational waves from compact binary coalescence, the coincident detection of an EM counterpart with a CBC event would add significant astrophysical information to our discoveries. Therefore, all-sky low-latency ($\sim$ 1 minute) CBC searches will provide rapid alerts for GW detections in order to enable observation of e.g., prompt x-ray and optical afterglows as well as to form a first response to gamma-ray bursts. Gamma-ray burst alerts will also trigger a deep, coherent CBC search targeting the GRB sky position, which will complete within ~ 2 hours and will allow us to test directly whether or not compact binary mergers are gamma ray burst progenitors. Low latency analysis activities include the development of low latency analysis pipelines, low-latency data quality assessment, low-latency significance estimation, sky localization for CBC sources, and joint gravitational wave/EM analyses in the advanced detector era.

   Once CBC sources have been detected, a significant amount of astrophysics can be extracted from the observed gravitational waveforms. The first few detections will allow us to make precise measurement of masses and spins to understand the properties of compact objects and their formation and to make accurate measurement of coalescence rates for CBC sources. After many detections we will constrain the neutron star equation of state, test the genuinely strong-field dynamics of space-time - a regime which can only be probed via direct gravitational wave detection - and conduct cosmological studies without the need for a cosmic distance ladder.

Achieving these goals requires LSC/Virgo scientists in the CBC group to prioritize: data quality, search pipeline development, rates and significance measurement, waveform development, and parameter estimation for detected sources.

2. **High priority**

   High priorities include expanding the CBC search for binary black holes beyond stellar mass, e.g. intermediate mass ratio inspirals, intermediate-mass binary black holes and eccentric binaries, which will necessitate the development of new data analysis algorithms and the implementation of associated template waveforms. Additionally, searching for neutron star binaries with significant component spin is also a high priority. Although neutron stars in binary systems have been observed to have small spin, some isolated neutron stars are known to spin significantly. If neutron stars with significant spins do exist in binary systems, then opportunities to detect them could be lost without a dedicated search.

3. **Additional priority**

   Building more accurate noise models for parameter estimation techniques can dramatically mitigate the effects of non-stationary, non-Gaussian noise on the fidelity of parameter inference. It is a priority to conduct a simulation campaign to study improved noise models for parameter estimation.

   The compact binary parameter space searched in higher priorities is not complete. It covers a plausible range of physical parameters based on observation and stellar evolution models with known detection techniques. However, there are other interesting but less plausible parameter spaces which would have a dramatic impact if discovered. Given additional resources, we would consider searching for compact objects below 1 solar mass. It is possible that neutron stars or black holes could exist with masses down to fractions of a solar mass and be in detectable binary systems. Additionally, the higher priority searches do not include template waveforms that take into account orbital precession. Currently, there are no complete methods to search for precessing binaries, however, work is ongoing to develop such a search and with additional resources we would conduct a precessing binary search in the future.

1.3 Searches for Continuous-wave Signals

The LSC/Virgo Continuous Waves (CW) Group aims to measure gravitational wave signals that are long-lived, nearly sinusoidal and extremely weak, believed to be emitted by rapidly rotating neutron stars in our galaxy. These stars can emit gravitational radiation through a variety of mechanisms, including elastic deformations, magnetic deformations, unstable r-mode oscillations, and free precession, all of which operate differently in accreting and non-accreting stars. Long-term simultaneous gravitational wave and electromagnetic observations of a galactic neutron star would support a rich astrophysical research program.

For known pulsars with measured spin frequencies, frequency derivatives and distances, energy conservation allows setting an upper limit on gravitational wave strain amplitude, known as the *spindown* limit, albeit with significant uncertainties due to poorly understood neutron star astrophysics. Previous searches in LIGO and Virgo data have obtained 95% confidence upper limits well below the spindown limits for the Crab and Vela pulsars. As interferometer sensitivities improve in the Advanced Detector Era, several dozen more known pulsars will become spindown-accessible, primarily at spin frequencies below 100 Hz. For suspected neutron stars with unknown spin frequencies, indirect upper limits based on estimated age or on estimated accretion rates can also be derived. Such indirect limits are more optimistic for non-accreting stars, but accreting neutron stars are more likely to be emitting near their limits.

Because there is so much astrophysical uncertainty in continuous gravitational wave emission and because electromagnetic astronomers have detected fewer than 2500 of the O(10^{8-9}) neutron stars believed to populate our galaxy, the CW group has established a broad program to search for gravitational wave emission from five distinct source categories, ordered below by decreasing *a priori* information known about the sources: 1) known pulsars with well measured timing; 2) other known or suspected isolated neutron stars with limited or no timing information; 3) known or suspected binary neutron star systems; 4) unknown isolated stars in any direction; and 5) unknown binary stars in any direction.

This ordering of categories corresponds to ordering by source strain sensitivity. Targeted searches using known ephemerides from radio, X-ray or γ-ray timing measurements can achieve strain sensitivities limited only by the intrinsic detector sensitivity and observation time spans with minimal trials factor corrections. Directed searches using known sky locations but having no *a priori* frequency information (e.g., *Cassiopeia A*) are degraded by trials factors that depend on the band size searched and on the assumed age of the source (which affects the number and range of higher-order spin derivatives to be searched). The sensitivity achievable with all-sky searches is still further limited by the need to make sky-location-dependent corrections for Doppler modulations of detected source frequency due to the Earth's motion (daily rotation and orbital motion). The number of sky points to search to maintain accurate demodulation grows rapidly with coherence time used in the search (time scale over which the signal is assumed to follow a precise phase model). The effect is severe enough to preclude all-sky searches using coherence times equal to the full observation spans of data runs. Adopting semi-coherent summing of data makes the computational problem tractable, but sacrifices additional sensitivity beyond that from the trials factor of exploring a larger parameter space. Directed searches for suspected neutron stars in binary systems with unknown source frequency must make similar sensitivity tradeoffs, and all-sky searches for sources in unknown binary systems define the current extreme in sensitivity tradeoff for tractability.

In the case of known objects, we have identified sources that seem to be the most promising, and should priorities need to be set because of limited resources (labor or computing), those sources will receive the highest priority. With these considerations in mind, the CW group plans a comprehensive search program in the Advanced Detector Era for all of these source categories, with the following priorities:

1. **Highest priority**
    - Targeted searches for the Crab and Vela pulsars and any other stars for which the spindown limit is likely to be beaten to within a factor of two. High-interest stars likely to fall in this category include PSR J0537–6910 and PSRJ1813–1246, among many others, as detector sensitivities improve. These analyses will include searching at the stellar spin frequency and twice that frequency.
    - Directed search for Cassiopeia A which is the youngest known neutron star in the galaxy, but for which the spin frequency is unknown.
    - Directed searches for the X-ray binaries Scorpius X–1, Cygnus X–3, PSR J1751–305 and 4U 1636-536. The first two are especially bright in X-rays, and in the torque-balance model, GW luminosity scales with X-ray luminosity, while there is evidence in the last two objects for sharp X-ray periodicities that may indicate an r-mode oscillation.
    - All-sky searches for unknown isolated stars. These searches necessarily suffer from degraded strain sensitivity relative to what can be achieved in the targeted and directed searches, but they cast a very wide net, offering a reasonable prospect of discovery.
2. **High priority**
    - Targeted searches for known pulsars for which the spindown limit is unlikely to be beaten, according to conventional theory, but which are extreme astrophysical objects of great interest.
    - Directed searches for young supernova remnants other than Cassiopeia A, including Supernova 1987A, for sources near the galactic center, for sources in nearby globular clusters and for unidentified γ-ray sources with pulsar-like spectra.
    - Directed searches for additional X-ray binaries.
3. **Additional priority**
    - All-sky searches for unknown binary stars. Because of the additional unknown orbital parameter space to search, these searches are most computationally demanding and must make the greatest tradeoffs in strain sensitivity for tractability.
    - All-sky searches for unknown isolated stars, using alternative algorithms.

For every type of search, the CW group supports at least two independent methods (pipelines). This redundancy provides greater robustness against incorrect assumptions in signal modeling and against non-optimum handling of instrumental artifacts. The robustness against incorrect signal modeling is especially important for accreting sources, such as Scorpius X–1, where the time span over which the coherence of the signal model can be safely assumed is uncertain. In fact, that time scale is likely to vary in response to fluctuations in accretion rate.

There is some overlap in the CW search space with searches carried out in the Burst and Stochastic working groups. Long-lived transients can be considered to be short-lived CW sources. A small joint subgroup with members from both the CW and Burst groups is carrying out work in this area. CW sources with deterministic but unknown phase evolution, such as from a neutron star in a binary system with uncertain parameters, may be detectable via the “radiometer” method in use by the Stochastic group. Tradeoffs among search methods for such sources are being explored in a joint CW/Stochastic mock data challenge focused on the search for Scorpius X-1.

1.4 Searches for Stochastic Backgrounds

The prime objective of the Stochastic Gravitational Wave Background (SGWB) group is to measure the stochastic background. A stochastic gravitational-wave background is formed from the superposition of many events or processes that are too weak and/or too numerous to be resolved individually, and which therefore combine to produce a stochastic background. A stochastic background can arise from cosmological sources such as inflation, cosmic strings, and pre-Big-Bang models. Alternatively, it can arise from astrophysical sources such as compact binary coalescences, supernovae, and neutron stars.

Comprehensive searches have been carried out using data from initial LIGO and Virgo. No signal was detected, but our results constrain the energy density of the stochastic background to be $\Omega_0 < 5.6 \times 10^{-6}$ at 95% confidence. Advanced detectors are expected to have about $10\times$ better strain sensitivity than the initial detectors, as well as to extend the sensitive band from 40 Hz down to 10 Hz. These improvements and wider bandwidth will enable breakthroughs in searches for the stochastic background, with a potential sensitivity of $\Omega_0 < 6 \times 10^{-10}$. The detection of a cosmological background would be a landmark discovery of enormous importance to the larger physics and astronomy community. Simulations studies show the detection of an astrophysical background is not unlikely, and it would be of great interest as a probe of the evolution of the Universe since the beginning of stellar activity.

The SGWB group has built on the cross-correlation infrastructure, originally designed to carry out searches for isotropic stochastic backgrounds, to diversify and to carry out a wide range of interesting analyses. The SGWB directional search provides a method of distinguishing between different stochastic sources using sky maps of gravitational-wave power; the narrowband radiometer has been used to search for gravitational waves from Scorpius X-1, the Galactic Center, and SN 1987A. The radiometer provides an important tool for gravitational-wave astronomy when there is significant uncertainty in the phase evolution of a neutron star signal (as is the case with the low-mass X-ray binary source, Scorpius X-1). The radiometer limits on Scorpius X-1 from initial LIGO remain the most constraining to date over a portion of the observing band, and the SGWB Group continues to develop the search, in collaboration with the Continuous Waves Group.

The SGWB group has developed a cross-correlation pipeline to search for very long-lived gravitational-wave transients lasting hours to weeks. It may be possible for neutron stars to emit transient gravitational waves on these time scales. Moreover, exotic models allow for the possibility of a seemingly persistent signal to start or stop during an observing run, also leading potentially to very long transient signals. An efficient very-long-transient detection algorithm will have other useful applications: it can establish if an apparently persistent source, e.g., observed in a stochastic background search, exhibits variability in time, and it can be used to understand the behavior of detector artifacts on timescales of days to weeks.

The SGWB group is actively involved in detector characterization efforts. Much of this work has overlap with both the Detector Characterization and SGWB groups. For example, the SGWB group uses Detector Characterization measurements of correlated magnetic noise in order to find solutions that minimize contamination in stochastic searches. Correlated noise, e.g., from Schumann resonances, can create coherence in widely separated detectors, which is possible to mistake for a stochastic background signal, thereby introducing a bias. The group is also developing a stochastic data-quality monitor to track search sensitivity in real time and to identify problematic sources of noise.

1. **Highest priority**

   The highest priorities of the SGWB group are the isotropic search, the directional search, and the search for very long transients. The isotropic analysis is the original *raison d'être* for the SGWB working group, and the detection of a stochastic background is the group's most compelling scientific deliverable. The directional search—which employs both a radiometer algorithm and a spherical harmonic decomposition algorithm—generates sky maps (and strain spectra), which can be used to identify cosmological or local anisotropies as well as point sources. This long-established analysis is an important tool for distinguishing between different sources of the stochastic background (e.g., isotropic signals vs. signals clustered in the galactic plane). While the directional search assesses the contribution to stochastic signals from different directions and from different frequency bins, the search for very long transients assesses the contribution from different times. The three searches together provide a complete understanding of the origin of any observed signal. We carry out a number of activities in support of these three searches including mock data challenges, modeling of different sources of the stochastic background, parameter estimation, folding of data into a sidereal day, detector characterization, and an extension to the isotropic search to look for non-standard polarization modes in the stochastic background.

2. **High priority**

   We designate as high-priority a program to follow up on CBC and burst detection candidates with a low-cost cross-correlation search. The search produces spectrograms showing the detection candidate in cross-correlated data and is designed to provide useful diagnostic tools for visualization and characterization of candidate events. For compact binary coalescence signals in particular, it offers a low-cost, independent verification of detections made by matched-filtering pipelines while filling in potential gaps caused by data-processing corner cases. In addition, this method offers a useful visualization of the GW signal in spectrograms.

3. **Additional priority**

   Additional priorities for the SGWB Group includes studies that are at a less mature stage than those listed above, such as measurements of non-Gaussianity of the stochastic background and a search for r-modes from neutron stars.

There is overlap in the SGWB group's search for very long-lived transients with searches being carried out in the Burst and Continuous Wave search groups. Continuous wave sources with deterministic but unknown phase evolution, such as from a neutron star in a binary system with uncertain parameters, may be detectable via the radiometer method in use by the SGWB group, or methods being developed in the continuous wave search group. Trade-offs among search methods for such sources are being explored in a joint Continuous Wave/Stochastic mock data challenge focused on the search for Scorpius X-1.

1.5 Characterization of the Detectors and their Data

Virgo

Noise mitigation, spectral lines identification, glitch reduction and data quality vetoes are the main tasks of the Virgo detector characterization group. Responsibilities include working with the commissioning team to track down any limitation to the detector's sensitivity, working with the calibration team to maintain the calibration and timing accuracy to an acceptable level for GW searches, and providing noise information and vetoes to the data analysis groups and commissioning team. During past science runs and commissioning periods, the Virgo detector characterization team has provided several investigation and monitoring tools, and data quality vetoes which impacted positively both commissioning activity and astrophysical searches.

Search Data Quality: A new Virgo data quality model has been developed and is currently implemented. This model defines workflows and procedures the group will follow to provide data quality products to searches. In particular, emphasis is made to produce and deliver search-specific data quality vetoes. On top of this, a new and ambitious online architecture is being implemented to provide vetoes to online search pipelines. We have developed with LIGO a common data quality segment database, to benefit Burst and CBC groups. It has been moved to production. Additional data quality needs specific to CW and Stochastic search groups include the identification of noise source contributions to spectral lines or non stationary and non linear features. For this, we use automatic spectral lines identification tools already well tested, and a line database.

Early AdvVirgo Characterization: The Virgo detector characterization team will begin noise and glitch studies on each commissioned sub-system as soon as they come online, in close collaboration with sub-system hardware coordinators and commissioners. A system of shifts has been organized. Periodically, a team of two shifters is on watch. They study transient and spectral noise using analysis tools developed by the group.

1. **Highest priority**
   The highest priority of the Virgo Detector Characterization is to find and mitigate the sources of noise and to provide data quality information to the LSC-Virgo search groups in order to reduce the impact of the remaining noises.

2. **High priority**
   Our current high priorities are the development of useful tools for commissioning and an early characterization of each sub-system of Advanced Virgo in order to reduce the need of vetoes in future searches. This will imply a coherent system of monitoring web pages, a spectral line database catalogue, identification of non stationary lines and a software infrastructure to provide useful online data quality information.

3. **Additional priority**
   Additional priorities for Virgo detector characterization are to develop improved methods to uncover the paths and the sources of the noise transients which most impact the searches, and to implement automated noise classification tools.

1.6 Data Calibration

LIGO

Calibration of the LIGO interferometer data is critical to the success of the searches and to the confidence in their results. This is a complex task that involves instrumental hardware measurements, detector modeling, computer programs, and extensive validation and review. Calibration is provided both in the frequency domain, as a frequency-indexed response function to be applied to the Fourier transform of the gravitational wave channel, and in the time-domain, as a derived digital time series representing strain as a function of time. The time domain calibrated data, along with an accompanying error budget, is the main calibration product. Early aLIGO critical calibration activities include:

- measurements of instrument transfer functions and calibration model parameters,
- development and improvement of instrumental measurements,
- estimation and reduction of the errors in the calibration data products,
- deployment and use of the photon calibrator as an independent cross-check of the calibration,
- development and improvement of time-domain data generation techniques, including use of gstlal and the aLIGO front-end system,
- development of pre-processed h(t) products, such as whitened, cleaned, and coherent data streams,
- development of on-line tools to monitor calibrated data quality, and
- a comprehensive review of entire calibration procedure.

The scope of the calibration team includes the timing of LIGO data. Traceable and closely monitored timing performance of the detectors is mission critical for reliable interferometer operation, astrophysical data analysis and discoveries. Critical timing tasks include:

- developing of injection techniques to determine accurate timing through direct test mass excitations,
- expanding the capabilities of data monitoring tools related to timing and phase calibration,
- enhancing the availability of timing diagnostics for various subsystems,
- measuring and documenting the timing performance of critical digital subsystems,
- measuring and documenting the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator, characterization of analog modules, etc.), and
- reviewing the physical/software implementation and documentation of the timing components of critical subsystems.

Virgo

During the Virgo science runs, the calibration measurements have been automated and extended to have some redundant data. It includes measurement of the absolute time of the Virgo data, measurement of the transfer function of the dark fringe photodiode readout electronics, measurement of the mirror and marionette actuation transfer functions and monitoring of the finesse of the arm cavities. The calibration output are then used (i) in the frequency-domain calibration, resulting in the Virgo sensitivity curve, (ii) in the time-domain calibration, resulting in the h(t) strain digital time series and (iii) for the hardware injections. Independent cross-check of the reconstruction has been done systematically during VSR4 using a photon calibrator.

The methods used for Virgo will still apply for AdV after some tuning for the new configuration. Simulations have been carried on for the a priori most challenging measurements, i.e. the measurement of the mirror actuation response. They confirm that the Virgo methods can still be applied, putting some constraints on the minimum force to be applied on the AdV arm mirrors. In parallel a conceptual design of the new photon calibrator to be developed for AdV is being finalized before the setup is built and then installed in 2015. Critical calibration activities are:

- development and improvement of instrumental measurements (in particular with the digital demodulation electronics of the photodiode readout),
- prototyping and installation of a photon calibrator,
- development of online tools to monitor the Virgo timing permanently,
- upgrade the h(t) reconstruction method after the study of the impact of some parameters that were neglected during the Virgo era.

1.7 Hardware Injections

Hardware injections are simulated gravitational wave signals added to LIGO and Virgo strain data by physically actuating on the test masses. They provide an end-to-end validation of our ability to detect gravitational waves: from the detector, through data analysis pipelines, to the interpretation of results. The hardware injection group is tasked with the development, testing, and maintenance of hardware injection infrastructure. This includes on-site software to carry out the injections at specified times. We also work with the search groups to maintain the software that generates gravitational waveforms suitable for injection.

Each data analysis group works with the hardware injection team, in different ways: Burst and CBC groups provide transient waveforms and determine suitable injection rates, the CW group selects the parameters for neutron star signals, which persist throughout the science run, and the SGWB group typically carries out one or two ≈ 10 min injections during each science run. The search groups analyze hardware injections during science and engineering runs to identify and solve problems as they come up, and the results of these studies are reported back to the hardware injection team so that adjustments can be made.

While most injections are known to the LSC, there are also blind injections, for a blind test of the analysis. Although blind injections are performed by a separate team, the hardware injection group is in charge of maintaining the blind injection infrastructure, nearly identical to the regular injection one, and provides training.

2 Search Plans for the Advanced Detector Era

This section collects the plans formulated by the astrophysics search groups for Advanced LIGO and Advanced Virgo. These plans connect the science case, analysis methods, criteria for result validation and publication, software development plans and resource needs for each search.

Contents

3.1 All-Sky Short-Duration Burst Search

3.1.1 Abstract

Supernovae, long gamma-ray bursts, and soft-gamma repeaters are among the variety of violent astrophysical phenomena that are expected to produce bursts of gravitational waves. These waves are ripples in spacetime, and carry the signature of the complicated physical processes that produced them. The detection of gravitational-wave bursts could therefore provide a direct probe of these relativistic systems, help us unravel the mechanism behind supernova explosions and reveal the inner workings of gamma-ray bursts.

The LSC and Virgo plan to undertake several Observing Runs in coming years and we present plans for searching for gravitational-wave bursts using data acquired by the advanced LIGO - advanced Virgo network during these Observing Runs (c.2015+). This search targets bursts of any form, without restriction to a particular signal type. We aim to cast the widest possible net for gravitational-wave transients and infer the properties of any observed signal which may originate from phenomena that are, at present, entirely unknown to science.

Parameter estimation algorithms will be used to estimate the morphology of possible gravitational-wave burst signals and their parameters. A key part of this search is the estimation of the position on the sky of possible gravitational-wave signals within minutes of data taking. This will allow partner telescopes to search for an electromagnetic counterpart, which could be vital to interpreting the nature of the source. The joint identification of a system in light and gravitational waves would be a watershed moment in multi-messenger astronomy.

3.1.2 Scientific Justification

Relativistic astrophysical systems such as supernovae, long gamma-ray bursts, and soft-gamma repeater giant flares are expected to produce bursts of gravitational waves. These systems are difficult to model, due in part to their complicated dynamics, and to the fact that the equation of state of matter at neutron-star densities is not known. Therefore, searches for gravitational-wave bursts from these systems cannot rely on an accurate signal model. Instead, methods must be developed that are capable of distinguishing weak gravitational wave (GW) signals from the cacophony of background noise fluctuations in our data *without* prior knowledge of the signal shape. This is in contrast to searches for GWs from, e.g., neutron star coalescences or isolated pulsars, which rely on precise models to separate signals from noise.

The all-sky short-duration burst search targets GW bursts of generic waveform in the advanced detector bandwidth, without assumptions on the signal's origin, direction, or time of arrival. While past all-sky burst searches have focused on signal durations of up to approximately one second, such searches in the advance detector era will extend to durations of 10 seconds which covers the parameter space of some gravitational wave emission models for supernovae and neutron star oscillations. Searches for unmodelled signals of durations between 10 seconds and 10 hours are covered in the long-duration Burst search plan. Potential sources of GW bursts include gamma-ray bursts, choked or low-luminosity gamma-ray bursts, core-collapse supernovae, soft-gamma repeater bursts, ringdowns of perturbed neutron stars or black holes, the coalescences of intermediate-mass black holes or highly eccentric binaries, and potentially as-yet-unknown systems. Burst search techniques are particularly useful for signals with compact time-frequency volume, where they are known to achieve sensitivities comparable to matched filters. They are also valuable for systems such as highly eccentric binaries where the number of templates required for a matched filter search is prohibitive. For the compact binary coalescence sources where burst methods are an efficient option, the work will proceed jointly with the CBC group; see §3.6.

In the absence of robust signal models, predictions of detection rates cannot be made with any certainty[1].

---

[1] We note that the rate of binary neutron star detections, which is the best-modelled GW source for the advanced detectors, is

The astrophysical impact of a positive result would greatly benefit from the reconstruction of the signal's characteristics (both waveform and source location), and the identification of counterparts in the electromagnetic or neutrino sector. The interpretation of null results will need to rely on some parametrisation of the burst signal, particularly the frequency range and total energy content. For example, burst searches of first-generation LIGO and Virgo data showed a typical sensitive range of $7\,\mathrm{Mpc}\,(E_{\mathrm{GW}}/10^{-2}\mathrm{M}_\odot)^{1/2}$ for standard sirens that emit energy E_{GW} isotropically in the 100-200 Hz band. This corresponds to a rate density limit of approximately $4 \times 10^{-4}\,\mathrm{Mpc}^{-3}\mathrm{y}^{-1}\,(E_{\mathrm{GW}}/10^{-2}\mathrm{M}_\odot)^{-3/2}$. This performance was shown to be robust across a variety of signal morphologies. [1]

In the advanced detector era, the chances of detection for the strongest GW sources ($E_{\mathrm{GW}} \geq 10^{-2}\mathrm{M}_\odot$) will scale as their detectable volume, since the average source density approaches homogeneity at the sensitive distance scales. For weaker sources, the chances of detection depend on the distribution of mass in the local universe. In particular, for sources whose initial LIGO and Virgo range was approximately 0.01 Mpc to 0.1 Mpc, the detection chances will increase less than the detectable volume due to the absence of significant additional mass on the 0.1 Mpc to 1 Mpc scale. The detection of GWs or their associated electromagnetic counterparts in the nearby universe may be aided by the use of galaxy catalogs as source priors. In addition, we note that the extension of the sensitive frequency band of the detectors to lower frequencies will also make new sources available, in particular higher-mass systems.

Historically, the burst group has performed many specialized searches targeting particular astrophysical transients or specific signal classes. In the advanced detector era, the all-sky short-duration Burst search will be the reference used to benchmark other burst searches for signals of similar duration. An assessment of the improvement offered by new approaches will be essential in deciding how or if such new searches will lead to LVC publications. The all-sky search results from low-latency analysis will also be used to trigger follow-up studies in electromagnetic channels, and as a first rapid check for GW counterparts to external triggers such as gamma-ray bursts.

3.1.3 Search Description

Searches for generic GW bursts require techniques that can distinguish signals from background noise fluctuations without prior knowledge of the signal waveform. The "gold standard" in this field is *coherent analysis*, in which the data from all the detectors in the network is combined using both amplitude and phase information before being searched. A GW signal in the data stream of each detector will build up coherently when those streams are combined, while coincident noise fluctuations can be eliminated based on their (non-)correlation between detectors. The analysis is typically done by performing a time-frequency decomposition of the data using wavelets or short Fourier transforms, and identifying candidate signals as clusters of "hot" pixels in the time-frequency maps.

Several independent pipelines based on these ideas have been used in the analysis of initial LIGO - Virgo data, including coherent waveburst (CWB) [2, 3, 4], X-PIPELINE [5], and STAMP [6]. Of these, only CWB has been used to date for all-sky short-duration Burst searches due to the computational cost of coherent analysis and the superior speed of CWB. However, development work continues on all three pipelines, and additional pipelines are under development. These include: the GSTLAL Excess Power (EXCESSPOWER) [7], which processes detector data streams separately ("incoherently") and looks for coincident transients; omicron+LAL Inference Burst (LIB), which also searches incoherently [8] to trigger a coherent MCMC followup; and BAYESWAVE [9], which performs Bayesian MCMC followups.

Given the unmodelled nature of GW bursts and the potentially large discovery space in the event of a detection, the all-sky search strategy aims to support multiple analyses that will offer both complementarity and redundancy. Multiple all-sky pipelines will allow for cross validation of search outputs and help the

---

uncertain to three orders of magnitude; see §3.2.

review our science results. In particular, by being subject to different systematics, additional all-sky analyses are expected to increase our confidence in the coverage for unexpected signal morphologies. We note further that, historically, the direct comparison of two or more pipelines has spurred improvements in each. Therefore, we will adopt at least two pipelines to perform all-sky searches. The baseline all-sky search pipeline will be coherent WaveBurst CWB2G (see §3.1.6).Additional pipelines will need to be reviewed by the start of each Observing Run.

In addition to these search pipelines, there are several pipelines that focus on the analysis of data from auxiliary and environmental channels to generate required veto and data quality information. Two such pipelines are KLEINE-WELLE and OMICRON. More details on pipeline development plans are given in §3.1.6.

For the coming Observing Runs, CWB2G, omicron+LIB and CWB2G+BAYESWAVE will be deployed to search all science data with two or more detectors operating at reasonable sensitivity. A minimum of two detectors is required to reject background noise fluctuations. Each search pipeline will actively work to improve their performance in preparation for improved detector sensitive in subsequent Observing Runs and tune their search parameters accordingly. The search will identify and rank candidate events, using GW correlation tests and data quality checks to reduce the noise background. The same time-lag method used in previous burst searches will be used to assess the background and thus the significance of candidate events. In this approach, search thresholds will initially be tuned using a small set of time lags before a larger, different set of time lags are used to estimate the background and, thus, the significance of zero-lag events. The significance estimated by each all-sky search pipeline will be combined to give a single, joint estimate, particularly for loud triggers which may be identified as candidate GW events. Should no GW events be identified, the combined significance estimate will be used to construct the rate upper limit.

Parameter estimation follow-ups of the loudest triggers will also be performed. BAYESWAVE will choose a CWB2G significance threshold so that it follows up background and zero-lag triggers at an average rate of once per week. LIB will choose a significance threshold so that it follow-up zero-lag triggers at an average rate of once per day. These analyses aim to estimate the signal properties for any detected gravitational-wave burst. Additionally, BAYESWAVE aims to characterise background trigger properties. Since parameter estimation work is a relatively new aspect of the all-sky search, outputs from all parameter estimation analyses, including CWB2G signal parameter estimates, will be compared to determine the operational characteristics of each algorithm.

As in S6/VSR2,3, an online minute-latency analysis will be used to provide alerts of significant events to electromagnetic observatories. The offline analyses will initially run on a small subset of all available livetime to identify any issues that many cause errors or delays in the computing time. This smaller run will also feed into data quality studies and data monitoring tasks. The definitive analysis will be performed offline using the best calibration, data quality, and other relevant information available.

3.1.4 Results validation plan

The coherent Waveburst search pipeline has had past experience analysing data from LIGO and Virgo. Therefore, its upgraded successor, coherent Waveburst 2G (CWB2G), will be considered the baseline analysis for the all-sky search. In addition to CWB2G, there will be two other measures of significance for each event: the results of the omicron+LIB pipeline and the BAYESWAVE follow-ups to CWB2G triggers. In the case that a pipeline sees a short duration transient, signficant events from one pipeline will be cross validated with events observed in the other pipelines. The additional validation pipelines will increase the robustness of the search by providing an independent analysis of the data. In particular, by being subject to different systematics, validation searches are expected to increase our confidence in the coverage for unexpected signal morphologies. It will also provide a significant opportunity for cross-checking and reviewing of results. Additionally, we will ensure the redundancy and cross-checks contribute to validation of any scientific results reported in an all-sky short-duration Burst search publication.

3.1.5 Publication plan

If the all-sky short-duration Burst search identify gravitational wave signals, the Burst group aim to have an article ready for publication within 3 months of the identification of any candidates. If no signals are found, the Burst group expect to publish an observational results paper for short-duration bursts following each advanced detector Observing Run. Based on the current run schedule [10], the observable four-volume VT (detection volume times observation time) for each successive run increases over that of the previous run by at least a factor of 2-3. [2] The rate limit set on a uniformly distributed population of sources would improve by this same factor, which we consider to be sufficiently interesting to merit publication even in the non-detection case.

In case of GW detection candidates, we consider $3\,\sigma$ to be a *minimum* significance threshold for an "evidence for" statement. Based on the initial LIGO and Virgo experience, we expect that confident detections will require improvement in our background rejection techniques. Currently these rely on single-detector characterization and data quality, and the tuning of coherent tests in network data analysis. Advanced methods for data cleaning and signal/background discrimination are under development, as described in §3.1.6.

Null results will be interpreted as typical sensitive distance and rate density limits for standard-candle sources. We will also compute limits based on standard-candle populations folded with galaxy catalogs when applicable, i.e., when the sensitive range is of order a few megaparsecs to a few tens of megaparsecs.

For each GW detection candidate, all burst pipelines report bulk information such as duration, bandwidth, and amplitude. In the absence of an *a priori* signal model, more refined statements about the properties of a candidate, such as reconstructed $h_{+,\times}(t)$, will rely on robust parameter estimation tools for generic bursts. Such tools are under active development, as discussed in §3.1.6. We will also explore schemes to classify GW candidates into pre-determined signal classes. This will allow us to set a statistical confidence conditional to a signal class of astrophysical interest. Finally, the inclusion in a publication of any multi-messenger information on candidates will depend on the nature of that information, and following the procedures set out in the relevant MoUs.

Our minimum goal is to complete the analysis of each run and the corresponding publication before the start of subsequent run. We will review in advance the analysis methods, as well as a few standard parts of the publication, to speed up the final internal review process. In the case of long observations, we aim to complete the offline analysis for a given data set within one month of fixing the calibration and data quality information. Our ultimate goal is to circulate a paper within the LSC and Virgo collaborations not more than 3 months after the calibration and data quality information is finalised.

---

[2] This also holds for the first advanced science run when compared to the total data collected by initial LIGO and Virgo.

3.1.6 Technical requirements and development plan

coherent WaveBurst:
CWB2G incorporates improvements upon the version run in the initial detector era. Development is now in the late stages. New features of particular interest include: a modular structure, allowing the plug-in of user-definable modules; an input data cleaning stage by regression of auxiliary channels; multi-resolution reconstruction of signal candidates using time-frequency pixels from different wavelet resolutions; and the ability to perform different post-processing using the same input trigger set (e.g. for specialised analyses targeting different polarization states).

An intermediate online version of CWB2G was tested in Engineering Run 6. The final version will be tested in ER7. The pipeline review is underway and should be completed by the start of the first Observing Run. Ongoing tests include the re-analysis of all initial LIGO-Virgo data. Other tests focus on questions including: the ability of regression to reject glitches which survive the standard data quality checks; how much signal reconstruction has improved; how much the confidence of candidates improves by classification into pre-determined signal and glitch classes; and optimal tuning using the false alarm rate density statistic [3].

Additional pipelines:
A feasibility study of all-sky analysis of initial LIGO data is currently underway, and advanced time-frequency clustering techniques are being explored [6]. EXCESSPOWER has been developed primarily for use in detector characterisation studies, but can be used for GW searches by applying coincidence criteria to triggers from different detectors using the infrastructure for the cosmic string analysis [11]. It has already been integrated into the existing online data transfer and analysis infrastructure and is partially reviewed as a consequence of the gstlal_inspiral review. The omicron+LIB pipeline has been tested by offline running over ER4,5 and S6/VSR2,3 data. BAYESWAVE has been applied to follow up background events identified by CWB2G as well as simulated signals, both in ER5 and S6/VSR2,3 data.

Mock Data Challenge (Simulation) infrastructure:
Burst simulations will be created using the lalsimulations infrastructure. The new code developed for generating burst injections has already been tested against the traditional burst simulation infrastructure, GRAVEN/BURSTMDC. GRAVEN/BURSTMDC will be be maintained in the near term as a reference with which to validate simulation engines until the lalsimulation infrastructure is fully deployed.

Glitch rejection and signal parameter estimation:
Improved glitch rejection is crucial for the all-sky search. In addition to the data cleaning and multi-resolution reconstruction efforts in CWB2G, the LIB and BAYESWAVE parameter estimation pipelines are being developed based on MCMC techniques. To quantify their performance, a parameter estimation challenge has been proposed[4]. Focusing on the all-sky background of the S6/VSR2,3 CWB search, this challenge will assess the glitch rejection and waveform reconstruction capabilities of the various pipelines.

Future searches:
The all-sky searches can also be adapted to perform future searches for signatures of alternative theories of gravity or bursts with memory. Alternative theories of gravity can be tested by modifying the all-sky burst search code to allow for polarization modes other than the plus and cross modes present in GR. Preliminary studies have focused on searching for scalar-mode GW bursts. Such research will be pursued more in the future, especially when a network with more GW detectors allows different polarization modes to be better separated.

---

[3] See https://dcc.ligo.org/LIGO-T1300869.

[4] See https://wiki.ligo.org/Bursts/AllSkyPE.

3.1.7 Resources

Detector Characterization:
All-sky searches have relied on data quality information to help reduce the noise background. While improved glitch-rejection capabilities are being developed, we expect that data quality will remain a key component of the all-sky search.

Calibration:
To meet our publication schedule, the offline search will require "final" calibrated $h(t)$ from three months after the beginning of each run (within two months for the 2015 run). All tests to date have shown that the calibration uncertainties will not have a significant impact on the analysis provided they are of the level already achieved in the S5-6 and VSR1-3 runs[5].

Review:
The CWB2G review has started, and is expected to complete before the start of O1. Given the relevance of the pipeline the review team is made up of 5 reviewers. The review of BAYESWAVE is underway, while the LIB review has been completed. The omicron review will begin soon though it is not expected to an onerous review since omicron shares many components with the Omega pipeline which has previously been reviewed.

---

[5] for detailed requirements, see `https://dcc.ligo.org/LIGO-T1300950`

3.2 Search For Binary Neutron Star Coalescences

3.2.1 Abstract

The coalescence of a binary neutron star (BNS) system is the most promising source of gravitational waves (GWs) for Advanced LIGO (aLIGO) and Advanced Virgo (AdV). Radio observations of double neutron star (NS) systems containing pulsars suggest a significant coalescence rate in the volume of the Universe accessible to the advanced detector network. The BNS gravitational waveform is well modeled in the most sensitive frequency band of the advanced detectors, allowing the use of matched filtering for detection. The detection of even a few BNS signals would greatly constrain the very broad uncertainties in the rate of BNS coalescence, thereby constraining the possible formation channels. Measuring the mass and spin distributions of BNS will inform stellar evolution, nuclear physics, and supernova physics. The presence of matter in the merger may give rise to a detectable electromagnetic (EM) counterpart, including a gamma-ray burst (GRB), an orphan afterglow or kilonova. BNS will provide a laboratory for measurement of the equation of state of neutron stars and for testing general relativity in the post-Newtonian approximation. BNS searches are crucial to the science potential of the LIGO-Virgo detector network.

3.2.2 Scientific Justification

Population synthesis models constrained by radio observations of double NS systems in the Milky Way provide an indirect estimate of the GW-driven BNS merger rate of $0.01 - 10\,\mathrm{Mpc}^{-3}\,\mathrm{Myr}^{-1}$, for an expected BNS detection rate of $0.4 - 1000$ per year in Advanced LIGO at design sensitivity [12]. In the most optimistic (but plausible) hypothesis, we could expect up to three detectable BNS mergers during a three-month run with early detector ranges of $60 \pm 20\,\mathrm{Mpc}$ in Advanced LIGO and $20\,\mathrm{Mpc}$ in Advanced Virgo; if no detections are made, we will at least be able to constrain the models which predict such rates.

The masses of known NSs are reported in a range from $1.0\,\mathrm{M}_{\odot}$ to $2.74\,\mathrm{M}_{\odot}$ [13], with mass greater than $0.9\,\mathrm{M}_{\odot}$ assumed. NSs in BNS systems have a narrow observed mass distribution, with masses reported from $1.0\,\mathrm{M}_{\odot}$ to $1.49\,\mathrm{M}_{\odot}$ [13], and are consistent with an underlying mass distribution of $(1.35 \pm 0.13)\,\mathrm{M}_{\odot}$ [14]. Theoretical models support a population of NSs formed in binaries through electron-capture collapse of O-Ne-Mg cores, and predict masses which are consistent with these observations [15, 16]. Lower mass Fe cores are predicted to lead to NSs with masses almost as low as $1\mathrm{M}_{\odot}$ [17].

Current astrophysical understanding indicates that the older NS in a binary system can be spun up through mass-transfer from its companion, which can increase the spindown timescale, although this process is not completely understood and it is not clear how efficient the spin-up process is. The observed dimensionless spins (J/m^2) for NSs in BNS systems (e.g., J0737-3039) are ≤ 0.04 [18], however the fastest known NS spin is 0.4 [19]. A search that neglects spin effects in the template waveforms will not incur an appreciable loss of efficiency if the NS spin is less than ~ 0.05 [20], whereas a search for more rapidly spinning systems would require the use of waveforms that capture the effect of spin [21, 22, 20]. BNS waveforms for systems with small spin are well modelled by post-Newtonian theory [23] at seventh order beyond the leading-order orbital phase [24].

Although it is unlikely that the first detected BNS GWs will be accompanied by an EM detection, there are several plausible EM counterparts to BNS mergers [25]. The detection and confident association of such a counterpart would enable a deeper study of compact object astrophysics than gravitational waves alone. The anticipated sky localization in the early advanced-detector network will constrain the source location to hundreds or thousands of square degrees, improving to ~ 10 square degrees as the network reaches full sensitivity [26, 27]. Low-latency pipelines [28, 29, 30, 31, 32, 33] will provide rapid-response triggers for joint EM searches.

With only a few BNS detections, we can constrain the merger event rate, leading to constraints on pulsar population and stellar evolution models. Source parameter estimation will enable us to obtain a census of masses and spins of NSs in binaries [34], which in turn will yield constraints on theoretical models that describe how these systems are formed.

NSs contain the highest densities of matter in the observable universe. The internal structure of NSs is constrained by nuclear experiments and astrophysical mass-radius measurements, which help to constrain the possible equation of state (EOS) of nuclear matter [35]. As binary NSs coalesce, the EOS will determine both tidal interactions during late inspiral and matter effects during merger. These effects are encoded in the gravitational waveform [36]. With a few tens of detections, it may be possible to significantly constrain the EOS by combining information from many coalescence events [37], although this requires the incorporation of waveform developments to better model effects beyond the level of current post-Newtonian templates to avoid systematic errors [38, 39]. In cases where an EM counterpart can be identified, further information can be used to understand the physics of the merger [40, 41].

Strongly relativistic bound systems have high orbital compactness $GM/(c^2R)$, with M the total mass and R the separation, and high orbital velocity v/c. From to-date radio observations of pulsars in binaries it can be seen that energy and angular momentum are lost through gravitational wave emission, but only at the level of the quadrupole formula [42]. Even the binary pulsar J0737-3039 [43] only has $GM/(c^2R) \sim 4.4 \times 10^{-6}$ and $v/c \sim 2 \times 10^{-3}$. By contrast, BNSs on the verge of merger will reach $GM/(c^2R) > 0.2$ and $v/c > 0.4$, with strong GW emission. Observing such events will give us access to the genuinely strong-field, relativistic regime of gravity, going well beyond leading-order effects. We will be able to probe the dynamical self-interaction of spacetime itself, for example the scattering of gravitational waves off the Schwarzschild curvature generated by the binary as a whole, an effect that enters the phase at 1.5PN [44, 45]. Such studies will provide stronger tests of general relativity than any that have been performed to date [46].

3.2.3 Search Description

During the initial-detector era we developed algorithms to search for BNS signals coincident in a detector network [47, 48, 49, 50, 51, 52]. The BNS search is conducted with matched filtering [48, 52, 30], where banks of template waveforms are constructed, and the search is parallelized across computers by splitting the template bank and/or by splitting the data to be filtered in time. Parameter estimation algorithms were developed and used to follow up candidate events [34]. Search results were published in a series of papers for the six LIGO and three Virgo "science runs": S1 [53], S2 [54], S3 and S4 [55], S5/VSR1 [56, 57, 58] and S6/VSR2,3 [59]; no detections were made, and an upper limit of $R_{90\%} \leq 130 \mathrm{Mpc}^{-3}\mathrm{Myr}^{-1}$ was placed on the rate of binary neutron star coalescence. This is about an order of magnitude above the most optimistic rate estimates. During the initial detector searches, the Compact Binary Coalescence (CBC) group developed a detection pipeline that addresses the challenges of analyzing data from real-world GW antennas, which often produce glitchy data with non-stationary and environment-dependent noise floors, and that experience unplanned data outages. Many of those techniques will go forward into the advanced-detector era, though some new techniques will be required. With improved low-frequency sensitivity, advanced-era detection pipelines will require longer waveform templates, and longer templates bring tighter orbital frequency discrimination leading to the need for a greater number of templates. The large number of long template waveforms poses new technical challenges almost all of which have been addressed during the time between the initial and advanced era science runs, however some tasks remain.

Template Mass. Historically, the BNS search has been performed in conjunction with searches for neutron star black hole binary (NSBH) and stellar-mass binary black hole (BBH) mergers [54, 55, 56, 57, 58, 59]. Although we will continue to pursue common search techniques for compact binaries of all masses, the physics of NSBH and BBH searches, particularly spin and merger effects, complicate the searches for these systems compared to a BNS search. We will search all valid observational strain data with two or more detectors operating in coincidence. The search will target BNS inspirals and mergers comprised of component masses spanning the plausible NS mass range of $1.0\,\mathrm{M}_\odot$ to $3\,\mathrm{M}_\odot$.

Template Spin. The spins of the NSs in all known BNS systems are quite low, below a dimensionless spin of 0.04 and much lower than the maximum spins observed in isolated NSs [19, 18]. Although such low spins are insignificant from a search point of view [20], it is prudent to allow for higher spins when choosing the parameter space to cover, and we will search for spins up to 0.4. In any case, all candidate events will be followed up with parameter estimation using waveforms with spin, which will allow us to estimate the spins of any candidates that are identified.

The optimal layout of templates in parameter space depends on the noise spectrum of the instrument. In Initial LIGO/Virgo searches, template banks were constructed every $\sim 40\,\mathrm{min}$ to track changing instrument noise. Fully exploiting the low frequency sensitivity of advanced detectors having a $10\,\mathrm{Hz}$ low frequency cutoff, covering the component mass and spin range described above, will require around 10^6 templates, the longest of which is about $35\,\mathrm{min}$ long. Such large banks will be regenerated less frequently with care to ensure a representative average of noise statistics does not lead to loss in sensitivity.

The science goals of the BNS search impose requirements on the way we will structure the search. The main components of the data analysis pipelines, illustrated in figure 1, are:

1. A low-latency BNS search, to issue alerts from the LSC-Virgo alert network for electromagnetic followup; two low-latency pipelines, GSTLAL-CBC [30, 31, 32, 33] and MULTI-BAND TEMPLATE ANALYSIS (MBTA) will run on distinct computing facilities to provide redundancy.
2. A traditional offline search pipeline, PYCBC, which takes advantage of final data quality cuts and calibration to provide the deepest search with full data quality information.
3. Parameter estimation is performed on candidates identified by the searches. Low-latency methods (e.g. BAYESTAR) will rapidly localize the sky position of the candidate, while higher latency methods (e.g. LALINFERENCE) will estimate all parameters including the masses and spins of the binary components using detailed waveform models.

The CBC group supports running three search pipelines to ensure the BNS science goals are met for the early aLIGO/AdV era. MBTA and GSTLAL-CBC will generate CBC triggers in low-latency for distribution to electromagnetic follow-up partners. The optimized, reproducible offline pipeline will run at longer latency (e.g. week latency) to incorporate additional data quality information not available in low-latency, to use final calibrated data, to span a larger region of parameter space, and to ensure that the search is robust to possible data dropouts or data quality issues that are possible in the low-latency environment. Additionally, the offline pipeline can generate high-statistics simulations to evaluate search sensitivity. In general, multiple pipelines help to increase confidence in results independent of code issues. The plans for validation of these pipelines are outlined in Section 3.4.7.

The detection confidence will be determined by measuring the false-alarm probability (FAP) of an event. In the past, false-alarm probabilities were measured using "time slides" to estimate the background. Event lists from separate instruments were shifted in time with respect to each other by amounts that do not correspond to physical propagation delays. Previous implementations proved computationally challenging for measuring the FAP of very rare events [60], therefore new techniques will be used to better measure the tails of the background distributions in all searches, and to provide rapid measurements of significance [61] in the online searches.

Environmental and internal disturbances of the interferometers can feed through to the $h(t)$ gravitational wave strain channel, producing intermittent glitches and reduced performance of the detectors. Online data quality channels will include information about the state of the detectors (e.g. science mode, calibration, low-level vetoes) and the event-by-event vetoes identifying glitches in the data. The low-latency data quality

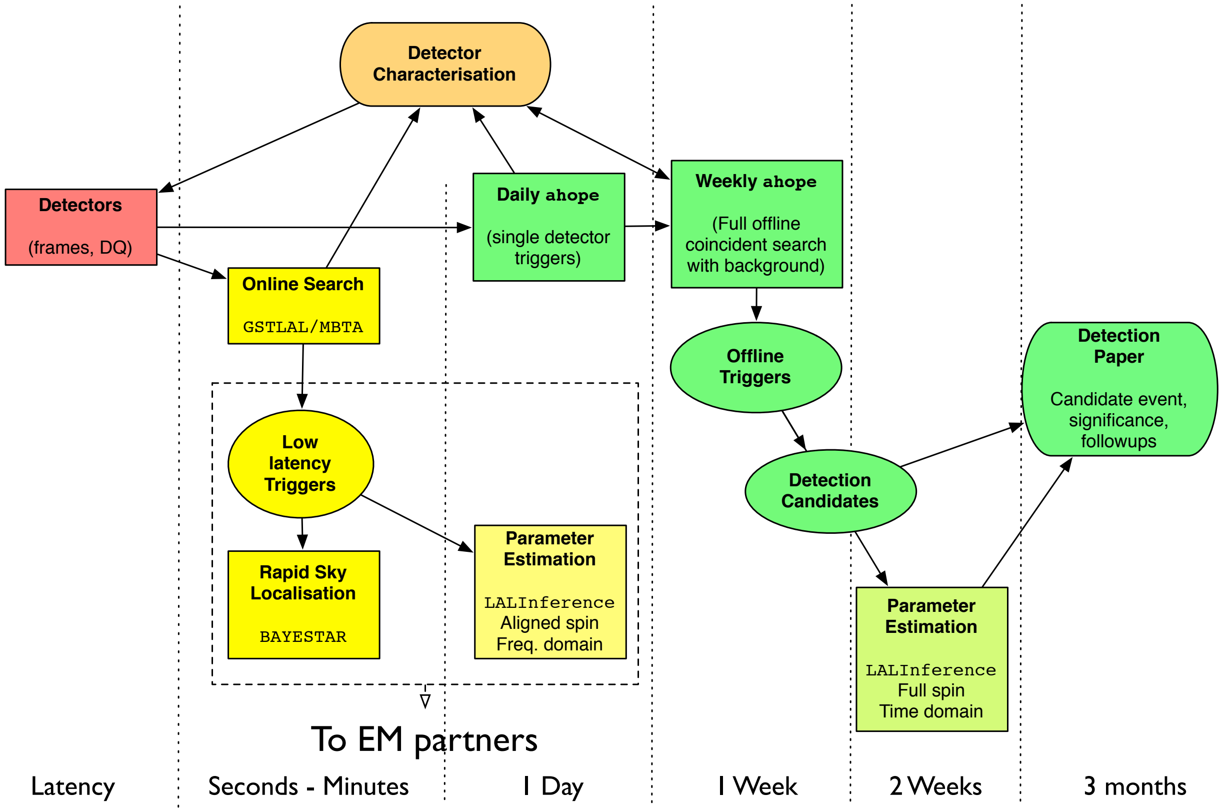

Figure 1: Overview of BNS search, identifying the core components for online and offline analyses, their latencies, and the interdependency between them. The low-latency components (in yellow) will provide rapid triggers for EM followup, whereas the higher latency part (in green) will provide final results for publication. The analysis pipelines are in typewriter font and are defined as follows: `GSTLALCBC`: gstreamer + LIGO Scientific Collaboration Analysis Library (LAL) - A software library for near real-time gravitational-wave analysis. `MBTA`: Multi-bank template analysis, a low latency compact binary search pipeline written and maintained by Virgo. `BAYESTAR` A rapid coherent sky localization pipeline that runs on the output of low-latency detection analyses. PYCBC-`ahope`, an advanced LIGO offline compact binary analysis pipeline. `LALInference`, a parameter estimation pipeline written with in the LAL library.

vector included with the online $h(t)$ data is required to provide data quality with seconds latency. Offline searches will use final revised data quality information available at higher latencies.

3.2.4 Validation plan

Any astrophysically significant candidate event will be validated using the multiple pipelines described above. A BNS pipeline comparison study has been performed to inform us of the variation between the output of pipelines when analysing the same data. In the event of a significant candidate being identified, we plan to follow these guidelines to validate it:

- Does the candidate have low false alarm rate in *at least one* pipeline?
- If a significant trigger is reported in more than one pipeline the responses of the pipelines should be consistent within expected differences. Differences in clustering may result in a different template being recorded as the trigger.
- If a significant trigger is reported in only one pipeline, we will act on a case-by-case basis, looking at the progress of a trigger through the pipelines with the aim of determining why differences are observed. If one pipeline is found to be able to rule out a candidate for reasons that were not taken into account by the triggering pipeline (e.g. updated data quality vetoes) then it should not be put forward as a detection candidate. If no reasons are found to rule out the candidate then it may progress through the detection process.
- Is a candidate found with consistent SNR when online pipelines are re-run offline?
- Is the candidate robust to slight changes in the analysis (PSD variation, choice of segment boundaries etc?)
- Does the candidate pass all checks in the CBC detection checklist?

In the case where no detection candidate is found, we will verify that multiple pipelines agree on estimates of sensitive volume within known differences in live time analysed and pipeline sensitivity.

3.2.5 Publication plan

In the past, we have published the results of BNS searches together with the results from NSBH and stellar-mass BBH [54, 55, 56, 57, 58, 59]. Going forward we will consider separate publications depending on whether we need to respond rapidly to a detection candidate, and the time-scale required to perform searches for systems containing black holes. NSBH and BBH searches require more computationally demanding search methods. Here we present some potential publication scenarios.

Publication time-lines will only be achievable if detector characterization efforts are well coupled to both the commissioning and analysis teams *before and during the run*, the effects of calibration uncertainties are understood, a sufficiently accurate and correct calibration is available, and the analysis pipelines (detection, rate measurement, and parameter estimation) are tested, compared, and reviewed prior to the start of the run.

Publication scenarios up to and including the first detection

Confident BNS detection In the case of confident detection, for example the S6 blind injection with false alarm probability $\sim 7 \times 10^{-5}$, the CBC group will aim to submit a paper for publication within 3 months of identifying the candidate, while allowing time for a full LSC/Virgo Collaboration (LVC) meeting to take place. LIGO-P1000146-v16 [60] provides a good model for what we expect from a first detection paper, including calculation of the event significance and estimation of the event parameters. Subsequent LSC-Virgo papers will provide more details about the event. An EM counterpart is not required to make a confident BNS detection. If a counterpart were to be used to elevate a

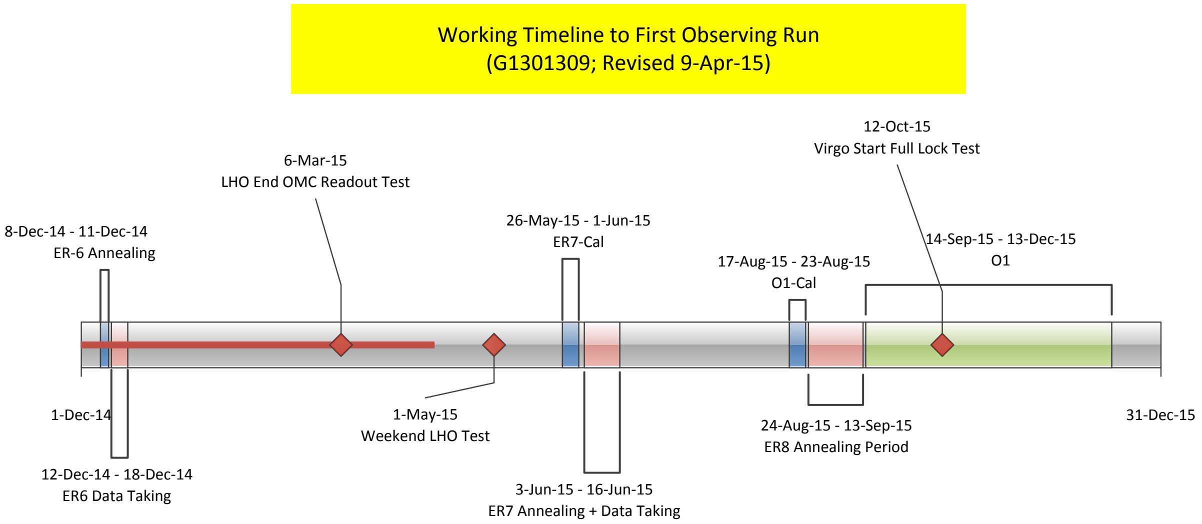

Figure 2: Proposed time line to first observing run (available as LIGO-G1301309 [63]).

marginal to confident detection, the procedure for doing so should be predetermined and vetted before observations begin, as in the GRB triggered search.

Marginal or no BNS detections If marginal detection candidates (for example with false alarm probability $> 1 \times 10^{-4}$) or no candidates are found, then the CBC group will use a ten-fold improvement in VT (the volume of space multiplied by the duration of its observation) to establish the publication cadence. Thus, the first publication in this scenario will require a factor of ten improvement over the BNS result published for LIGO's sixth science run and Virgo's science runs 2 and 3 [59]. If NSBH and/or BBH searches are completed on a similar time scale, the CBC group will submit a paper covering the updated merger rate estimates for the BNS, NSBH and stellar mass BBH searches within three months of reaching the ten-fold VT improvement.

Post-first-detection publication scenarios

As the sensitivity of the detectors improves, we expect to make multiple detections. After the first 4 detections [62], we do not foresee publishing a separate paper with each new signal unless there is a compelling scientific reason to do so. Instead we will publish results at the end of short observing runs (< 6 months), or at predetermined intervals throughout the run. Once astronomical alerts are made public, the initial publication of events may be via Astronomical Telegrams (ATELs) when appropriate.

3.2.6 Technical requirements and development plan

The LSC and Virgo proposal for early observing run scenarios is described in Ref. [64]. The CBC group will work with the commissioning team to update and refine this proposal to ensure the fastest possible route to the detection of a BNS with aLIGO/AdV. As aLIGO/AdV sub-systems are commissioned, the CBC group will collaborate with the Detector Characterization and Commissioning groups to ensure the best data quality possible for the BNS search. In preparation for aLIGO/AdV, the CBC group has adopted two tools: Mock Data Challenges (MDCs) and Engineering Runs (ERs).

- **Engineering runs (ERs)** will engage the CBC group with detector characterization, commissioning, and data analysis infrastructure development. Previous searches have demonstrated that early and close coupling between the search group and the detector group is critical for a timely completion of the analysis [59]. ERs have begun to use real data from detector sub-systems, and we expect to get closer to using $h(t)$ data as Advanced LIGO and Advanced Virgo are commissioned. In addition to detector characterization, engineering runs provide a platform for end-to-end testing of the online and offline analysis infrastructure.
- **Mock Data Challenges (MDCs)** provide a platform for large-scale testing and validation of pipelines using offline or simulated data. Performance in the MDCs is a required milestone to demonstrate that search and parameter estimation pipelines have reached the appropriate sensitivity and accuracy for BNS science, and for comparisons between pipelines with overlapping science targets prior to the first science runs.

We have divided technical activities into a number of sub-groups, which we describe below.

Waveforms. Search and parameter estimation pipelines require an accurate waveform. As the BNS signals are dominated by the long inspiral phase, a BNS detection search can use the TaylorF2 3.5 post-Newtonian stationary phase approximation waveform families which describes the inspiral part of the coalescence. These approximants are already used in non-spinning template banks, which fully capture the physics of BNS systems with spins less than $\chi \leq 0.05$ [20]. TaylorF2 models with leading order spin corrections can be used in template banks with component spins aligned with the orbital angular momentum, which fully capture the physics of BNS systems having spins less than $\chi \leq 0.4$ [20]. Full parameter estimation requires the inclusion of spin effects in the waveforms and template banks using models such as SpinTaylorT4. Required actions include:

- a SpinMDC that is currently being conducted using recolored S6 data to test the efficacy of aligned-spin BNS searches. This serves as a followup to [20]. Results indicate that by using an aligned-spin template bank, a significant increase in sensitive distance is achieved for highly spinning injections $(0.2 \leq \chi \leq 0.4)$ while no significant change is observed for astrophysically motivated BNS injections [65]. After each pipeline confirms these results, a decision must be made by the CBC group on whether to implement spinning BNS searches.
- improvements in the generation speed of time domain waveforms which capture the full range of physics, including spins, enabling parameter estimation to take full advantage of frequencies below 40 Hz within the timeline envisioned for publication. Such optimizations are being coordinated with tasks within the Parameter Estimation sub-group.

Pipelines. All three CBC search pipelines, GSTLAL-CBC, MBTA and the offline pipeline, are being actively maintained and are planned to be available in the advanced detector era. The initial LIGO offline pipeline, IHOPE [52], is well-tested and understood and could be run on the early data without further development. However, this code evolved over the initial LIGO era, and it is not flexible enough to meet the needs of the advanced detector era. A new workflow using the PYCBC code suite [66] is being created based on the experience of initial LIGO. The IHOPE pipeline will be maintained as a fall-back for early science runs, however current development of the offline pipeline is focused on PYCBC.

MBTA and GSTLAL-CBC are specifically designed for advanced detector BNS searches. The MBTA review has begun with a review committee convened in January 2015. GSTLAL-CBC review has begun with a review committee convened in January 2014. The PYCBC review has begun with a review committee convened in July 2014.

MDCs will be used to confirm the validity of the three search pipelines, and to ensure that both low latency triggers and final results are generated unambiguously: either by selecting a single pipeline for each or by developing a method for combining results across pipelines. Since early science runs are only expected to have astrophysically interesting sensitivity above about 30Hz, this will be our first target for tuning, followed by tuning for detectors that approach the full low-frequency sensitivity.

A common MDC is currently being finalized to develop and test our strategies for the validation of astrophysical results when running multiple pipelines as requested by the multiple pipelines policy [67]. Results show that the pipelines have differences in recovered SNR and FAR estimates [68]. Work is ongoing to follow up (rare) simulation candidates with apparent FAR inconsistencies. In such cases, we seek to identify the cause of the discrepancy and show that pipeline behaviour is consistent with the presence of claimed signal. Thus validation will depend on the outcome of these followups. Required actions include:

- understanding the consistency and relative sensitivities of the pipelines in the common MDC and developing a strategy and methods for validation of results. Additionally, use ER7 and ER8 to practice and test these validation methods.
- implement the proposal that has been submitted for how to combine results from the BNS, NSBH and BBH offline pycbc searches in the early aLIGO/AdV era [69]. A similar proposal exists for the online pipelines [70]. Work is ongoing to develop the template banks for these MDCs and perform the analyses.

Data quality. Interaction with the detector characterization (DetChar) and commissioning groups cuts across CBC searches. This effort will be performed within the CBC data quality (DQ) group using ER data. Details of data quality information in aLIGO are provided in [71].

The DetChar group has developed a basic plan for delivering online calibration, state, and injection DQ information together with online DQ triggers from the veto generator called iDQ [71]. Progress has been made in creating an Online Detector Characterization (ODC) system to provide real-time information on the state of the interferometers to low latency searches. The ODC-MASTER channel at each site will have a bit associated with science mode and other bits for the DQ of the instrument.

Regarding the veto system for offline searches, the DetChar group is currently building the infrastructure to support a veto category system similar to the system used in S5/S6. The offline searches will begin to test this as soon as it is ready. Required actions include:

- testing of the online DQ infrastructure in ER7 and ER8 provided that the ODC channel definitions have been finalized and the iDQ pipeline is running in low latency.
- testing the veto system for offline searches beginning before ER7. The CBC group should collaborate with the DetChar and Burst groups to design a sensible veto system.
- deciding how to handle different types of glitches. Do we veto triggers, down-rank them, or excise glitches from h(t)? Investigations are ongoing regarding the procedure for particular types of glitches.

Rates and significance. We will ensure that we have in hand methods to evaluate background rates down to the required level and to use search results to generate rate limits or intervals. Required actions include:

- understanding the impact of removing zero-lag coincidences from the background versus leaving them in. Preliminary results of the Hamlet test MDC can be found in [72]. A paper and conclusions are expected before O1.
- ongoing investigations into the uncertainties in estimates of false alarm probability and their implications for detection confidence. These investigations include the Hamlet test MDC [72] and the MDCs being performed for pipeline comparison.
- implementing tools for rate estimation with multiple populations. The reviewed loudest-event statistic method will remain available to compute rate upper limits; however, rather than further developing it to obtain nonzero rate intervals in the detection case, we plan to use a newer and more powerful

technique for rate estimation with multiple populations [73] which is currently undergoing testing and validation; review of any new methods and code will be required preparatory to O1.
- understanding systematic and statistical uncertainties in astrophysical rate estimation methods and their behaviour in the presence of one or more signals. Investigations include the multiple populations technique described in [73].

EM followup. Gravitational wave candidates will be collected in a central database (GraceDB), which will allow categorisation of triggers according to their originating pipeline, their type, etc. GraceDB also allows additional data produced by followup investigations of events to be associated with the candidates in an easily accessible way. Significant triggers will be processed rapidly to issue alerts for EM observers.

- To avoid collisions between different event streams submitting to the GraceDB database, three fields must be specified: (`Group`, `Pipeline`, `Search`). The values of these fields need to be finalized, particularly for the `Search` field.
- There are two aspects of releasing information on candidate events: the VOEvent message that the LVC sends to the Gamma-ray Coordinates Network (GCN) and the BAYESTAR HEALPix sky localisation maps. The format of the latter has been been finalized as FITS files while the format of the former is not yet finalized.

Parameter Estimation (PE) and Testing GR. To accurately and rapidly extract the parameters of a signal, two Bayesian inference packages called LALINFERENCE [74] and BAYESTAR [75] have been developed, the former with the goal of providing full information about the source parameters and the evidence integral for a model, the latter with the specific aim of returning sky localisation information in quasi-real time. Both packages have been reviewed, and are kept up-to-date as new features are introduced. An earlier version of LALINFERENCE was successfully used in previous science runs [34], and we have performed cross-validation of the BAYESTAR and LALINFERENCE methods through an extended BNS sky localization MDC [76, 77]. LALINFERENCE implements multiple independent sampling algorithms to allow for cross-checking of its results [74], which is part of the rolling review process for this code.

Further developments are required to fully exploit the advanced detector data, and are being pursued by the PE sub-group. Required actions include:

- Parameter estimation performance must be increased to allow use of improved sensitivity below 40 Hz. There are several efforts underway which address this. Advances in algorithm approaches are centered on two fronts: (i) A massively parallel ensemble sampler that is currently in the testing phase and has shown efficiency savings of a factor ≈ 100 [78]; (ii) A highly parallelizable approach based on a spherical harmonic mode decomposition to represent each physically distinct source through which one can efficiently evaluate the likelihood for generic source positions and orientations, independent of waveform length or generation time. In addition, by integrating over all extrinsic parameters and by using a purely Monte Carlo integration strategy, the calculation can be efficiently parallelized over the intrinsic and extrinsic space [79, 80]. Work is also in progress to deploy a “multi-banding" approach to the likelihood calculation, in which waveforms are sampled at at different rates depending on the instantaneous frequency of the signal.
- For estimation of detection candidates we must ensure robustness against calibration errors and noise artifacts. The methods developed in [81, 82] have been partially included and reviewed in LALINFERENCE; they will be fully included by the end of the year.
- Systematic effects coming from calibration uncertainties and limitations of the waveform models need to be addressed. Regarding calibration, the necessary pieces of code have been developed and tested, and the goal is to have them fully operational by O1. Effect of systematics coming from waveform uncertainties are the focus of intense development, but will not be ready by O1.

- GR tests will follow using a population or loud signals to test different models. The TIGER pipeline has been tested on S6/VSR2-3 recolored data using non-GR injections. The background distribution does not show any particular features that would hinder our detection capabilities. The TIGER pipeline review is ongoing.

3.2.7 Resources

Calibration. We will require calibrated data over the bandwidth of the inspiral signals (from 30 Hz to 2 kHz in O1). The most stringent calibration requirements are for parameter estimation, where the data is analysed coherently. This has been studied using simulated calibration error curves in [83, 84], which found that the most significant bias occurs for sky localisation accuracy when amplitude calibration errors have opposite signs in each detector. The study found that a maximum 1σ error of 10% in amplitude and $5\,\mathrm{deg}$ in phase calibration would be a suitable requirement for parameter estimation work.

The effect of calibration errors on detection was studied in [85, 83], with the conclusion that at the level outlined above the bias in Optimal SNR should be less than 3%, however the matched filter SNR and therefore distance estimation will depend linearly on the calibration error, affecting the estimated sensitive volume of any search. This will have an impact on the inferred rates, with a 10% error in amplitude giving around 30% error in the estimated volume and event rate.

We would like to have calibration to an accuracy of 10% in amplitude and $5\,\mathrm{deg}$ in phase prior to the start of any science run, and to have the calibration accuracy confirmed after a detection is made.

Detector Characterization. The BNS and CBC searches will work with the detector characterization and commissioning teams to identify and eliminate sources of transient noise. We will quantify the effect of these transients on our search sensitivity. One of the primary engineering run goals will be to develop the tools and techniques to effectively identify and mitigate glitches in the data. During data taking runs, we will examine the data on a daily and weekly basis for detector characterization. A daily CBC detector characterization pipeline will be run at the observatories to provide a “first look” at the data quality for detector characterization. The CBC group was able to provide at least weekly analysis of BNS data quality during S6 by engaging a team of around ten people involved in running the analysis, resolving any issues that arise, and regularly looking at data quality. Providing DQ information with a latency of less than one day will require coordination with the detchar and commissioning groups, and a discussion of the role of on-site scientists in aLIGO/AdV. To achieve our DQ goals, we strongly encourage engagement from the commissioning group in the CBC search, and extended visits by CBC scientists to the observatories.

Review We will need review teams for the following components: Search algorithms (3 pipelines), Waveforms, Rate and significance estimation, Parameter estimation. The review of all critical components must complete before ER8. Depending on the anticipated length of the reviews, some will need to begin up to a year before that. In addition, the following will be reviewed when early data becomes available: Data Quality, Calibration.

Person-power. At present 57 individuals have committed to work on aspects of BNS search development. In addition, a team of approximately ten people (assuming some overlap with the DQ team) will be required to maintain searches.

3.3 Search for Stellar-Mass Binary Black Hole Coalescences

3.3.1 Abstract

Coalescing stellar-mass binary black holes (BBH) are one of the most promising sources of gravitational waves (GWs) to be observed by LIGO and Virgo, with up to tens of BBH mergers per year at advanced design sensitivity [12]. The first detection of a stellar-mass BBH would prove the existence of a so far unobserved class of binary, constituting a major discovery in astrophysics. Measurement of the distribution of mass parameters in these binaries could distinguish between astrophysical models of compact object populations. As the knowledge of merger rates from current astrophysical observations of black holes and theoretical modelling is extremely uncertain, improved upper limits resulting from non-detection can constrain how massive stars and black holes evolve.

Some of the questions we aim to answer by searching for the gravitational waves emitted by binary black holes are: What happens when massive stars collapse to form black holes? How much of their mass escapes, and how much falls into black holes? Are black holes in binaries spinning, and if so how fast? How do massive binary star systems evolve? Is material transferred between the two stars? How does the collapse of one star affect the other? Do black hole binaries form mainly from existing binary systems of massive stars, or from random encounters between previously formed black holes? Does Einstein's theory of general relativity give an accurate description of merging black holes?

3.3.2 Scientific Justification

Stellar-mass BBHs are predicted from population synthesis simulations of the formation and evolution of massive binary stars [86] and through dynamical modeling of dense stellar clusters [87]. The estimated BBH coalescence rates arising from such predictions are very uncertain: the rate estimates for a $(10{+}10)\mathrm{M}_\odot$ BBH range from 10^{-4} to $0.3\,\mathrm{Mpc}^{-3}\mathrm{Myr}^{-1}$ [12]. Current rate upper limits based on LIGO and Virgo observations are only a factor $\sim$4 above the upper end of this range [59]. This means that even non-detection of these systems in the advanced observing runs will begin to contrain astrophysical models of their formation, and if such systems exist in this range of predicted rates, we have a realistic chance of observing them with LIGO and Virgo.

Both the lower and upper bounds of possible black hole masses are uncertain. The mass distribution of Galactic stellar-mass black holes is estimated in [88, 89, 90]; X-ray observations indicate black hole masses in the range $5 \leq M_\bullet/\mathrm{M}_\odot \leq 20$, confirmed with dynamical mass measurements for 16 black holes. An apparent lack of black holes ("mass gap") in the range 3–5 $\mathrm{M}_\odot$ [88, 89, 91] has been ascribed to the supernova explosion mechanism [92, 93]. However, objects initially forming as neutron stars may accrete material during a common envelope phase and collapse to low-mass (i.e., $< 5\mathrm{M}_\odot$), near-maximally spinning black holes [94]. The highest well-determined neutron star masses are close to $2\,\mathrm{M}_\odot$ [95, 96]; thus BBHs may contain component black holes as small as $2\,\mathrm{M}_\odot$.

The most massive stellar mass black holes are observed in extragalactic high-mass X-ray binaries, IC10 X-1 and NGC300 X-1, with black hole masses of $20{-}30\,\mathrm{M}_\odot$ and with Wolf-Rayet star companions [97, 98]. These systems are likely binary stars that formed in low-metallicity environments. Population synthesis based on recent stellar wind models allows for isolated black hole masses up to $\sim 80\,\mathrm{M}_\odot$ [93, 99]. Stellar black holes with masses above $100\,\mathrm{M}_\odot$ are conceivable [100]. However, recent population synthesis models estimate that the masses of black holes in stellar binaries are most likely to be $\lesssim 50\,\mathrm{M}_\odot$ due to common envelope binary evolution [101, 102]. There is no direct observational limit on BBH mass ratios.

X-ray observations of accreting black holes indicate a fairly uniform distribution of spins over the entire range allowed by general relativity, $0 \leq \chi \leq 1$, where $\chi = S/m^2$ is the dimensionless spin [103, 104, 105, 106, 107, 108, 109]; both low (χ $\sim$0.1) [110] and high ($\chi > 0.85$) values [111] are represented. The microquasar XTE J1550-564 [112] and population synthesis models [113] indicate small spin-orbit misalignment in field binaries. For massive field binary progenitors, the common envelope phase and mass transfer [114] are expected to cause strong correlations between spins and masses of the two black holes in field binaries [115]. However, no such correlations are expected for dynamically formed BBH.

The first few confident BBH detections will begin to distinguish astrophysical models of compact object populations. At least one intrinsic source parameter will be measured accurately (chirp mass, in the low-mass limit). For parameter ranges where detections are made, merger rates could be estimated to within a factor of a few; the absence of detections in other ranges will significantly tighten upper limits. Thus, classes of models can be ruled out entirely and parameters in population synthesis models can be constrained [116]. A detection with moderately high signal-to-noise ratio (SNR) may yield strong evidence for at least one spinning component, independently establishing the existence of black hole spin in General Relativity.

Several detections will begin to establish an observational distribution over chirp mass, and, possibly, spin or mass ratio, further constraining models of BBH populations [117]. Some of the louder signals could give indications for or against precessing component spins, and/or allow component masses to be well estimated, allowing us to test competing models of binary formation, and possibly test the "mass gap". The mass distribution depends on many aspects of the input physics, and ~ 10 detections may allow us to exclude many alternatives [118, 119] and provide information on star formation (metallicity), stellar winds, mechanisms of stellar collapse, binary evolution (supernova kicks, mass transfer episodes) and possibly on kicks from BBH mergers for dynamically formed binaries.

A more speculative aim is to test the validity of our waveform models based on post-Newtonian expansion, numerical relativity and black hole perturbation theory. Establishing a deviation from the predictions of General Relativity, for example the Kerr description of black hole, would be a major discovery.

3.3.3 Search Procedure

Search Pipelines In initial LIGO/Virgo the "ihope" pipeline [52] was developed and used to search for signals from non-spinning BBHs in GW detector data. To cover the range of possible masses and spins, a "bank" containing large numbers of template waveforms is used to matched filter the data. In prior searches the BBH parameter space was split in two: a "low-mass" pipeline, which used inspiral-only templates with total masses $M \leq 25\,\mathrm{M}_{\odot}$ [56, 57, 58, 59], and a "high-mass" pipeline using inspiral-merger-ringdown (IMR) templates to search for signals with $25 \leq M \leq 100\,\mathrm{M}_{\odot}$ [120, 121]. These searches yielded upper limits that were a factor 4 higher than optimistic rate predictions.

We plan to use two pipelines to search for stellar-mass BBHs: a workflow using the `PyCBC` code suite [66], and a workflow using the `gstLAL` infrastructure [30]. Both pipelines use matched filters to search the data for compact binary coalescences (CBCs). The `PyCBC` pipeline is a new implementation of the frequency-domain based `ihope` pipeline that was used to analyze initial LIGO data. It is run in an "offline" mode, with data analyzed in ~ 1 week batches. The `gstLAL` pipeline is a time-domain based search, originally developed to operate in an "online", low-latency mode for fast electromagnetic-followup of binary neutron star (BNS) and neutron star-black hole (NSBH) candidates. Although no electromagnetic counterpart is expected for stellar-mass BBHs, `gstLAL` can be run in an "offline" mode. Further, each pipeline uses different signal-based vetoes and ranking statistics, allowing for cross checks on potential gravitational-wave candidates. We therefore plan to use both pipelines to search for BBHs. How to compare and combine results between the two is currently being studied.

Mass and spin ranges As discussed above, BBHs may have a very wide range of component masses. We will search for BBHs with templates having a minimum component mass of $2\,\mathrm{M}_{\odot}$; this will allow us to test the possible "mass gap" between neutron stars and black holes, and to be sensitive to possible highly-spinning low-mass black hole formed via collapse from an accreting neutron star.

There is no fundamental distinction between what we call here "stellar-mass" binary black holes (BBH) and the more massive class of binaries referred to as "intermediate mass" black hole binaries (IMBH). Instead, the division of parameter between the two searches will be determined by the regime of efficacy of the matched filters and of signal-based vetoes. Based on past experience [59, 121], we expect existing signal-based vetoes to adequately mitigate noise transients for lower-mass, longer-duration signals, allowing us to produce results for these systems within a shorter time frame. We therefore plan to search for stellar-mass BBHs with $M \leq 50\,\mathrm{M}_\odot$ and minimum component mass $m \geq 2\,\mathrm{M}_\odot$ as our highest-priority search. We leave the higher mass systems for the IMBH search to occur on a more relaxed time scale in order to test and implement more sophisticated noise-mitigation techniques. We may revise the mass limits of these searches as the low-frequency ($\lesssim 30\,\mathrm{Hz}$) noise of the detectors improves in later observing runs.

Possible black-hole spins cover the whole range $0 \leq |\chi| \leq 1$. Recently, it has become possible to generate a template bank that covers the full range of non-precessing ("aligned") spins and masses of the stellar-mass BBH parameter space. It has been shown that this bank is significantly more sensitive to BBH signals than a non-spinning bank, particularly to signals with spins > 0.6 and component masses $\lesssim 20\,\mathrm{M}_\odot$ [122]. If BBH spins are misaligned with respect to the orbital angular momentum, the orbit will precess about the total angular momentum. Some precessing systems may be badly recovered even by aligned-spin templates; detecting them would require a pipeline specifically tuned to precessing signals. Developing such a pipeline would be a substantial longer-term project (see *Technical requirements and development plan* below), hence we do not envisage a precessing search until later observing runs approaching aLIGO design sensitivity. Until then, we plan to use the aligned-spin search.

Search procedure and latency Based on experience in S6/VSR2-3, we aim to run initial analyses in ~1-week batches within about a week of data taking. Each week, we will perform checks on the efficiency and correctness of data quality (DQ) flags, on the limiting background of the search, and on the ability of the search to recover simulated signals. We will use these checks to tune the pipeline to account for changes in data quality if necessary, prior to examining the search results containing potential GW candidate events.

If one or more events with low false alarm probability are then found, additional followup work on detector operation, data quality, candidate significance, estimation of source parameters, and rate inference will be needed. Initial studies may take 1-3 weeks, and are included within the 3 month timeline to public release for a first detection claim. Bayesian likelihood sampling algorithms will be used to determine the likely source parameters and their uncertainties, and to assess the evidence for the presence of spinning components and the (mis)alignment of their spins. Multiple waveform models (if available) will be used to check for consistency.

3.3.4 Validation plan

Any astrophysically significant candidate event will be validated using the multiple pipelines described above. A BBH pipeline comparison study should be performed, similar to that for BNS, to inform us of the variation between the output of pipelines when analysing the same data. In the event of a significant candidate being identified, we plan to follow these guidelines to validate it:

- Does the candidate have low false alarm rate in *at least one* pipeline?
- If a significant trigger is reported in more than one pipeline the responses of the pipelines should be consistent within expected differences. Differences in clustering may result in a different template being recorded as the trigger.
- If a significant trigger is reported in only one pipeline, we will act on a case-by-case basis, looking at the progress of a trigger through the pipelines with the aim of determining why differences are

observed. If one pipeline is found to be able to rule out a candidate for reasons that were not taken into account by the triggering pipeline (e.g. updated data quality vetoes) then it should not be put forward as a detection candidate. If no reasons are found to rule out the candidate then it may progress through the detection process.

- Is the candidate robust to slight changes in the analysis (PSD variation, choice of segment boundaries etc?)
- Does the candidate pass all checks in the CBC detection checklist?

In the case where no detection candidate is found, we will verify that multiple pipelines agree on estimates of sensitive volume within known differences in live time analysed and pipeline sensitivity.

3.3.5 Publication plan

Publication scenarios up to and including the first detection

Confident BBH detection Based on experience in analyzing simulated signals in the initial LIGO era [60, 123], we will consider events with significance $\gtrsim 4\sigma$ (false alarm probability $\lesssim 6 \times 10^{-5}$) to be detection candidates. The first such candidate that passes all followup checks will warrant its own publication, with the language describing the candidate based on its false alarm probability (e.g., a 5σ event may be considered a confident detection). Any BBH event with significance $\gtrsim 4\sigma$ and passing initial followup checks will be presented to the LIGO and Virgo collaborations as a detection candidate. Detailed followup will then be performed, including checks on: the analysis results; the data quality at the candidate time and over the whole of the analysis; the detector behaviour and possible environmental disturbances; the overall consistency of the event with gravitational-wave signals; any other investigations that the Detection Committee or the collaborations consider necessary. At the same time a short paper will be drafted describing the search and the candidate, and discussing possible astrophysical implications.

In order to meet a schedule for public information release within 3 months of identifying a detection candidate in the offline analysis, detector characterization and calibration should be active before the start of a observing run. The analysis methods (codes and run procedures) used for the search, significance calculation and source parameter estimation must also be officially reviewed before application.

For a candidate event happening early in an observing run, additional data may be needed to assess detector operation, the noise background of the search, and astrophysical source rates, to a level necessary to support a detection claim. In exceptional cases, if new and relevant information appears towards the end of the run, it may be necessary to extend the 3 month schedule.

Marginal candidates If one or more events that pass initial followup are found with significance $\lesssim 4\sigma$ but $\gtrsim 3\sigma$ (false alarm probability $\text{few} \times 10^{-3}$), the events will be logged to support a possible joint detection in combination with events from other analysis times (and possibly other compact binary searches). If no strong candidates subsequently appear we expect to wait until the end of a data-taking period before presenting a possible case for detection.

Non-detection and upper limits If the analysis results are consistent with non-detection, we aim to submit a paper describing updated rate estimates for the BNS, NSBH and BBH searches whenever there is a tenfold increase in VT, where V is the volume of the Universe within which a signal could have been detected over the search time T.

For the BBH search, a relatively small increase in detector sensitivity may put nontrivial bounds on population synthesis models in the non-detection case. For example, Fig. 3 shows that “Submodel A”

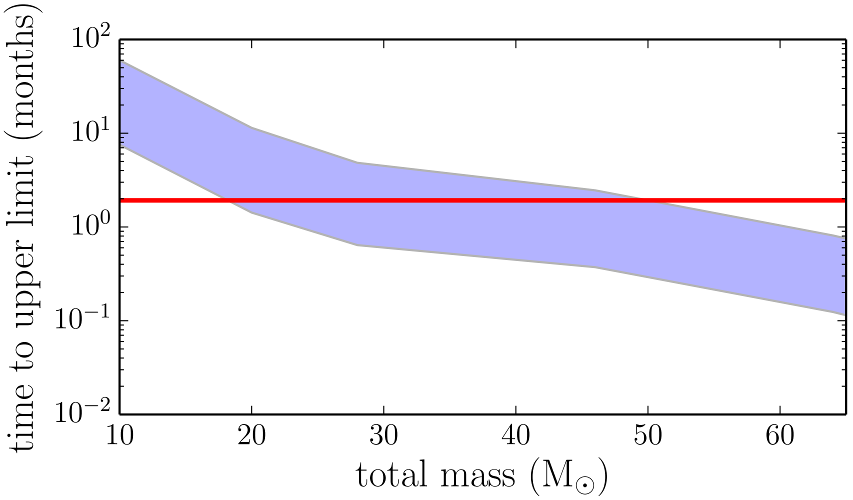

Figure 3: Amount of search time needed in early advanced LIGO to place upper limits that would constrain optimistic population synthesis models ("Submodel A" in [101]) of the rate of BBH coalescence, as a function of total mass. The red-line region indicates the estimated search time in the first observing run (O1) assuming it is 3 months long with a coincident duty cycle between Hanford and Livingston of 64% [124]. The blue-shaded region indicates the time to upper-limit assuming the optimistic (lower line) and pessimistic (upper line) "Early" PSDs in [124]. For instance, if the detectors have the optimistic noise PSD, Submodel A will be ruled out at the 90% confidence level for systems with total masses $\gtrsim 20\,\mathrm{M}_\odot$.

described in [101] would be ruled out for certain equal-mass, non-spinning systems with just 3 months of run time predicted for the first aLIGO observing run (O1).

Post first-detection At the end of observing runs or periods of data-taking of order 6 months–1 year we will prepare longer papers with a detailed summary of analysis results and discussing all interesting events found. Such summary papers may be joint between different compact binary searches, and will include updated source rate estimates (upper limits for searches where there is no detection). This publication schedule is also appropriate in the post-detection era when multiple signals may be seen in over an observing run or data-taking period. If a well-defined BBH signal population is seen it may also be appropriate to publish separate papers on features specific to BBHs.

3.3.6 Technical requirements and development plan

Waveforms Accurate and computationally efficient waveform models for BBH coalescences are crucial to the scientific output of the search. The waveforms are used as matched filters for detection, as simulated signals for tuning the pipeline, estimating the source parameters of a detected coalescence signal, and for interpreting search results in terms of the astrophysical coalescence rate. Waveform models require significant work, partly done outside the LSC, to develop and implement reliably, as well as thorough review including comparison between different models.

Including the merger and ringdown phases of the coalescence in search templates is important for BBH signals. We wish to use the SEOBNRv2 waveform model for our templates, as it models these effects for all masses and mass-ratios that we plan to search for, and includes non-precessing spin effects covering nearly the entire range of possible spin magnitudes (specifically, up to spins of 0.99) [125]. The implementation of this waveform in LAL has been reviewed. However, the generation time of SEOBNRv2 is too slow to be used in a search. Instead, we plan to use the double-spin reduced order model (ROM) of SEOBNRv2 [126]. This surrogate is substantially faster than the full model, and it has been shown to be effectual to SEOBNRv2 over the entire parameter space [122]. Finishing the review of the double-spin SEOBNRv2 ROM before the start of O1 is therefore a priority.

Since our templates model only effects from aligned spins, but real astrophysical BBHs may have misaligned spins, we additionally require models of precessing IMR signals in order to properly interpret our serarch results. The IMRPhenomP model gives a fast approximate representation of IMR signals from single-spin precessing binaries [127], and is already implemented in LAL, but requires review. In the longer term, SEOBNRv3, a model for double-spin precessing BBH, will be implemented and reviewed, and surrogate models will be developed for SEOBNRv3 that make its computation efficient for high-statistic sensitivity analysis. This model will provide a cross-check against conclusions drawn from IMRPhenomP studies, and may substantially widen the mass and spin parameter space over which such conclusions can be made.

Analysis pipeline The data analysis pipeline begins with the layout of matched filter templates over the parameter space to be searched. Templates are laid out such that largest possible loss in SNR due to the discreetness of the bank is less than a predetermined value (usually 3%). A bank is "effectual" if it is found to meet this requirement for a set of simulated signals. Two techniques exist for placing templates: a "geometric" algorithm in which templates are placed on a grid using a post-Newtonian metric [20], and a "stochastic" algorithm which uses randomly drawn points to cover the parameter space to the desired match [128]. It has been shown [122] that by using a combination of these two techniques we can produce an effectual template bank to non-precessing signals with spins $-0.99 \leq \chi \leq 0.99$, component masses $m_{1,2} \in [2, 48]\,\mathrm{M}_{\odot}$, and total masses $M \leq 50\,\mathrm{M}_{\odot}$. This bank is significantly more sensitive than the non-spinning bank to low-mass, high-spin ($m_{1,2} \lesssim 20\,\mathrm{M}_{\odot}$ and $\chi \gtrsim 0.6$) systems, with a sensitive volume that is 10 times or larger than the non-spinning bank to these signals. Over the entire parameter space we expect the aligned-spin bank to be $\sim 25\%$ more sensitive than the non-spinning bank. We therefore plan to use this bank in the first observing run of advanced LIGO (O1).

Assuming that in O1 we will not be sensitive to signals below 30 Hz, we will need $\sim 65\,000$ templates to carry out an aligned-spin search. We estimate that this number will grow to $\sim 240\,000$ and $\sim 580\,000$ templates in the second (O2) and third (O3) observing runs, respectively. There is, however, significant overlap between low-mass BBH templates and high-mass NSBH and BNS templates. By generating a single bank that covers the entirety of the BNS, NSBH, and BBH parameter spaces we may be able to reduce these template counts. We are currently studying the efficacy of combining searches, as well as any complications that may arise.

Studies on BNS and NSBH systems indicate that a search using aligned-spin waveforms will also perform better at detecting precessing binary black hole systems than a non-spinning search [20, 129]. Even so, it may be necessary to develop a search that includes precessional effects in the search templates. Doing so is challenging, as these signals depend on more parameters: the spin components in the binary orbital plane, as well as the binary's orientation and polarization. It may be possible to reduce the parameter space to a smaller set of dominant physical parameters, but this remains to be tested. Even with this reduction, a conservative estimate for the number of templates needed to perform a search in aLIGO is $O(10^6)$. The additional templates will inevitably increase the noise level of the search; the sensitivity of the search to possible astrophysical signal distributions must therefore be taken into account in order to determine if, and for which spin orientations, a precessing search is advantageous.

As a first step we will investigate the sensitivity of the aligned-spin bank for precessing inspiral-only and/or IMR waveforms (initially, IMRPhenomP [127], and SEOBNRv3 [130] when available — expected late 2015). Due to these uncertainties and technical challenges we do not plan to deploy a precessing search in the first advanced LIGO observing run.

We plan to use waveforms that only include the dominant mode. It has been shown that including higher harmonics in templates does not improve search sensitivity to non-spinning signals with masses and mass ratios covered by our search [131, 132]. The effect of higher harmonics in spinning signals currently cannot be quantified due to the lack of suitable IMR waveform models.

Pipeline Past searches regenerated (non-spinning) template banks approximately every 40 minutes. This was done to capture changes in the sensitivity of the instrument over this timescale. Generation of an aligned-spin bank is much more computationally expensive; it is not feasible to regenerate the bank on the same time scale. However, using S6 data, it has been shown that using a fixed bank generated with a PSD averaged over ~ 1 week time scale results in no loss in sensitivity. We therefore plan to generate a bank once per week-long run in O1. It may be more advantageous to regenerate the bank when some characteristic of the noise has changed. We plan to study other indicators for when to regenerate the bank for future observing runs.

We demand that candidate events be coincident across multiple detectors. To determine coincidence, past searches used “ethinca” [51] which allows template of different masses to be coincident, accounting for parameter correlations. However, MDC studies indicate that “exact” match — which requires the same template in all detectors — has a slight sensitivity advantage over ethinca. This is particularly true for high masses where the post-Newtonian metric becomes increasingly inaccurate. Furthermore, ethinca does not currently account for component spins, making it difficult to use in an aligned-spin search. We therefore plan to use “exact” match for determining coincidence across detectors.

Data Quality The BBH search is generally more susceptible to non-Gaussian transients (glitches) than the BNS search, on account of the shorter templates present in the BBH search; for example, while the shortest BNS templates coalesce at $\sim 1000\,\mathrm{Hz}$, the shortest BBH templates ring down at $\sim 100\,\mathrm{Hz}$. Thus, removing or down-ranking bad data is of greater importance in performing an effective search. The task is complicated by the broad parameter space covered by the BBH search, as templates in the lowest-mass part of the search will have very different glitch responses than those in the highest-mass part of the space. The primary product of detector characterization in advanced runs is expected to be a set of “DQ triggers” that indicate significant excess noise in environmental and auxiliary channels, along with other information, such as the duration and characteristic frequency of the triggers. In order to generate appropriate sub-standard data, these triggers should be correlated with non-Gaussian (high-SNR) triggers in the BBH search, and criteria for the efficiency of the resulting segments should be determined.

Rates and significance A challenge for any GW search is identifying a detection candidate with very low false alarm rate. Background estimation in the `PyCBC` pipeline is based on the rate of coincidence between triggers in different detectors after shifting each detectors’ data relative to each other by an amount greater than the light travel time between the detectors. A large number of such time shifts is performed to produce independent background samples, which are then used to measure the significance of a coincident trigger in the non-time-shifted analysis. The `gstLAL` pipeline estimates event significance using a “slideless” method, which is more computationally efficient and may reach smaller false alarm probabilities over a given data set. The method works by multiplying together the probability density functions of single detector triggers to estimate the probability of getting a multiple-detector trigger with a given set of parameters [133, 134].

Both background estimation methods are known to produce answers which depend strongly on whether the event being ranked is kept in the list of events used to estimate the background [59]. Furthermore, we not yet know whether these two methods always give comparable results, particularly at high masses where the noise distribution and coincidence test are harder to model. Both of these issues are currently under investigation.

We are also investigating the optimal strategy for binning the background, whose characteristics can vary strongly in the high mass regime considered in this search. Previous searches have used simple, hand-tuned functions of SNR and χ^2 to separate signal from background. However, these will be far from optimal at high masses, where signal-based χ^2 tests [49] become ineffective. Variation of the background over the parameter space has so far been dealt with by dividing triggers into a small number of bins. However, efficiency and background estimation should be improved if the variation can be modelled by a smooth function.

One interesting output of a search is an inference on source merger rates: currently this takes the form of upper limits calculated via the Bayesian loudest-event statistic [135] using estimated background and signal distributions. Statistical errors in this calculation need to be quantified and controlled, and changes in implementation may be needed to deal with the first few detections. In the limit of large numbers of signals a "cut-and-count" method is simpler. However, it is unclear when this condition is satisfied. Developments here should also be applicable to BNS and NSBH.

Parameter Estimation The information gained from accurate estimates of source parameters will be crucial to answer astrophysical questions. Parameter estimation pipelines have already demonstrated the ability to analyze candidate events in the BBH mass range in initial LIGO and Virgo data [136]. If we are able to measure BBH component masses, the possible "mass gap" between $3 - 5\,\mathrm{M}_\odot$ can be investigated. Distinguishing between precessing and spin-aligned BBH will also have strong astrophysical implications.

Systematic bias from waveform uncertainty will always be a concern for parameter estimation efforts. This can be somewhat mitigated by implementing improved waveform models (see above). Additionally, parameter estimation pipelines will be run with different waveform models to check for consistency and characterize the level of any potential bias. Studies are in progress to determine the expected accuracy of such measurements in the advanced era, and point to areas where further development is needed.

Waveform generation is the dominant cost of parameter estimation, rising strongly as either the mass or the minimum sensitive frequency decreases. Computational costs are estimated using spinning inspiral-only (namely, SpinTaylorT4) waveforms. Though estimating BBH parameters requires merger and ringdown, the IMRPhenom$X$ models can be generated more quickly than SpinTaylorT4. SEOBNR waveforms are too computational expensive to use in parameter estimation. However, efforts are underway to use reduced order modeling to create computationally-efficient surrogate models [137, 138, 139, 140, 141, 142, 126]. To be usable within the same computational constraints, surrogate SEOBNR models should be comparable to SpinTaylorT4 in the cost of generation.

Testing General Relativity Bayesian methods for testing whether the waveforms of any detected signals are consistent with the expectations from General Relativity have been developed [143] and are demonstrated for BNS systems [46]. In the next few years it is planned to implement these methods for BBH systems, including fully spinning waveforms, and address the many additional challenges arising for such systems including uncertainty of waveform models and the effects of non-ideal data quality. A mature BBH analysis pipeline is not expected by O1 (late 2015), but will be run later in the advanced observing runs once the feasibility of testing General Relativity with BBH systems is demonstrated. In order to confidently detect a deviation from General Relativity waveforms, a large number of "background" (i.e., simulated General Relativity waveform) sample results may be needed, with a computing cost of order(1000) cores running over months. Feasibility studies will have similar computing loads.

3.3.7 Resources

Calibration. We will require calibrated data over the bandwidth of the inspiral signals (from $30\,\mathrm{Hz}$ to 2 kHz in O1). The most stringent calibration requirements are for parameter estimation, where the data is analysed coherently. This has been studied using simulated calibration error curves in [83, 84], which found that the most significant bias occurs for sky localisation accuracy when amplitude calibration errors have opposite signs in each detector. The study found that a maximum 1σ error of 10% in amplitude and $5\,\mathrm{deg}$ in phase calibration would be a suitable requirement for parameter estimation work.

The effect of calibration errors on detection was studied in [85, 83], with the conclusion that at the level outlined above the bias in Optimal SNR should be less than 3%, however the matched filter SNR and therefore distance estimation will depend linearly on the calibration error, affecting the estimated sensitive volume of any search. This will have an impact on the inferred rates, with a 10% error in amplitude giving around 30% error in the estimated volume and event rate.

We would like to have calibration to an accuracy of 10% in amplitude and $5\,\mathrm{deg}$ in phase prior to the start of any science run, and to have the calibration accuracy confirmed after a detection is made.

with well-quantified uncertainties before the start of any observing run.

Detector Characterization & Data Quality Reliable knowledge of detector operation – specifically, science, calibration and analyzable data segment information – is a prerequisite for our analyses. Additionally, non-Gaussian noise transients in strain data have a significant impact on the sensitivity of BBH searches, particularly at higher masses where templates are shorter and easier for glitches to mimic. We will work with the detector characterization and commissioning teams to identify and eliminate sources of transient noise. A primary goal is to perform automated single-detector BBH analyses (which may be run together with BNS/NSBH) with a latency of one day, and to develop tools to identify and mitigate glitches by comparing BBH triggers with other low-latency trigger generators and establishing correlations with environmental and auxiliary channels. This effort will require a team of ~ 5 people, some from within the Detector Characterization team and some from within the compact binary DQ team, to develop and maintain DQ analyses and liaise with commissioners and LSC Fellows at the detector sites. Methods and codes should also be applicable to BNS, NSBH and/or IMBH searches; thus effort may be "shared" between sources. We also strongly encourage engagement from commissioning groups and LSC Fellows in compact binary searches, facilitated by CBC-specific figures of merit for detector operation.

Review Methods and codes used for the following tasks (many of them common to several CBC searches) will need review before the first advanced science data: Search algorithms; waveforms; significance and rate estimation; parameter estimation. The level of review required will range from checking patches and upgrades relative to previously existing code, to complete reviews for new methods and/or codes. Reviews of code used for the search pipelines that we expect to use (pycbc and gstlal) are currently on going; we expect these reviews to complete in the Summer of 2015. Methods used for Data Quality and Calibration will be reviewed when early advanced strain data becomes available.

Search and detection A team of approximately ten people will be required to maintain BBH searches. If strong candidates are identified the preparation for a possible detection claim will additionally call on CBC review team(s), on the DAC and Detection Committee, and on other expertise in the collaborations, in particular for calibration, detector characterization and instrumental operation.

3.4 Search For Neutron Star – Black Hole Coalescences

3.4.1 Abstract

To date, no binary system consisting of a neutron-star and a black-hole (NSBH) has been observed. At design sensitivity, Advanced LIGO and Advanced Virgo will be able to observe the merger of a $1.4\,\mathrm{M}_\odot$ neutron-star and a $10\,\mathrm{M}_\odot$ black hole at a maximum distance of $\sim\ 1\,\mathrm{Gpc}$. Current best estimates predict that 0.2-300 NSBH systems merge every year within the predicted sensitive volume of the Advanced LIGO and Advanced Virgo network. There are significant uncertainties in these measurements and gravitational wave observations may be the only way to obtain much tighter constraints on the rate of NSBH mergers. Observations of NSBH systems will allow scientists to explore the distribution of black-hole masses, black-hole spin magnitudes and the orientation of that spin with respect to the orbit. This will help to answer important astrophysical questions, such as whether there is a "mass gap" between the most massive neutron stars and least massive black holes [91], and whether black holes preferentially have large spins [111]. For certain configurations the neutron star will tidally disrupt during merger and this can give rise to an observable electromagnetic counterpart [144]. Jointly observing neutron-star–black-hole mergers in the electromagnetic and gravitational-wave spectra will allow a much better understanding of how, and in what circumstances, neutron stars disrupt during the merger. This can provide information about the underlying neutron star equation-of-state [145].

Achieving this sensitivity to NSBH systems with Advanced LIGO and Advanced Virgo requires highly accurate models of the gravitational-wave signals emitted and required matched-filtering search techniques utilizing very large banks of filter waveforms. Even a single neutron-star black-hole observation potentially yields much astrophysical insight due to the importance of the black hole's spin and precessional effects, the effect of the neutron star equation-of-state on the inspiral and merger, and the potential tidal disruption in the observed gravitational wave signal.

3.4.2 Scientific Justification

The first observation of gravitational waves (GWs) from the merger of a NSBH system would provide a direct measurement of the masses and spins of the system. Due to degeneracies between these parameters, only the chirp mass will be recovered accurately if the system does not exhibit precessional effects [146]. However, if the system does exhibit modulation due to precession in the gravitational-wave signal it may be possible to measure the masses and spins to a precision of $< 10\%$ [147, 148] or better [149].

With ~ 10 NSBH observations it will be possible to more closely probe the underlying physics [150]. One of the first pieces of information that we will have access to is the observed rate of NSBH coalescences. We will also directly have access to a measurement of the mass and the spin distribution of both the BH and neutron star (NS). Extracting information about the masses and spins of NSBH systems will allow us to probe different formation models by investigating how well the observed distributions agree with the theoretical distributions. Measuring the spin distribution of BHs in NSBH systems would provide information about the formation of NSBH field binaries. The distribution of the BH spin magnitude and orientation would provide information about the size of the kick imparted on the compact objects during their formation [151, 152, 153, 154, 155] and could help to distinguish whether NSBH systems are predominantly formed by dynamical capture or stellar evolution of field binaries [156, 157, 158, 159]. Additionally, the magnitude of the spin of the BH would provide information on the amount of accretion experienced by the BH [110, 94, 160]. All marginal and confident detection candidates will be examined to see if any deviations from general relativity can be observed from any single signal or from a collection of signals.

Fully general-relativistic numerical simulations of NSBH systems have been performed (for e.g. [161, 162, 163, 164, 165, 166, 144]) and show that certain combinations of mass, spin, and NS equation of state (EOS) parameters will cause the neutron star to tidally disrupt before coalescence. These systems could power the central engines of short gamma ray bursts (GRBs) or produce other types of prompt or delayed electromagnetic (EM) counterparts [167]. The identification of an electromagnetic counterpart could provide several insights into the physics of NSBH systems. The estimated parameters from many of these systems, along with EM counterparts or lack-there-of, would allow us to test the predictions of numerical relativity for which configurations of the BH mass, spin, and NS EOS produce accretion disks that power short GRBs [168, 167, 169]. If NSBH binary mergers produce short GRBs, a GW measurement of the inclination angle of the binary may help in understanding the structure and geometry of the jet [170, 171]. Additionally, galaxy host identification of EM counterparts can allow us to better constrain their formation process, as with (the possibly identical) short GRB events [172]. For short GRBs, host galaxy information has been used to constrain their typical age; host galaxies also let us constrain the size of the kicks imparted on the binary from supernovae [154, 153].

As with binary-neutron-star (BNS) systems, finite size effects may become measurable with collections of observations. This would aid in making a statistical NS radius and EOS measurement [173, 174, 37, 39]. Investigation is needed to understand if this kind of study is better facilitated with NSBH signals than with BNS signals.

NSBH systems differ fundamentally from binary black hole (BBH) systems as they have the possibility of tidally disrupting the neutron star leading to a waveform that differs substantially from those of BBH in the merger and ringdown phase. This tidal disruption also entails the possibility of electromagentic (EM) signals, making NSBH systems a target for rapid EM follow-up and inclusion in low-latency searches. Unlike BNS systems, NSBH systmes are likely to have larger total masses, placing the late inspiral part of the signal closer to the most sensitive bucket region of Advanced LIGO and Advanced Virgo detectors. This requires special attention that the template waveforms contain all relevant relativistic effects and also separate tuning of data quality vetoes. Since the black hole may have close to maximal spin, it is important that spin effects are included in NSBH searches.

In the absence of a detection of NSBH GW signals, astrophysically relevant constraints can be placed on the rate of NSBH coalescences. The absence of NSBH detections will also constrain the fraction of short GRBs powered by NSBH mergers.

3.4.3 Source Parameters

NSBH systems are thought to be efficiently formed in one of two ways: either through the stellar evolution of field binaries or through dynamical capture of a neutron star by a black hole [156, 157, 158, 159]. Though no NSBH systems are known to exist, one likely progenitor has been observed, Cyg X-3 [175]. Rates for the coalescence of NSBH systems are not well known, however a “realistic” estimate from population synthesis of field binaries is given as 0.03 $\mathrm{Mpc}^{-3}\ \mathrm{Myr}^{-1}$ [176]. A “pessimistic” estimate is given as 6×10^{-4} $\mathrm{Mpc}^{-3}\ \mathrm{Myr}^{-1}$ and an “optimistic” estimate as 1 $\mathrm{Mpc}^{-3}\ \mathrm{Myr}^{-1}$ [176]. These yield observation rates for Advanced LIGO and Advanced Virgo of 0.2 - 300 yr^{-1}.

The mass distribution of NSBH systems is not well constrained. However, it is possible to place estimates on the mass ranges by using the masses of neutron stars and black holes observed in other systems. The masses of known neutron stars are reported to be in the range $0.7\,\mathrm{M}_\odot$ to $2.7\,\mathrm{M}_\odot$ with a mean mass of $\sim 1.4\,\mathrm{M}_\odot$ [14, 177], though the lower value, $0.7\,\mathrm{M}_\odot$, comes from an imprecise measurement of a single system that is also consistent with a higher mass. Neutron stars in binary neutron star systems have a more narrow observed mass distribution of $(1.35\pm0.14)\,\mathrm{M}_\odot$ [14, 177]. Theoretical models support the production of a population of neutron stars formed in binaries through electron-capture collapse of O-Ne-Mg cores, and predict masses which are consistent with these observations [178, 179]. Lower mass Fe cores are predicted to lead to neutron stars with masses almost as low as $1\,\mathrm{M}_\odot$ [17].

The mass distribution of Galactic stellar mass black holes is estimated in [180, 88, 181], and X-ray observations yield black hole (BH) masses $5 \leq M_\bullet/\mathrm{M}_\odot \leq 20$, confirmed with dynamical mass measurements for 16 BH. An apparent lack of BH masses in the range 3–5 $\mathrm{M}_\odot$ [180, 91] has been ascribed to the supernova explosion mechanism [92, 93]. However, black holes formed from stellar evolution may exist with masses down to $2\,\mathrm{M}_\odot$ especially if they are formed from matter accreted onto neutron stars [94]. The most massive stellar mass BHs are observed in extragalactic high-mass X-ray binaries, IC10 X-1 and NGC300 X-1, with BH masses of $20 - 30\,\mathrm{M}_\odot$ and with Wolf-Rayet star companions [97, 182]. These systems are likely field stars that formed in low-metallicity environments. Population synthesis based on recent stellar wind models allows for isolated black hole masses up to $\sim 80\,\mathrm{M}_\odot$ [93, 183]. Common envelope binary evolution may reduce the maximum expected component mass and total mass to $\lesssim 100\,\mathrm{M}_\odot$ [101], however stellar BH with mass above $100\,\mathrm{M}_\odot$ are conceivable [100], overlapping the range associated with intermediate-mass black holes formed by repeated mergers.

X-ray observations of accreting black holes indicate a fairly uniform distribution of black hole spin [103, 104, 105, 106, 107, 108, 109]. Some black holes observed in X-ray binaries have very large dimensionless spins (> 0.7), while others could have much lower spins (~ 0.1) [110]. Measured black hole spins in high-mass X-ray binary systems have large values (> 0.85), and these systems are more likely to be progenitors of NSBH binaries [111]. Isolated black hole spins are only constrained by the relativistic Kerr bound $\chi = S/m^2 \leq 1$. The microquasar XTE J1550-564 [112] and population synthesis models [113] indicate small spin-orbit misalignment in field binaries. Dynamically formed NSBH systems, in contrast, are expected to have no correlation between the spins and the orbit.

At birth, the spin period of a neutron star is believed to be in the range of 10 to 140 ms, which corresponds to a dimensionless spin of 0.04 [150], depending on the equation-of-state (EOS) of the NS. These natal spin periods are expected to decrease significantly in the long time between the formation of the neutron star and the merger of the two objects. The observed dimensionless spins for neutron stars in binary-neutron-star systems (e.g., J0737-3039) are ≤ 0.04 [184] It is possible for neutron stars to be spun up through accretion to much higher spins that will persist until merger [19], for example a 1 ms pulsar has a dimensionless spin of 0.4. However, it is unlikely for a field NSBH system to be spun up by accretion as the BH would normally form first. It is therefore plausible that the neutron-star spin will be negligible in NSBH systems formed from field binaries. However, dynamically formed NSBH systems may have neutron star spin values up to 0.4.

3.4.4 Search Description

During the initial-detector era, algorithms were developed for searching for NSBH mergers, along with BNS and binary black hole (BBH) mergers, and estimating their parameters in a network of detectors [185, 48, 186, 187, 188, 136]. These searches did not include the effect of the black hole's spin angular momentum in their search templates. Search methods which incorporated spin effects were considered, but were found to increase the false alarm rate which resulted in a less sensitive search [189, 190, 191]. Search results were published in a series of papers: S3 and S4 [192], S5/VSR1 [193, 57, 194] and S6/VSR2,3 [59]. No observations were made, and an upper limit of $R_{90\%} \leq 36\mathrm{Mpc}^{-3}\mathrm{Myr}^{-1}$ was placed on the rate of neutron-star–black-hole coalescence, which is almost 2 orders of magnitude above the most optimistic rate estimates. A number of the techniques used in the initial detector era, including the matched-filter maximization algorithm, the chi-squared glitch-rejection tests, significance calculations and parameter estimation techniques can be directly used in the advanced-detector era. Significant refinements are needed for advanced detectors, including the inclusion of spin effects in search templates and extending search and parameter estimation pipelines to incorporate the longer waveform templates required [195].

To date, the NSBH search has been performed in conjunction with CBC searches for BNS and stellar-mass BBH [192, 193, 57, 194, 59]. Although we will continue to pursue common search techniques for the three sources, the physics of NSBH and BBH searches, particularly spin and merger effects, complicate the searches for these systems compared to a BNS search. Therefore we may perform a dedicated NSBH search in the advanced-detector era.

The search will be performed with matched filtering. This requires covering the parameter space of interest with a template bank of filter waveforms. In the initial detector era, template banks were created using methods described in [185]. However, this method was limited to waveforms that do not include spin effects and do not include the most recent terms in the frequency-domain TaylorF2 waveform expansion [23]. It has been shown that neglecting spin effects in the filter waveforms in the advanced-detector era will result in the observation rate of NSBH mergers being reduced to 65% of the optimal rate of $0.2 - 300\text{yr}^{-1}$ [196]. This assumes an isotropic in direction and uniform in magnitude spin distribution and assumes Gaussian noise. More refined studies have been performed in realistic non-Gaussian noise that confirm that a significant number of NSBH mergers could missed if the effects of spin are not included [195, 66]. For specific system configurations, especially those with large spin magnitudes and high mass-ratio, a non-spinning search would achieve a detection rate of only 20% of observable NSBH systems [195]. High spin and high mass-ratio systems can give waveforms with very unusual features, for example the sense of orientation reverting as the orbital plane precesses by more than 90 degrees [197]. Those systems will be rare, but if we are able to identify them, we will be able to extract particularly much information about the source [147, 149]. Even a single such detection would have potentially dramatic inferences on astrophysics. However, because only a small region of parameter space shows such features, this will not show up on parameter-averaged efficiency studies.

Recent work has made it possible for template banks to be placed including the effects of spin *aligned* with the orbital angular momentum, and including the latest terms in waveform expansions [198, 196, 195, 66]. We will use these developments for placing template banks for NSBH searches with Advanced LIGO (aLIGO) and Advanced Virgo (AdV) from O1. The placement of waveforms in these template banks depends on the noise spectrum of the instrument. Historically, template banks have been constructed on an hourly basis to track changing instrument noise. In the advanced detector era the size of the template banks, especially with the inclusion of spin, may make this impractical. NSBH template banks are typically the largest of all sources in the CBC group. Investigations are underway to test if template banks can be used for longer timescales without unacceptable reductions in efficiency or increases to the number of templates.

NSBH signals will be searched for in all available science-quality data where at least two detectors are operating. To facilitate optimal performance of the matched-filtering algorithms across large-scale computing clusters the search is parallelized by splitting the template bank and/or by splitting the data to be filtered in time. Currently, we only have search algorithms capable of using search templates where the component spins are aligned with the orbital angular momentum. This will cause some loss in search efficiency as precessional effects, which will be present in real NSBH signals, will not be included in the filter waveforms, however these losses are not judged to be critical [196, 66]. Previously tested algorithms incorporating precession increased false alarm rates and thus did not show an overall increase in search sensitivity [191, 199]. Therefore, our initial search proposal is for a search using only aligned-spin template waveforms (that is. not including precessional effects), while we develop and assess potential new methods for efficiently incorporating precession. The technology needed for an aligned spin search has already been demonstrated in recolored data and engineering runs [195] and will be used in O1.

The search will target NSBH inspirals and mergers comprised of component masses spanning the entire plausible NS mass range of $1\text{M}_\odot$ to $3\text{M}_\odot$. Note that while the error bars on observations of NS masses extend as low as $0.7\text{M}_\odot$ [14], there is no known mechanism by which such low-mass NSs could be formed. The range of BH masses is currently uncertain, we will use a lower limit on the BH mass of $2\,\text{M}_\odot$. We plan to use an upper limit of $50\text{M}_\odot$ for the total mass of the NSBH source. Above this point are so-called "IMRI" sources, which are described in a later search plan. This boundary is motivated both by the lack of observations of black holes with mass greater than $50\text{M}_\odot$ and the degrading performance of signal-based consistency tests at higher masses. We will span the full range of (aligned) black hole spins. The expected range of neutron star spins (< 0.05) is negligible and does not need to be considered [196].

The CBC group supports running three search pipelines to ensure the NSBH science goals are met for the early aLIGO/AdV era. MBTA and GSTLAL-CBC are low-latency online searches and `PyCBC` will be used for the offline search. MBTA and GSTLAL-CBC will generate CBC triggers in low-latency for distribution to electromagnetic follow-up partners. The optimized, reproducible `PyCBC` based offline pipeline will run at longer latency (e.g. week latency) to incorporate additional data quality information not available in low-latency, to use final calibrated data, to span a larger region of parameter space, and to ensure that the search is robust to possible data dropouts or data quality issues that are possible in the low-latency environment. Additionally, the offline pipeline can generate high-statistics simulations to evaluate search sensitivity. In general, multiple pipelines help to increase confidence in results. The plans for validation of these pipelines are outlined in Section 3.4.7.

The offline search takes advantage of data quality and calibration information that is produced with latencies much longer than a few minutes and will be used to make final statements about search sensitivity and detection significance in publications. It is expected that this search will be performed on Max Planck Institute computing facilities in Germany. In addition, because NSBH mergers are potentially visible electromagnetically, a low-latency NSBH search will be conducted to participate in the LSC-Virgo alert network. To effect this, two low-latency pipelines, Virgo-based `MBTA` [28] and LSC-based GSTLAL-CBC [200, 31, 32, 33], have been developed. It is expected that in the first advanced-detector era observing runs GSTLAL-CBC will run on the LIGO compute facility at Caltech and `MBTA` will run on the Virgo compute facility at Cascina. Because timeliness is essential for low-latency searches, two-site redundancy is important as it allows for computer facility maintenance to occur without risking the loss of coincident gravitational-wave and electromagnetic observations.

To determine the detection confidence of events observed by the search pipelines we will calculate a false-alarm probability (FAP) for each event. Historically, false-alarm probabilities have been measured using “time slides” — event lists from separate instruments are shifted in time with respect to each other by amounts that do not correspond to physical propagation delays, and coincidence rates measured. These time slide techniques have been demonstrated to be capable of measuring the FAP of very rare events [59], however, in the implementation of [59], can be computationally inefficient. Therefore a new implementation to calculate FAP using time-slides in a highly computationally efficient manner is being developed in tested in the `PyCBC` search. GSTLAL-CBC has developed a “slideless” technique that has been demonstrated in both engineering runs and MDCs [201, 134]. MBTA calculates the false alarm rate without using time-slides from the measured trigger rate by assuming that the single detector trigger rate is constant over the time period studied.

Detection candidates arising from any of the search pipelines will be promptly followed up by parameter estimation pipelines. These parameter estimation pipelines are described in [136]. To check for consistency, at least two methods will be used to sample the high dimensional parameter space: Markov Chain Monte Carlo [202, 203] and nested sampling [204]. Additionally, several waveform models will be used, including the precessing post-Newtonian models used in [196, 205], the best-available spinning effective one-body (EOB) [206, 207, 130], and so-called “phenomenological” spinning inspiral-merger-ringdown (IMR) waveforms [208, 127]. This redundancy will check for consistency among waveform models and sampling methods, as well as quantify any potential systematic biases due to waveform model uncertainty. Waveform development is on-going and the best available waveform models at the time of a detection will be used for parameter estimation.

Both the low-latency and the final offline data quality information will be supplied to the pipelines by the Detector Characterization group [71]. This will include information about the state of the detectors (e.g. science mode, reliability of the calibration, hardware injections) and veto segments for non-standard operation. Data quality triggers will identify probable times and parameters of glitches in the data. The CBC Data Quality subgroup will work with the Detector Characterization group to decide how to best apply this information in the search pipeline in order to maximize detection efficiency, and to identify what data quality issues most warrant instrumental investigation and intervention.

3.4.5 Validation plan

Any astrophysically significant candidate event will be validated using the multiple pipelines described above. An NSBH pipeline comparison study should be performedi, similar to that for BNS, to inform us of the variation between the output of pipelines when analysing the same data. In the event of a significant candidate being identified, we plan to follow these guidelines to validate it:

- Does the candidate have low false alarm rate in *at least one* pipeline?
- If a significant trigger is reported in more than one pipeline the responses of the pipelines should be consistent within expected differences. Differences in clustering may result in a different template being recorded as the trigger.
- If a significant trigger is reported in only one pipeline, we will act on a case-by-case basis, looking at the progress of a trigger through the pipelines with the aim of determining why differences are observed. If one pipeline is found to be able to rule out a candidate for reasons that were not taken into account by the triggering pipeline (e.g. updated data quality vetoes) then it should not be put forward as a detection candidate. If no reasons are found to rule out the candidate then it may progress through the detection process.
- Is a candidate found with consistent SNR when online pipelines are re-run offline?
- Is the candidate robust to slight changes in the analysis (PSD variation, choice of segment boundaries etc?)
- Does the candidate pass all checks in the CBC detection checklist?

In the case where no detection candidate is found, we will verify that multiple pipelines agree on estimates of sensitive volume within known differences in live time analysed and pipeline sensitivity.

3.4.6 Publication plan

In the past, we have published non-detection papers of NSBH searches together with the results from BNS and BBH searches [192, 193, 57, 194, 59]. In the advanced-detector era we will consider separate publications depending on whether or not we have a detection candidate, and the time-scale of the BNS and BBH searches. We will not hold up a BNS or BBH detection paper, or paper that places astrophysically interesting bounds on merger rates, to get the results of the NSBH search, which may require more development. The target journal for these publications is to be determined in consultation with the wider group and collaborations. The addition of electromagnetic counterparts to a NSBH detection may also change the publication plan according to the external memoranda of understanding, although an electromagnetic counterpart to the first gravitational-wave detection in the first few years of operation is unlikely [124]. Here we present some potential publication scenarios.

Publication time-lines will only be achievable if detector characterization efforts are well coupled to both the commissioning and analysis teams *before and during the run*, the effects of calibration uncertainties are understood, a sufficiently accurate and correct calibration is available, and the analysis pipelines (detection, rate measurement, and parameter estimation) are tested, compared, and reviewed prior to the start of the run. It is important for members of the NSBH search team to liaise with these various groups within the collaboration to ensure that these requirements are met.

Publication scenarios up to and including the first NSBH detection

Confident NSBH detection In the case of a clear detection, the CBC group will aim to submit a paper for publication within 3 months of identifying the candidate, while allowing time for a full LVC meeting to take place. An example of a confident detection candidate is the "big dog" blind hardware injection, which occurred near the end of the LIGO's S6 and Virgo's VSR3 [59, 60]. At the same time a short paper will be drafted describing the search and the candidate(s) and discussing possible astrophysical implications based on their physical parameters and inferences on the source population. LIGO-P1000146-v16 [60] provides a good model for what we expect from a first detection paper, including calculation of the event significance and estimation of the event parameters. Subsequent LSC-Virgo papers will provide more details about the event. Note: An electromagnetic counterpart is not required, and is not expected, to make the first confident NSBH detection. Additional constraints and considerations imposed by the memoranda of understanding with external observatories may impact publications if applicable.

Marginal NSBH detection It is possible that a detection candidate will be "marginal"; where there is strong evidence for a detection, but the null hypothesis cannot be firmly ruled out. An example of this is the "equinox event" hardware injection from LIGO's S5 [123, 209]. The publication of even a marginal NSBH detection should be prioritized *if a clear gravitational-wave detection has already been made*. In this case the schedule of the confident detection should be followed, but the marginal significance should be clearly stated in the paper's title, abstract and conclusions. If a marginal detection occurs *before* a confident first detection has been published, we will wait until the end of an observing run. If no clear detection is found before the end of the observing run, we will publish the details of any marginal NSBH detections. We would consider a joint publication with BNS and/or BBH if they also have marginal candidates but no single clear detection.

No NSBH detections If no NSBH observations occur, as is expected for the first observation runs, then the NSBH search team will use a ten-fold improvement in VT (the volume of space multiplied by the duration of its observation) to establish the minimum publication rate. If BNS and/or BBH searches are completed on a similar time scale, the CBC group may submit a paper covering the updated merger rate upper limits for the BNS, NSBH and stellar mass BBH searches within six months of reaching the ten-fold VT improvement.

Post-first-detection NSBH publication scenarios

As the sensitivity of the detectors improves, it is possible that we will make multiple NSBH detections. After the first detections, we do not foresee publishing a separate paper with each new signal unless there is a compelling reason to do so (that is., first discovery with observed precession, or unexpected masses, etc.). Instead we will publish results at the end of short observation runs ($<$ 6 months), or at predetermined intervals throughout the run. Once astronomical alerts are made public, the initial publication of events may be via Astronomical Telegrams when appropriate.

3.4.7 Technical requirements and development plan

A proposed roadmap for the development of the aLIGO and AdV detectors towards their design sensitivity is given in [124]. Updates to this roadmap should be driven by maximizing the chances of a BNS detection, as BNS signals are considered the most likely candidate for a first detection. Nevertheless, searches for NSBH signals may lead to a detection in the early observation runs and we should be ready for this possibility. As aLIGO/AdV sub-systems are commissioned, the compact binary coalescence (CBC) group will collaborate with the Detector Characterization and Commissioning groups to ensure the best data quality possible for the BNS search. This data-quality information will be directly applicable to NSBH searches, including some information from data-quality for BBH searches, at the higher range of masses.

To help prepare for aLIGO/AdV, the CBC group has two tools at its disposal: Mock Data Challenges (MDCs) and Engineering Runs (ERs). Previous searches have demonstrated that early and close coupling between the search group and the detector group is critical to our ability to complete analyses in a timely way. The purpose of ERs is to engage the CBC group with detector characterization, commissioning, and data analysis infrastructure development. ERs are using real detector data from the two LIGO instruments. In addition to detector characterization, engineering runs provide a platform for end-to-end testing of the low-latency and offline analysis infrastructure. MDCs provide a platform for large-scale testing and validation of pipelines using offline or simulated data. Performance in the MDCs is a crucial milestone to demonstrate that search and parameter estimation pipelines have reached the appropriate sensitivity and accuracy for aLIGO/AdV NSBH science, and for comparisons between search pipelines with overlapping science targets. Participation in the MDCs will be and has been a critical component of testing prior to the first observation runs.

NSBH binaries are possible progenitors of short GRBs, if the NS tidally disrupts before merger, which makes them interesting candidates for multi-messenger astronomy. There are plans to follow up GRBs detected electromagnetically with a targeted search for gravitational waves, but this is considered a separate search effort and is described in another document. It is also possible that an NSBH gravitational-wave signal could be detected first by this search, and used to trigger searches for an associated electromagnetic afterglow. Therefore, we will promptly issue alerts for interesting candidate events with as much low-latency information on sky localization and parameter estimation as we can provide using the tools developed for the low-latency search. However, we emphasize that an electromagnetic counterpart is not necessary for a confident gravitational-wave detection, and that the sky localization of any events is likely to be poor [124].

As much of the ongoing development is shared with the BBH and BNS searches there will be significant overlap, cooperation and collaboration between the teams developing and running the NSBH, BNS and BBH searches. This is being done through regular teleconferences between the search teams and through the use of a common code base. Within the CBC group, we have divided technical activities into a number of sub-groups, described below. We have identified the critical and desirable tasks and their time estimates. The main tasks are:

Waveforms: Search and parameter estimation pipelines require accurate waveforms. The currently available post-Newtonian waveform models show significant discrepancies between the various models for NSBH waveforms with asymmetric masses ($m_1/m_2 \gtrsim 6$) and large spin magnitudes ($\|\vec{S}_{\rm BH}\|/m_{\rm BH}^2 \gtrsim 0.5$) [205]. From that we infer that a similar discrepancy exists between these models and waveforms produced by nature [205]. This will reduce the detection rate of our search pipelines and will introduce a bias in attempts to extract the parameters of detected systems. Improving the waveform models should be seen as a high priority for NSBH searches. Deriving additional terms in the Post-Newtonian expansions of the waveform models, incorporating field theory methods or using information from numerical relativity are potential ways to do this. The development of precessing inspiral-merger-ringdown (IMR) waveforms valid in the NSBH region of parameter space is also a priority to detect NSBH systems where the black hole is at the high end of the expected mass distribution. Much of this effort will take place in the broader numerical and analytical relativity communities, but any such waveform improvements should be promptly implemented into analysis software. In particular, efforts are underway to implement the precessing IMR models “SEOBNRv3” [130] and “IMRPhenomP” [127] into the software. These waveform models should be implemented and reviewed before the first aLIGO observing run.

Specific Tasks Planned for Waveform Development

- Finish implementation and review of the SEOBNRv3 precessing IMR waveform models.
- Begin development of a highly-efficient SEOBNRv3 reduced-order model approximant.
- Finish implementation of the PhenomD IMR waveform model
- Develop a precessing version of PhenomD using the infrastructure used to create PhenomP.
- Determine which waveform family to use as detection templates with consideration for both accuracy and computational cost. All available spinning waveform families will likely be used for parameter estimation, provided they are not so expensive as to make parameter estimation computationally intractable.

Search Pipelines: The CBC group has three search pipelines in development: `PyCBC-ahope` (a new implementation of the traditional offline search pipeline, `ihope`, allowing for longer template waveforms and greatly increased flexibility); `MBTA` (multi-band template analysis; used for low-latency search of S6/VSR3); and `GSTLAL-CBC` However, the angular-momentum (spin) of the black hole should not be neglected in searches for NSBH systems with aLIGO and AdV. Each of these search pipelines is currently testing and demonstrating their ability to run with aligned-spin template waveforms. A careful comparison of detection rates between non-spinning and aligned-spin template banks in real data and as a function of detector sensitivity is also being conducted. There are also a number of potential methods for searching with precessing waveforms as templates [189, 210, 191, 199], though such methods have yet to demonstrate an improvement in sensitivity on initial LIGO and Virgo data [191, 199]. These methods should be applied and tested on the ERs and MDCs, but aligned-spin search development is the priority.

Search Pipeline Validation MDCs will be used to confirm the validity of the three pipelines, and to ensure that both low latency triggers and final results are generated unambiguously: either by selecting a single pipeline for each or by developing a method for combining results across pipelines. Since early science runs are only expected to have astrophysically interesting sensitivity above about 30Hz, this will be our first target for tuning, followed by tuning for detectors that approach the full low-frequency sensitivity. This will be a priority issue for the BNS group, and the NSBH group will utilize developments in that context.

A common MDC is currently being finalized to develop and test our strategies for the validation of astrophysical results when running multiple pipelines as requested by the multiple pipelines policy [67]. Results from BNS studies show that the pipelines have differences in recovered SNR and FAR estimates and this is expected to also be true for NSBH. Work is ongoing to follow up (rare) simulation candidates with apparent FAR inconsistencies. In such cases, we seek to identify the cause of the discrepancy and show that pipeline behaviour is consistent with the presence of a claimed signal. Thus validation will depend on the outcome of these followups.

Specific Tasks Planned for Search Pipeline Development

- Use MDCs containing precessing gravitational waveforms to demonstrate that an aligned-spin template search has better efficiency at fixed false alarm rate than a non-spinning template search. If any of the pipelines cannot demonstrate such an improvement before the start of an observing run they will use non-spinning templates for that search.

Data quality: It is vital that the commissioning team, Detector Characterization and CBC groups communicate with each other during and leading up to the early observation runs. The CBC data-quality group will attempt to help early commissioning work as much as possible by conducting careful data-quality studies on data from the various aLIGO and AdV sub-systems as they come online. This will especially be conducted through the engineering runs. CBC data quality efforts should focus on the needs of the BNS search, as

this is the most likely candidate for the first detections. However, the NSBH signals and search methods are similar enough that these data quality vetoes and techniques should also be applicable to the NSBH search.
Rates and significance: Accurately evaluating the significance of potential gravitational-wave signals and inferring rates of CBC mergers from gravitational-wave observations (or lack thereof) will be a high priority task for the BNS sub-group. The NSBH sub-group will directly apply methods developed in that context. Here there is significant overlap between the BNS and NSBH sub-groups.
EM follow-up: The low-latency pipelines will report high significance triggers to an alert system for EM observers. Again, this development will be a higher priority for the BNS group, so work here should focus on those systems, but also allow for rapid follow-up of NSBH system in the case such a detection is made.
Parameter Estimation: Parameter estimation will be run on all marginal and confident detection candidates arriving in the search pipeline. Parameter estimation pipelines have already demonstrated the ability to analyze candidate events in initial LIGO and Virgo data [136]. This included analyzing simulated NSBH signals with precessing post-Newtonian models. Systematic bias from waveform uncertainty will always be a concern for parameter estimation efforts. We can mitigate this somewhat by implementing improved waveform models as described above. Additionally, we will run parameter estimation pipelines with different waveform models to check for consistency and characterize the level of any potential bias. Waveform generation is the dominant cost of parameter estimation and rises strongly as either the mass or the minimum sensitive frequency decrease. Depending on the waveform model, parameter estimation of BNS signals in initial LIGO can be performed in roughly 8 hours to 2 weeks using 16 CPU cores. Computational cost is not an immediate barrier for NSBH parameter estimation with any of the currently available waveform families. If the early advanced detectors have sensitivity down to ~ 30 Hz, the computational cost of NSBH parameter estimation will be less than the cost of BNS parameter estimation for the initial detectors. If aLIGO and AdV approach their design sensitivity at $\sim 10 - 15$ Hz, computational cost will rise by a factor of ~ 20 relative to the cost for a minimum frequency of 30 Hz and may become a significant barrier to NSBH parameter estimation, but this will not be an immediate focus of the NSBH group. Of particular interest for parameter estimation are *precessing* IMR models (because we wish to include all available physical effects), ideally accurate for all mass ratios $m_1/m_2 \lesssim 15$ (as most NSBH binaries will have very asymmetric masses). Efforts are underway to review and use reduced order modeling of IMR waveforms to create computationally-efficient surrogate models [137, 138, 139, 140, 141, 142, 126] for expensive waveform families such as SEOBNRv2, thus enabling parameter estimation with these waveform families.

Specific Tasks Planned for Parameter Estimation

- Develop accurate, computationally-efficient waveform models (EOB-based and phenomenological) for use in parameter estimation. In particular we require precessing IMR waveforms valid for asymmetric masses.
- Explore the use of surrogate waveform models or alternative approaches to parameter estimation so that slower waveform models may be used in parameter estimation.

Testing general relativity: All marginal and confident detection candidates will be examined to see if any deviations from general relativity can be observed from any single signal or from a collection of signals. Bayesian methods for testing whether the waveforms of any detected signals are consistent with the expectations from general relativity (GR) have been developed [211] and are demonstrated for BNS systems [46]. Implementation for NSBH systems, including fully spinning waveforms, and addressing the many additional challenges arising for such systems, is expected to be undertaken over the next few years; a NSBH analysis is not expected to be available and reviewed by the first Advanced LIGO observing run in late 2015. However, if the feasibility of testing GR with BBH systems can be demonstrated, we expect that such a pipeline will be run later in Advanced LIGO. In order to confidently detect a deviation from GR waveforms, a large number of ‘background’ (i.e. simulated GR waveform) sample results may be needed, with a computing cost of order(1000) cores running over months. Feasibility studies will have similar computing loads.

3.4.8 Resources

Calibration. We will require calibrated data over the bandwidth of the inspiral signals (from 30 Hz to 2 kHz in O1). The most stringent calibration requirements are for parameter estimation, where the data is analysed coherently. This has been studied using simulated calibration error curves in [83, 84], which found that the most significant bias occurs for sky localisation accuracy when amplitude calibration errors have opposite signs in each detector. The study found that a maximum 1σ error of 10% in amplitude and $5\,\mathrm{deg}$ in phase calibration would be a suitable requirement for parameter estimation work.

The effect of calibration errors on detection was studied in [85, 83], with the conclusion that at the level outlined above the bias in Optimal SNR should be less than 3%, however the matched filter SNR and therefore distance estimation will depend linearly on the calibration error, affecting the estimated sensitive volume of any search. This will have an impact on the inferred rates, with a 10% error in amplitude giving around 30% error in the estimated volume and event rate.

We would like to have calibration to an accuracy of 10% in amplitude and $5\,\mathrm{deg}$ in phase prior to the start of any science run, and to have the calibration accuracy confirmed after a detection is made.

Detector Characterization & Data Quality Reliable knowledge of detector operation (for example science/calibration/analysis segment information) is a prerequisite for NSBH analyses. Additionally, noise transients in the detector data will have a significant impact on the sensitivity of the search. We will work with the detector characterization and commissioning teams to identify and eliminate sources of transient noise. The CBC group’s daily detector characterization pipeline (included in the BNS search computation cost) will be used to provide rapid feedback on detector performance affecting astrophysical search sensitivity. We will quantify the effect of these transients on our search sensitivity. One of the primary engineering run goals will be to develop the tools and techniques to effectively identify and mitigate glitches in the data. During data taking runs, we will examine the data on a daily and weekly basis for detector characterization. The CBC group was able to provide at least weekly analysis of data quality during S6 by engaging a team of around ten people involved in running the analysis, resolving any issues that arose, and regularly looking at data quality. Providing data-quality information with a latency of less than one day will require coordination with the detchar and commissioning groups. To achieve our data-quality goals, we strongly encourage engagement from the commissioning group in the CBC search, and extended visits by CBC scientists to the observatories.

Computational Requirements Further details about the computational requirements for the NSBH searches can be found in LSC Computing Plan Summary April 2015, LIGO-T1500118. The NSBH portion of the offline CBC searches comprises about 26% of the computational costs for the LSC highest priorities in O3. This corresponds to approximately 70 million service units as defined in LIGO-T1500118. An additional 18 million service units are required for the low-latency search, which includes some coverage of the NSBH parameter region.

It is possible that future observational evidence between now and 2018 may restrict the parameter space that needs to be searched, which would decrease the computational cost. As an example, if we were to search a restricted range with the black hole mass $\in (7.0, 12.0)\,\mathrm{M}_{\odot}$ and neutron star mass $\in (1.0, 3.0)\,\mathrm{M}_{\odot}$, would require approximately a third less templates and reduce the computational cost correspondingly.

Computational Infrastructure and Services The following computational infrastructure and services must be available, tested and working to successfully perform searches for NSBH mergers:

- The calibrated offline h(t) data must be available at remote sites where analysis will be performed with a latency of no more than a day.
- A service must be available to easily locate detector data files on any cluster.
- A service must be available from which to obtain the state of the detector at any time, with a latency of no more than seconds after the data is taken.
- A service must be available to obtain data quality segments and/or vetoes at any time. The final version of this information should be available on a $\sim$ week timescale.
- Services must be available to allow analysis workflow to be planned and executed on the sites where analyses will be performed.
- Compute nodes must be POSIX compatible with at least 4GB of RAM per CPU core.
- Clusters should have a pool of at least 50 "performance" nodes with at least 10GB of RAM per CPU for memory intensive post-processing jobs.
- Compute nodes must have local "scratch" space available with at least 100GB of available disk space per core.
- Clusters must have a large capacity storage area, which is accessible from all compute nodes.
- Compute nodes must have access to the lscsoft package bundle, including, but not limited to, lalsuite, gstlal, pycbc, matplotlib, numpy, scipy.
- Compute clusters must have ligo.org authenticated web servers for display of results. These should support symlinks to the storage area.
- A service should be available for archiving analysis products and to easily locate old results and data products.
- Compute clusters must be able to support jobs that run using multi-threading and using MPI.
- Compute clusters should have ability for setting up database servers, such as MySQL and Postgres for storing large number of triggers. Initial estimates predict that 10 TB per 6 calendar months of 3-ifo data would be sufficient storage space.

We also consider a possible search using precessing templates. As discussed earlier, an effective search using precessing templates does not yet exist, but is being investigated and developed. Therefore we can only provide an estimate of computational cost using one of the search techniques proposed for searching for precessing binaries [189]. The precessing search has a significantly larger computational cost due to two factors: The first factor is the increase in the size of the intrinsic template parameter space due to the effect of precession on the gravitational-wave phasing of the signal. Ref. [189] made an order-of-magnitude estimate of 76000 templates for the Initial LIGO noise curve for the region of parameter space with the black hole mass $\in (7.0, 12.0)\,\mathrm{M}_{\odot}$ and neutron star mass $\in (1.0, 3.0)\,\mathrm{M}_{\odot}$. Scaling this up to Advanced LIGO (increasing the number of templates by an order of magnitude due to the improved low-frequency sensitivity of Advanced LIGO) predicts that $760,000$ templates would be needed to perform a precessing search in this region. The number of non-precessing spin templates needed to cover the same region of parameter space is $\sim 670,000$ suggesting that precession causes a $\sim 13\%$ increase in the size of the template parameter space. Using this to scale up the full NSBH parameter space results in approximately $2,300,000$ templates for the precessing bank. Significantly compounding this is the fact that the filters described in

Ref. [189] require *five* times as much computational power as the aligned spin search, as for each template the data is projected onto ten orthogonal basis vectors, rather than two, in order to analytically maximize over further precessional effects in the waveform amplitude and phase. Consequently a search using the methods described in Ref. [189] would dominate the computational cost requirements of the CBC searches. However, we note that such a precessing NSBH search would cover the aligned spin search as a sub-set of the parameter space, so either an aligned-spin search or an precessing search would be performed, but not both.

A search with precessing templates will not be possible during the earliest aLIGO observing runs, as considerable development is needed to implement a search pipeline. However, if such a search is to be ready for later aLIGO and AdV observing runs of circa 2018, it will require the allocation of resources for development in 2015-2018, beyond those used for the spin-aligned or non-spinning search effort. A precessing pipeline would need to be run on $\sim$ months of real or simulated data to assess its performance relative to other search pipelines. It is almost certain several iterations would be needed to tune the precessing search and re-evaluate its performance. Thus, we estimate that developing a precessing search would require enough CPU cores allocated over the course of 3 years to run an experimental precessing search pipeline on a total of $\sim$ 1 year of real or simulated data.

We plan to follow-up the interesting event candidates with parameter estimation pipelines. Running a Markov Chain Monte Carlo (MCMC) parameter estimation pipeline with the precessing SpinTaylorT4 waveform model takes roughly 10 days on 16 CPU cores. We would analyze each event with several different waveform families including, but not necessarily limited to, the post-Newtonian SpinTaylorT2 and SpinTaylorF2, the phenomenological IMRPhenomP and PhenSpin models, and ideally one or more variant of spinning EOB models. Most of these waveform families have a comparable or smaller cost than SpinTaylorT4, while the EOB models are considerably more expensive. Therefore, we multiply the cost of a single pipeline by 10 to allow for the use of several waveform families. Additionally, we would use multiple sampling methods for parameter estimation, including MCMC, nested sampling and possibly a third method such as a machine learning-based technique.

Review The review of several search components will overlap between source groups: i.e., search algorithms (3 pipelines), waveforms, rates and significance estimation and parameter estimation. However, NSBH specific detection algorithms will require additional review before they are deployed for production searches. It would be prudent to start the review of these elements as they become tested during the MDCs and ERs. Review of some components has already started (in overlap with BNS) or will start soon. This will allow enough time for review of components that are developed later to be performed prior to the start of the first observation run. In addition, the following will be reviewed when early data becomes available: Data Quality, Calibration.

Search A team of approximately ten people will be required to maintain the various NSBH searches. There will be considerable overlap between the NSBH, BNS and BBH search groups, so some parts of this (e.g running the online pipelines and maintaining the electromagnetic alert algorithms) will not be unique to NSBH. We will aim to have electromagnetic alerts sent within $\mathcal{O}(1) minute$ of data being taken, and detection/non-detection statements within 2 weeks of taking data. Parameter estimation of candidate events should be available within a month of the detection of an event.

Candidate verification and follow-up To get from the observation of a clear detection candidate in the search to being ready to submit a paper will require much extra work, checks and confirmation from many groups within the collaborations. This is expected to include detailed checks on the analysis results; on data quality at the candidate time(s) and over the whole of the analysis; on detector behaviour and possible environmental disturbances; on the overall consistency of the events with GW signals; and any other investigations that the Detection Committee or the collaborations consider necessary.

3.5 Search for GRB Sources of Transient Gravitational Waves

3.5.1 Abstract

Extremely energetic bursts of gamma-rays from cosmological sources are observed by orbiting satellite detectors at a rate of about one per day. These extra-galactic events are generally referred to as GRBs. Astrophysical evidence has led to the hypothesis that GRBs herald the creation of a compact object (a black hole or neutron star) by way of two distinct pathways, corresponding to two phenomenologically recognized GRB categories: short-duration (< 2 s) bursts with generally harder spectra, and long-duration (> 2 s) bursts with generally softer spectra. Both progenitor categories, collapsars for long GRBs and mergers of compact binary systems for short GRBs, are expected to be sources of transient gravitational waves (GWs). The detection of a GW signal in coincidence with a GRB would provide tremendous insight in the astrophysics of these systems. A GW signal associated to a long GRB would give new astrophysical insight into long GRB progenitors, which in general are not expected to be efficient GW radiators, but some models do predict significant GW emission. A merger signal associated to a short GRB would confirm the compact binary merger nature of the engine and allow for measurements of the binary components masses and spins, as well as constraints on the beaming angles and the neutron star equation of state. A collection of joint short GRBs with redshift and GW measurements will enable a relatively systematics-free measurement of the Hubble parameter at low redshift, which would provide constraints on cosmological models.

Since the GWs would be observable within seconds of the onset of the gamma-ray detections, we plan to perform sensitive *triggered searches* with two distinct algorithms. Optimal searches with GW waveform templates will be sensitive to short GRBs out to ~ 400 Mpc for NS-NS mergers and ~ 1 Gpc for NS-BH systems in advanced detectors at full sensitivity. Because of the astrophysical uncertainty of the long GRB mechanisms, the GW emission cannot yet be modeled *a priori*, so a *burst* algorithm is also employed for GW searches triggered by both long and short GRBs. The space and time constraints for mergers which are associated with a GRB allow for a *coherent* version of both types of searches. We plan to run archival searches on the refined data samples available after a few weeks, as well as a more computationally costly search focused on long-duration GW signals associated with long GRBs.

3.5.2 Scientific Justification

GRBs are generally associated with systems which are also expected to be GW sources: compact binary coalescences for short GRBs, with gamma-ray duration $< 2\,\mathrm{s}$ and harder spectra, and collapsars for long GRBs, which typically last $> 2\,\mathrm{s}$ and have softer spectra. The cosmological distance of GRBs remains the principal challenge for detecting GWs associated with GRBs with advanced detectors.

Cosmological short GRBs likely result from mergers of binary systems consisting of two neutron stars (NS-NS) or a neutron star and stellar mass ($< 10\,\mathrm{M}_\odot$) black hole (NS-BH). At full sensitivity, they are potentially observable by advanced GW detectors to $\sim 400\,\mathrm{Mpc}$ for NS-NS or $\sim 1\,\mathrm{Gpc}$ for NS-BH at a rate of $\sim 1\,\mathrm{yr}^{-1}$ for each [212, 213]. Short GRBs can also result from soft gamma-ray repeaters (SGRs) of galactic or near-galactic origin (*e.g.* the likely progenitors of GRB 070201 and GRB 051103).

It is widely believed that binary mergers are the progenitors of short GRBs. The recent detection of a possible kilonova associated with a short GRB [214] has further supported this association. However, it is only the observation of a GW signal that will conclusively show that the progenitor is a binary merger. While the merger scenario is preferred and at least one NS is required, both NS-NS and NS-BH progenitors are possible; GW observations will allow us to determine which of these it is. A population of GW-GRB observations will allow us to measure the fraction of GRBs associated with each progenitor type. The degeneracy between distance and inclination angle means that it will be difficult to measure the GRB beaming angle based on measurement of a single system. However, with observations of a population of binary merger sources, with and without GRB counterparts, we can constrain the average opening angle.

Long GRBs are associated with the gravitational collapse of massive stars, also termed collapsars, but perhaps with multiple sub-pathways and astrophysical details which are far from being fully understood. They are detected more frequently than short GRBs, are typically brighter in gamma rays, and are observed to larger distances. Their observable distance in GWs is unknown, but is unlikely to be as large as short GRBs. For basic stellar collapse, the GW range is likely to be galactic or near-galactic. Collapse leading to rotational instabilities or core fragmentation could be observable to 200 Mpc. Long GRBs associated with accretion disk related mechanisms could produce long-duration GW emissions, perhaps observable to tens of Mpc. The estimated rate for long GRBs within 100 Mpc is $\sim 1\,\mathrm{yr}^{-1}$ [215, 216]. The hypothesized population of low-luminosity long GRBs [217] would also have a rate of $\sim 1\,\mathrm{yr}^{-1}$.

Long-duration GRBs are expected to be dominantly associated with collapsar models – rapidly rotating massive stars which collapse to BHs (or NSs). The long GRBs reveal themselves with a wide range of observable properties, leading to speculation that there may be sub-classes involving different mechanisms. In addition, some models predict GW emission associated with the accretion disk itself, or with a post-collapse proto-neutron star, which would give rise to long-duration (> 1 s) GW emission. The observation of X-ray "plateaus" following the GRB on timescales of tens of minutes to hours after the main burst, has suggested that GRB central engines may live longer (1000 s) than previously thought. Given the mysteries associated with long GRBs, it is clear that any significant GW detection would greatly contribute to our understanding of the underlying astrophysics. Finally, a low-luminosity population of long GRBs would benefit from lower-threshold gamma-ray triggering.

The association between GWs and GRBs also impacts fundamental tests and cosmology. For example, the speed of gravity can be probed to $\sim 10^{-16}$, and "dark matter emulator" models, which allow MOND-like models to fully eliminate the need for dark matter, can be effectively ruled out with the first observation of a coincident GRB-GW event [218].

After some number of years of running near design sensitivity, it may be possible to accumulate a sufficient number of GW detections in association with short GRBs to make meaningful statements about cosmology. A well-measured binary inspiral waveform can determine the luminosity distance D_L which is absolutely calibrated, up to GW detector calibration uncertainties, and independent of the electromagnetic distance ladder [219]. However, to probe the cosmological expansion, one also needs the redshift, thus a joint observation of gravitational waves and electromagnetic afterglow for the GRB.

The potential science outcomes of such joint observation include testing the EM cosmological distance ladder for z up to 0.1 or 0.2 and setting dark energy constraints. With enough short GRBs a few % measurement of the Hubble constant H_0 will be possible [220]. A precise measurement of H_0, along with the CMB, tightly constrains models of cosmological expansion and addresses the question of continued acceleration into the current epoch. The dark energy figure of merit [221] is used to quantify the worth of future constraints, by testing deviations from a pure cosmological constant: even after the next generation of BAO measurements from WFIRST, a 1% measurement of H_0 will improve the figure of merit by $\sim$ 40% [221].

3.5.3 Search Description

The nature of GRB progenitors enables a GW *triggered search*: a search restricted to short on-source time windows that are defined by the gamma-ray detections. The source position derived from gamma-ray observations, constrains the possible delays in the GW arrival time in the LIGO and Virgo observatories. Together, knowledge of the time and location of the event are powerful handles for the reduction of accidental background in GW searches, which translates to a larger maximum distance ("horizon distance") for a GW detection [222, 213]

In addition, the existence of a GRB trigger significantly reduces the search parameter space, for example with the source orientation, spin, or binary mass ratio, which allows us to deploy a fully coherent analysis restricted to (near-)circularly polarized signals that would is be possible in the all-sky searches for bursts, NS-NS, or NS-BH.

A number of different searches will be deployed in the search for gravitational wave signals associated with GRBs. We will run an un-modelled burst search, a templated inspiral search and a search for long duration transients. For the templated search, we will cover a similar parameter space as for the NS-NS and NS-BH searches combined, as we require at least one neutron star in the system to produce a GRB. Numerical models can predict the range of masses and spins of the black hole in an NS-BH system that may lead to a GRB; if the black hole is too massive, the NS is swallowed whole and no GRB can be formed [223]. Taking this into account allows us to reduce the number of templates searched.

Both searches — for short and long GRBs — will be promptly initiated, within about two to thirty minutes following a GRB detection notice, and results from the GW searches should be available within a few hours. The results that will be provided, in order of increasing latency, are:

- As soon as possible after the announcement of a GRB, we will provide details on which of the detectors were operational, as well as an estimate of the network sensitivity in the direction of the GRB.

- Low latency analyses will look for coincidences between events in the online burst, BNS and NSBH searches (described in those plans) and GRBs. These have the advantage of being very low latency, and of requiring a little additional computational cost over the existing all sky searches. The results from these searches will be ready in minutes after the GRB and will give a preliminary indication of any candidate event and associated significance.

- We will run two fully coherent analyses around the time of the GRBs, as was done for the S5-6 and VSR1-4 searches: the X-PIPELINE burst algorithm on all GRBs [224], long or short, and the PYGRB binary merger search on any GRBs that are identified as short, or possibly short [225, 226] These searches will be initiated as soon as possible after the GRB alert, and we will endeavour to produce first results from this analysis within hours of the GRB alert, so that these can be shared with the broader astronomical community. For both analyses, in the initial run we will calculate a minimum of 1,000 background trials to enable us to quote a 1 in 1,000 false alarm probability for any significant event. Subsequently, we will refine the search results by performing more background trials, if necessary, to better estimate the significance of any event and by performing injections to test the sensitivity of the search. These follow-up runs will be performed in the offline/archival framework. The final results will be available within a few days.

- We will run an archival search with a latency of around a month. This will be a rerun of the data with the final data quality, calibration, etc. It will be the result that is used in non-detection publications, much as has been done in the past. Additionally, at this time we will run a search for longer duration signals with the STAMP pipeline.

3.5.4 Results validation plan

Validation searches for the different GRB analyses will be done by using the all-sky triggers produced by others plans:

- cWB for the long GRBs analyzed with X-PIPELINE
- pyCBC for the short GRBs analyzed with pyGRB
- the long duration transient search (Stampas, Sphrad,...) for the long transient GRBs

The coincidence search will de bone with the Raven pipeline currently developed for the low latency follow-up of GRB events.

3.5.5 Publication Plan

As a general rule, we plan to share anything of interest in low latency with MOU partners in the GW-Electromagnetic Follow-up program. Combined GW–gamma-ray sky maps are an example of such information. We also plan to release public GCN notices or GCN circulars, similar to the one reported by IceCube on GRB130427A [227], where we acknowledge there was an event of interest and we provide basic detector status information (on/off, antenna response, etc.). The Fermi GBM team has volunteered to include such status information in their GCN notices. They would presumably read the LIGO-Virgo status from a web page or equivalent.

Detection confidence for the GRB triggered searches will be increased if the associated GW sky map overlaps the GRB error box. Otherwise, we expect detection confidence to be analogous to that of the all-sky CBC or burst analysis, as the case applies. The publication plan will depend on the detection and astrophysical scenario as follows:

- **A significant non-detection.** (like GRB 070201 [228]) We plan to publish with a goal of $\sim$2 weeks after the GRB. This would apply to (nearby) GRBs for which there is a reasonable expectation of GW detection. For a short GRB this can be well defined, *e.g.* when the distance is known and is within the current NS-BH horizon. For a long GRB, we have no expectation, except perhaps for a galactic or local neighborhood GRB. We must keep in mind that most GRBs do not have an associated redshift and we know that the gamma-ray brightness is not a good indicator for distance. For this quick turnaround, we need pre-reviewed pipelines and pre-reviewed paper templates.

- **Early detection.** If one or more of the first several detections is a GRB, we would then plan to publish one paper per event.

- **End of run upper limits.** If the average exclusion distance upper limit is at least twice that of the current published result (currently S6/VSR2-3), we will then publish a paper including all vetted GRBs from Fermi, Swift, and IPN on a time scale of 3 to 6 months.

- **Sub-prime GRB search.** A search focused on a population of low-luminosity GRBs would benefit from a gamma-ray trigger sample obtained with a lower than usual threshold. If no significant GW detection results from this search, the results can be included in the end of run paper.

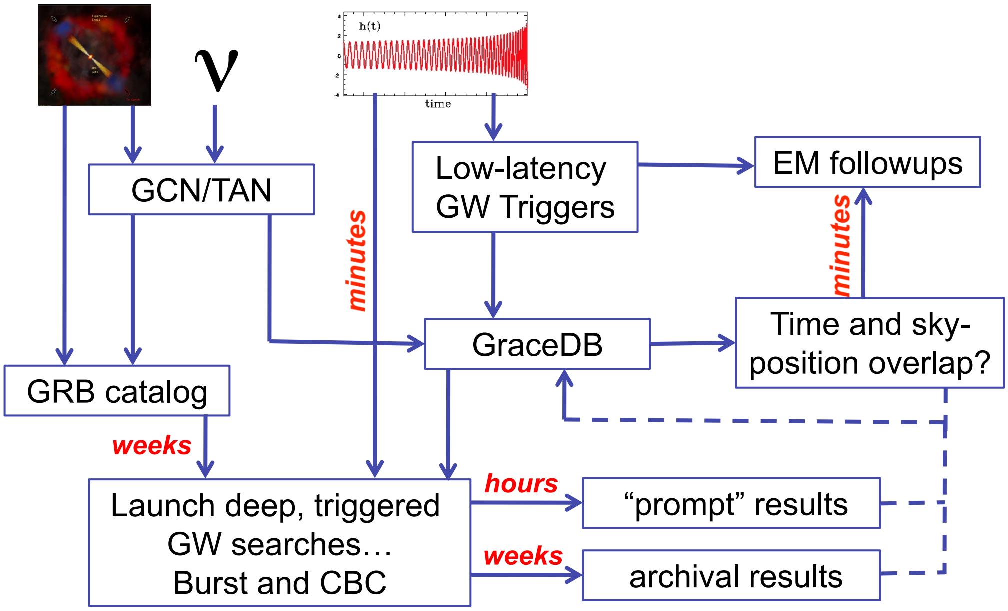

Figure 4: Flow of gamma-ray and GW data for the GRB searches. There are three main time scales, each having an associated search: Low-latency (minutes), prompt triggered (hours), archival (weeks). The ν indicates that a galactic supernova with a corresponding SNEWS alert could be handled similarly to the GRB bursts search.

3.5.6 Technical Requirements and Development Plan

The GRB searches involve 5 pipelines and 3 analysis timescales, as shown in Fig. 4.

Low latency analysis. The low-latency ($\sim$minutes) all-sky pipelines are developed by the Burst and CBC groups. Candidates from these analyses can be compared to the GRB triggers to provide additional evidence by way of time coincidence and spatial overlap, using GW sky maps and GRB error boxes.

Triggered GW pipelines. ($\sim$hour, or *prompt*, and $\sim$week, or *archival* timelines)

- **Bursts.** The X-PIPELINE [?] is a coherent search which has been used for almost all of the GRB searches in the initial detectors era. The code is stable, known, reviewed, and can be used *as is*. A steady program of updates was applied to improve it, including efforts to speed up the code. A recent version using multivariate analysis, for higher efficiency and lower computational cost, is in development but has not been reviewed yet. Its performance on the next engineering run will determine which version will be used in the first science run and whether a review of the modifications is needed.

- **CBC.** PyGRB is a fully-coherent search employing matched filtering to target NS-NS and NS-BH waveforms in association with GRBs [225, 226]. This pipeline was used in S6/VSR23 GCN searches and in S56/VSR123 IPN searches. The pipeline searches for GW signatures in a 6 s window around the time of the GRB and uses O(1000 s) of data around the central time to estimate the significance of events seen in the 6 s window. A number of enhancements were added to the code since S6. For example, the ability to do time slides has been added, the NS-BH template bank has been reduced by excluding target sources in which the NS cannot be tidally disrupted, and the low-frequency cutoff has been reduced to 15 Hz. In addition, the pipeline generator has been rewritten within the PyCBC framework, although the underlying coherent search code is largely unchanged. The code is now able to run on GRBs where only short stretches of data are available, and to use longer stretches of data when these are available.

- **Promp triggered search infrastructure.** The coherent Bursts (X-PIPELINE) and CBC (PYGRB) GRB pipelines will be run in prompt triggered mode, to give relatively quick feedback from our most sensitive searches, on ∼hour timescale. In ER5, both pipelines were operating in immediate triggered mode and sending triggers to GRACEDB.
- **Long Duration GW bursts.** The STAMP pipeline targets GW signals lasting 10–1000 s. It has been applied to initial LIGO data to constrain extreme models of stellar collapse [229]. The pipeline is reviewed and can be applied "as is" to advanced detector data. More recent work has focused on the development of the seedless clustering algorithm **stochtrack**, which can improve the reach of STAMP by as much as a factor of 2 [230]. Stochtrack is not at present reviewed, but it is a relatively small subroutine called by the otherwise reviewed STAMP pipeline.
- **GRB plateaus.** An analysis technique is in development to target secular bar-mode signals of duration $\sim 1000\,\mathrm{s}$, triggered by GRBs that show evidence for longer-lived energy injection (*plateaus*). The plan is to develop a search with a sensitivity that is intermediate between matched filtering and un-modeled searches. This analysis could complement the burst search with tools that extract astrophysics information from the results of an un-modeled search.

3.5.7 Resources

Number of events

The computing needs of prompt and archival triggered searches are driven by the number of GRBs in the samples. We will assume that the GRB rate for the early ADE will be the same as that of the S6/VSR2-3 runs, in which the rate was dominated by the Fermi-GBM, with significant contributions by the Swift-BAT and the other IPN detectors. As the IPN detections are not conveyed promptly via the GCN, they are only considered for the archival searches. Table 3 summarizes these rates.

Table 3: S6/VSR2-3 detected GRB rates

Fermi + Swift (GCN) GRBs	350/year
short GRB fraction of above	17%
IPN GRBs not in GCN sample	85/year
short GRB fraction of above	10%

We will ask to retrieve the low latency alerts from GCN through a Web Marshall (GRACEDB) getting triggers to start analyses with a medium latency ($\sim$ 30 minutes). Past analyses were run on 50% of the received GRB triggers, asking for at least 2 IFO and sufficient data around the trigger time. On a longer timescale we need to insure getting all triggers from offline catalogs with the vetting procedure described below.

Data needs

The prompt triggered search aims to launch an analysis within $\sim$ 10 minutes of receiving the GCN trigger notification. Thus, online frames with $h(t)$ calibrated data, gathered from shared memory, will need to be used with a circular buffer of at least 24 hours. Additional needs include:

- The archival searches using X-PIPELINE, PYGRB or STAMP will use standard offline frames.
- The infrastructure will need to retrieve the data on the cluster, i.e. a metadata service that maps the scientific description of detector data to frame file locations (diskcacheAPI, ligo_data_find).
- The low-latency search only requires an analysis of the data products of the low-latency, all sky searches; it does not require frame data.
- All the pipelines use the typical job submission system and therefore will then need the workflow planning and execution environment (Pegasus, DAGMan, Condor).
- The results of the analyses will be gathered and shared with the LIGO.ORG authenticated web severs (auth project, LDG).
- Low and medium latency analyses plan to store results in the Web Marshall system (GRACEDB).

Detector Characterization:

We will need to retrieve detector status and data quality information before the search start in order to determine which periods are to be excluded around the GRB time. This will require: infrastructure to generate detector state and data quality information (ODC, DMT); metadata service that provides information about when the detector data is analyzable (DQSEGDB, ligolw_segment_query_dqsegdb, segments_from_cats) within 15 min; method to locate veto trigger files from detector characterization.

Calibration:

If subsequent modifications to the offline calibration result in a significant change of results, as defined by the criteria for the corresponding all-sky analyses, a full rerun will be performed on all of the GRBs.

Review:

- For the X-PIPELINE, the expected review will depend on the configuration selected as the multivariate configuration, which has not been reviewed yet. Otherwise, most of the modifications done recently have been reviewed for different analysis projects.
- For PYGRB, review will be needed on the different upgrades foreseen, see *Technical Requirements and Development Plan*.
- The core of the STAMP pipeline has already been reviewed for previous searches. Some parts will also be reviewed for on-going projects on the burst and stochastic sides. We anticipate short review times for this project.
- As we plan to have a short time scale to publish a significant detection, we plan to have paper templates reviewed before the start of the first run of the Advanced Detectors.

Documentation:
This will be one of the main projects, as we would like to have any member of the collaboration able to act as GRB *advocate* of a GW analysis, or to check the online results. We then need to insure that all the proper documentation will be available to start the analysis, monitor the processing and interpret the results.

3.6 Search for Intermediate Mass Black Hole Binary Coalescences

3.6.1 Abstract

Intermediate Mass Black Holes (IMBHs) are conjectured to occupy the mass space between stellar-mass and massive black holes, roughly 100 to $10^5\,\mathrm{M}_\odot$. The coalescence and merger of IMBH binaries with masses of a few hundred solar masses is one of the promising sources of gravitational waves for Advanced LIGO and Advanced Virgo. A single detection of an IMBH binary merger would provide the first unambiguous proof of IMBH existence. Searches for these sources have already been conducted on data collected by the initial GW detectors during the S5/S6 and VSR1-VSR3 runs; no IMBH binaries were discovered. The established upper limits on the IMBH binary merger rates, $\sim 10^{-7}$ Mpc^{-3} yr^{-1} for binaries with $\sim\ 100\,\mathrm{M}_\odot$ companions, are a few orders of magnitude higher than the expected optimistic rates, $\sim 10^{-9} - 10^{-10}$ Mpc^{-3} yr^{-1}. Preliminary studies show that detection of a few IMBH binary mergers is plausible at advanced detector sensitivities.

3.6.2 Scientific Justification

Stellar-mass black holes, originating from core collapse of massive stars, have been observed in the mass range up to $\sim$ 30–50 $\mathrm{M}_\odot$. Massive black holes, exceeding $10^5\,\mathrm{M}_\odot$, appear to be generic in galactic centers. Intermediate-mass black holes (IMBHs) occupy the mass range between these two. IMBHs with a mass of a few hundred solar masses may generically exist in globular clusters [231, 232]. These IMBHs may form binaries, either when two or more IMBHs are formed in the same cluster [233], or as a result of a merger of two clusters each of which contains an IMBH in the suitable mass range [234]. A large number of IMBH mergers may be a generic feature of some mechanisms of structure formation, although these are likely to occur at high redshifts [235]. IMBH binaries could also form as a result of evolution of isolated binaries with very high initial stellar masses [236].

No IMBHs in the mass range of interest ($\lesssim 1000\,\mathrm{M}_\odot$ for advanced detectors) have been detected so far. Thus, a single detection will be revolutionary, as it will prove unambiguously that black holes exist in the mass range between stellar-mass and massive black holes. IMBH binary detections will also serve as probes of globular cluster dynamics, and, potentially, as probes of structure formation and growth of massive black holes. IMBH binary measurements could also inform our understanding of the formation and evolution of the most massive stars. On the other hand, the lack of confident IMBH observations to date and the complexity of modeling the evolution of very massive stars mean that little is known about these objects. It is impossible to quote lower limits on the IMBH binary merger rate, which may, in fact, be zero.

If IMBHs in this mass range do merge in binaries, little is known about their mass distribution. However, we do expect that merger and ringdown will make a very significant contribution to the signal-to-noise ratio of observed gravitational waves from such systems, because most or all of the inspiral for massive systems will occur at frequencies below the detector band. Therefore, waveforms that include merger and ringdown phases, such as (S)EOBNR [206], will be necessary to accurately model IMBH binaries. Similarly, little is known about the spin distribution of IMBHs; they may have high spins, and because of the dynamical interactions likely involved in IMBH binary formation, the spins are likely to be misaligned, leading to precession.

We can very crudely estimate optimistic IMBH binary merger rate predictions as follows. For IMBH binary mergers in globular clusters, it is very unlikely to have more than $O(1)$ merger per history of globular cluster. The space density of globular clusters is approximately 3 Mpc^{-3}, and a typical cluster is about 10 Gyr old, leading to an upper limit on the IMBH binary merger rate of 3×10^{-10} Mpc^{-3} yr^{-1}. IMBH binary formation from very massive isolated stellar binaries in galactic fields could yield rates a few times higher than this, but there are many uncertainties. This optimistic rate is a factor of several hundred lower than the upper limits obtained in previous LIGO-Virgo searches [237].

A search for IMBH binaries would thus have a chance of yielding a detection once $\sim 3 \times 10^9$ Mpc$^3 \cdot$ yr (comoving volume $\cdot$ surveyed time). Assuming the $\sim 3-$month O1 run operates at the high-sensitivity version of the early aLIGO noise spectrum, it will reach approximately this sensitivity, providing a chance of making the first IMBH detection. IMBH binaries with a "redshifted" mass of $M(1+z) \sim 260\,\mathrm{M}_\odot$ could be detected to a luminosity horizon distance of ~ 4.8 Gpc. Larger detection volumes are possible for spin-aligned systems.

Previous searches covering the IMBH binary parameter space included the coherent-WaveBurst search [238, 237] and the CBC ringdown search [239]. This search will have some overlap in mass space with the CBC search for binary black hole systems. The Burst all-sky search will also be sensitive to some IMBH binary mergers.

3.6.3 Search Description

IMBH binaries (IMBHBs) can be effectively searched for with both template (CBC) and unmodeled (Burst) techniques. We propose to carry out both types of searches over all data with two or more detectors operating in coincidence, starting with the O1 run, and combine the results of several pipelines into joint publications.

Due to the large binary mass ($\mathrm{M_T} > 50\,\mathrm{M}_\odot$), the IMBHBs merge at low frequencies ($f_{merger} \sim 1/\mathrm{M_T}$ < 200 Hz) and the IMBHB waveforms are dominated by the late inspiral, merger and ringdown phases. The IMBHB signals detectable by second generation GW detectors are relatively short, lasting less than a few seconds in band, which allows the use of un-modeled or weakly-modeled excess power (Burst) searches. The Burst searches do not require templates and provide a robust IMBHB detection, which is insensitive to the waveform uncertainties. Matched-filtering searches should be optimal under the assumption that the source waveforms are known perfectly. The requirements to both searches are discussed in §3.6.6.

The coherent WaveBurst (cWB) burst search is well established [3]. Two IMBHB burst searches have been performed with data from initial LIGO and Virgo [238, 237]. For the IMBHB search with Advanced Detector data the upgraded cWB pipeline (called CWB2G) will be used. The efficiency of the CWB2G pipeline to the IMBHB sources is expected to be comparable to the original cWB pipeline, with significant improvements in background rejection and computational performance. The CWB2G upgrade is nearly complete, the first stable version is available.

The CBC search has previously covered the total mass range above $50\,\mathrm{M}_\odot$ with ringdown-only template searches [239], while binaries with a total mass up to $100\,\mathrm{M}_\odot$ were searches for with full inspiral-merger-ringdown searches [120, 240]. Also, during Engineering Run 3, a successful search over the IMBHB mass parameter space was carried out with the inspiral-merger-ringdown templates via the streaming matched-filtering gstlal pipeline [241]. The gstlal pipeline is particularly relevant for advanced detector BNS searches, as it has been developed with long-duration templates and low latencies in mind. These features are not critical for the IMBHB search, where templates are relatively short and the low-latency detection is not an issue; however, other features of gstlal, such as accurate estimation of false alarm probability without the need for large numbers of time slides, are still relevant. Also, the ability of gstlal to search for binary black holes with aligned spins has been successfully demonstrated [242]. However, further work will be necessary (see §3.6.6) to tune the gstlal pipeline for IMBHB sources.

The results of these two searches will be combined by using the False Alarm Density (FAD) statistic [243], which has been already used in the burst IMBHB search [237]. This search will cover IMBH binaries with a total mass $\geq 50\,\mathrm{M}_\odot$ and a mass ratio $0.1 \leq q \leq 1$, where q is the ratio of the smaller component mass to the larger one.

3.6.4 Result validation plan

We will have two searches, the CBC gstlal search and the Burst cWB search, running on the same data. These are completely independent searches, and as such will provide the robust means for an independent validation of search results. While we do not require that both search pipelines find the candidate to be significant, the preferred signal interpretation (i.e., inferred signal parameters) and our studies of the relative sensitivity of the two searches will make it possible to determine whether the results are consistent with an astrophysical interpretation. Similarly, parameter estimation follow-up with the LALInference pipeline [74], a parameter estimation pipeline written with in the LAL library (see 3.2), will provide an additional check on the candidate events.

3.6.5 Publication Plan

Any IMBH binary detection will be very exciting, since they are likely to be the first definitive proof of the existence of intermediate-mass black holes. Upper limits will, in general, be less informative than for other binary types, because of the uncertainties in formation scenarios. Therefore, the focus of this search is on detections. Rather than committing to a paper for every science run, we will publish a paper either when there is a single confident detection, an accumulation of one or more marginal detections before a break in detector operations, or a significant improvement over existing upper limits, as specified below.

Confident detection. In the case of a clear detection, such as the S6 blind injection with false alarm probability $\sim 7 \times 10^{-5}$, we aim to submit observational papers within 6 months of the first detection, and within 3 months for subsequent detections. Based on past experience, this timeline is contingent on the technical and resource requirements outlined in §3.6.6 and §3.6.7. The S6 blind injection manuscript, LIGO-P1000146-v16 [60], is a good model for what we expect from a first detection paper, including calculation of the event significance and estimation of the event parameters. Where necessary, subsequent collaboration papers may provide more details about the events, such as tests of general relativity carried out on these events.

Marginal detections. If one or more significant but not outstanding detection candidates, based on their false alarm probability, are found prior to a break in data collection, we will submit a paper describing the analysis and providing the details of any marginal candidates. This will include follow-up studies on these candidates. The goal is to submit this paper within 6 months of the end of data collection.

No detections. If there is a break in detector operations and the accumulated data allows us to significantly improve existing upper limits, either by lowering upper limits in previously covered regions of the parameter space by an order of magnitude or more, or by covering a previously unexplored region of parameter space, we will submit a paper with updated rate upper limits within 6 months of the end of data collection. This will be reduced to 3 months after the first joint upper limits paper.

In all cases, we will submit a single paper combining the results of all searches covering the IMBH parameter space and described in Section 3.6.3. Because of limited prior experience with such joint publications, we are budgeting for somewhat longer times for the first paper, to be reduced for future publications. Upper limit papers may include an overlap in the covered parameter space with other searches (e.g. the CBC BBH search); this would provide a confirmation of the sensitivity of different techniques.

3.6.6 Technical Requirements and Development Plan

Pipelines. Both the CWB2G and gstlal pipelines should satisfy general IMBHB analysis requirements: detect IMBHB sources with high efficiency comparable to an optimal matched filter and establish significance of detected events at the FAR level of 1 event per several thousands years. Both pipelines should be ready for the first aLIGO data taking runs as outlined in the Milestones, Table **??**. The gstlal algorithm requires validation that its template bank is covering the IMBHB mass parameter space and that there is no significant loss of efficiency due to the BH spins and spin precession. The CWB2G algorithm is much less dependent on the model of the IMBHB sources and provide a robust detection in the full IMBHB parameter space accessible by aLIGO. However, for both searches the waveform uncertainties may bias the estimation of the search sensitivity (range), source parameters, upper limits and measured astrophysical rates. The bias and the corresponding systematic errors due to the model uncertainties have to be estimated with the different waveform families. These studies should be performed as part of the mock data challenge (MDC), which will provide a platform for the end-to-end testing of the CWB2G and gstlal analysis infrastructure. The MDCs will also validate the application of the FAD statistic to combining the results of the two searches. Low-latency analyses and EM-follow-up are not of a high priority for the IMBHB search, since EM counterparts are not expected to accompany IMBH binary mergers. Therefore, the main analysis mode is offline. However, both pipelines are designed for the low-latency analysis and can be run online.

Waveforms. In general, the burst search does not require very accurate waveforms: the existing family of EOBNR waveforms is expected to be adequate for the burst IMBHB search. The CWB2G pipeline will allow for a robust (model-independent) estimation of the source parameters, such as sky location, polarization and total mass (by reconstructing the ringdown phase of the IMBHB merger). However, the CBC gstlal search and modeled parameter estimation require accurate waveforms. To estimate the CWB2G and gstlal detection efficiency and systematic errors due to waveform uncertainties, several waveform families will be used in the analysis, including SEOBNRv2, IMRPhenomP, and ROM surrogates.

Mock Data Challenges. To prepare for aLIGO/AdV runs, the Burst and CBC groups are using Mock Data Challenges (MDC), MDCs will be a critical component of testing prior to the first science runs. IMBHB injections into S5/S6/VSR1-3 noise recolored to aLIGO/AdV noise spectra are used to test pipeline performance and data analysis infrastructure. The pipeline performance is evaluated based on the False Alarm rate Density (FAD) figure of merit [243]. MDC studies serve to validate the performance of search pipelines on IMBHB injections and the FAD method for combining pipeline outputs.

Data quality. Data Quality is a critical-path issue for the IMBHB search. The IMBHB signal is expected at low frequencies ($\lesssim$ 100 Hz), which are strongly affected by non-stationary instrumental and environmental noises. The IMBHB signal may have just a few cycles in-band, which makes it difficult to distinguish from noise glitches. A significant fraction of glitches can be rejected with the coherent network analysis (performed by CWB2G) or with template matching and signal-based vetoes (CBC analyses), but background rejection provided by the pipelines usually is not sufficient. For the initial detector data, the false alarm rate of the burst pipeline was dominated by the spike glitches — sine-Gaussian-like signals at frequencies below 200Hz with FAR of $\sim$ 1 event per 30 years. The source of these glitches is not understood and it is likely that similar glitches will be present in the aLIGO data as well, affecting confident detection of GW signals. Therefore, interaction with the DetChar and commissioning groups and development of advanced detector characterization methods is critical for the IMBHB searches. As aLIGO/AdV sub-systems are commissioned, the Burst and CBC group will collaborate with the Detector Characterization and Commissioning groups to ensure the best data quality possible for the IMBH search.

3.6.7 Resources

Development and Simulation studies 3 FTEs are required for the development of the IMBHB searches and for the MDC studies. Of these, 2 FTEs are needed for simulations. This includes generating MDC data sets, running CWB2G and gstlal jobs, and analyzing the resulting triggers . This also includes post-processing the CWB2G and gstlal triggers and finalizing the statistical procedure for the estimation of trigger significance. The development needs are 0.1 FTE for CWB2G (studies of IMBHB-specific constraints) and 0.9 FTE for gstlal (finalizing the IMBH template bank and search statistic).

Analysis of aLIGO data 3 FTEs will be required to maintain and run the searches.

Detector Characterization & Data Quality As discussed in 3.6.6, data quality, particularly at low frequencies, is critical for the IMBHB search. We will work with the detector characterization and commissioning teams to identify and eliminate sources of transient noise at low frequencies. We will quantify the effect of these transients on our search sensitivity. One of the primary engineering run goals will be to develop the tools and techniques to effectively identify and mitigate glitches in the data. During data taking runs, we will examine the data on a daily and weekly basis for detector characterization. The IMBHB analysis will require 2 FTEs to carry out data quality studies.

Calibration requirements We will require calibrated data over the sensitive band of the detectors (from the low-frequency wall to 1kHz). Both Burst and CBC IMBHB search groups plan work together with the CBC waveform sub-group on simulating calibration errors to test calibration requirements.

Review We will need review teams for the following search components: IMBH-specific details of the search pipelines, post-production code for combining their results, waveforms (where not reviewed elsewhere) and parameter estimation (aspects specific to IMBHB parameter estimation – LALInference code has been reviewed and upgrades are continuously reviewed). Most of the review effort is not IMBHB specific; we do not anticipate that IMBHB-specific aspects of the review to require more than 0.5 FTE years. We expect that Data Quality and Calibration will be reviewed as part of the aLIGO baseline CBC BNS and Burst All-sky analyses.

3.7 Search for Eccentric Binary Black Holes

3.7.1 Abstract

Compact binary coalescences (CBC), containing some combination of neutron stars (NS) and black holes (BH), are the most promising sources for first detection with advanced gravitational wave (GW) detectors. Compact binaries formed from stellar binary progenitors, which have historically been considered the most probable sources, are expected to circularize due to gravitational-wave emission prior to reaching the sensitive band of advanced detectors. However, other types of dynamically formed CBC sources covering a large range of component masses, spins and eccentricities are also possible. For example, dynamically formed compact binaries may retain significant residual eccentricity when they enter the sensitive band of Advanced LIGO and Advanced Virgo. The inspiral and merger-ringdown (IMR) of these eccentric binary black holes (eBBH) may therefore be a promising candidate for gravitational-wave detection. A detection of these sources would provide information regarding the viability of several proposed dynamical formation mechanisms, and a measurement of the eccentricity would help discriminate among those scenarios. However, standard CBC searches using quasi-circular IMR waveforms from stellar-mass binaries will not detect these systems for eccentricities $e >\sim 0.05$ [244], so a dedicated search for these potential sources is required. The expected event rate for eccentric binary sources varies wildly, from a negligibly small event rate up to rates that exceed the expected rates for quasi-circular binaries. This uncertainty illustrates the underlying uncertainty in the astrophysics governing the various proposed formation mechanisms, and serves to justify a dedicated eBBH search exploring the CBC parameter space not covered by existing CBC searches.

3.7.2 Scientific Justification

The focus of previous gravitational-wave searches for compact-object (CO) binaries has centered on quasicircular systems, since gravitational waves are known to circularize binaries, and there is ample time to circularize if the system was formed from a stellar binary progenitor. However, more recent theoretical work has suggested that galactic nuclei and globular clusters may be promising settings for the formation of dynamical capture binaries. Since these systems can form with large eccentricities and very small initial separations, there is good reason to expect that a significant amount of eccentricity will remain when the binaries evolve into the Advanced LIGO (aLIGO) band. In galactic centers, mass segregation around the central massive BH can lead to large densities of stellar mass BHs. The Fokker-Planck model used in [245] suggests that our galactic nucleus should have ~ 2000 BHs and ~ 400 NSs in the central 0.1 pc. In [246, 247], the event rate for the formation of BH-BH binaries from GW capture in this setting was estimated to be between 0.01 and 1.0 yr^{-1} Gpc^{-3}, with corresponding Advanced LIGO detection rates of $\approx 1 - 10^2$ yr^{-1}. The formation of BH-NS binaries is estimated to be $\sim 1\%$ of this rate [246].

Dynamical capture binaries may also form in globular clusters (GCs) that undergo core collapse [248, 249]. In [250], it was estimated that binary formation through tidal capture would result in a NS-NS tidal capture rate that would peak at ~ 50 yr^{-1} Gpc^{-3} at $z = 0.7$, falling to ~ 30 yr^{-1} Gpc^{-3} by $z = 0$. They also provide a scaling to BH-NS and BH-BH mergers which gives rates that peak at ~ 70 yr^{-1}Gpc^{-3} and ~ 20 yr^{-1}Gpc^{-3} for BH-NS and BH-BH mergers, respectively.

There is also the possibility that eccentric mergers could result from hierarchical triples through the Kozai mechanism. This has been suggested to occur in BH-BH mergers in GCs [251, 252, 253] and CO mergers around supermassive BHs in galactic nuclei [254], as well as in coevolved or dynamically formed BH-NS or NS-NS binaries [255]. Efforts to understand this mechanism in the general-relativistic regime are ongoing (see e.g. [256]), and the event rates of these systems are not well known (though see [257]).

To estimate the fraction of dynamical capture binaries that retain high eccentricity, we first note that the relationship between impact parameter b and pericenter distance r_p is $r_p \approx b^2 v^2 / 2M$. In other words, the cross section $\sigma \propto b^2$ scales *linearly* with r_p, rather than quadratically as one might expect. If the initial periapse is $r_{p,i}$, and we consider the repeated burst phase to end at a periapse of $r_{p,f}$ with eccentricity e_f, then from [258], $r_{p,i} \approx 0.57 r_{p,f}(1+e_f)e_f^{-12/19}[1+O(e_f^2)]$. For example, if a binary with $e_f > 0.1$ by $r_{p,f} = 10M$ can be considered to have a significantly eccentric inspiral phase, then this corresponds to all systems with $r_{p,i} < 27M$. In galactic nuclei, this is between 60% and 80% for mass ratio $q = 1 - 0.1$, so the majority of systems from the aforementioned rate estimates will have significant eccentricity with a repeated burst phase occurring in band. This fraction is significantly lower in globular clusters due to the smaller velocity dispersion, such that more systems will form with large periapses.

For all of these scenarios, the event rates for Advanced LIGO are very uncertain, and range from effectively zero to exceeding the predicted event rate for quasicircular binaries. A null result will significantly constrain the efficiency of the aforementioned mechanisms, and the fact that a very large event rate remains a viable possibility necessitates a concerted effort to search for these specific signals.

3.7.3 Search Description

The purpose of the eBBH search is to explore the CBC parameter space not covered by the CBC searches planned for aLIGO. One possible approach would be to extend the CBC template banks to cover the corresponding part of the eBBH parameter space. Such template bank should cover not only eccentric waveforms, but spinning waveforms as well, because many anticipated sources are binary black holes where the spin effects can be significant. This is a very challenging task, which is unlikely to be realized in the immediate future. Therefore, the eBBH search will be conducted with the cWB2G burst pipeline, tuned for eBBH sources. This tuning is required for improving the cWB2G performance for low mass eBBH sources, which could be formed dynamically in galactic nuclei. Typical total masses for such binaries are expected to be in the range of 5 – 100 $M_{\odot}$. GW signals from more massive BH-BH systems are dominated by the merger and the ring-down phases where the effects of eccentricity are less important and the detection of such sources is covered by the IMBBH burst searches. The cWB2G IMBBH search should efficiently detect high mass systems even if they retain residual eccentricity.

After the cWB search pipeline identifies and reconstructs candidate signals, we will also apply parameter estimation algorithms to characterize the source parameters. Such source characterization is critical in the eBBH search to understand the eBBH formation mechanism. For example, if the event followup analysis reveals that there are repeating bursts before the BH-BH merger and/or multi-chirp structure (high-order modes), this would be a signature of a highly eccentric binary formed dynamically in a galactic nucleus. The cWB2G pipeline already has a chirp reconstruction algorithm implemented, which needs to be extended for reconstruction of multi-chirp signatures.

Apart from the minor algorithmic details and different tuning strategy, the eBBH analysis uses the same search algorithm as the All-sky and IMBBH searches, which are extensively validated with the other search pipeline: Omicron+LIB and GSTLAL respectively. Also, for validation of the eBBH candidates and better parameter estimation the BayesWave pipeline will be used. BayesWave coherently models transient events in LIGO data as a linear combination of wavelet basis functions [259]. The number of basis functions needed in the analysis is a parameter of the model and is determined through Bayesian model selection. BayesWave calculates relative probabilities that excess power in the data is due to a gravitational wave or an instrument artifact, In addition, for the signal model, BayesWave provides parameter estimates. The most basic implementation of BayesWave uses a uniform prior in the time-frequency plane for the location of basis functions and has shown promise in standard burst analyses in S6/VSR2 and ER5 data. The sensitivity of BayesWave can be improved by adopting more informative priors in the time-frequency location of wavelets, which are motivated by a particular signal morphology, such as eBBH signals. Numerical evolutions of eBBH waveform models provide predictions for the time-frequency locations of GW bursts emitted at pericenter passages. While the simulations are not accurate enough to construct templates for matched filtering, they are ideally suited as priors for signal recovery in BayesWave [260]. Ultimately, we will use the $f(t)$ predicted by these simulations to guide where in time-frequency new wavelets can be added to the signal model resulting in a “happy medium” between an eyes-wide-open burst search and a strict template-based matched filtering analysis, but initially, we will use leading order expressions for $f(t)$ [258]. The $f(t)$ track itself is a function of the eBBH system parameters. By simultaneously varying the track parameters along with the basis functions in the model, we will simultaneously produce estimates for the physical parameters of the system such as mass, spin, and eccentricity. A proof of concept study using $f(t)$ priors from post-Newtonian waveforms to perform dedicated analysis of CBC signals and provide estimates of the binary mass and spin parameters using BayesWave is under way [261] and is naturally extensible to eBBH searches.

3.7.4 Publication Plan

The eBBH search covers the same mass parameter space as the high-mass CBC BH-BH search. However, it will be sensitive to a wider class of BH-BH sources with arbitrary spins and eccentricities. The eBBH publication plan is described below and depends on the possible outcomes of both searches.

Confident detection. In the case of a clear detection with a false alarm probability of 10^{-4} or lower, we aim to submit a paper within 6 months of the first detection, and within 3 months for subsequent detections. It will be an observational paper, if the event is not observed (or observed with low significance) by the CBC BH-BH search. In the event that the same circular BH-BH events are observed by both the eBBH and CBC BH-BH searches, we aim to write a followup paper to the CBC detection. Based on past experience, this timeline is contingent on the technical and resource requirements outlined in §3.7.5 and §3.7.7. Where it is necessary, subsequent collaboration papers may provide more details about the events and their refined astrophysical interpretation.

Marginal detections. If one or more significant, but not outstanding, detection candidates (false alarm probability $\sim 10^{-3}$) are found prior to a break in data collection, we will submit a paper describing the analysis and providing the follow-up studies of any marginal candidates. The goal is to submit this paper within 6 months of the end of data collection.

No detections. If there is a break in detector operations and the accumulated data allows us to significantly improve the existing upper limits, we will submit a paper with the updated rate upper limit within 6 months of the end of data collection.

3.7.5 Technical Requirements and Development Plan

There are two main technical requirements for a successful execution of the eBBH search: a) confident detection of BH-BH sources with a wide range of source eccentricities and spins and b) reconstruction of BH-BH sources and unambiguous identification of eBBH signatures.

The first requirement is satisfied by using the cWB2G pipeline, which is proven to be sensitive to eBBH sources. However, additional optimization of the cWB2G eBBH analysis will be necessary to improve the detection of eBBH sources. To capture long eBBH waveforms, the cWB2G pipeline extracts principal components from the multiresolution time-frequency data. Currently, the extraction of principle components is sub-optimal and needs improvement. Also, the reconstruction of the binary chirp mass provides a powerful selection cut for background rejection. The current reconstruction procedure needs to be updated to account for the higher chirp modes (multi-chirp) in the eBBH signal.

The second requirement is important to understand the astrophysical origin of detected events. Both the cWB2G and BayesWave need to be updated to be able to identify the multi-chirp structure of detected waveforms, provide measures of eccentricity and spins, and identify the trailing sequence of periodic bursts typical for dynamically formed binaries with large in-band eccentricities. In the case of BayesWave, a prior on $f(t)$ that uses leading-order expressions for eccentric binaries will be implemented and tested, to estimate the size of parameter biases that result from using a lower-order prior to model injected signals that contain higher order contributions. This study will be completed in advance of O1. Beyond O1, development will be required to implement our knowledge of higher order contributions to $f(t)$ in the prior calculation performed in BayesWave.

Implementing a search that can address all of the aforementioned astrophysical questions will require improvements to the existing model for binary black-hole systems on orbits with arbitrary eccentricity. The model that has recently been implemented within lalsimulation [262] is a frequency domain model that is capable of generating faithful frequency-domain inspiral waveforms for BNS sources with $e < 0.4$ at 10 Hz. The range of eccentricity over which this model can provide faithful templates increases with the total mass of the binary, so that, for instance, a $(6+6)\mathrm{M}_{\odot}$ BH-BH binary can be modeled faithfully for $e < 0.6$ at 10 Hz. This version of the code will be stress tested through simulations across the relevant parameter space and comparison with existing models in the appropriate limits, and then it will need to undergo code review. The code will then be used for generation of eccentric injections to assess the sensitivity and performance of cWB and BayesWave (as well as any other existing search pipelines) to the eccentric regime.

For BH-BH sources with larger masses, the late inspiral and merger-ringdown portions of the waveform contribute nonnegligibly to the amount of available match-filtered SNR within the band. For this reason, a key area of model development will be the development and implementation of a frequency-domain merger-ringdown waveform. The merger-ringdown model will be a frequency-domain version of a separate, time-domain model for sources with large eccentricities that has already been used to test and tune the cWB2G pipeline [260]. By improving the accuracy of the inspiral model at ever larger eccentricities through the inclusion of higher order post-Newtonian effects, and the addition of an accurate merger-ringdown model that can be attached to this (or, in principle, to any other) inspiral model, we plan to maximize the range of source eccentricities and masses that can be faithfully approximated with a single waveform model.

The suggested waveform model has been optimized for speed and it is a natural candidate for use as a template generator for CBC search algorithms like gstlal. Also we will investigate other approaches for EBBH detection, such as nested sampling [263]. However, nested sampling and the development of eBBH templates that span the full range of possible eccentricities and masses are the longer term projects with an uncertain time scale.

3.7.6 Milestones

Both cWB2G and BayesWave will be ready for the O1 all-sky analysis. The same software is used for the eBBH analysis. However, our plan is to conduct the full eBBH analysis starting with the O2 run. There are two reasons for that: a) for the two-detector Livingston/Hanford network there are small differences between the ALLSKY and the EBBH cWB2G settings. These searches become significantly different when the Virgo detector joins the run, and b) the eBBH simulation work may not be finished before O1.

3.7.7 Resources

Development and Simulation studies 2 FTEs are required for the development of the eBBH search and for the MDC studies: 1 FTE to run cWB2G analysis and 1 FTE to run BayesWave. This includes generating MDC data sets, running CWB2G and BayesWave jobs, and analyzing the resulting triggers . The development of CWB2G is complete. The development of BayesWave is nearly complete. The remaining development items are reconstruction of eBBH waveforms in BayesWave, demonstration of reconstruction of the eBBH signals with BayesWave, and review of both algorithms will be finished before the O1 run. An additional 1 FTE is required for development of eBBH templates that can be used by the gstlal pipeline.

Analysis of aLIGO data 2 FTEs will be required to maintain and run the eBBH cWB2G search and the BayesWave followup.

Detector Characterization & Data Quality As with other burst searches, the data quality (DQ) is critical for the eBBH search. Most of the DQ work will be conducted in the scope of the all-sky and IMBBH searches. No additional DQ resources are required for the eBBH search.

Calibration requirements We will require calibrated data over the sensitive band of the detectors (from the low-frequency wall to 1kHz). The eBBH calibration requirements are the same as for the cWB IMBBH search and they are investigated as part of the IMBBH search proposal.

Review The analysis software used in the eBBH search is reviewed as part of the all-sky and IMBBH proposals. The review of the cWB2G pipeline is in progress and will be completed in early 2015. The BayesWave review has start in the end of 2014. We expect that the majority of the analysis will be reviewed prior to the beginning of the O1 run. The additional review items include the eBBH-specific details of the analysis (both in cWB2G and BayesWave) and the review of the existing frequency domain inspiral eBBH waveforms. It is envisioned that the review of all of the additional components could begin in the second half of 2015. We expect that Data Quality and Calibration will be reviewed as part of the aLIGO baseline CBC BNS and Burst All-sky analyses.

3.8 Search for Intermediate-mass-ratio Coalescences

3.8.1 Abstract

There is some observational evidence that intermediate mass black holes (IMBHs), with mass in the range $100–1000 M_{\odot}$ may be present in globular clusters. In these dense stellar environments, encounters between the IMBH and other stellar objects in the cluster could lead to the capture of a compact object and the formation of a binary with mass ratio of $1 : 10$ or smaller. The subsequent gravitational-wave driven inspiral and merger of the compact object with the IMBH could generate gravitational waves detectable by Advanced LIGO and Advanced Virgo. These systems are usually referred to as intermediate-mass-ratio inspirals (IMRIs) in the literature, and we will refer to them as IMRIs in the following, but we have used "coalescence" in the search title to emphasise that the merger and ringdown phases could also contribute significantly in the LIGO/Virgo band. The event rate for IMRIs is highly uncertain, given the uncertainty in the number density of IMBHs in particular. However, the IMRI rate could be comparable to or higher than the rate of IMBH-IMBH mergers. Detection of a single IMRI could provide the first direct proof of the existence of IMBHs and constraints on their number density. Any observations would have profound implications for our understanding of black holes and the evolution of globular clusters, since these systems are currently highly speculative and virtually unconstrained observationally. These systems also have strong potential as probes of fundamental physics. A search for these sources poses various complications, which is why they require a dedicated analysis. They lie outside the usual parameter range of current searches, the sources sweep through a relatively narrow range of frequencies, and the systems lie in a regime where accurate modelling of the gravitational waveforms is particularly difficult. However, several approximate models of IMRI signals have been proposed in recent years which can now be used to assess the efficiency of pipelines. Given the uncertainty in waveform modelling, we propose to use both burst and template based search techniques. We will use the coherent WaveBurst [3] with tuning for the appropriate mass range and a template based search with gstlal [241]. The performance of both techniques will be assessed using the best available models in the IMRI mass range.

3.8.2 Scientific Justification

There is some observational evidence that IMBHs of a few hundred solar masses may exist generically in globular clusters [231, 232]. Various scenarios have been suggested for IMBH formation in these environments, including the runaway collision of massive stars, mergers between stellar-mass black holes that sink to the centre of globular clusters through mass segregation, gas accretion onto a black hole seed early in the cluster's lifetime or direct collapse of low-metallicity pop-III stars. Numerical simulations suggest that in the dense environments of globular clusters, IMBHs could merge with many compact objects over their lifetime. After dynamical capture, an IMBH-compact object binary is hardened through three-body interactions with other stars in the cluster until it becomes sufficiently hard that gravitational-wave emission takes over and drives the binary to coalesce [264]. This gradual hardening mechanism means that the binary is likely to have sufficient time to circularize. In [264], an upper limit on the Advanced LIGO IMRI rate of $\sim 30\text{yr}^{-1}$ was estimated, with more realistic rates of ~ 1 every 3 years if the inspiraling compact objects are primarily $\sim 1.4 M_{\odot}$ neutron stars or ~ 10 per year if the typical compact object is a $\sim 10 M_{\odot}$ black hole. More recently, the discovery of several very massive stars in the Large Magellanic Cloud has indicated that IMBHs do not necessarily need a dense stellar environment to form [236], opening up a possible new channel for IMRI formation.

The observation of one or more IMRIs would have significant astrophysical implications. The detection of an IMBH in an IMRI (or IMBH-IMBH binary) would provide the first direct evidence for the existence of black holes in the intermediate-mass range. This would have profound implications for our understanding of stellar evolution in the dense environments of globular clusters, which are the most likely hosts of such black holes. The rates of IMRI events and the parameters of the observed systems encode information about the physical processes occurring in globular clusters and the relative efficiency of the sequence of segregation, capture and hardening that most likely gives rise to IMRI events. In addition, IMRI sources can be used as probes of fundamental physics, in particular to test the no-hair property of the IMBH [265, 266]. A $1.4 + 100 M_\odot$ non-spinning IMRI system generates approximately $340/85/5$ waveform cycles between a frequency of $10/20/40$Hz and plunge, throughout which time the compact object is within $8/5/3.5$ Schwarzschild radii of the IMBH [267]. These strong-field waveform cycles encode a map of the strong-field space-time structure just outside the horizon of the IMBH which can be used to verify that the multipole structure of the spacetime obeys the no-hair theorem expected for Kerr black holes in general relativity. We note that, although the detection of an IMRI event is extremely uncertain, due to the very uncertain number density of IMBHs in the Universe, the scientific pay-off both for astrophysics and for fundamental physics of just one IMRI observation would be immense.

There are several reasons for having a dedicated IMRI search, rather than subsuming this search into one of the existing CBC or burst searches. IMRIs are outside of the parameter space included in conventional inspiral searches. The mass ratios and chirp masses are much smaller than usually considered in the CBC high-mass search and in the proposed IMBH-IMBH binary search. The templates are also much longer than IMBH-binary templates due to the lower mass ratio. Spin effects are likely to be more important for IMRIs as all of the observed cycles will be in the strong-field regime where spin effects are most important. The size of precessional effects will be suppressed due to the low mass ratio, but these will accumulate over the large number of observed waveform cycles and so it might also be necessary to use precessing templates. Therefore, an IMRI CBC search will require the construction of a dedicated template bank. An additional complication for the IMRI search is that IMRI waveforms are difficult to model, lying between near-equal mass systems that can be described by PN and numerical waveforms and extreme-mass ratio systems that can be described by perturbation theory. An IMRI search will have to tolerate greater uncertainty in the templates than other searches. A dedicated IMRI burst search will also be required since the necessary background rejection strategy will be different for IMRIs than existing burst searches for higher chirp mass systems.

Given that both the prospects for IMRI detection and the scientific pay-off of any IMRI observations are significantly enhanced by a lower low-frequency cut-off and that significant effort is required to realise this search we propose to have this search ready for O2 [124]. We plan to run a version of the search, with all the most up-to-date reviewed code available, in O1 as a detection search only. If an event candidate is identified it will be followed up, but if not we will not plan to publish an upper limit since any pipeline used in O1 will not be optimised.

3.8.3 Search Description

We will use both template based searches (gstlal) and burst search techniques (cWB) to identify candidate signals. These will be followed up using parameter estimation codes in LALInference for IMRI waveform models. The template based search will employ a dedicated IMRI template bank, covering the relevant parameter range and tuned to be effective for IMRIs. This search will cover binaries with a total mass $\geq 20\,\mathrm{M}_\odot$ and a mass ratio $0.1 \geq q$, where q is the ratio of the smaller component mass to the larger one; the lowest companion mass we will include is $1\,\mathrm{M}_\odot$, while the upper mass cutoff will depend on the sensitivity of the detectors. The lower cut off in total mass was chosen to ensure a sufficient overlap with other CBC searches so that no events are missed, but not too much overlap. While a $20\,\mathrm{M}_\odot$ black hole is not an IMBH, if such a black hole merges with a neutron star the mass ratio is ~ 0.05, which is in the IMRI range where no numerical simulations are available and therefore is subject to the same modelling uncertainties that will be the focus of much of the IMRI search development. It is highly likely that it will be necessary to include spin effects in the templates. Theoretical work suggests that in the IMRI range the spin of the larger object is very important and, while the effect of the spin of the smaller object is instantaneously suppressed due to the small mass ratio, the effect of precession accumulates over the large number of observed waveform cycles and therefore precessing templates may be needed [268, 269]. The template bank will therefore ultimately need to employ non-precessing spinning waveform models to increase coverage and it may even be necessary to use a precessing waveform model. Nonetheless, non-spinning waveform models should provide some sensitivity in the IMRI range and since such models are already reviewed and available we will use these initially. As more complex waveform models become available these will be used to improve the search and assess the effectiveness of the simpler templates. To cover the IMRI parameter range using TaylorF2 3.5PN templates with a low frequency cut-off of 10 Hz requires $\sim 50,000$ templates (estimated using $pycbc\_geom\_nonspinbank$, though stochastic template bank placement will probably be used for the actual search). Using this template bank with EOBNRv2 non-spinning templates and injections has been shown to be effectual. However, this number will change as investigations into the necessity of including spin effects are completed. Nonetheless, the IMRI parameter space should be coverable with $O(10^5 - 10^6)$ templates, which could be pre-computed if necessary. In addition, the templates will be considerably shorter than binary-neutron star templates due to the higher primary mass which will make the evaluation of this bank computationally much simpler. The IMRI template search may also require the use of a more relaxed chi-squared veto to allow for waveform uncertainties. Understanding these issues is part of the development plan described below.

The burst search will be needed to pick up events that the gstlal search misses, for instance due to the use of poor waveform models. The burst search will use cWB to identify triggers in the usual way, but background rejection will be tuned for the IMRI parameter range using the best available IMRI models for injections (see waveform discussion in “Technical Requirements” section). The additional costs of the IMRI search with cWB are small as it will borrow from the all-sky and IMBH binary searches (see computational requirements below). We will use both burst and inspiral pipelines to improve coverage and handle uncertainty. The short duration of IMRI signals for the highest mass systems in the early Advanced LIGO era will make a burst search quite sensitive to these systems, but they will be less sensitive to the longer inspirals expected for lower mass IMBHs and IMRIs observed in later science runs. The current uncertainty in IMRI waveforms also means that the template based search will potentially not be as effective, as the templates could have moderately large mismatches with astrophysical signals. The results of these two searches will be combined by using the False Alarm Density (FAD) statistic [243], which has been already used in the burst IMBHB search in initial-detector data [237], and is being developed as a technique for combining Burst and CBC searches as part of the joint CBC and Burst IMBHB search.

Parameter estimation follow-up will use the LALInference parameter-estimation pipeline and IMRI waveform models included in LAL. Robustly identifying an event as being an IMRI will require measurement of the mass ratio. This parameter should be measured more precisely in the IMRI range than it is for more comparable mass systems due to the large number of observed strong-field waveform cycles [270]. IMBHB parameter estimation has been demonstrated to mass ratios of 1:10, on the edge of the IMRI range [271], and there have been some (unpublished) efforts to extend this further into the IMRI range. However, this was only preliminary work for a very small number of mass ratios, spins and other parameters, so further work will be required to explore the efficiency of parameter estimation codes and the likely precision of parameter estimation over the full parameter space of IMRIs. Current work has also ignored the effects of waveform uncertainty which must be understood. This will be handled by performing inference with multiple waveform models or using uncertainty-marginalisation techniques currently under development.

The search for IMRI-like triggers will be applied to all data. Neither the cWB search nor the inspiral search is expected to be particularly computationally demanding so this should be possible in practice.

3.8.4 Result validation plan

We will have two searches, the CBC gstlal search and the Burst cWB search, running on the same data. These are completely independent searches, and as such will provide the robust means for an independent validation of search results. While we do not require that both search pipelines find the candidate to be significant, the preferred signal interpretation (i.e., inferred signal parameters) and our studies of the relative sensitivity of the two searches will make it possible to determine whether the results are consistent with an astrophysical interpretation. Similarly, parameter estimation follow-up with the LALInference pipeline [74], a parameter estimation pipeline written with in the LAL library (see 3.2), will provide an additional check on the candidate events.

3.8.5 Publication Plan

We aim to have an IMRI search operating during O1 on a best-effort-available basis, to facilitate serendipitous discoveries. No upper limits will be published until all elements of the search have been fully developed and the impact of waveform uncertainties in particular has been properly understood. We aim to have a full IMRI search operating in O2.

Confident detection. If we have a confident detection of a system with IMRI-like mass and mass ratio, which has been followed up using parameter estimation codes, this will be a significant astrophysical result. In the case of a clear detection, such as the S6 blind injection with false alarm probability $\sim 7 \times 10^{-5}$, we aim to submit observational papers within 6 months of the first detection, and within 3 months for subsequent detections. We will publish a detection statement and parameter measurements, including both statistical and systematic uncertainties. The former include any parameter degeneracies intrinsic to these systems, while the latter arise from waveform model uncertainties.

Marginal or no detections. If only marginal candidates (defined as events with false alarm probability above some threshold, say 10^{-4}) or no candidates are found then we we will submit a paper describing the analysis, providing the details of any marginal candidates and an upper limit. The first paper will be based on a minimum of 6 months of data and the analysis will be performed only after the IMRI search has been fully implemented and reviewed. There will be an additional requirement that the sensitive volume of the IMRI search is at least a factor of ten better than the part of the IMRI parameter space covered previously (e.g., [120]). The paper will be submitted within 6 months of the search review being completed. Subsequent improved upper limits will be published if there is a break in detector operations and the accumulated data allows us to significantly improve existing upper limits, either by lowering upper limits in previously covered regions of the parameter space by an order of magnitude or more, or by covering a previously unexplored region of parameter space. For these subsequent papers, we will submit a paper with updated rate upper limits within 3 months of the end of data collection (or 6 months if previously unexplored regions of parameter space have been added). In the absence of detections from the IMBH binary searches, we would consider a joint upper limit paper from the two searches since both contribute statements about the number density of IMBHs and their environments.

3.8.6 Technical Requirements and Development Plan

The development of an IMRI search will have various elements. The inspiral search and any parameter estimation follow-up will rely on the existence of templates for IMRI waveforms. IMRIs are very difficult to model, as they spend many cycles in the strong-field regime, which makes post-Newtonian and numerical techniques inappropriate, but have mass ratios that are too large for perturbation theory to be applicable [272]. Several waveform models are now available which make predictions in the IMRI range — EOBNR(v2) [273, 274], SEOBNR(v2) [125], Huerta & Gair [270], Callister & Gair [275], etc. While none of these models are likely to be correct, they will capture the main features of IMRI waveforms and the differences between the models will be characteristic of the size of the difference between astrophysical IMRI waveforms and any one of these models. A necessary first-step to understanding the effect of waveform uncertainties on IMRI searches is to have multiple models available for injections and templates, extending the analysis of [276] to explore the impact of waveform uncertainty on the χ^2 veto as well as the effectiveness of parameter-space coverage. Implementation of these models in LAL has begun; in particular, Callister & Gair waveforms can now be used in LAL.

The only developments required for the cWB IMRI search are in the tuning of post-processing of triggers to reject background efficiently in the IMRI regime. This will be done using each of the IMRI models to understand the impact of waveform systematics.

For the template-based search in gstlal several developments are needed:

- Construct template banks in the IMRI range. Test their efficiency for IMRI detection.
- Explore which waveform models are needed. In particular, is it necessary to use a precessing-waveform template bank? Can template banks constructed for one of the IMRI waveform models detect injected IMRIs described by one of the other models?
- How much do standard vetoes reduce our search efficiency in the IMRI range and how do we modify them to recover the lost efficiency? For instance, an uncertain waveform may have a poor effective SNR or fail a χ^2 test as the model will not perfectly match the data. Tuning of the detection statistic may be required to optimise search efficiency.
- Tune the gstlal pipeline in the IMRI range to maximise the detection efficiency.

For parameter-estimation of IMRI systems, studies of the effectiveness of existing PE codes within LALInference for the characterisation of IMRIs are needed. Tuning of theses pipelines for the IMRI parameter space may be required. The best way to include systematic uncertainties from waveform differences in parameter estimates must also be investigated. This is particularly important for IMRIs, given the uncertainties that exist in our understanding of IMRI models. The development of techniques for folding such uncertainties into parameter estimates is underway and will be used here. The utility of template-free parameter estimation techniques, such as STAMP or BayesWave, to characterise IMRI signals, will also be explored.

3.8.7 Resources

Some of the challenges of developing the joint matched filtering and unmodeled search for IMRIs are the same as those for the IMBH binary coalescence search. Since that search is at a further stage of development at the time of writing, we rely on the Intermediate Mass Black Hole (IMBHB) search preparation to address some of the common challenges, such as using the False Alarm Density statistic to combine Burst and CBC search results. The resources required here are therefore only additional resources *on top* of those requested in the IMBHB search plan.

Development and Simulation studies The biggest challenge for the IMRI search is producing credible waveforms in the IMRI regime and assessing the accuracy of existing waveforms.

Detector Characterization & Data Quality Data quality, particularly at low frequencies, is critical for the IMRI search. As discussed in 3.6.6, we will work with detector characterization and commissioning teams to identify and eliminate sources of transient noise at low frequencies.

Calibration requirements We will require calibrated data over the sensitive band of the detectors (from the low-frequency wall to 1kHz). We will need to confirm that IMBHB calibration requirements are sufficient for IMRI searches. Following this, we will provide more detailed feedback to the calibration team on the level of calibration accuracy required for the IMRI search. The desired calibration should be available for the analysis shortly (1-3 months) after the beginning of each science run.

Review We will need review teams for the following search components: IMRI-specific details of the search pipelines, post-production code for combining their results, waveforms (where not reviewed elsewhere) and parameter estimation (aspects specific to IMRI parameter estimation – LALInference code has been reviewed and upgrades are continuously reviewed). We expect that Data Quality and Calibration will be reviewed as part of the aLIGO baseline CBC BNS and Burst All-sky analyses.

3.9 High-energy Neutrino Multimessenger Analysis

3.9.1 Abstract

Some dynamical processes with strong gravitational-wave (GW) emission, such as compact binary mergers or stellar core-collapse with rapidly rotating cores, can drive relativistic outflows that result in the emission of high-energy neutrinos. Detecting both messengers from a common source would provide the unique opportunity to develop and fine-tune our understanding of the connection between the central engine, its surroundings, and the nature of relativistic outflows. A joint search also increases the sensitivity compared to GW-only or neutrino-only searches, and can be especially interesting for sources that are difficult to detect electromagnetically.

We present the plans for the multimessenger search for GWs and high-energy neutrinos using the advanced LIGO - advanced Virgo network (c.2015+) as well as the IceCube detector. This search targets (i) GW bursts of any form, without restriction to a particular signal type, as well as (ii) GWs from compact binary mergers, both in coincidence with one of more high-energy neutrinos. The presented search plan closely follows the GW+neutrino multimessenger search performed for initial detectors. It will combine the time of arrival, significance and point spread function from neutrinos and GWs, along with a blue-luminosity-based weight and direction of galaxies in the local universe, to determine the significance of multimessenger event candidates.

3.9.2 Scientific Justification

Many of the most violent and interesting phenomena producing GW transients are also thought to be sources of high-energy neutrinos [277, 278, 279, 280, 281, 282]. These non-thermal, GeV-PeV neutrinos neutrinos are thought be produced in relativistic outflows driven by central engines also responsible for GW emission [283, 284, 285, 286, 287, 288, 289, 290]. The progenitors of both long and short gamma-ray bursts (GRBs), core-collapse supernovae with rapidly rotating cores, and highly magnetized neutron stars (magnetars) are thought to produce GWs and high-energy neutrinos that may be detectable out to relevant distances [282]. A particularly interesting development is the recent detection of astrophysical PeV-energy neutrinos [291], which will be an important science target for multimessenger observations.

Studies indicate that multimessenger searches with advanced detectors will be able to probe the characteristic parameter space for GRBs [281]. For compact binary mergers, which are the likely progenitors of short GRBs, we expect GW emission detectable out to $\sim 450\,$Mpc for optimal direction and orientation, and even farther for BH-NS mergers, with advanced detectors at design sensitivity. Core-collapse supernovae with rapidly-rotating cores, which are the likely progenitors of long-GRBs, low-luminosity GRBs and choked GRBs, may be detectable out to $\sim 200\,$Mpc upon the emission of $10^{-2}\mathrm{M}_\odot\mathrm{c}^2$ via, e.g., non-axisymmetric instabilities of the millisecond proto-NS that forms upon core-collapse. Neutrino emission is uncertain, but is expected to be comparable to the observed gamma-ray emission of $\sim 10^{51}$ erg for some sub-photospheric and collisionless shock models.

Search results can be used to constrain joint emission models [281]. Upon non-detection, the source rate will be constrained to a rate comparable to or below the core-collapse supernova rate of $\sim 10^5 - 10^6\,\mathrm{Gpc}^{-3}\mathrm{yr}^{-1}$ for characteristic emission parameters (above). The expected rate of joint GW+neutrino sources will probably be between the core-collapse supernova rate and the rate of long GRBs.

There are multiple scientific benefits of simultaneously observing GWs and high-energy neutrinos from a common source: (i) The combined information from GW and high-energy neutrino observatories can greatly enhances our confidence in a joint detection [277, 278, 279, 280, 281, 282, 292, 293]. (ii) GWs and high-energy neutrinos both carry information from the depth of their source that is, to a large extent, complementary to the information carried by electromagnetic radiation. While the GW signature of cosmic events is characteristic of the dynamics of their central engine, a high-energy neutrino flux is reflective of the presence of hadrons in the relativistic outflow generated and driven by the central engine. (iii) The emission of high-energy neutrinos is tightly connected to the creation of high-energy photons (gamma rays) by the outflow. There are specific cases in which the source optical thickness is large and prevents the gamma-rays from leaving the source. One of the most interesting prospects of joint GW - high-energy neutrino searches are common sources that are dark in gamma rays. Prominent sources of this type are choked GRBs [285, 294, 295] or low-luminosity GRBs. These sources are difficult to detect with electromagnetic observatories, and hence provide an exciting opportunity to joint GW+high-energy neutrinos searches that can discover them and/or constrain their population [280, 281, 293].

The IceCube neutrino detector [296], from which reconstructed neutrinos will be taken for the joint search, is already under operation, and will be in operation during ADE. Additionally, IceCube is planned to be extended during early ADE, and will be operating with further increased sensitivity.

The ANTARES neutrino detector operates in its complete configuration since mid 2008. It will be in operation during the currently planned period for O1 (end of 2015). It will cease operation in 2016 and will be gradually superseded by KM3NeT. KM3NeT is currently in its construction phase 1 aiming at 24 and 7 lines installed close to Sicily and Toulon respectively. Phase 1 should be completed in 2016. Few lines of KM3NeT could be operating already end of 2015. The ANTARES/KM3NeT detectors will also be used for joint searches.

In short, GW+high-energy neutrino observational results have already proved to produce exciting scientific results [281, 292], while the projected constraints [281] and expectations (e.g., [297]) suggest that multimessenger GW+high-energy neutrino searches will be a fruitful direction of research during ADE.

3.9.3 Search Description

The multimessenger GWHEN pipeline aims to identify GWs and neutrinos emitted from a common transient source. It uses GW event candidates identified by searches in LIGO-Virgo data, and neutrinos identified by IceCube or other neutrino telescopes. The goal is to find GWs and neutrinos of common origin. In the initial detector era, a multimessenger search pipeline was developed to search for temporally and directionally coincident GW and neutrino signals [298, 280, 293]. The search pipeline additionally allows for the use of galaxy catalogs, as many of the target phenomena are expected to be occurring from within or near galaxies other than the Milky Way. Other astrophysical catalogs or distributions also present an interesting option, while general all-sky searches are also important. The search pipeline was developed to be directly applicable to multimessenger searches in ADE. The corresponding code, developed and used in initial searches, will be used for ADE searches. While it makes use of available directional information, it is also effective with limited or no directional information. This scenario can be particularly useful

1. in the early ADE when only 2 detectors are available
2. with neutrino cascade events that are poorly localized
3. other cases in which neutrino information have no localization

The joint analysis uses a test statistic for GW+neutrino event candidates, which combines the significance and directional distribution of GW and astrophysical neutrino event candidates, and optionally galaxies. We only consider a GW and neutrino for analysis if they arrive within a time window of 500 s. For GWs, we measure the significance of an event candidate by adopting the test statistic of the used GW pipeline, and comparing it to the background distribution. The background distribution is estimated as standard for the used GW pipelines, using time-shifted data for cWB2G and other techniques not requiring time shifts for CBC. To estimate sensitivity, the search will adopt standard simulated signals for which test statistics are determined by the GW pipelines. For high-energy neutrinos, we calculate the significance of a neutrino by comparing the reconstructed energy of the neutrino to the distribution of the reconstructed energies of all detected neutrinos (implicitly assuming that most of them are background events). The background distribution for neutrino events is estimated as standard for neutrino searches, using the reconstructed direction and energy of detected neutrinos. We calculate the significance of a galaxy using its blue luminosity and its distance, assuming that the source rate in a given galaxy is proportional to its absolute blue luminosity. We then combine the point spread functions of GWs, neutrinos and galaxies. For GWs, we take the point spread function, the so-called skymap, from the used GW pipeline. For neutrinos, we calculate the point spread function form the reconstructed neutrino direction and directional uncertainty. We finally combine the GW, neutrino and galaxy significances, as well as the p-value derived from the combination of the point spread functions, using Fisher's method, i.e. we consider the logarithm of the product of the p-values. The background distribution of the GW-neutrino joint test statistic is obtained by time-shifting the data streams of GW detectors relative to each other and by randomly permuting the neutrino times of arrival, while keeping the neutrinos' directions relative to the detector and energy fixed.

The search will receive GW triggers with high false alarm rate threshold from an all-sky GW search that is being run on ADE data. We will adopt a FAR of 10/day or the FAR provided by the GW pipeline. GW signal injections will be used to determine the search sensitivity and to set upper limits in the case of non-detection. The pipeline will take a GW test statistic, time of arrival, and directional information (skymap) for each trigger. For more detailed information see [298].

The search will receive neutrino events from a neutrino transient search developed by the respective neutrino collaborations (IceCube, ANTARES, KM3NeT). The pipeline will take a neutrino test statistic, time of arrival, and directional information (skymap) for each trigger. For more detailed information see [298].

We will carry out two GW+neutrino searches, one for exceptional events and one for regular events. An event will be considered exceptional if it has sufficient significance such that it could be beneficial to follow the event up with electromagnetic of other observations, or if the event is a possible detection by itself. The specific significance threshold for considering an event exceptional will be determined based on available follow-up capabilities, latency from GW and neutrino searches, and the events likelihood of being an astrophysical event, or if this can be determined with additional observations. All other events that do not satisfy this exceptionality criterion will be considered regular.

Both analyses for exceptional and regular events will be carried out by the GWHEN group using the method described in [298] for all joint triggers. All data that is analyzed by the all-sky searches will be analyzed in the GWHEN search. Data quality will be taken into account by the GW searches themselves. The GWHEN pipeline will adopt the output of GW pipelines, therefore it will not additionally consider data quality.

The GWHEN search will use GW event candidates from the cWB2G pipeline, as well as the standard CBC pipeline. cWB2G will provide triggers covering unmodeled transients expected, e.g., from the progenitors of long GRBs, while the standard CBC search will provide triggers covering compact binary mergers, the likely progenitors of short GRBs. For the CBC-neutrino analysis, both binary NS and NS-BH templates are relevant as potential short GRB progenitors. Both types of sources are interesting for joint GW+neutrino detections.

We will carry out a search that identifies exceptional GW+neutrino events. These events will be a subject to low-latency analysis on an event-by-event basis. Some of these events may be subjected to EM follow-up (the latency of these will be determined based on the GW and neutrino latencies. The latency due to GWHEN will be much smaller.). We will communicate with the EM follow-up group to arrange a proper follow-up of the identified GWHEN events (e.g., by email).

Additionally, we will carry out an analysis for regular GW+neutrino events. These events will be analyzed in ensemble (i.e. the results will be interpreted for all of the regular GW+neutrino events within a measurement duration, similarly to the initial searches).

Low-latency joint GW+high-energy neutrino searches will constitute an interesting new direction for the advanced detector era. Both GW and high-energy neutrino detectors and their implemented event reconstruction algorithms will be able to provide low-latency events that in turn can be used in low-latency joint searches. As both GWs and high-energy neutrinos can arrive prior to the onset of electromagnetic emission from sources such as GRBs, joint GW+high-energy neutrino events may be primary targets for electromagnetic follow-up searches. Additionally, due to the sub-degree direction reconstruction available with neutrino detectors, joint events will exhibit significantly improved localization compared to GW-only triggers, further aiding EM-followup searches.

3.9.4 Publication plan

For GWHEN detection candidates, we consider the "$3\,\sigma$" threshold to be a *minimum* threshold on the significance to be considered for a possible "evidence for" statement. This $3\,\sigma$ is calculated for the whole analysis period. In the case of a clear first detection ($> 3\,\sigma$) of a multimessenger GW - high-energy neutrino event, we will aim to publish results within 3 months of identifying the candidate.

As the sensitivity of the detectors improve, we expect more detections. Upon regular detections, we foresee the publication of a short detection summary after each GW science run. We will aim to publish results within 3 months following the end of a science run.

Upon no detection or marginal detection, we expect to publish a multimessenger GW and high-energy neutrino search results following each ADE science run. This will involve the determination of source rate upper limits. Similarly to the expected increase of the expected observable number of sources between each science run, we require a source rate upper limit improvement of at least 2-3 in order for the results to merit a separate publication. We will aim to publish results within 3 months following the end of a science run.

3.9.5 Technical requirements and development plan

The ADE search will use the multimessenger search pipeline developed for joint searches with IceCube [298]. The same pipeline can be used for ANTARES and KM3NeT searches as well. The development of the pipeline is finished and reviewed, and it is ready to be used for ADE searches.

As the pipeline is built on the foundations of utilizing GW triggers from all-sky GW searches, its adaptation to ADE will be done on the all-sky-search level, and no adaptation is needed in the pipeline itself. The search will use the online triggers generated by the cWB2G GW pipeline, as well as CBC triggers by processing the time, test statistic and skymap from these events by the GWHEN pipeline. GW triggers from cWB2G will be acquired through direct access to the LDAS server. The same triggers will be used as in the online CBC and cWB2G analyses. The data stored on the server is of essentially identical to the data used in the reviewed GWHEN analysis, therefore no additional review is necessary. Low-latency trigger notifications will be received via gracedb with an expected latency of $\sim$ minutes. The received triggers will be automatically received by a special-purpose "alert code", and will be automatically processed by the reviewed pipeline. CBC triggers will be accessed via gracedb by the alert code. The alert code will be a standard software that will automatically receive incoming GW and neutrino trigger alerts, and feeds these alerts to the GWHEN pipeline.

The IceCube collaboration, as well as ANTARES and KM3NeT will provide readily usable triggers for the search, for which no development is needed. Low-latency trigger notifications will be received via gracedb by the alert code with an expected latency of $\sim$ minutes. The received triggers will be automatically processed by the reviewed pipeline.

The low-latency joint search pipeline will be tested using GW and neutrino signals generated by the GW pipelines (engineering run) and IceCube (real observations) in an identical data format as for ADE, and will be accessed identically to the access during ADE.

3.9.6 Resources

The multimessenger search requires low-latency GW triggers from an all-sky search pipeline (along with low-latency neutrino triggers). Beyond this, there is no requirement for the search. Running the pipeline is computationally cheap. The low-latency reception of GW triggers with direction reconstruction is important.

The multimessenger search requires 0.5 FTE for running and maintaining the search and its infrastructure, as well as the organization of publication efforts.

For the search the MOU between IceCube and the LVC needs to be renewed. Similarly, the MOU between ANTARES and the LVC needs to be renewed. An MOU needs to be signed between KM3NeT and the LVC. IceCube is already generating the data products required for this search.

Personnel: IceCube GWHEN search development/testing will need up to 0.1 FTE, while running will need up to 0.1 FTE. ANTARES/KM3NeT GWHEN search development/testing will need up to 0.5 FTE, while running will need up to 0.1 FTE.

3.10 Search for transients from Cosmic Strings

3.10.1 Abstract

Powerful bursts of gravitational waves (GWs) are expected to be produced by cosmic string cusps and kinks. The search for GW signals from cusps was performed on initial detectors' data and was published in 2014. We propose to conduct the same search using the data of Advanced LIGO and Advanced Virgo detectors. In addition, we plan to search for kink signals as well.

3.10.2 Scientific Justification

A cosmic network of strings may form as a result of phase transitions in the early Universe [299]. When a U(1) symmetry is broken in multiple causally disconnected spacetime regions, one-dimensional topological defects, i.e., strings, are expected to form [300]. For a long time, cosmic strings were considered candidate sources for structure formation in the early Universe [301]. Cosmic microwave background (CMB) experiments, however, have shown that cosmic strings can only contribute up to a few percent of the overall anisotropies observed [302, 303, 304, 305, 306]. More recently it was realized that strings can also be produced within the framework of string theory inspired cosmological models and grow to cosmic scales [307, 308, 309, 310, 311]. Cosmic strings produced in string theory motivated models (dubbed "cosmic superstrings") have received much attention since they could provide observational signatures of string theory [312, 313].

Observational constraints on cosmic string models are often given as bounds on the string tension $G\mu$ ($c = 1$), where G is Newton's constant and μ the mass per unit length. Such constraints have been derived from direct searches for line discontinuities in the CMB temperature maps [314, 315, 316] and from simulations of string-sourced CMB anisotropies [302, 303, 304, 305, 317, 318]. These analyses, based on various assumptions about the string network, set upper limits on $G\mu$ in the range of 10^{-7}–10^{-6}. The recent results from the Planck mission [306] constrain $G\mu$ to be lower than 1.5×10^{-7} and 3.2×10^{-7} for Nambu-Goto and Abelian-Higgs strings, respectively.

A promising way of detecting the presence of cosmic strings and superstrings is the gravitational wave (GW) emission from loops [319, 320]. When two string segments meet, they exchange partners or intercommute with a probability p. For superstrings, the reconnection probability can be less than unity ($10^{-4} < p < 1$ [321]) while field theory simulations show that topological strings will essentially always reconnect. This is partly due to the fact that fundamental strings interact probabilistically. Furthermore, superstring models have extra spatial dimensions so that even though two strings may meet in three dimensions, they miss each other in the extra dimensions. When a string intercommutes with itself, a closed loop breaks off. The loop oscillates, radiates gravitationally, and decays (see [322] for a review of the dynamics of cosmic string loops).

Special points on the cosmic string loop play important roles – cusps and kinks. Cusps are points along the string with large Lorentz boosts. Kinks are loop discontinuities that forms in particular every time inter-commuting occurs. Both produce powerful bursts of gravitational radiation [323].

We propose to perform searches for such GW transient signatures from cosmic strings in data from the advanced LIGO and Virgo gravitational wave detectors.

The GW emission by cosmic strings depends on the loop size, which is often written as a fraction of the horizon at the time of formation $l = \alpha t$, where t is the cosmic time. Early simulations such as Ref. [324] suggested that the size of loops is set by gravitational backreaction and $\alpha \leq \Gamma G\mu$, where $\Gamma \sim 50$ [300]. More recent simulations favor cosmic string networks where the size of loops is dictated by the large scale dynamics of the network, in which case $\alpha \lesssim 1$ [325, 326]. We parametrize $\alpha = \varepsilon \Gamma G\mu$ with $\varepsilon < 1$ following the convention of Ref. [319].

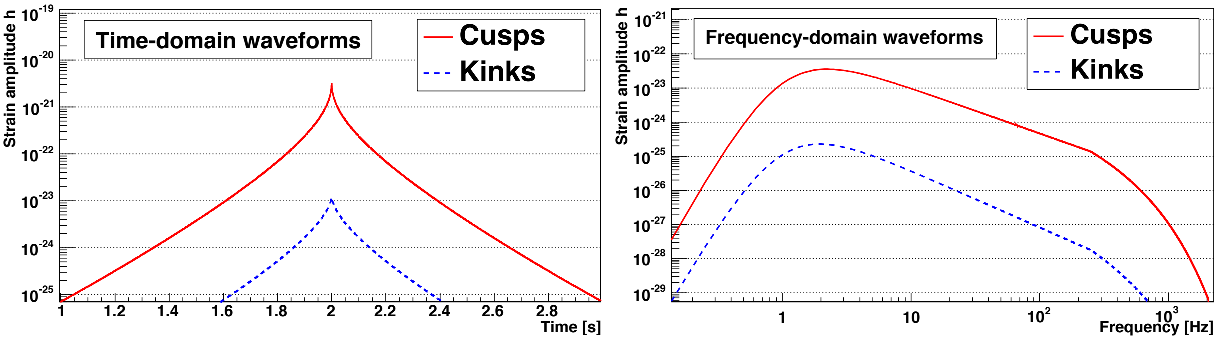

Figure 5: Time-domain (left) and frequency-domain (right) waveforms for gravitational waves produced by cusps (red plain curve) and kinks (blue dashed curve). These waveforms have been computed with the same set of cosmic string parameters: $G\mu = 10^{-8}$, $\varepsilon = 10^{-5}$, $z = 1$ and $f_h = 250$ Hz.

The possibility of direct detection of GW bursts from cosmic string cusps was first suggested in 2000 by Berezinsky *et al.* [327]. Shortly after, Damour and Vilenkin showed that the stochastic GW background generated by oscillating loops is strongly non-Gaussian [323]. Occasional sharp bursts of GWs produced by cusps are expected to stand out above the stochastic background [323, 328, 319]. Damour and Vilenkin predict that the GW burst signal produced by cusps is linearly polarized and the expected waveform in the frequency domain (using logarithmic Fourier transform) is $h_{cusp}(f) = A_{cusp} f^{-4/3}$ ($h_{kink}(f) = A_{kink} f^{-5/3}$) with an exponential decay that sets on at some high-frequency cutoff f_h. Figure 5 shows an example of such a waveform for a given set of parameters. The signal amplitude A_{cusp} is determined by the string tension, the loop size, and the propagation distance. Direct searches for these signatures such as the one proposed here explore the parameter space $(G\mu, \varepsilon, p)$.

Constraints on cosmic string parameters obtained from the search for GWs produced by cusps have been published in 2014 [11]. The next generation of ground-based GW detectors will probe the cosmic string parameter space further. The improved sensitivity of Advanced LIGO [329] and Advanced Virgo [330] will eventually allow us to search for cosmic strings with an order of magnitude lower tension.

Recent studies shows that GWs produced by cosmic string kinks might also be of interest for LIGO and Virgo. The signal amplitude is smaller than for cusps (see Fig. 5) but a proliferation mechanism could offer a production rate several order of magnitude larger [331] see Fig. 6). In the future, it is planned to search for GWs associated to both cusps and kinks and to publish combined results.

3.10.3 Search Description

The search for GW bursts from cosmic strings begins with a matched-filter analysis of strain data from each detector separately [332]. It consists of projecting the whitened data onto an overpopulated [6] template bank defined by a set of cusp waveforms with a high-frequency cutoff spanning from 75 up to 8192 Hz. This procedure results in a time series for the signal-to-noise ratio (SNR) for each template. An event is identified when the SNR > 4 and only the template with the largest SNR is retained when several templates are triggered at the same time. A set of five variables is used to characterize an event. The event time t_e and the SNR ρ are determined by the point where the SNR time series is maximum. The triggered template provides the high-frequency cutoff f_h and the amplitude A. In addition, a χ^2 parameter can be computed to characterize the match between the event and the signal waveform in the time domain [186].

To discriminate true signals from background events, we apply the multivariate technique described in Ref. [333], which uses a set of simulated GW events and typical noise events to statistically infer the

---

[6] the maximal mismatch between two consecutive templates is 0.001.

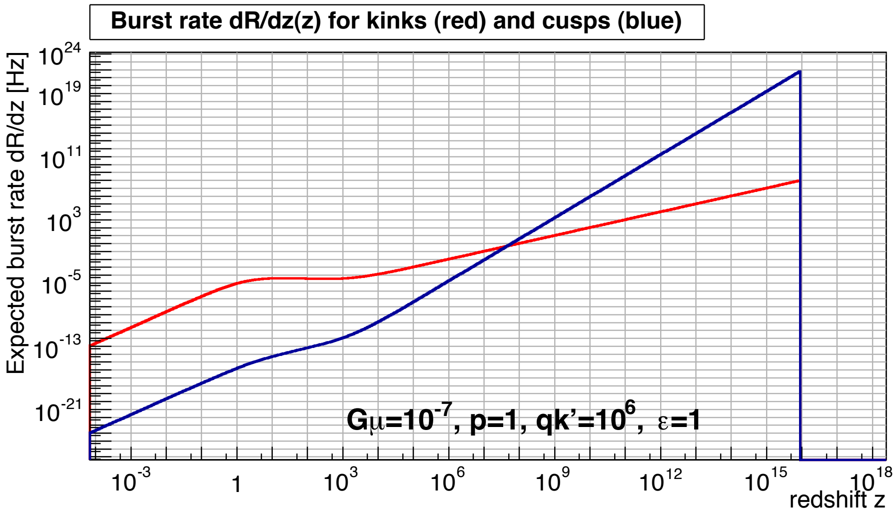

Figure 6: Predicted rate of GW bursts produced by cosmic string cusps and kinks. The proliferation mechanism, parameterized by $qk' \gg 1$, favors the kink event rate over a wide range of redshifts.

probability for a candidate to be signal or noise.

Figure 7 displays the region of the cosmic string parameter space that is excluded by the analysis of initial detectors' data [11]. For comparison, we also show limits, fixing p at 10^{-3}, derived from constraints on the GW stochastic background spectrum. These limits were computed adopting the same cosmic string model and using the same parameters $(G\mu, \varepsilon, p)$. Our result improves the indirect CMB bound [334, 335] by a factor 3 for intermediate ε values. It nicely complements existing limits provided by pulsar timing experiments for large ε [336, 305] and by the LIGO stochastic search in the very small loop regime [335].

3.10.4 Results validation plan

If an event is found to significantly differ from what is expected from background, we will immediately trigger the following procedure:

- we will ping the detector characterization groups to investiagte detector noise at the time of the event.
- We will compare the results of our search with the untriggered all-sky searches (CWB2G, EXCESSPOWER, and omicron+LIB).

3.10.5 Publication plan

We will run the cosmic string pipeline over O1 data to search for both cusp and kink signals. O1 data will be analyzed in a detection-only mode. Even if the O1 sensitivity is improved by some factor, this will not impact the upper-limits significantly enough. The O1-only analysis is indeed penalized by the low livetime and the presence of only 2 detectors. The O1 analysis will lead to a publication only if a detection is made. In this case, a paper for LVC consideration should be ready within 6 months after the run ends.

In case of non detection, the plan is to publish combined cusp and kink upper limit results of O1/2/3.

3.10.6 Resources

Detector characterization Studies performed over S5 data showed that data quality vetoes had a low impact on the events of this search as the background distribution was found to be almost Gaussian. In

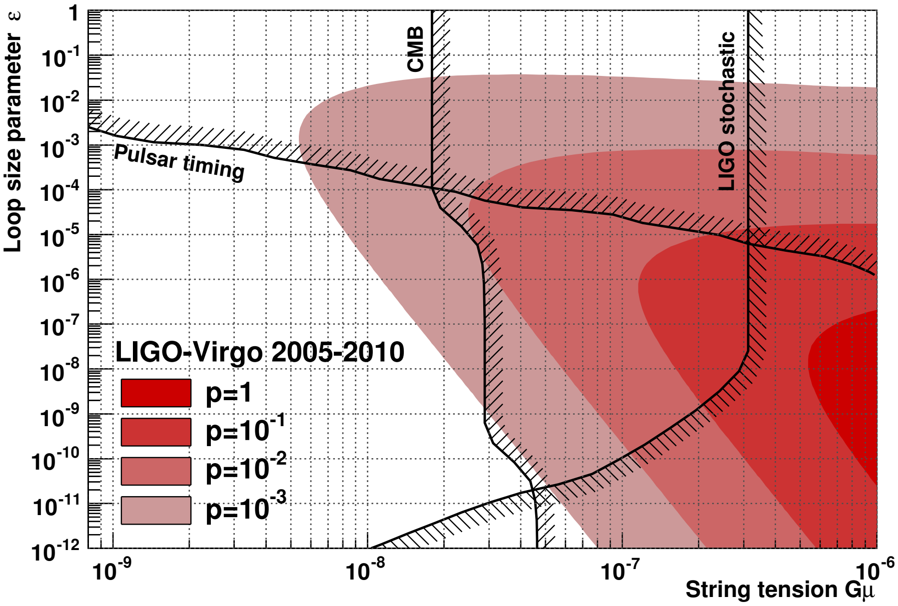

Figure 7: This plot presents existing constraints on cosmic string parameters: the string tension $G\mu$, the loop size parameter ε and the probability p that two string segments interact when they meet. The S5-6/VSR1-2-3 analysis is able to reject the regions filled in red. For comparison, other constraints derived from searches of a GW background from cosmic strings (pulsar / CMB / LIGO stochastic) are given fixing p at 10^{-3}.

the future, it is planned to carefully select useful vetoes for this type of search and not to apply the vetoes blindly.

Calibration The cosmic string search is sensitive mostly in the bucket of the detectors. Special care should be taken to insure the best calibration uncertainties at these frequencies.

Review The cosmic string search pipeline was already reviewed. The review process of the furture analysis is expected to be straight-forward and to require little manpower

3.11 Search plan for long-lived gravitational-wave transients

3.11.1 Abstract

Unmodeled long-lived gravitational-wave transients (lasting $\gtrsim$ 10–$10,000\,\mathrm{s}$) are an interesting class of signals for advanced detectors. Such long-lived transients have been linked to the death of massive stars. In one class of models, gravitational waves are emitted by a rapidly spinning protoneutron star, which may be spun up through fallback accretion. In another class of models, the signal comes from the motion of clumps in an accretion disk. In either case, the signals are long-lived, narrowband, and may occur with a sufficiently high rate so as to be observed with advanced detectors. Other possible scenarios for long-lived gravitational-wave emission include protoneutron star convection, rotational instabilities in merger remnants, r-mode instabilities associated with glitching pulsars and type I bursts from accreting pulsars, eccentric binary systems, and, of course, surprises. One targeted long-transient search with initial LIGO data is published and a second all-sky search is nearing completion. Future searches for long-lived transients will require coordination with various LIGO-Virgo subgroups and search plans. This search plan describes the scientific justification for long-transient searches and the technical description and requirements for such searches. These searches will carried out using the STAMP-AS pipeline, X-PIPELINE, and CWB2G. All three pipelines are sensitive to gravitational-wave signals in this long-duration regime, and their results will contribute to both this search effort and verification if a signal is detected.

3.11.2 Scientific Justification

Sources of long-lived (lasting $\gtrsim$ 10–$10000\,\mathrm{s}$) gravitational-wave transients are reviewed in [337]. The most promising sources fall into two categories, both of which begin with the death of a massive star.

Protoneutron stars. The first scenario relies on the formation of a protoneutron star. If the protoneutron star is born spinning rapidly, it may develop an instability (e.g., a bar mode), leading to the strong emission of long-lived, narrowband gravitational waves [338]. A protoneutron star can also be spun up through accretion of stellar remnant fallback such that an instability sets in [339]. The resulting gravitational-wave emission can last for $\approx$ 40–$3100\,\mathrm{s}$ [340]. Advanced LIGO might detect rotational instability signals from protoneutron stars out to to distances of up to $\approx 40\,\mathrm{Mpc}$ [341, 342]. The rate of observed supernovae in this volume is on the order of $\approx$ 10–$30\,\mathrm{yr}^{-1}$ [340], though the fraction of these stellar explosions that might result in an accretion fallback signal is currently unknown. Nonetheless, a single such detection would provide an unparalleled glimpse into the moments following stellar collapse and the birth of a neutron star or black hole.

Accretion disks. The second scenario relies on the formation of an accretion disk following stellar collapse. A central spinning black hole drives turbulence in the accretion torus, which leads to the formation of clumps. The motion of these clumps emits long-lived narrowband gravitational waves [343, 344, 345], on the time-scale reaching 10^3 sec. Alternatively, clumps may form through accretion disk fragmentation, also leading to the production of gravitational waves [346]. The rate and energy budget of accretion disk instability signals are debated. However, we estimate that advanced LIGO can observe accretion disk instability events out to distances of $540\,\mathrm{Mpc}\,(E_{\mathrm{gw}}/0.1M_{\odot})^{1/2}$ where E_{gw} is the gravitational-wave energy budget [341, 342]. We note that long gamma ray-bursts are observed at a rate of about $0.3\,\mathrm{yr}^{-1}$ within this radius (and many are likely to be missed due to beaming) [347]. A single detection would provide unprecedented information about the environment following the collapse of a massive star and could shed light on the mechanics of long gamma-ray bursts.

Other sources. Other scenarios for the production of long-lived gravitational waves include protoneutron star convection [348], rotational instabilities in merger remnants [337], and eccentric binary systems [349]. While these sources are associated with signals on the time-scale of 1 minute to 1 hour, much longer transient signals (on the time-scale of days) are also possible, for example in glitching pulsars or in accreting fast millisecond pulsars. The subfield of theoretical investigations into long-lived transients is fluid, and it is prudent to have a search dedicated to whatever long-lived transient signals may be awaiting discovery: predicted models, yet-to-be-predicted models, and total surprises.

Relationship to other searches. This search plan complements the burst all-sky search for short-duration signals [350] by extending the parameter space to longer durations. We will coordinate with the short-duration search teams to determine the complementarity of the two searches and the overlap in sensitivity of the pipelines involved. For example, recent benchmark tests have indicated that STAMP-AS, CWB2G (with a long-burst search configuration), and X-PIPELINE have comparable sensitivities for signals with duration 10–200 s while CWB2G (with a short-burst search configuration) has better sensitivity below 10 s. This search also overlaps with triggered searches for long-lived transients described in the GRB search plan 3.5 and the neutron star transient search plan 3.12. We will work with these other ongoing efforts to study long-lived transients, and collaborate to identify the advantages and complementarities of different pipelines.

This search also overlaps with efforts in the CW group to look for long-lived transient signals from neutron stars. We will coordinate with the CW group to identify overlapping interests. Finally, this work overlaps ongoing work in the Stochastic group in which STAMP is used to study the time dependence of a stochastic background signal and, in doing so, to identify very long-lived ($\gtrsim$ 10 hours) signals [351]. The two projects are complementary as they focus on different time scales, and the stochastic search is primarily concerned with the effect of transient phenomena (astrophysical or due to noise) on the stochastic background search, whereas this search is concerned with studying astrophysical long-lived transient signals.

3.11.3 Search Description

This search proposal describes using advanced LIGO and advanced Virgo data to conduct an all-sky, all-time search for long-lived signals from 10–10^4 s. We plan to carry out these searches using STAMP-AS, X-PIPELINE, and CWB2G. Each of these pipelines uses a distinct approach to the analysis, providing for robustness and cross-validation of the searches in this relatively unexplored region of parameter space.

STAMP-AS

STAMP-AS [337, 352, 347, 341, 342] is an all-sky extension of the STAMP matlab based library[7] developed initially by the Stochastic group and also used for detchar purposes. The pipeline cross-correlates the output of any combination of detectors. While fully coherent the analysis remains computationally affordable thanks to a series of compromise (reduced sky grid for instance) and specific developments (pixel based clustering adapted for hundreds of second long signals). The S5 and S6 LIGO data have been analyzed using STAMP-AS to search for transient signals of duration up to 600 sec in the 40–1000 Hz frequency band. Some of the outputs include: managing the computational cost of the search, understanding the background and minimizing the effect of noise transients to a level approaching Gaussian noise, and investigating the applicability of the search to eccentric binary waveforms [354, 355]. The pipeline has also been verified via a comprehensive review.

We plan to have a hierarchical strategy to select the most significant events without wasting computional resources. We will run the STAMP-AS standard configuration with 200 time slides to reveal the presence of significant events at 2.8-σ level for a single search. We will then determine if it can be explained in terms of a detector artifact. If the candidate is promising, we will present the results to the burst group and estimate its significance running more time slides using a quicker time slide code (semi coherent method) and using alternative analytical methods to time slides **??**.

X-PIPELINE

---

[7] STAMP-AS resides in the Matapps svn repository [353]

X-PIPELINE [224, 222] performs coherent burst searches of data from arbitrary networks of detectors and has been used in many externally triggered searches [356, 357, 358, 359, 360, 361]. Recently, X-PIPELINE has been extended to an all-sky pipeline by replacing the sky-grid based analysis by a spherical radiometer analysis [362, 363]; this technique allows correlations between detectors to be computed quickly for signals with durations longer than $\sim$1 s. The spherical radiometer formulation also provides a fundamentally different approach to all-sky searches than the traditional grid-based approach used by other burst pipelines.

Our search strategy is to perform an initial analysis with enough time slides to tune for maximum sensitivity at the 3-σ detection level for a single pipeline. In the case of a loud on-source event, additional time slides will be run to estimate significance to the 3-σ level with a trials factor of 3 (accounting for the CWB2G and STAMP-AS pipelines also participating in the search). See [364] and Sec. **??** for details of the associated CPU costs.

CWB2G

The CWB2G pipeline is described in detail in the short-duration burst search plan [350]. The extension of CWB2G to the long duration region of the burst parameter space is achieved by running the baseline all-sky analysis with one more additional time-frequency resolution (1 s) optimised for long-duration transients at no significant increase in the computational cost (<10%). Therefore, there is no need to run a separate long-duration CWB2G analysis and the cost of this search is already accounted for in the short-duration search plan. The extended CWB2G pipeline can efficiently detect signals with the durations up to several hundred seconds. However, it is currently not capable of detecting signals longer than 1000 s. Additional exploratory work is necessary to adapt the CWB2G pipeline to longer signals.

For all three pipelines, all science data with two or more detectors operating at reasonable sensitivity will be searched, to maximise chances of a detection. Each search will identify and rank candidate events, using GW correlation tests and data quality checks to reduce the noise background. The traditional time-slide method will be used to assess the background and thus the significance of candidate events, as well as quicker alternative methods. Pipeline sensitivities will be evaluated for a number of test waveform morphologies[8]. It is foreseen that the offline analyses will initially be run on a small subset of all available livetime or with a small number of time slides to identify any issues; the results will also feed into data quality studies and data monitoring tasks. The definitive analysis will be performed offline using the best calibration, data quality, and other relevant information available.

3.11.4 Result Validation Plan

The long duration transient search will probe a regime where some types of signals are predicted, but much uncertainty remains. In the case that a pipeline sees a long duration transient the best confirmation strategy will be to examine how the signal appears in the searches by the two other pipelines. We note that each pipeline is tuned and run independently, and rely on a radically different algorithmic implementation, so each provides an independent confirmation of the results of the others. Given the event properties and the known relative sensitivities of the pipelines, we expect to be able to understand why the event is or is not seen by each pipeline with a given significance. This will require more comprehensive comparisons to understand the consistency and relative sensitivities of the three pipelines. These comparisons are a high priority activity in preparation for the first science runs.

3.11.5 Publication plan

In the event of a detection, we plan to publish as soon as possible, and ideally within six months of the event. A detection paper would focus on the attributes of the observed signal: the reconstructed time-

---

[8] See https://wiki.ligo.org/Main/STAMPAllSkyWaveforms for the waveforms used in the 2005-10 STAMP-AS all-sky search.

frequency evolution, the time and sky location of the event, etc. The discovery of even one long-lived gravitational-wave transient that does not conform to a compact binary coalescence model would be of enormous interest to the broad astronomical community. Thus, we would recommend submitting a detection paper to a high-visibility journal. We will also investigate the utility of burst parameter estimation pipelines [350, 365, 366] for reconstructing long-duration signals.

In the event that no candidate event is observed, we will set limits on the distance and/or rate to long-lived sources given different signal models. Waveforms for protoneutron star models and accretion disk instability models are available and the latter have been successfully used in an initial LIGO analysis [347]. Other ad-hoc waveforms will be also considered. We recommend publishing a paper with updated limits if/when the constraints on the distance to a long-transient source with a fixed energy budget improve by a factor of at least two (corresponding to a factor of eight improvement on the rate) *and* there is a convenient break in data-taking (the end of a run). Papers reporting on improved limits can be submitted to journals such as PRD or CQG. It may also be suitable to include these results in a broader paper (e.g., jointly with the short-duration burst searches).

3.11.6 Technical requirements and development plan

STAMP-AS

The baseline plan is to carry out the analysis with the current stable version of STAMP-AS. Most of the STAMP-AS code has been reviewed for the published initial LIGO targeted analysis [347] and the all-sky search [355]. A handful of improvements are currently under test. Among them, one could cite: the use of the new stochtrack seedless clustering algorithm [341, 342], an alternative SNR pixel calculation to improve the sensitivity to monochromatic signals. The seedless clustering code has been tested and benchmarked with both Monte Carlo and recolored noise [341, 342]. Furthermore we also plan to continue to work on the optimisation of the code in order to speed up the computational time for the background estimation. STAMP-AS is a fully coherent search that is costly when doing time slides. We are investigating pipeline modification to reduce the cost and study alternative methods to time slides [367]. Besides, the current implementation of STAMP-AS allows us to search for signal as long as $\sim 1000\,\mathrm{s}$. Beyond, STAMP-AS is facing memory and computational issues and changes are mandatory. In coordination with similar developments carried out by the Stochastic group we plan to increase the signal duration coverage for signals longer than 1000 s after 01.

Significant coding effort has been made to take advantage of highly-parallel architecture. For instance, the seedless clustering algorithm stochtrack runs most efficiently on GPUs and multi-core CPUs. Note that we plan to adapt the STAMP-AS pipeline to run on Virgo clusters (Lyon and Bologna). This development should not happen before O2.

In order to catch problems early and facilitate a smooth analysis, we plan to start analyzing O1 data immediately as it becomes available with a goal of estimating the background and sensitivity on an approximately biweekly basis. We will coordinate with potentially overlapping efforts to maintain a mock dataset with long-lived injections in order to compare the domain of utility of different searches/algorithms.

X-PIPELINE

The baseline plan is to carry out the analysis with the spherical radiometer version of X-PIPELINE. Most of the matlab and python wrapper codes have already been reviewed as part of the GRB and supernova searches, so review effort will focus on the fast cross-correlator codes [362]. This search has been approved for review and a review team is being assembled.

Current effort is focused on tuning the analysis to optimise the sensitivity to signals with durations in the 10 s-1000 s range. This study is being done through a comprehensive analysis of the S6/VSR2,3 data set [9]. In principle, the current implementation of X-PIPELINE can detect signals of arbitrarily long duration provided

---

[9] https://wiki.ligo.org/Bursts/AllSkySecondPipeline?rev=4

they are non-monochromatic. However, the longest signals tested to date are $\sim 250\,\mathrm{s}$; characterisation of the sensitivity to longer signals is an outstanding priority. Longer-term development will also include testing multi-variate analysis methods for signal/background discrimination in X-PIPELINE triggers, and further code optimisation to improve speed and memory usage.

CWB2G

The CWB2G pipeline is currently under review in preparation for O1; details of development and review milestones can be found in the short-duration burst search plan [350].

3.11.7 Resources

Detector Characterization

We do not anticipate that the long-transient searches will require significant detector characterization work beyond what is already planned for other searches. During initial LIGO, STAMP-AS searches were shown to produce well-behaved noise distributions (similar to stationary Gaussian noise) through the application of notches for instrumental lines, basic data quality flags, and a glitch identification flag [352, 347]. This is probably attributable to the long time scales probed by the search as well as the fact that STAMP-AS uses cross-correlation [352]. However, the loudest background events of each search will be scrutinized to understand the main non-Gaussian sources of background events and provide feedback to the detector characterisation group.

Calibration

There are no special requirements for calibration beyond what has already been requested by the four data-analysis group [368]. Quoting the short duration all-sky burst search plan: "To meet our publication schedule, the offline search will require "final" calibrated $h(t)$ from three months after the beginning of each run (within two months for the 2015 run). All tests to date have shown that the calibration uncertainties will not have a significant impact on the analysis provided they are of the level already achieved in the S5-6 and VSR1-3 runs."

No specific long-duration hardware injections are requested.

Review

The STAMP-AS review has started in July 2014 and its completion for the S5/S6 search is expected by mid 2015. The X-PIPELINE review has just started. The CWB2G review is on-going; see the all-sky short duration plan for details.

3.12 Searches For Transient Gravitational Waves From Isolated Neutron Stars

3.12.1 Abstract

Violent phenomena associated with neutron stars (NSs), such as flaring activity in magnetars and binary coalescence, may result in the excitation of various oscillatory modes which leads to transient gravitational wave emission. Although the strength and likelihood of any putative signal across a variety of sources ranges from being completely unknown to likely weak and only detectable for galactic or very local galaxies, it is thought that these signals convey unique information on the neutron star equation of state. The detection and characterization of gravitational waves (GWs) associated with NS oscillations holds the potential for GW neutron astereoseismology, whereby NS oscillation mode identification and characterization leads to constraints on the equation of state (EOS). Here, we provide details for a plan of readiness for opportunistic searches for transient GWs from neutron star oscillations for *extraordinary* events, such as a hyperflare from a galactic magnetar or the detection of a binary neutron star inspiral GW signal.

3.12.2 Scientific Justification

Galactic Magnetar Hyperflares

Isolated NSs exhibit a variety of violently energetic and unpredictable phenomena. For example, isolated NSs with powerful magnetic fields ($B \sim 10^{15}$ G) known as *magnetars*, undergo episodic X- and gamma-ray burst activity which manifest as anomalous X-ray pulsars (AXPs) and soft-gamma repeaters (SGRs). Typical isotropic burst energies are usually lower than 10^{42} erg, although three SGR "hyperflares", with energies 10^{44}-10^{46} erg have been observed in the past 30 years [369, 370, 371]. Indeed, it is likely that a fraction ($\sim 15\%$) of short gamma-ray bursts are, in fact, extragalactic magnetar hyperflares.

While these phenomena are not well understood, it is reasonable to expect that such a sudden and localized energy release could excite non-radial oscillation modes in the star, leading to GW emission. The most efficient GW-emission mode is the f-mode, since this is damped almost entirely via GW emission with no competing damping time-scales. The signal thus emitted would be expected to resemble a damped sinusoid lasting O(100) ms and oscillating principally around 1–3 kHz, depending on the NS EOS [372]. However, if the mechanism for a violent event like a magnetar flare is confined to the surface of the star, it may be that, rather than global f-modes, torsional oscillation modes are excited, potentially leading to GW emission at ∼10–2000 Hz. The quasi-periodic oscillations (QPOs) observed in the pulsating tail of giant flares from soft-gamma repeaters such as the 2004 event in SGR 1806-20 are believed to be associated with these modes [373, 374]. While torsional modes themselves are not likely to be particularly efficient GW emitters, if they couple to modes in the NS core, one may expect some form of damped GW signal at similar frequencies to the observed electromagnetic QPOs. Finally, and more speculatively, dynamical or secular instabilities (see e.g., [375]) in the NS may become active following some violent event, leading to transient, but long-lived GW emission.

Previous LIGO/Virgo analyses have generally searched for generic short-duration bursts associated with violent events such as magnetar flaring activity [376, 377, 378], pulsar glitches [379] and potentially nearby ($\sim$ 1 Mpc) short γ-ray bursts where the burst may be an extragalactic magnetar hyperflare [380, 359]. To date, there has only been one galactic hyperflare, the 2004 event in SGR 1806-20, which coincided with a science run and an analysis was performed to search for GWs associated with quasi-periodic oscillations observed in the X-ray light curve for that extraordinary event [381]. The magnetar analyses in particular were met with interest in the astrophysical community and, given the proximity and energetics of the potential source, the science potential of identifying and characterizing NS mode frequencies, as well as the ability to perform an externally-triggered search with a relatively well-defined event time and most likely, sky-location, we propose to continue this legacy in O1 and beyond. However, the uncertainties in the GW detectability and even in the likelihood of mode excitation suggest that we should focus attention only on the most extraordinary events in isolated neutron stars.

Detection prospects for GW signals from magnetars range from unknown to unlikely, depending on which oscillation modes are excited (i.e., the frequency of the GW emission) and the dominant damping mechanism (i.e., the stability and duration of the GW emission). *We therefore propose only to analyze extraordinary events from magnetars, where there is evidence of a hyper-flare and the source can be confirmed as being within our own Galaxy, such as the 2004 flare in SGR 1806-20.* For a known galactic magnetar, the null-result would, as in past analyses, take the form of GW energy upper limits at fiducial frequencies using ad hoc signal injections. The astrophysical significance of the result would depend on the EM energy of the hyperflare and the uncertainty in the source distance: we only consider a null-result to be astrophysically interesting if the GW energy upper limit is within a factor of unity of the measured electromagnetic energy emission.

Post-merger Signals From Binary Neutron Star Coalescence

The inspiral and coalescence of binary neutron star systems is an extremely rich source of astrophysical phenomena, ranging from strong GW emission to short gamma-ray bursts and kilonovae. It is also another potential source of GWs from NS oscillations. It has become apparent in the past few years that the likely outcome of a binary neutron star merger is the formation of a stable, or quasi-stable, neutron star remnant, hereafter referred to as a post-merger neutron star (PMNS) (e.g., [382, 383, 384, 385]. The stability of the PMNS depends both on the NS EOS and the mass of the system: soft EOSs and/or high-masses will result in either prompt or delayed collapse to a black hole (where the delay arises as the object is temporarily supported by centrifugal forces which eventually diminish as it spins down). Transient non-axisymmetric deformations in the surviving postmerger remnant lead to a short duration ($\sim 10-100\,\mathrm{ms}$) burst of GWs with rich high frequency content, dominated by emission from f-mode oscillations at $\sim 2-4\,\mathrm{kHz}$ and generally lower-frequency sub-dominant peaks from nonlinear couplings between certain oscillation modes [386]. A number of studies [385, 387, 388], have identified and confirmed a correlation between the dominant postmerger oscillation frequency (i.e., half the peak GW emission frequency) and the radius of a fiducial cold, non-rotating neutron star. Due to the short-duration and high-frequency nature of the merger/post-merger burst this source is likely to be detectable in aLIGO only for nearby events (10–50 Mpc; see [389] for a comprehensive detectability study of the short burst following the merger) even at design sensitivity.

The significance of the detection, however, would be profound as EOS constraints may be made by the mere detection of the signal and measurement of the dominant oscillation frequency (a relatively easily measured parameter).

If the PMNS remains stable, additional higher-amplitude, longer-duration signals may also be emitted. If, for example, the merger results in a stable magnetar due to magnetic field amplification, the star may be distorted into a prolate spheroid leading to a slowly down-chirping GW signal which sweeps down from $\sim$ kHz, lasting $\mathcal{O}(10^3-10^4)$s [390]. Such a signal may be detectable to $\sim 100\,\mathrm{Mpc}$ (assuming aLIGO design sensitivity). Given the high rotational energies involved, it may also be possible that the dynamical bar-mode instability may activate and survive for a substantial length of time in a stable post-merger remnant, again leading to strong GW emission.

Given the potential for high-impact astrophysics in the event of detecting any post-merger signal, we propose targeted follow-ups of detections of binary neutron star coalescence signals for both the short, merger/post-merger burst, as well as the longer duration transient from a stable PMNS

We note, however, that a null-result would be insignificant: the strength of the post-merger emission is uncertain but likely to be weak and high-frequency, it is not clear that the distance to the source would be well known (introducing an associated uncertainty into an energy upper limit) and there is the potential scenario that the merger results in prompt-collapse to a stellar mass black hole, whose ring-down is far less likely to be detectable than the post-merger oscillations of a surviving NS. This distance uncertainty, together with the degeneracy between a weak/distant source and the prompt-collapse scenario would likely render upper limits somewhat uninteresting.

3.12.3 Search Description

This search proposal discusses searches for GWs triggered by two distinct events: hyperflares in Galactic magnetars and post-merger oscillations in NS remnants which survive binary neutron star coalescence. Many aspects of the proposed search methods and potential GW signals, however, are commmon to both. In particular, both events have the potential to result in:

1. A short (10's of ms), high-frequency burst of GWs associated with f-mode oscillations.
2. Longer-duration (100-1000 s) transients associated with phenomena such as torsional oscillations, instabilities or magnetic-field induced quadrupole moments.

As a preliminary minimal effort and low-latency search, triggers arising from the online burst all-sky search should immediately be examined to look for any temporal and spacial correlation with either class of trigger. A galactic hyperflare would likely manifest as an extraordinarily bright and unusually soft-spectrum short gamma-ray burst and would generate mass interest in the astronomy community. We will manually monitor (e.g.,) the GCN notices and other notification systems such as Atel for indications of Galactic hyperflares. A binary neutron star coalescence which is near enough to stand a hope of observing an associated burst signal will, of course, be detected by the low-latency GSTLAL/MBTA BNS search (see 3.2). A nearby merger would likely also be detected by the online burst analysis. In both cases, there will be a well defined event time and, potentially, a sky localisation with which to correlate online search triggers.

Galactic Magnetar Hyperflares

Following this low-latency check, and in the case of a confirmed galactic magnetar hyperflare, a targeted analysis will be performed using X-PIPELINE [5, 222], a pipeline designed for GW follow-up of astrophysical triggers with known sky-location and a specified time window in which a GW signal may be expected (the "on-source"). X-PIPELINE has been used in a number of LSC-Virgo observational results papers in GRB and neutrino-triggered searches [356, 357, 359, 358, 360, **?**]. In its standard configuration, X-PIPELINE is designed to target short-duration, low frequency bursts in the most sensitive regime of LIGO and Virgo. For magnetar analyses, however, we require sensitivity up to f-mode frequencies (i.e., $\sim 1 - 4\,\mathrm{Kz}$) and long-duration ($\sim 100 - 1000\,\mathrm{s}$) transients. We will, therefore, aim to extend the X-PIPELINE analysis to at least 4 kHz. Studies are currently underway to optimize the selection of time-frequency resolutions used in the X-PIPELINE analysis in order to provide some overlap in sensitivity to long-duration signals with a separate analysis targeting long-duration bursts. The main tool for targeting these long-duration signals will be the STAMP-GRB pipeline [337] which has previously been used to search for long-lived GW signals associated with GRBs [347]. No significant modification to the STAMP analysis is expected to be required.

Detection candidates for short-duration signals associated with extraordinary magnetar events will be followed up with parameter estimation tools currently under development such as BAYESWAVE and LAL-INFERENCE BURST; see Sect. 3.1. This will allow reconstruction of the impinging waveform and basic spectral analysis, which will be important for mode identification and astrophysical interpretation. Basic spectral analysis and reconstruction of the frequency evolution of long duration transients can and will be addressed by the STAMP analysis.

In summary, we propose the following in the event of a hyperflare from a galactic magnetar:

- **Low-latency analysis (∼minutes–hours):** Initial check for temporal and spatial coincidence of triggers arising from the online burst analysis with the electromagnetic or GW trigger.

- **Higher latency analysis (∼week–month):** Upon confirmation of the galactic and extraordinary nature of a magnetar event, dedicated X-PIPELINE and STAMP analyses will be performed *manually*[10]. X-PIPELINE will be used to target short-duration bursts from 10's of Hz up to $\sim 4\,\text{kHz}$ and will be tuned to maintain sensitivity to longer duration bursts and provide some degree of overlap in sensitivity with the STAMP-based tools which will be used to target exclusively long-duration bursts. Short-duration candidates to be followed-up and characterized with burst parameter estimation (PE) tools.

Post-merger Signals From Binary Neutron Star Coalescence
In the event that the *offline* weekly ahope CBC analysis confirms a BNS inspiral detection initially made by the online GSTLAL/MBTA CBC analysis, we will initiate a burst-led follow-up to characterize the post-merger scenario and detect signals from a surviving NS remnant.

The procedure to target post-merger signals following a BNS inspiral detection is very similar to the proposed magnetar analysis. The notable difference in this case is that we will already have a GW detection (the inspiral) and we wish to infer the late-time behaviour of the same system (prompt vs delayed collapse, measurement of post-merger oscillation frequencies) from the evolution of the coalescence signal. In the case of a BNS inspiral, therefore, the burst parameter estimation tools will be deployed as a follow-up to the inspiral detection and *not* as a follow-up to a short-duration burst detection via e.g., X-PIPELINE. Long-duration bursts arising from instabilities or quadrupole deformations in stable post-merger remnants are somewhat separate from the initial coalescence, have a less certain start time and are not easily targeted by the parameter estimation tools currently in development. As with the magnetar event then, long-duration bursts from stable remnants will be targeted by the STAMP-GRB pipeline.

In summary, we propose the following in the event of a hyperflare from a galactic magnetar and/or a BNS inspiral detection:

- **Low-latency analysis (∼minutes–hours):** Initial check for temporal and spatial coincidence of triggers arising from the online burst analysis with the electromagnetic trigger in the case of a Galactic magnetar hyperflare, or coincidence with a low-latency GSTLAL/MBTA GW trigger in the case of binary neutron star coalescence.
- **Higher latency analysis (∼week–month):** Burst PE tools will be used to characterise the short-duration, high-frequency merger/post-merger signal. The STAMP-pipeline will be used to target any later long-duration bursts from stable remnants.

3.12.4 Result Validation Plan

Galactic Magnetar Hyperflare Search

The first stage in the search for GWs associated with an extraordinary Galactic magnetar event, prior even to launching the offline STAMP (targeting longer duration signals) and X-PIPELINE (targeting shorter duration signals) analyses, will be examination of any available all-sky burst triggers from around the time of the event. If the all-sky analysis provides a detection candidate, we certainly expect the *targeted* analyses to see the same event; we therefore expect consistency between the basic parameters of the signal (e.g., characteristic times and frequencies), insofar as the time-frequency resolution and basis functions permit.

In the event the targeted analyses provide a candidate, but all-sky analyses with sensitivity in overlapping parameter space (as determined from the characteristics of the targeted search trigger) do not,

---

[10] Given the rare nature of the targeted events, we do not feel that investing in automated infrastructure is necessiated

the significance of the candidate will be assesed to determine whether the lack of a matching all-sky trigger is consistent with the plausible amplitude of the signal and the thresholds and tuning choices applied in each pipeline.

Finally, in the event that a trigger arises from a targeted analysis with no possible overlap with existing all-sky analysis (for example, the high-frequency all-sky analysis may run at higher latency than the targeted high-frequency magnetar analysis), cross-validation will take place via studying the results of following up the targeted analysis trigger with parameter estimation tools in the form of a dedicated BayesWave follow-up and/or the production of a CWB coherent event display[11].

In all cases, a coherent event display will be used to more thoroughly study the trigger and reconstructed signal.

Post-merger Signals From Binary Neutron Star Coalescence

The validation procedure for any long-duration burst counterpart to a BNS merger will be identical to that proposed above; checking for consistency with the significance and time-frequency content of candidate events between all-sky triggers and coherent event display follow-ups.

The results of the parameter estimation follow-up study targeting the immediate short-duration burst at and after the merger will be compared with a) available high-frequency all-sky triggers and b) the output of a CWB coherent event display. There should be consistency in the bulk features (e.g., root-sum-squared amplitudes in each detector, characteristic times & frequencies) between the parameter estimation output and CWB results.

3.12.5 Publication Plan

Galactic Magnetar Hyperflare Search: The analysis of GW data near to the time of a galactic hyperflare will be of high interest even in the event of a null result. Our goal, therefore, will be to have a complete paper draft within 3 months of the event. We provide specific plans for the cases of confident, marginal and no detections below.

Evidence for detection Evidence for detection would be comprised of a temporally and spatially coincident (with the hyperflare) trigger with significance estimated at $> 3\sigma$, accompanied by a robust waveform reconstruction (i.e., narrow posterior measurements, free of artefacts) and spectral analysis. The target significance for a confident detection will be the usual 5σ and efforts will be made to estimate to this level for triggers at $> 3\sigma$.

Marginal evidence for detection Marginal detection would be comprised of a detection candidate with significance at 2–3σ. Reconstructed waveform polarizations (and frequency content) would be presented but with less emphasis than for a confident detection.

No detections. In the event of no detection, the publication plan would be contingent on an GW energy upper limit:

- For GW energy upper limits comparable (i.e., within a factor $\sim$unity) to the isotropic equivalent electromagnetic energy, the result would be considered a significant non-detection, given the extraordinary nature of the event and the precendent for interesting magnetar analyses.
- If the GW energy upper limit does not compare favorably to the electromagnetic emission, the result would not necessitate its own publication and a simpler non-detection statement in either an end-of-run all-sky or GRB-specific publication would suffice.

---

[11] the web page produced by CWB detailing reconstructed parameters, time–frequency maps, reconstructed detector responses, likelihood time–frequency maps and reconstructed parameter skymaps

Post-BNS Follow-up: Any detection of high-frequency power associated with a BNS inspiral signal would be a major discovery. Papers reporting confident or marginal detections must be made on the same timescale as the deep CBC parameter estimation studies. The following criteria refer to the short-duration, high-frequency burst associated with the immediate post-merger regime of the coalescence signal:

Confident detection. Confident detection of a BNS post-merger signal would be comprised of clear evidence for a high-frequency oscillation component to the BNS merger in the 1–4 kHz band. 'Evidence' here would be comprised of a combination of Bayesian model selection statistics (i.e., an odds ratio, whose significance threshold is TBD from on-going parameter estimation studies), signal amplitude posteriors inconsistent with zero amplitude *and* a clear dominant mode in the signal frequency posterior probability density function. This scenario would consitute extremely high profile science and would allow accurate spectral analysis of the post-merger oscillation with potentially important implications for the neutron star EOS.

Marginal detections. Significance measures for marginal detection here are comprised of smaller values for Bayesian model selection statistics, amplitude posteriors peaked away from, but not inconsistent with, zero and some indication of at least one distinct mode in the signal frequency posterior PDF.

Even a marginal detection of high-frequency GW power in close temporal coincidence with a BNS inspiral signal would be an extremely significant discovery, confirming that the NS EOS is sufficiently stiff to, at least temporarily, survive gravitational collapse. In addition, it it is possible that estimation of the *dominant* post-merger oscillation frequency (also the most robust, Universal characteristic of these signals) will still be made with reasonable accuracy at low SNR. We do, therefore, plan for publication in the event of marginal detection.

No detections. With regards to the immediate short burst following a BNS signal, the significance of a null detection of a post-merger signal is marginal. We do not, therefore, envisage a dedicated publication in the event that this signal is not observed. GW amplitude or energy upper limits, however, may be of more interest in light of the potential for stable magnetar formation. If one assumes that a stable magnetar is formed, it may be possible to constrain some models for the magnetic field-induced quadrupole moment of a nascent magnetar. This would be particularly interesting in the event of an accurate distance measurement from e.g., host galaxy identification[12]. A dedicated publication reporting the non-detection of a short duration post-merger burst and the GW energy upper limit for the long-duration signal is envisioned. We do, however, expect that a remark would be made in a dedicated inspiral parameter estimation publication reporting that this analysis was conducted but found no evidence of a post-merger signal.

The publication plans and criteria associated with the search for a long-duration transient following a BNS inspiral detection are almost identical to those for long-duration transients in magnetar hyperflares. The exception is that the publication plan for the null detection scenario is identical to that for the short-duration burst: a remark in an inspiral-specific publication that the search was conducted but no evidence for a signal was found.

3.12.6 Technical Requirements and Development Plan

These searches will use X-PIPELINE in a similar configuration as for SNEWS-alert analyses, with an extended frequency range and some additional choice of time-frequency resolution; a modest amount of testing and validation is required to sanity check the performance of the algorithm at higher frequencies

---

[12] This would allow for a reasonable constraint on the GW energy emitted by the source and not simply the GW fluence at the Earth

and additional time-frequency resolutions (expected to be one additional resolution). The STAMP-based analyses are not expected to require further development.

The burst parameter estimation tools currently under development and described in section 3.1 will be used to follow up candidate detections for short-duration bursts. In addition to the on-going development and performance characterization[13] of these parameter estimation pipelines, some additional tuning will be desirable to target the high-frequency and late-time evolution of the BNS signal and, potentially, restrict the sky area searched to that which is consistent with the inspiral detection.

The bulk of development work is expected to be the development and use of appropriate waveform simulation infrastructure. For the short-duration post-merger bursts (where the detector antenna patter does not change significantly over the duration of the signal), the codes and infrastructure developed for the NINJA analyses should prove sufficient for assessing the performance of the analysis. The set of short-duration merger/post-merger waveform simulations used in the study in [389] is available for use in these studies and further simulations from other groups and collaborators may become available in the coming months. Longer duration signals where the antenna pattern changes significantly over the course of the signal will require some additional development and validation. Simulations of these longer duration signals will be based on astrophysically motivated analytic models such as those described in [390] and [391], as well as more ad hoc simulations to explore and characterize the sensitivity to non-stationary frequency content and signals whose frequency content tracks observed QPOs.

3.12.7 Resources

The expected number of events of direct interest to the NS group (Galactic hyperflares and BNS inspirals) during O1 is likely to be < 1. As stated, the analyses for both events is very similar: targeted short-duration burst analysis with X-PIPELINE and a long-duration burst analysis with STAMP with some overlap in sensitivity. Targeted burst parameter estimation analyses will be performed for both sources where appropriate.

Review Review requirements are minimal. The proposed pipelines (X-PIPELINE, STAMP) are reviewed and mature. Only configuration changes to the default used in X-PIPELINE will be required.

Similarly, the parameter estimation follow-up tools, BAYESWAVE and/or LIB are already to be reviewed for various other analyses; again, the only difference between those analyses and this would be in the configuration since it is possible that there may be a high frequency detection candidate.

Development No significant technical development should be necessary. The only changes from other, mature analyses will be in the time window and frequency range analysed. Table **??** details the development milestones and timelines involved.

Detector characterization: Neither of these proposed searches has any specific detector characterization / quality requirements above those of other triggered burst searches.

Calibration: It should be noted, however, that we do expect to analyze frequencies up to at least 4 kHz; *we therefore make the explicit request for calibrated data up to at least 4096 Hz with uncertainties matching those achieved for previous magnetar analyses in S5 onwards*. These searches are not expected to run at low-latency so this calibration requirement does not seem likely to pose additional overhead to the planned second pipeline all-sky analysis (which extends up to 4 kHz).

EM alerts: There is no immediate plan or obvious requirement that the results of these low-latency need to be communicated to electromagnetic partner observatories: both Galactic magnetar hyperflares and GW detections from binary neutron star coalescence are rare and will be of sufficient intrin-

---

[13] See https://wiki.ligo.org/Bursts/AllSkyPE

sic astrophysical importance that the results of the analyses proposed here will not contribute strongly to a decision to study the source further.

Finally, it is worth emphasising that we do not perceive any conflict or overlap with other externally triggered analyses. While it is likely that a Galactic hyperflare will initially trigger the online GRB burst analysis, such an event is unlikely to be included in a final GRB-specific publication. Furthermore, the analyses proposed here cover a wider and more appropriate signal space for magnetar triggers. Similarly, in the case of other related triggers arising from e.g., neutrino detectors or radio observatories, we do not propose any modifications to the search plans in order to maximise the signal space coverage.

3.13 Gravitational-Waves from Galactic and Near-Extragalactic Core-Collapse Supernovae

3.13.1 Abstract

The next supernova to explode in the Milky Way will be the most important astronomical event of the century. With their advanced-generation detectors, the LIGO Scientific Collaboration and the Virgo collaboration will for the first time have a good chance of picking up the gravitational waves (GWs) emitted if the supernova involved the collapse of a massive star and is a so-called *core-collapse supernova*.

Once a massive star (about $10\times$ the mass of the Sun or more) exhausts its fuel, its core collapses to a hot proto-neutron star. The proto-neutron star cools by emitting neutrinos, a kind of very weakly interacting elementary particle. A shock wave is launched from the proto-neutron star and plows through the stellar mantle. When it breaks out of the star's surface, it lights up the star in a supernova explosion. It is theorized that the shock wave is powered by the absorption of a small fraction ($\sim$10%) of the emitted neutrinos within the first second of the proto-neutron star's life. The neutrinos heat up the material above the proto-neutron star. This drives convection (think: hot buoyant bubbles rise, cold bubbles sink). This convection – a bulk motion of large amounts of material – leads to the emission of GWs. If the star's core was spinning, then also a burst of GWs that lasts for a few milliseconds is emitted when the proto-neutron star is first formed.

Neutrinos were detected from supernova SN 1987A, which exploded in the Large Magellanic Cloud, a neighbor galaxy of the Milky Way. Back then, no detector with sufficient sensitivity to detect GWs from a supernova was operational. Theoretical predictions of the GW emission suggest that the Advanced LIGO detectors at their planned sensitivity for the O1 science data run will have a chance of detecting the GWs from a core-collapse supernova in the Milky Way. The projected improved sensitivities in the 2016-17 (O2) and 2017-18 (O3) runs will possibly allow detection out to the Magellanic Clouds and the establishment of more stringent upper limits on the strength of GW mission for supernovae occurring in the Andromeda galaxy and even more distant galaxies in the Virgo cluster of galaxies. The chance of detection for a supernova in the Milky Way will be particularly high, because the neutrinos from the supernova will be caught in neutrino detectors. They will provide a very clear "timestamp" for when to look in the Advanced LIGO data.

Observing and characterizing GWs from the next nearby core-collapse supernova will open up a completely new observational window. It will allow us to measure, for example, how rapidly the progenitor star's core was spinning and how strong convection was right before the explosion. GW detection will thus help elucidate much of the still unknown physics that is going on as a star transitions from collapse to explosion.

In the search we propose here, we plan to rapidly follow up supernova neutrino event candidates sent by the SuperNova Early Warning System (SNEWS, `http://snews.bnl.gov/`) with model-independent GW searches that use constraints on GW emission time and sky location. Within days of a possible detection, we will characterize the GW signal and extract astrophysical information.

3.13.2 Scientific Justification

The next galactic (core-collapse) supernova ([CC]SN) will be the most important astronomical event of the century and the world will look to the LIGO Scientific Collaboration and the Virgo Collaboration for statements about its GW emission. Multimessenger observations in photons, neutrinos, and gravitational waves of the next nearby CCSN will revolutionize our understanding of massive star structure and angular momentum distribution, of core collapse dynamics, of the still uncertain CCSN explosion mechanism, of explosive nucleosynthesis and mixing of synthesized elements in the explosion, and of fundamental physics such as the equation of state of nuclear matter and neutrino interactions.

GWs are produced by bulk aspherical accelerated motion of mass-energy and in the CCSN context are thus a direct probe of the uncertain degree of asymmetry of the supernova engine. GWs are expected to be emitted by a broad range of processes (e.g., [392, 393]) many of which can be directly associated with particular explosion mechanisms. For example, the GW signal from neutrino-driven convection can be connected to the neutrino mechanism or the strong signal from rotating core collapse can be linked to magnetorotational explosions (e.g., [394, 395]). GWs from rotating core collapse can be used to measure the angular momentum of the collapsing core [396, 397]. GWs from neutrino-driven convection and the standing accretion shock instability can be used (1) to infer the moment of the onset of explosion and (2) to constrain the structure of the nascent neutron star and in this way put constraints on the nuclear equation of state (in combination with neutrino information; e.g., [398, 399, 400]). An abrupt end of GW and neutrino emission would unambiguously herald the formation of a black hole (e.g., [401, 402]).

Much work has been directed toward understanding GW emission from CCSNe and the detectability of the waves at various levels of sophistication: theoretical estimates of optimal or angle-averaged signal to noise ratios (SNRs; e.g., [393]), sensitivity estimates based on real re-colored noise with projected O1 sensitivity and a two-detector network [403], and preliminary detection upper limits from the ongoing search for GWs from distant CCSNe in S5/A5/S6 data [404]. These studies show that a galactic CCSN at a distance of $1 - \mathrm{few\,kpc}$ may be detectable by an L1–H1 O1 detector configuration. In the planned O2 and O3 science runs, improved detector sensitivity may allow detection throughout the Milky Way even if pessimistic emission scenarios are considered. A non-detection of a confirmed core collapse event, even during O1, would put highly significant constraints on the degree of asymmetry of the supernova engine. Detection of a CCSN at Small/Large Magellanic Cloud (SMC/LMC) distance will likely require projected the sensitivity projected for later science runs and optimistic emission scenarios such as rotating core collapse, but will still yield astrophysically highly interesting upper limits in the case of a non-detection.

MeV-energy neutrino and GW emission in a CCSN will set in within a few milliseconds of each other (e.g., [405]). A galactic (extragalactic) CCSN at $10\,\mathrm{kpc}$ ($100\,\mathrm{kpc}$) will result in $\mathcal{O}(10,000)$ ($\mathcal{O}(100)$) neutrino events in worldwide detectors. Limited directional information, at an error level of $5^\circ - 20^\circ$, will be available from neutrino observations [406]. Information about these neutrino detections will be disseminated with low latency by the SuperNova Early Warning System (SNEWS, `http://snews.bnl.gov/`, [407]) and can be used to set a tight ($\sim 4\,\mathrm{min}$, [403]) on-source window for a rapid-follow-up GW search. Such a short on-source window dramatically reduces the background of all-sky/all-time search and searches relying on information from electromagnetic detection of the CCSN.

We propose an online SNEWS-triggered low-latency, rapid-follow-up search for GWs from nearby CCSNe with X-PIPELINE [5]. The motiviation for the online search is to provide alerts to the scientific community containing the GW-estimated sky position, ahead of the light emission of SN that is expected up to a day later. Subsequent to a SNEWS trigger with or without GW detection candidate from the online search, we will carry out deeper offline analyses with X-PIPELINE and Coherent WaveBurst 2G (CWB2G, [408]). Upon detection, effort will be directed toward waveform reconstruction and parameter estimation to determine baseline signal properties such as central frequency, time-frequency evolution, and to extract astrophysical parameters such as the most likely explosion mechanism and the amount of angular momentum present in the collapsing core.

This search is motivated and informed by the ongoing S5/A5/S6 optically-triggered search for GWs from distant CCSNe [404] and by the sensitivity study of [403]. CCSNe are prime galactic science for Advanced LIGO and Advanced Virgo. CCSNe occuring at distances beyond the Milky Way are unlikely to be detectable with projected two-detector O1 sensitivity on the basis of current GW signal predictions from detailed multi-dimensional simulations (e.g., [399, 409, 393]). However, extreme emission scenarios associated with longer-lasting ($\sim$ one to two seconds) barmode instabilities (e.g., [410]) or fragmenting accretion disks [346] could be constrained for more distant CCSNe [403].

In addition to the SNEWS-triggered search in O1 and subsequent science runs, we are also working on two additional projects:
(**1**) the development of a search for setting upper limits for GW emission from extragalactic CCSNe exploding within $10 - 20\,\mathrm{Mpc}$ of Earth, using optical triggers. This search will be carried out on data from O1, O2,

and O3. Its motivation is to exclude the most extreme GW emission scenarios for CCSNe at much higher sensitivity than possible with S5/A5/S6 data [404].
(**2**) the development of a deep search for coincident sub-threshold GW and neutrino events from potential (dark) local-group CCSNe. This search will first be carried out on archival data/triggers from S5/A5/S6 LIGO/Virgo data and triggers from the IceCube [411, 412], LVD [413, 414], and Borexino [415, 416] neutrino detector collaborations.

3.13.3 Search Description

SNEWS-triggered search: The goal of the SNEWS-triggered analysis is to provide a directed online search for a GW burst associated with a galactic or SMC/LMC CCSN. The directed online search will be carried out with X-PIPELINE [5] and will complement the all-sky online analysis with Coherent WaveBurst 2G (described in the all-sky burst search proposal, [350]). After a SNEWS trigger, we will carry out deeper offline analyses with both X-PIPELINE and CWB2G, which will complete within a few days. The detection statements of the two pipelines will be combined as in the initial LIGO CCSN search [404]. An important issue for this search is that there is a significant probability that at the time of a SNEWS trigger only one detector will be online. Extreme care will be required to vet detection candidates and to take full statistical advantage of the information provided by SNEWS. It is understood that poor data quality around the SNEWS trigger may prevent a strong detection statement. Single-detector searches with and without external triggers have been carried out before by Virgo [417, 418] and we will build upon the experience gained by this previous work. A preliminary analysis using multivariate analysis and X-Pipeline on a single detector trigger example demonstrates some improvement in the efficiencies, and work will continue to charaterize this effect [419].

A nearby CCSN will produce a prominent signal in the global array of neutrino detectors such as Super-Kamiokande [420, 421], Borexino [415, 416], and LVD [413, 414]. In preparation for such an event, the neutrino community has an established alert system known as SNEWS [407]. SNEWS will provide an automated email alert of "GOLD" events to registered users with an estimated latency of five minutes or less.[14] The best pointing accuracy will be approximately $5^\circ - 20^\circ$ from Super-Kamiokande [406], but this information may not be immediately available at the time of the alert. We therefore envision an all-sky scan for GWs at the rapid-follow-up stage and incorporate directional information in a subsequent offline search.

The SNEWS-triggered online search will be conducted online using X-PIPELINE [5, 222] and offline with X-PIPELINE and CWB2G. Both pipelines can operate in a single-detector mode in the case only one detector is operational at the time of the SNEWS trigger. For the single-detector case, we will incorporate the lessons learned and experience gained from previous single-detector searches by Virgo [417, 418]. X-PIPELINE [5, 222] is a pipeline designed for GW follow-up of astrophysical triggers and has been used in a number of LSC-Virgo observational results papers in searches triggered by gamma-ray bursts and high-energy neutrinos [356, 357, 359, 358, 360, **?**]. The X-PIPELINE analysis will be launched automatically by the same online monitoring software used for the GRB search. It will analyze a 4-minute block of data around the alert time for associated GW bursts up to $\sim$4 kHz [392, 393]. Background estimation will be performed using all data within three hours of the time of the alert; this will be sufficient to provide up to $\sim 5 \times 10^4$ background trials. Half of these are used to tune the background rejection tests on the local background to maximize sensitivity. The other half will allow us to estimate the significance of any candidates to greater than 3σ. The X-PIPELINE analysis is fully automated, including tuning and candidate identification, and can run without any human intervention.

In the case of a SNEWS alert, the X-PIPELINE analysis will be rerun offline as needed after any significant changes to calibrations, data quality, or as further information on the supernova is received. In

---

[14] The false alarm rate for SNEWS alerts is estimated to be less than one per century [407].

particular, we forsee re-running the analysis as improved sky position information becomes available and a more sensitive search (perhaps by up to a factor of two, e.g., [224]) becomes possible.

For the directed CWB2G offline search, we will use the 4 minutes of data to perform the search and 3 hours of data before the SNEWS trigger to estimate the background. We will present a detection statement with 4-sigma confidence within two days of the SNEWS trigger. The actual duration of individual analysis jobs will be less than 15 minutes (estimated from the online speed of CWB2G, [408]), but preparation and checking will require human input and intervention. The CWB2G offline search will be run in multiple configurations using parameter settings optimized to best resolve the different expected GW emission mechanisms. The base CWB2G code will be run with a sky mask that will take advantage of the almost exact knowledge of the sky location of the CCSN at the time of the offline analysis. In the unlikely case no or only poor direction information is available at the time of the CWB2G offline analysis, the sky mask will be dropped or tuned accordingly. CWB2G without and with sky mask is expected to be fully reviewed by the end of the 2015 calendar year. The CWB2G team is considering additional developments for the directed galactic/near-extragalactic CCSN search proposed here: (**1**) A ring-shaped sky mask including all the pixels that have time delays between the interferometers consistent with the information on the CCSN sky location. For this improvement, code freeze is expected to happen at the end of June 2015. (**2**) A CWB2G plug-in that implements the noise reduction procedure of [422] under development at the University of Texas at Brownsville. This procedure employs a method of noise reduction known as Harmonic Regeneration Noise Reduction (HRNR) algorithm known from audio processing (e.g., [423]). If mature, a review of this procedure will be requested at the end of June 2015. In fall 2015 and spring 2016, we will explore two more modifications, namely using BayesWave as a follow-up to CWB2G and modifying the clustering procedure of triggers in the time frequency plane so that the search becomes more sensitive to waveforms with partially disconnected time frequency islands. The goal is to have these developments reviewed before the start of O2 if they turn out to yield an improvement.

Parameter Estimation Follow-Up: Upon an identification of a detection candidate by X-PIPELINE (or in the offline-only CWB2G analysis), we intend to carry out an array of rapid follow-up analyses to infer general signal properties and physics/astrophysics from the detection candidate.

- We will characterize the detection candidate waveform's central frequency and time-frequency content and reconstruct the waveform using tuned configurations of the BAYESWAVE [365], and CWB2G [408] algorithms. We will also utilize the LALInference_Burst (LIB, [366]) pipeline for parameter estimation, using the sine-Gaussian signal model. It is expected that the targeted capabilities of LIB, BAYESWAVE, and CWB2G will be reviewed by the start of O1.

- We will employ the Supernova Model Evidence Extractor (SMEE, [395]) to find the CCSN explosion mechanism that is best fit by the detection candidate. Improvements to SMEE are currently being implemented, but a full review until the start of O1 is unlikely. Should there be a detection candidate, the burst group will carry out an ad-hoc review of SMEE and expedite a more detailed review for subsequent follow-up papers. We expect a formal review of the SMEE pipeline to start at the end of O1 and be completed by the start of O2.

- We will use a multivariate regression model for waveform prediction by Engels *et al.* [397] in combination with MCMC sampling to estimate the total angular momentum and the degree of differential rotation in the progenitor star's inner core. We will use the GW signal catalog of Abdikamalov *et al.* [396] as the bases for our waveform model. This method has not been reviewed and a review until O1 is unlikely. Should there be a detection candidate that SMEE identifies as rotating core collapse, the burst group will conduct an ad-hoc review of this approach and expedited a more detailed review for subsequent follow-up papers.

- We will employ a novel time-frequency (spectrogram-based) analysis to infer physical parameters such as the postbounce accretion rate, the structure of the proto-neutron star, and the strength of neutrino-driven turbulence. This analysis is presently under development and is unlikely to be mature by the start of O1.

EM-triggered search to establish upper limits for extragalactic CCSNe: This search will be along the lines of the current S5/A6/S6 search [404], but will be carried out with CWB2G and with the same noise reduction techniques implemented for the SNEWS-triggered search. We will select CCSNe that have well established on-source windows from late optical non-detections / early optical detections by supernova surveys. This search is at the planning stage and more details will be worked out by 2016-17 as O3 draws closer. We do not plan a parameter estimation component for this search. In the 2015-2016 time frame covered by this White Paper, we are conducting a sensitivity study for extragalactic CCSNe using O1/O2/O3 expected sensitivity curves.

Neutrino-triggered search for (dark) local-group sub-threshold GW bursts from CCSNe: A new effort aiming to perform a joint search for GWs and low-energy neutrinos from CCSNe is under development. This search intends to use neutrino data from individual neutrino detectors (and thus not available as SNEWS alerts) in order to improve the search sensitivity (in terms of distance reach and overall model parameters). This is achieved by lowering thresholds of individual detectors. The effort looks forward in establishing a working relation with the low-energy neutrino detector community, preparing us for the next nearby CCSN and developing joint analysis strategies that will maximize the science output of such astrophysical events. A proposal –and a corresponding Memorandum of Understanding– for such a joint search was approved in 2014 by the LIGO-Virgo collaborations and the Borexino [415, 416], LVD [413, 414], and IceCube [411, 412] neutrino detector collaborations. The effort so far has been focusing on developing common simulation engines for the three neutrino detectors in order to establish their joint sensitivities. Adding GW simulations into this framework as well as analyzing archival data from the 2005-2010 science as a prototypical search are the next steps in this. We expect this effort to ultimately lead during later observing runs of Advanced LIGO-Virgo to a low-latency search for CCSNe that will involve both the neutrino and GW data.

3.13.4 Results Validation Plan

In the event a galactic CCSN occurs and the X-PIPELINE search is triggered by a SNEWS alert, we will immediately follow up the low-latency analysis to validate its results:

- We will carry out offline analyses with X-PIPELINE and CWB2G that take information about the location of the CCSN into account. This offline follow-up is integral part of our search plan (see §3.13.3).
- We will compare the results of our search with the untriggered all-sky CWB2G, EXCESSPOWER, and omicron+LIB.
- We already plan to employ multiple parameter estimation approaches (see Search Description, §3.13.3). These generally complement each other, but also have some overlap (e.g., in signal reconstruction). We will use overlapping capabilities for cross-validation. The SMEE pipeline exists in multiple implementations and we will cross-validate results of each independent implementation.

3.13.5 Publication Plan

A galactic/near extragalactic CCSN will be an extremely important astrophysical event. A GW analysis will be of very high interest even if no signal was detected. Our goal will be to have a complete paper draft within

3 months of the event. The target journal would be a high-profile journals such as *Physical Review Letters* or *The Astrophysical Journal Letters*. We envision a more extendend report within 6 months of the event in a subsequent paper with more details on parameter estimation and constraints on astrophysical parameters and fundamental physics. The target journal for this will be *The Astrophysical Journal* or *Physical Review D*.

3.13.6 Technical Requirements and Development Plan (O1 SNEWS-triggered search)

SNEWS-triggered online search: X-PIPELINE will be used. The core pipeline and the online triggering system were used to do fully autonomous GRB-triggered searches in S6/VSR2,3. Only small modifications are expected to be needed for the SNEWS triggered search, with minimal review burden. The major milestones for the search already accomplished:

- We have enabled GraceDB to receive and parse SNEWS alerts.
- In terms of parameters, we anticipate applying the whole range of X-PIPELINE likelihoods and 7 different FFT lengths. In order to obtain a 3 sigma detection in a timely matter, the calculations of some of the likelihoods have been optimized to run in 5 percent of the previous time. In addition, we are maximizing the use of internal time slides and minimizing the use of external time slides in an attempt to get the initial all-sky runtime below 3 hours. We are not quite there, but we expect to get there within Q3 2015 (before 01).
- We have implemented code which appropriately calculates the H1-L1 time delays that will cover all potential source locations at the initial alert with no sky-position information.
- We have investigated what X-PIPELINE parameters are optimal to use in the case of single detector data.

The major X-PIPELINE milestones yet to be accomplished:

- Implement and test sending alerts of detected GW events for EM follow-up, with emphasis on the estimated sky position.
- Determine if multivariate analysis (MVA) improves the sensitivity to supernova waveforms. The MVA extension of X-PIPELINE has been determined to be useful for GRBs and is already part of the GRB search plan [424].
- Quantify the sensitivity loss from the single detector runs and the all-sky run compared to the optical search.

Offline follow-up: Both X-PIPELINE and CWB2G will be used. No further modifications of X-PIPELINE are required. The version of CWB2G to be run will be very similar to the reviewed all-sky version and similar search methodology as in the optically-triggered S5/A5/S6 CCSN search will be used to combine detection or upper limit statements of the two pipelines [404]. The main difference on the CWB2G side compared to the S5/A5/S6 optical search will be a shorter on-source window and the use of data before the on-source window for background estimation. Development in CWB2G for the dedicated follow-up will include the following:

- The implementation and testing of a ring-shaped sky mask. This new plug in of CWB2G has already been implemented and testing has started. The case of a SNEWS-triggered event is trivial to handle, but an evolving sky mask in a long on-source window as expected for the O1/O2/O3 EM-triggered CCSN search will require more development work. We do not do not expect a full review of the latter feature to complete until O2.

- The characterization and optimization of CWB2G for the case of a single-detector event (which is quite likely in O1 and O2). CWB2G has been run in this mode, but no statistical studies have been performed to determined how many σ in detection confidence could be obtained for this kind of search (in combination with a neutrino trigger). We expect this to be a key area of work in Q3/Q4 2015. We will incorporate lessons learned and experience gained from single-detector searches by Virgo [417, 418].

- The implementation and testing of a CWB2G module for the noise reduction procedure developed by the University of Texas at Brownsville group [422]. This method involves a two step noise reduction approach based on accurate spectral estimation. The Matlab/C++ codes are fully modular and can be integrated with the pipeline with minimal intervention. The goal of this analysis is to enhance the efficiency of the CWB2G search for GWs from CCSNe. This procedure has been investigated for 2 months and it appears that while the estimated signal to noise of the gravitational waves is improved, this is also the case for the non-Gaussian transients (glitches) in the background noise. We have derived receiver operating curves for several combinations of the internal CWB2G parameters and Q3/Q4 2015 will be used to understand if there are combinations that allow to consistently improve sensitivity with respect to the CWB2G version without noise reduction.

Parameter Estimation Follow-Up: At the time of writing, the LIB, BAYESWAVE, SMEE parameter estimation / model selection, and waveform reconstruction with BAYESWAVE and CWB2G have reached partial maturity. Some additional development and simulation work is needed to ready these parameter estimation pipelines for follow-up of a potential SN trigger.

The BAYESWAVE pipeline currently makes assumptions about the polarization of the detected signals and the SMEE pipeline still requires further development (inclusion of multiple detectors, different kinds of principle components and signal models etc.).

Significant work will be required to ready multivariate/MCMC angular momentum estimation [396, 397], signal reconstructions/characterization, and advanced parameter estimation using GW spectrograms that will go after progenitor parameters such as the postbounce accretion rate or the compactness parameter [425]. The goal is (1) to run preliminary versions of these algorithms and ad-hoc review them if there is detection in O1 and (2) bring these algorithms to full maturity and review them by O2 (or O3 in some cases).

The major milestones for parameter estimation are:

- Further test, bring to full maturity, and review basic signal characteristic extraction and signal reconstruction with the optimal set of the LIB, BAYESWAVE, and CWB2G codes.

- Deploy a multi-detector SMEE pipeline that can be run on X-PIPELINE and CWB2G detection candidates. Compare two independent implementations (Matlab, Python) of SMEE and complete the second SMEE methods study. We expect multi-detector SMEE to be mature but not reviewed by O1. It will be run only if there is a detection candidate in O1 and will then receive an ad-hoc review. A full review of SMEE is expected to start at the end of O1 and be completed by O2.

- Ready a multivariate regression / MCMC pipeline for angular momentum estimation [397, 396] that can be run on X-PIPELINE and CWB2G detection candidates. We expect a preliminary version of this pipeline to be ready to be run in case of a detection in O1 (and to be reviewed in an ad-hoc review) and a reviewed mature version to be available by O2.

- Develop and test an advanced parameter estimation pipeline based on GW spectrograms to infer progenitor physics such as postbounce accretion rate and compactness from neutrino-driven CCSNe. We plan for an initial version of this pipeline to be ready by O2 and a fully mature and reviewed one to be completed by O3.

3.13.7 Resources

A galactic core collapse supernova would be an event so rare and so important that it would be imperative for the LSC and Virgo to make available all necessary resources for a complete and rigorous analysis; the scientific world would be watching. In particular, we would expect that the requested computing resources would be made available *even if this requires temporarily pre-empting non-SNEWS analyses*. However, time-critical low-latency searches should be allowed to continue. Also, computational requirements for additional supernova analyses not already discussed in this search plan will need to be justified and will require an ad-hoc review. This is to ensure that other searches are not unreasonably pre-empted by unplanned supernova work.

The various parameter estimation and waveform reconstruction pipelines that would be used in a follow-up analysis of a SN trigger are still under development and require computing resources for testing and simulation work in the 2015-2016 time frame.

The O1/O2/O3 and the neutrino-triggered searches are still under development. The O1/O2/O3 sensitivity study using X-PIPELINE for an optically/SNEWS triggered SN event will require computing resources for simulations. The neutrino-triggered (extragalactic) search will need minimal resources in the 2015-2016 timeframe covered by this White Paper.

3.14 Search for transients in coincidence with Fast Radio Bursts

3.14.1 Abstract

Since the publication in summer 2013 of four Fast Radio Bursts (FRBs) identified in Parkes Telescope data [426] there has been considerable scientific interest in these millisecond-scale radio transients which, based on their observed dispersion measures, appear to occur at cosmological distance scales. Currently, while numerous papers have suggested plausible sources for these radio transients, their origin is unclear. While not all plausible mechanisms for emission of FRBs are likely to result in simultaneous gravitational wave (GW) emission at detectable frequencies, there are several mechanisms which may result in coincident signals. The High Time Resolution Universe (HTRU) Collaboration has negotiated an MoU (see M1400007) with the LSC and Virgo in order to share trigger times, location and other information on currently unpublished Fast Radio Bursts. We have performed searches using GEO and Virgo data for several FRBs already under this agreement. This search plan describes plans to continue analysis of FRBs during O1 as per the MoU signed with HTRU.

3.14.2 Scientific Justification

Little has been definitively determined about the source of fast radio bursts, so we have elected to perform a “burst” search to look for gravitational waves as broadly as possible, in hopes of providing insight into the mechanism causing FRB emission. If FRBs result from stellar collapse scenarios [427] or if they are actually terrestrial “perytons” rather than a true astrophysical signal then we would not expect a GW signal. There are, however, a number of astrophysical phenomena that may plausibly produce gravitational waves in close coincidence with radio frequency emission matching the characteristics of FRBs. The discussion presented here is not an exhaustive list of source candidates, rather we focus on three general classes of sources which may produce both GWs and FRBs with frequency and duration suitable to the instruments being used in this analysis.

Cosmic strings - Cosmic strings, formed during symmetry breaking in the early universe, are topological defects thought to be capable of emitting large amounts of energy from their cusps or kinks [428] and have been suggested as a possible emission mechanism for FRBs [429, 430]. A cosmic string cusp may emit gravitational waves with a $f^{-4/3}$ frequency dependence up to a cutoff frequency [431], potentially at frequencies and amplitudes detectable by ground-based interferometers [332, 11].

This class of sources is particularly promising for purposes of an FRB related search since the distance scales on which GW signals may be observable for cosmic strings are consistent with the cosmological scales suggested by current FRB observations.

Binary neutron star coalescence - There are several models for radio emission in coincidence with a compact binary coalescence GW signal. This may be pulsar-like radio emission, either from the reactivation of the dormant pulsar emission in one of the neutron stars through interactions prior to merger [432] or by a hypermassive neutron star, which may sometimes result as an intermediate result of a merger before collapsing to a black hole, emitting at radio frequencies through a pulsar mechansism [433]. A third possible mechanism is the radiation at radio frequencies as a result of magnetospheric interactions [434].

Given an appropriate density in the surrounding environment, the gravitational waves emitted by a compact binary coalescence may induce electromagnetic radiation through magnetohydrodynamic interactions. While this interaction would directly produce radiation at the same relatively low frequencies as the GWs themselves, upconversion through inverse compton radiation may result in emission at radio frequencies [435]. This particular magnetohydrodynamic mechanism does not necessarily require neutron star coalescence as the mechanism for production of the GWs, but this class of source is likely to be able to produce GWs of suitable amplitude and may be surrounded by an environment suitable to this mechanism [436].

Single neutron stars - Most models of gravitational-wave emission resulting from single neutron stars would most likely produce signals too weak to detect at the distance scales suggested by the dispersion measures under consideration. However, if FRBs result from extreme SGR events as has been suggested [427], if future observations result in a lower dispersion measure event, or if some subset of FRBs is much closer than estimates due to the bulk of the dispersion happening close to the source rather than in the intergalactic medium, this class of models is still worth considering.

Transient gravitational wave emission can occur when a temporary deformation of a rapidly rotating neutron star creates a quadrupolar moment. Typically, this is believed to happen as a result of crust cracking from magnetic, gravitational or superfluid forces, dubbed a starquake [437], or from other asteroseismic phenomena resulting in shifting of the neutron star's crust [438]. While asteroseismology may result in several distinct types of quasinormal oscillatory modes of the neutron star, the f-mode is the most promising for the purpose of gravitational wave detection. Gravitational emission resulting from f-mode oscillation typically peaks around 2 kHz, although the exact emission depends on several factors, including the neutron star equation of state and the mass of the emitting neutron star [372]. The amplitude of the GW emission even in optimistic cases, however, is small enough that sensitivity to this type of source will be limited to our own Galaxy even in the advanced detector era.

Radio pulsars result from beamed emission from the poles of a rapidly rotating, highly magnetized neutron star sweeping past the Earth, producing reliably periodic radio signals. The asteroseismic events described above may result in a distinct increase in the rotation rates of these neutron stars, typically followed by a gradual return to their original period. This phenomenon, called a pulsar glitch, has been observed across a large number of pulsars, especially younger ones (see e.g. [439] and references therein). A search for gravitational wave emission from quasinormal modes in coincidence with the observed glitching of pulsar was the subject of a previous LIGO publication [379]. Models for neutron star asteroseismic phenomena similar to those under discussion have also motivated previous gravitational wave searches in coincidence with SGR flares [440].

The standard indication of an asteroseismic event in an isolated neutron star is a pulsar glitch, but there are are plausible mechanisms that could result in the observation of a transient radio pulse. This could simply be through the pulsar radio emission coming into view from the Earth as the pulsar's orbit shifts slightly, but there is also some evidence that pulsar-like radio emission can be "switched on" in coincidence with with a glitching mechanism [441, 442, 443]. We therefore consider single neutron stars as a possible source of coincident GW and radio transient events.

3.14.3 Search Description

The GW/FRB coincidence search will be conducted as a "target of opportunity" search if and when the timing and spatial information of known FRBs are provided by the HTRU collaboration. Based on previous observation rates, the rate of observation is expected to be less than 1 FRB per month.

Since the search for transient GWs of unknown morphology in coincidence with a known electromagnetic signal necessitates an approach similar to established burst GRB searches, the FRB-coincident analysis described in this search plan uses a "GRB-type" X-Pipeline analysis. The analysis uses existing X-Pipeline code and does not require separate code review, however the parameters of the analysis have been adjusted to reflect the astrophysical models under consideration for this particular search. Relative to the standard X-Pipeline based burst GRB search, the FRB coincident search utilizes a shorter (4 minute) duration on-source window. This reflects our expectation that joint radio and GW emission should generally occur within seconds of each other for the set of viable models we consider, based on a case-by-case analysis of the sources considered in the "scientific justification" section. We also use tighter spatial localization given Parkes' relatively narrow field of view, as well as a different set of simulated waveform software injections used to tune background rejection cuts that are consistent with the sources described in the "scientific justification" section above. Specifically, two compact binary coalescence, two cosmic string cusp, two sine gaussian, two chirplet sine gaussian and two damped sinusoids of different frequencies are used to represent a broad distribution of plausible emission mechanisms. A subset of these waveforms overlaps the set used for standard GRB X-Pipeline searches, including the NS-NS (1.4 solar mass) and NS-BH (1.4 and 10 solar mass, respectively) waveform types. Since single neutron star emission mechanisms are under consideration as a plausible source of joint emission, we also truncate the analysis at 3 kHz rather than 2 kHz where higher frequency analysis is permitted by the calibration of the interferometers.

As with similar externally triggered searches, a single trial false alarm probability of 0.01 or lower would trigger some follow-up activities as outlined in the X-Pipeline wiki pages. If the candidate survives initial scrutiny, further follow-up procedures as laid out on the burst "detection checklist" would be initiated as per, for example, the "Big Dog" blind injection event [59]. The candidate may be rejected before the entire list is executed. In the event of a viable detection candidate, parameter estimation may be a useful means of distinguishing between various emission scenarios, especially as a diverse group of possible sources are considered in this analysis.

3.14.4 Publication Plan

No publications are planned regarding specific FRBs or sets of FRBs in O1 regarding upper limits set on FRB triggers. The expected trigger rate is much lower than other similar analyses and the coupling between radio and GW signals, while very plausible for reasons described in the search plan, is not expected with the same degree of confidence as in, for example, GRB searches, so upper-limit based papers for these radio transients are not expected to be produced with the same frequency. "Evidence" or "detection" of a GW, according to standards agreed to by LIGO and Virgo, would merit a publication, however. Members of the HTRU collaboration would be included as authors in this paper. If detections were made in both an FRB triggered search and another analysis over the same data, a single "detection" paper would be written in collaboration with the other searches. The collaborations involved would reach a decision about dedicated follow-up papers focused on implications of a specific analysis.

It should be noted that since the source of FRBs is still an ongoing debate the burden of proof to establish "evidence" or "detection" may be higher than in externally triggered searches where the joint emission mechanism is more clearly established (such as GRB searches). This proposal is concerned with a search plan for O1, but continued plans to follow-up FRBs as external triggers for gravitational waves in later observation runs may be affected by the evolving science around FRBs. In particular, if the broader astrophysics community is able to come to a consensus regarding the origin of these bursts then the motivation for pursuing FRB-GW analyses will need to be revisited.

A collaboration paper regarding archival initial detector era burst searches for GWs in coincidence with short duration radio transients is currently under preparation. The bulk of radio triggers analyzed for this paper are from the Green Bank telescope drift-scan survey [444] and appear to occur at Galactic distance scales based on observed dispersion measures. The nature and scope of this analysis is roughly laid out in [445]. The paper will contain discussion of Fast Radio Bursts, including description of plans to continue monitoring for GW signals in coincidence with FRBs in O1 and beyond. However, this paper will not refer to specific FRBs coincident with advanced LIGO/Virgo (O1) data.

3.14.5 Results validation plan

X-Pipeline has been used in numerous previous externally triggered searches [446, 357, 447], including initial LIGO analyses with nearly identical parameters to those used in the planned radio transient search and is therefore considered a dependable search method. While there is no plan to run a second targeted pipeline on radio transients, all-sky analyses over the same stretch of data will serve as a useful cross-check in order to validate the results. As the primary burst all-sky [448] pipeline, coherent WaveBurst [408] is expected to be the main tool used for this, although in the event of an interesting trigger, omicron+LIB and possibly CBC pipelines [58, 59] would be utilized as well.

3.14.6 Resources

Computing needs - Requirements on a per-trigger basis are similar to the burst analysis of a Swift GRB. We estimate that 9500 SUs would be required for each FRB event based on previous events occuring in initial LIGO/Virgo data. Since these FRBs are identified by Parkes infrequently (less than one event per month based on identified FRBs [449] and data sets with known non-detections) the overall use of computing resources is negligible relative to the most computationally intensive all-sky all-time searches. While we anticipate conducting these FRB-coincident searches with turnaround times on the order of days rather than months as for archival searches, in the case that computational resources were required for an externally triggered event at a higher designated priority level, the FRB analysis could be temporarily delayed upon request without a significant loss of scientific value.

Detector Characterization and Calibration - As the FRB search primarily utilizes X-Pipeline and other existing low-latency infrastructure no additional effort should be required for purposes of detector characterization or calibration. (The possible exception to this statement would be a statistically significant event which would require exercising the burst detection checklist.)

Review - Most of the code utilized was already reviewed in the context of S6/VSR2-3 GRB searches [450]. Supplemental review for minor code adjustments and analysis procedures specific to radio-coincident analyses was performed as part of the Green Bank analysis (see G1200800, M1100024) . Therefore no additional code review is required for ongoing FRB searches.

3.15 All-sky Searches for Isolated Spinning Neutron Stars

3.15.1 Abstract

Gravitational waves (GWs) are small ripples in the geometry of space-time, predicted to exist by Einstein's Theory of General Relativity, and propagating at the speed of light. Rapidly spinning neutron stars in our galaxy may emit such waves if they are not perfectly symmetric around their spin axis. Their deviation from such axisymmetry is usually characterized by a parameter called *ellipticity* which, for example, might be as large as 10^{-6} for a broad, 1-cm-high *bulge* across the surface of the 20-km-diameter star. A comprehensive search has been carried out in data from the LIGO and Virgo gravitational wave interferometers for signals from such spinning stars. No such signal was detected, allowing scientists to rule out rapidly spinning neutron stars with ellipticities greater than 10^{-6} anywhere within 2000 light years of the Earth. Those limits are more stringent (10^{-7}) for neutron stars within 200 light years. We describe here the plans to extend the all-sky search for isolated spinning neutron stars to the next generation of ground based gravitational wave detectors. New data to be taken in the coming year after improvements to the LIGO and Virgo detectors is expected to be still more sensitive, allowing scientists to probe deeper into the galaxy and down to smaller ellipticities, in the hope of detecting these expected continuous waves for the first time.

3.15.2 Scientific Justification

Rapidly rotating neutron stars (NS) are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band. These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [451, 452, 453], magnetic deformations [454, 455], unstable r-mode oscillations [451, 456, 457], and free precession [458]. A review of these emission mechanisms can be found in [459]. Here, we focus on the all-sky search for unknown, isolated neutron stars. The number of undiscovered, electromagnetically quiet neutron stars within 5 kpc can be estimated to be $O(10^6 - 10^7)$ from the neutron star birth rate [460], although it is likely that only a tiny fraction would both be rotating fast enough to be accessible to LIGO [461] and remain bound to the galaxy over the age of the galaxy [462]. Only ~ 2000 radio or x-ray galactic pulsars have been discovered so far [463]. It has been argued, based on the observed supernova rate and inferred population of neutron stars in the galaxy, that the indirect limit on the strongest signal from this population is no more than

$$h_{\mathrm{IL}} = 4 \times 10^{-24} \left(\frac{30\,\mathrm{yr}}{\tau}\right)^{1/2}, \tag{1}$$

where τ is the mean time between supernovae in the Galaxy. The latest and most detailed derivation of this upper limit is given in [459]. Note, however, that a later simulation analysis found significantly lower expectations that depend on the assumed source frequency and ellipticity [464]. Moreover, these calculations make the optimistic assumption that a sizable fraction of neutron stars are gravitars, *i.e.*, their spindown is dominated by gravitational wave energy loss. In the end, the first detection of a previously unknown source may come down to a statistical fluctuation of a star's being especially near the Earth.

The absence of a signal detection does not place strong limits on neutron star physics, unfortunately. While exotic equations of state (EOS) *permit* relatively large ellipticities, accessible to the initial and advanced detectors, those EOS *do not require* that such ellipticities be realized in stars. The actual ellipticity of a star will presumably depend on its initial conditions and perhaps on its environment via post-birth accretion. Hence we simply do not know how close we might be to a first detection. On the other hand, observation of a CW signal would potentially reveal a great deal about NS EOS, especially once a corresponding electromagnetic counterpart were identified, ideally, in multiple bands.

3.15.3 Search Description

Given our ignorance of the physics of neutron stars (which may well be quark stars or hybrid stars) and given the tiny fraction observed to date ($<10^{-5}$ most likely), it seems prudent to cast a wide net in searches for these exotic yet pervasive objects, wide both spatially and in assumed phase evolution.

Below we describe five search pipelines, all of which are computationally bound, that make different tradeoffs in intrinsic strain sensitivity *vs.* robustness against deviations from assumed phase modeling. Four are "quick-look" pipelines, and one (Einstein@Home) aims at both deeper and broader searches, exploiting the substantial resources available via distributed computing with O(100K) users. With these pipelines we intend to search a GW frequency band from as low as 10 Hz (detector sensitivity permitting) to as high as 2000 Hz. The fastest known pulsar rotates at 716 Hz, implying GW radiation at 1432 Hz for a rotating-ellipsoid model; to be conservative, we plan to search well above this observational limit. The spindown range over which to search is affected by available computing resources, by the assumed minimum age of the stars, by the assumed ellipticity distribution of the stars and by the assumed maximum distance to the stars. As a minimum, we will search the spindown range -10^{-8} Hz s^{-1} $< \dot{f}_{\mathrm{GW}} <$ 10^{-9} Hz s^{-1}, where a slightly positive frequency derivative upper bound is chosen in the spirit of keeping one's eyes wide open and mindful that a star in a long-period binary system could appear to have a spin-up instead of a spin-down. Choosing rectangular coverage in the f–$\dot{f}$ plane is somewhat arbitrary in its simplicity, but has the virtue of automatically searching younger stellar ages at low frequencies, where young pulsars are predominant. We will likely also search for still higher spindown magnitudes and perhaps higher frequencies in lengthier searches, as discussed below. For now we plan to search these nominal frequency and spindown ranges with no special preference given to any subset (e.g., via longer coherence or observation times), but we may refine our search strategy if future detailed studies based on galactic neutron star modeling suggest that doing so can substantially improve detection probability.

An all-sky search faces formidable computing challenges. The parameter space over which one must search includes source right ascension and declination, source frequency and at least one frequency derivative, along with source inclination and polarization angles. For coherent integrations, the computational cost scales as the sixth power of the coherence time (if a single frequency derivative suffices) and as the cube of the upper limit on search band frequency. To exploit the power of long observation times requires semi-coherent search methods with detection statistics created from multiple coherent observation times.

Quick-Look Searches

The quick-look search pipelines will be run on LVC clusters and may be in contention for limited resources in the months immediately following the conclusion of a data run. To assist in allocating those resources, if needed, each pipeline team will be asked to run on a relatively narrow band ($\sim$100 Hz) of the run's data and report preliminary results promptly. The sensitivity, robustness and timeliness of those results may play a role in resource allocation, depending upon demand and supply.

PowerFlux [465, 466] applies a stack-slide approach [467] to compute from many thousands of 30-minute Short Fourier Transforms (SFTs) an average strain power corrected for antenna pattern and with weighting based on detector noise and antenna pattern. Its signal estimator is direct excess strain power for an assumed source direction and polarization, using one circular polarization projection and four linear polarization projections. PowerFlux corrects explicitly for Doppler modulations of apparent source frequency due to the Earth's rotation and its orbital motion around the Solar System Barycenter. Source frequencies are searched with $\sim$0.56 mHz spacing and limits presented separately for 0.25 Hz bands. Interesting outliers are followed up with the loose-coherence zooming technique [468, 469].

The Sky Hough method [470, 471, 472] uses short coherence times (30 minutes) to compute spectral powers. These values are compared against a threshold, with noise weights summed together over the course of a run, to create a semi-coherent detection statistic. These statistics are histogrammed in bins of sky location, corresponding to different assumed f and $\dot{f}$ templates, with outliers indicated by high final counts. A recent improvement was the addition of a χ^2-statistic computed over subsets of the data run, which has proven useful in vetoing certain detector artifacts [473].

The hierarchical Frequency Hough method [474, 475, 476, 477] uses relatively short coherence times (8000 s up to 128 Hz and 1000 s at higher frequency). Significant peaks in the equalized spectra are selected and used as input to the frequency Hough transform, which, for each direction in the sky, connects the data time/frequency plane to the source frequency/spin-down plane. Candidates are selected using coarse and refined grids in the parameter space and are then subject to coincidences, validation and follow-up.

The time-domain $\mathcal{F}$-statistic method [478, 479] computes the $\mathcal{F}$-statistic over coherence times of $\sim$2 days and then searches for coincidences among candidates over the course of a data run with consistent source parameters. Extensive line and transient cleaning is used in data preparation. Work is under way to parallelize the code using the Message Passing Interface (MPI) to run on very large computing facilities. There are plans to implement a follow-up procedure based on global correlations, as in the Einstein@Home search and to exploit GPUs.

Deep and Broad Search:

The **Einstein@Home** [471] distributed computing resources enable the use of longer coherence times ($\sim$day in recent searches) in an all-sky search with smaller mismatch than can presently be supported by the LIGO and Virgo computing clusters for those coherence times. The detection statistic for each coherence time is the well known $\mathcal{F}$-statistic for a particular choice of sky location, frequency, frequency derivative, marginalized over unknown source inclination and polarization. Individual $\mathcal{F}$-statistic values are used in a power-summing algorithm [471] to derive a semi-coherent detection statistic over a full data run. Recent improvements in the Einstein@Home infrastructure include (1) lower effective thresholds by moving more computing to remote hosts, (2) global correlations in search parameters to make semi-coherent combination more efficient, and (3) Bayesian line-veto inference to suppress single-detector artifacts.

All pipelines, with the exception of Einstein@Home, share comparable sensitivity and computational cost; the mock data challenge described in §3.15.6 will yield a quantitative comparison amongst them. We do not foresee significant improvements in the performance of these pipelines before the Advanced Detector Era, although there is opportunity in some pipelines for improved optimization.

Two additional search methods are under development aiming at sensitivity improvement while maintaining robustness against uncertainty in the source model: 1) a “loosely coherent” method [468], which builds on top of the PowerFlux infrastructure; and 2) a cross-correlation method [480] which bridges between semi-coherent and coherent methods, with the possibility of parameter tuning. The computational viability of these two methods for all-sky searches remains uncertain and will not be discussed further here. Note, however, that loose coherence allows a small patch of sky to be probed more deeply (“spotlight search”), a technique being applied in a recently completed S6 search of two spurs of the Orion galactic spiral arms.

Despite our wish to field a broad suite of approaches, the present proliferation of pipelines seems excessive. We are developing a standard set of performance metrics via a multi-stage mock data challenge, discussed in §3.15.6, that will allow us to identify redundancies; as a result, we expect the all-sky pipelines to undergo a period of consolidation and stabilization. If computational resources become too limited to support running each pipeline over the full parameter space, these metrics should aid in setting priorities. The goal is not to converge on a single pipeline, but we do want to understand clearly the justification for each one that remains. The justifications can include:

- Best sensitivity
- Ability to cover (with astrophysically interesting – if not best – sensitivity) the largest region of parameter space
- Best robustness against signal deviations from assumed phase model
- Fastest pipeline for quick looks at data
- Deliberate redundancy (using an independent software base) for safety in this critical search

The consolidation of pipelines will also be accompanied by stabilization, both in the testing and freezing of software and in building up a team of pipeline users, to avoid present reliance on single individuals.

This search will also benefit from two follow-up pipelines for "zooming in" on interesting sources, including the "loose coherence" approach [468, 469] used in the full-S5 PowerFlux paper [481] and the "Nomad" approach [482] used in the full-S5 Einstein@Home paper [483]. In the case of an isolated neutron star detection, this will allow us to focus promptly on pinning down the source parameters via zooming. There are several important considerations to bear in mind with respect to such follow-ups. First, the SNR increase from zooming can be an order of magnitude or more, allowing relatively precise parameter estimation compared to what is found upon first detection. Second, one can then go back to older, less sensitive data and still find a source with very good SNR that had been missed previously; moreover, that expanded time baseline itself offers improved parameter precision, at the cost of requiring an expanded search range to account for extrapolation errors due to uncertainties in the initial (discovery) parameters. Finally, although the semi-coherent methods used in all-sky searches have sensitivities that improve only as the fourth root of observation time, once a true signal has been pinpointed, then the SNR and parameter estimation from coherent followups improve as the square root of observation time (assuming stable detector noise during the data run).

Electromagnetic follow-ups play a critical role in assessing the scientific impact of a detection in this search. Before publishing a CW discovery, we wish to consult electromagnetic partners with access to radio, X-ray and γ-ray telescopes on what they can see in the direction of the source. With a year's observation time and good SNR, we expect to achieve an angular resolution of O(arcsec), which should suffice for excellent electromagnetic follow-up. We already have partnerships in the radio, X-ray and γ-ray communities for obtaining ephemerides for our targeted pulsar searches. A discovery publication might include co-authors from those communities who carry out follow-ups at our request (whether or not they confirm the source), a blanket request we envision going out through our existing liaisons in those communities[15]. Inclusion of these authors in the discovery paper (as opposed to later papers or companion papers), would depend on the timeliness and astrophysical importance of their results. We fully expect results from some observers that would be rapid and enhance the astrophysical impact of the discovery paper.

Assuming that pulsations are indeed observed in one or more electromagnetic bands in follow-up, there is immediate astrophysical payoff. If the gravitational wave frequency is twice that of the rotation frequency, with no GW emission observed at the rotation frequency, then a rotating triaxial ellipsoid model is favored. On the other hand, if the two frequencies agree, then precession probably plays an important role. If the GW frequency is approximately 4/3 the rotation frequency, then r-modes are strongly favored, yielding direct information on the interior fluid motion of a neutron star. In the conventional scenario of $f_{\rm GW} = 2f_{\rm Rot}$, the

---

[15] We envision a very lightweight MOU arrangement with electromagnetic observers. In the event of a GW discovery, observers would be offered the opportunity to sign an MOU setting out publications ground rules, in order to obtain the GW source coordinates and ephemeris. While leaks of such parameters to MOU non-signers are likely, the pulsar community has established an etiquette for protecting priority on discoveries that makes such leaks benign, we believe.

phase of the GW emission relative to the phase of the electromagnetic pulse(s) gives insight into relation of the stellar quadrupole asymmetry to the magnetic poles. Similarly, the inclination angle and transverse polarization angle of the star's spin axis can be inferred from the GW emission, and if deep electromagnetic observations detect a pulsar wind nebula around the star, it may be possible to compare the GW- and EM-inferred angles. More generally, one can envision a sequence of beyond-discovery collaboration publications, as we and electromagnetic astronomers continue studying a source that (we hope) remains detectable throughout the ADE.

Another interesting follow-up possibility after a discovery is to exploit advanced interferometer narrow-banding to improve parameter estimation. Whether the potentially improved SNR from narrow-banding is worth the cost in sensitivity in other bands and in interferometer downtime is difficult to judge in advance, however. Unless all interferometers are narrow-banded together, the SNR gain from narrow-banding over broadband operation may be modest, and given current astrophysical knowledge, there is no compelling reason to choose narrow-banding in an all-sky isolated-star search. We can't even say with confidence that detection of an unknown high-frequency millisecond star is more or less likely than detection of a young, low-frequency star. (Even for low-mass X-ray binaries, such as Scorpius X-1, which were often cited in the past as targets for narrow-banding, it is not at all clear that narrow-banding is truly sensible.)

3.15.4 Result Validation Plan

As discussed in section 3.15.3, "zooming" in on all-sky candidates offers the potential for an order of magnitude increase in signal-to-noise ratio for a source that follows the signal model of an isolated star with constant spindown over the observation span. With SNRs of 50-100 or more, confirmed by the three independent targeted-search pipelines described elsewhere, there will be no uncertainty as to whether a detection has been made.

If there is no detection, the fact that there are up to five independent all-sky pipelines of comparable sensitivity will give us confidence we have not missed an easily detectable signal. An ongoing mock data challenge using injections into S6 data will give us strong guidance as to relative pipeline sensitivities prior to O1. In addition, there will be a "real data challenge" at the end of O1, in which all pipelines will be asked to produce timely results on a 100-Hz test band of the O1 data, to assess further their relative performance. In the event of a shortage of computing resources, this comparison could lead to re-prioritizations of search jobs.

3.15.5 Publication Plan

Following the observing scenarios document [484], we assume here a 3-month run in 2015, a 6-month run in 2016-2017, and a 9-month run in 2017-2018. In each run we assume an intrinsic strain sensitivity significantly improved over previous runs. Hence we envision at least one all-sky publication from each run and perhaps more if deeper or broader searches require more computational time. Roughly speaking, all-sky upper limits on strain amplitude scale with the noise floor of the most sensitive instrument operating and with the fourth root of observation time. Steadily improving detector sensitivity will argue for expedited analysis of each data set, including a pragmatic approach to detector artifacts and unpromising outliers[16] For each of the first three runs, we plan the following all-sky publications:

1. A quick-look analysis, based on one or more pipelines that use short coherence times (30 minutes or less) and that report final results promptly. Ideally, the publication for each run will be completed

---

[16] This approach contrasts with that used for the S6 run, where a long hiatus of new data argued for digging aggressively into the noise to try teasing out signals, which necessitated following up millions of outliers. This development of systematic follow-up procedures should be helpful in advanced detector searches. The substantial non-Gaussianity of the S6 data also required developing new techniques to avoid excluding large swaths of the search band.

within a few months of the end of data collection, before the next run begins.

2. A deeper and broader analysis, based on one or more pipelines that use longer coherence times [O(day)]. Based on past experience, these results are likely to be delayed by as much as an extra year. Nonetheless, with a 2-year proprietary period in the open data era, deep-search results should be published on each data set before it is made public. Again, pragmatic choices on follow-up procedures can ensure such timelines are kept, if necessary. At current pipeline sensitivities, the primary advantage of Einstein@Home lies in the breadth of its parameter space coverage, but refinements in the outlier followup could lead to significant further gains in sensitivity from lowering 1st-stage SNR thresholds (to be evaluated with mock data challenges described below). That said, large delays in producing results, combined with availability of newer, more sensitive data, could make publication contingent upon broad parameter space coverage. Hence we envision increasing spindown range by an order of magnitude (with perhaps an increase in frequency range) while maintaining or improving upon the strain sensitivity of the quick-look searches.

Although we implicitly assume here that published results from the first three runs will be upper limits on CW sources, we may well detect a signal from the all-sky searches. In that event we would, of course, publish one or more papers on the detected source and would collaborate with radio, optical, X-ray and γ-ray astronomers on deriving astrophysical insight from joint observations, as discussed in §3.15.3.

If detector sensitivity does not improve significantly from one year to the next, *e.g.*, 2016 running is comparable to that in 2015, then the searches will nonetheless be carried out, but the journal (or perhaps conference proceedings) to which results are submitted for publication may be adjusted accordingly.

Pipeline	**Published observational results to date**	**Search & review status**
PowerFlux	*PRD* 77 (2008) 022001 (S4) *PRL* 102 (2009) 111102 (S5) *PRD* 85 (2012) 022001 (S5)	S6 search under review
Einstein@home	*PRD* 79 (2009) 022001 (S4) *PRD* 80 (2009) 042003 (S5) *PRD* 87 (2013) 042001 (S5)	S6 search under way
Sky Hough	*PRD* 72 (2005) 102004 (S2) *PRD* 77 (2008) 022001 (S4) *CQG* 31 (2014) 085014 (S5)	S5 search review completed
Frequency Hough		VSR2/4 search under review
Time Domain $\mathcal{F}$-statistic	*CQG* 31 (2014) 165014	VSR1 search review completed

Table 4: Summary of CW all-sky pipelines, including published observational results from previous data runs and the current search and review status.

3.15.6 Technical requirements and development plan

Pipelines status and review

PowerFlux has been used in published searches on LIGO S4 data [465] and S5 data [485, 481], with significant improvements from one publication to the next. Preliminary results for the LIGO S6 data have been presented to the collaborations, and outlier follow-up is under way. The multitude of spectral artifacts in S6 data (far higher than in earlier LIGO runs) has led to a longer-than-expected follow-up analysis and to the development of a novel method of setting upper limits in non-Gaussian data [486], based on a "universal statistic." Review of the S6 results is well under way, with publication expected in summer 2015. The tagged pipeline from this review serves as a reviewed baseline for the ADE.

The **Sky Hough** pipeline has been used in published or soon-to-be-published searches in LIGO S2, S4 and S5 data [487, 465, 472]. This reviewed pipeline can be frozen for ADE use, as is.

The review of the **Frequency Hough** method and analysis results began in late 2011 and is expected to finish in spring 2015, with the subsequent publication of all-sky upper limits on sources in the Virgo VSR2/4 data with frequencies below 128 Hz. The pipeline used in this search will be tagged for default use in the ADE, although further pipeline development may occur (which would require additional review).

The review of the **time-domain $\mathcal{F}$-statistic** method began in fall 2012 and completed in fall 2013. The method has been applied so far to the VSR1 data set, and an observational paper was published in 2014. The pipeline has been tagged for default ADE use. The future addition of global correlations component would require additional review.

Einstein@Home has been used in published searches on LIGO S4 data [471] and S5 data [488, 483], with significant improvements from one publication to the next. Recent production runs on S6 data have included the global correlations and line-veto improvements, but post-processing to produce upper limits and follow up outliers is still under way. Review of the most recent full E@H all-sky pipeline is planned for completion in 2015.

Mock Data Challenge The performances of the various pipelines used in this search are being evaluated via a standardized set of software signal injections into a parallel stream of real data from the LIGO and Virgo interferometers (S6 and VSR2/4). Two sets of 3000 injections each of isolated stars over the sky and frequency band have been created, and a multi-stage mock data challenge is well under way. The stages of the challenge include:

- Stage 1 (Concluded Fall 2012) – 20 injections (upper limits and yes/no detection success)
- Stage 2 (Concluded Fall 2013) – 200 injections (same as stage 1 plus efficiency curves)
- Stage 3 (Concluding Summer 2015) – 1500 injections (same as stage 2, with better statistics)
- Stage 4 (Concluding Summer 2015) – 1500 injections (same as stage 3 but with *blind* injections)

In fall 2012, the Stage 1 subset of about 20 injections was used for initial testing of all-sky pipelines and of the challenge itself. Only one all-sky pipeline (PowerFlux) was ready to participate in the Stage 1 challenge, while two pipelines (PowerFlux and the multi-IFO Sky Hough) participated in Stage 2. Other pipelines have been urged to join the Stages 3 and 4 challenges, which will conclude in 2015. Preliminary results have been produced by all pipelines for Stage 3, but not with consistent criteria for detection and the setting of upper limits. Search teams have been informed that reviewing resources and substantial LVC computing resources will be supported for only those ADE pipelines that participate, and that performance in the challenge will be reported at collaboration meetings.

By means of standardized figures of merit, the mock data challenge will allow us to assess, for each pipeline, the sensitivity in real data, the robustness against detector artifacts, the coverage of parameter space in frequency and spindown, and the computational cost. We expect the mock data challenge to quantify differences in intrinsic sensitivities among the all-sky pipelines, giving guidance on which pipeline(s) to treat as highest-priority in access to computing and reviewing resources. That does *not* mean that continued work on the other pipelines will be discouraged. As discussed above, redundancy provided by independently written pipelines (*e.g.*, red team *vs.*blue team) is prudent and a path to innovation. Nonetheless, the presentation of results in a publication will focus on those from the most sensitive pipeline, with only brief mention of corroborating results from alternative searches.

3.16 Targeted Searches for Gravitational Waves from Known Pulsars

3.16.1 Abstract

Gravitational waves (GWs) are small ripples in the geometry of space-time, propagating at the speed of light, predicted to exist by Einstein's General Theory of Relativity. Rapidly spinning neutron stars in our galaxy may emit such waves if they are not perfectly symmetric around their spin axis. Their deviation from such axisymmetry is usually characterized by a parameter called the *ellipticity* of the neutron star which, for example, might be as large as 10^{-6} for a broad, 1-cm-high *bulge* across the surface of the 20-km-diameter star.

Searches have been carried out in data from the LIGO and Virgo gravitational wave detectors for signals from known radio and X-ray pulsars. No GW signals have been detected from these stars, leading to upper limits on their ellipticities reaching as low as 10^{-8}. These limits provide tests of neutron star structure and of fundamental nuclear physics in a realm inaccessible to terrestrial experiments. We describe here the plans to extend targeted searches for gravitational waves from known pulsars to the next generation of ground based detectors: Advanced LIGO and Advanced Virgo.

3.16.2 Scientific Justification

Rapidly rotating neutron stars are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band. These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [451, 452, 453], magnetic deformations [454, 455], unstable r-mode oscillations [456, 451, 457], and free precession [458]. A review of these emission mechanisms can be found in [459].

Here, we focus on the targeted search for gravitational waves from all known neutron stars are also radio, X-ray or gamma ray pulsars. Given our ignorance of the physics of neutron stars (which may well be quark stars or hybrid stars), and the relatively small fraction of their full population that are currently known, it is prudent to target these known stars with deep search techniques that are tailored to exploit the knowledge of their rotation and spindown rates to maximise our sensitivity to gravitational waves.

Our pipelines target a subset of sources for which pulsations are observed in radio, X-ray, or other electromagnetic radiation, giving timing with sufficient precision to merit a deep, usually single-template, search. Electromagnetic (EM) astronomy can tell us precisely the sky positions, frequencies, and frequency evolution of these objects, so that targeted analyses need search only a small parameter space and are not computationally limited. EM astronomy also sets an upper limit on the gravitational wave strain we could see from a known pulsar, assuming that all of the observed spindown is due to gravitational wave emission. In terms of the distance D, gravitational wave frequency $f_{\rm gw}$ and its time derivative $\dot{f}_{\rm gw}$, this indirect limit is [459]

$$h_{\rm IL} = 8.1 \times 10^{-25} \left(\frac{1\,{\rm kpc}}{D}\right) \left(\frac{-\dot{f}_{\rm gw}}{10^{-10}\,{\rm Hz/s}} \frac{100\,{\rm Hz}}{f_{\rm gw}}\right)^{1/2} \left(\frac{I}{10^{38}{\rm kg\,m}^2}\right)^{1/2}. \tag{2}$$

Here I is the star's moment of inertia (as estimated by theory as these are yet to be directly observed) and could be higher than the fiducial value by a factor of up to three. For most pulsars the distance D is determined by combining their observed radio dispersion measure with a model of the galactic HII distribution and is uncertain to at least 20%. Analysis of the LIGO full S5/S6 data and the Virgo VSR2/VSR4 data has improved on this indirect "spindown limit" by a factor of 10 for the Crab pulsar (at 59.45 Hz) and by a factor of 3 for the Vela pulsar (22.38 Hz) [489]. Other pulsars for which the spindown limit was approached in S6/VSR2/VSR4 include PSRs J0205+6449 (30.45 Hz), J1833−1034 (32.33 Hz), J1813−1749 (44.74 Hz), J1913+1011 (50.59 Hz), J1952+3252 (55.69 Hz), J0737−3039A (88.11 Hz) and J0537−6910 (123.95 Hz) [490].

The discussion above assumes gravitational wave emission from a triaxial neutron star, with the electromagnetic and gravitational wave components rotating as one unit. Detecting such emission would represent the first ever measurement of the difference between the two (equatorial) components of the inertia tensor. This would provide important information on the strength and strain profile of the solid phase of the star (the crust, or possibly a solid core) and/or information on the nature of the internal magnetic field.

While this form of gravitational wave emission is the simplest and most plausible, it is by no means the only possible wave generation mechanism. Alternatives include free precession, excited modes of oscillation of the fluid, and the spindown of a multi-component star. The identification of such mechanisms would provide information on asymmetries in the inertia tensor, the viscosity of the fluid, and the internal structure of the star. However, the observational challenge is correspondingly greater, as the gravitational wave emission no longer occurs at twice the spin frequency, and special care should be given to the choice of search range in parameter space (i.e., the wave frequency and its time derivatives) that is both astrophysically reasonable and computationally feasible. Such a search has already been carried out for the Crab pulsar, concentrating on a small patch of parameter space centred on (twice) the observed spin frequency [491]. A more comprehensive search over an enlarged parameter space and for more pulsars is needed to fully exploit the science potential of targeted searches.

Targeted searches look for gravitational wave emission from pulsars of known position, rotation frequency, spindown rate, and binary orbital parameters where necessary. This information greatly reduces the size of the search space, and allows us to perform a fully coherent search over the full dataset. Timing accuracy is sufficient to maintain coherence both during and between science runs, and the relative phasing of the interferometers is also sufficiently well determined for us to be able to combine data from all runs and all detectors coherently, resulting in the best signal sensitivities achievable by LIGO and Virgo.

3.16.3 Search Description

Targeted searches require extreme precision and are not as resilient to ephemerides and coding mistakes as multi-template searches. Our strategy is to use multiple pipelines for these searches, cross-validated on hardware and software injections; for a more robust result, these pipelines use complementary techniques to reject out-of-band signals and for statistical interpretation. In principle we can use *all* the data we have available, both historical and current, and data from detectors with any sensitivity. However, in practice the more up-to-date data from a single IFO usually dominates the overall sensitivity. The derived parameters define the orientation of the neutron star and the strain amplitude at Earth.

Time domain Bayesian Pipeline. The time-domain Bayesian method has been applied successfully to data from the first six LSC science runs [492, 493, 494, 495, 496, 489] and to the Virgo VSR2/VSR4 runs [497, 489]. The method is described in detail in [498]; the inclusion of binary system parameters is described in [499]. This pipeline is designed to carry out robust signal extraction and optimal parameter estimation, rather than search over a large parameter space. Its primary purposes are therefore (1) to perform searches for signals from known pulsars and, (2) to determine the astrophysical parameters of candidate sources.

The method comprises a heterodyne and filtering stage to extract interferometer data in a tracked 1/60 Hz band centered on the expected gravitational wave signal, followed by a Bayesian parameter estimation stage. This second stage delivers an upper limit to the strain amplitude of any signal and an estimation of the signal's parameters, should one be detected, in the form of marginal posterior probability distributions for the signal parameters. The method has successfully determined the parameters of the injected signals in all our science runs. The most computationally intensive part of the search is the heterodyning and down-sampling of the raw data. Currently this takes about 25 min per pulsar per detector per day of data. We have developed a new method of computing the tracked 1/60th Hz band by interpolating between frequency bins in one-minute short Fourier transforms (SFTs) of the data rather than by heterodyning the timeseries. This gives a very significant speed-up when processing multiple targets and is suitable for all isolated and some binary targets. For these targets, the computational load for this preprocessing step becomes trivial.

The parameter estimation stage uses a Markov-chain Monte-Carlo (MCMC) algorithm to explore the unknown parameter space, combining the data with priors on pulsar parameters (including spin axis orientation) to return a joint posterior distribution for the parameters of each pulsar. From this posterior distribution we form marginalised posteriors on each parameter and can infer an upper limit on the gravitational wave amplitude. In addition to this, a new parameter estimation code has been developed, based on the `lalinference` functions within `lalsuite` (these are also being used as the main parameter estimation tools in the CBC searches), which uses nested sampling to explore the parameter space. As well as the posterior distribution this algorithm also returns the “evidence” for the signal model. This evidence, when compared to the evidence that the data consists of noise alone (or contains a different signal model), can be used as a detection statistic. The code has also been designed to be more flexible at including different pulsar emission models and allowing searches over expanded parameter spaces. In the advanced detector era (ADE) we also plan a fully automated version of the pipeline, producing preliminary data products and results a few days after data has been collected.

Time domain matched-filter method using the $\mathcal{F}$ and G statistics. Here, as above, we assume we know the pulsar’s position, frequency and frequency derivatives so that the standard $\mathcal{F}$-statistic [478], used in all-sky searches, can be applied to a single template corresponding to the target in question. If in addition the orientation of the spin axis of the pulsar is known, the signal can be detected using the recently derived $\mathcal{G}$-statistic [500]. Once the signal is detected, a calculation of the $\mathcal{F}$-statistic enables estimation of the amplitude, phase, polarization and inclination angles, whereas application of the $\mathcal{G}$-statistic results in estimation of the amplitude and phase. If the computed value of the $\mathcal{F}$-statistic or $\mathcal{G}$-statistic is not significant, we can derive an upper limit on the gravitational wave signal using a standard frequentist approach, that is by injecting signals into the data with random parameters. The input data for the calculation of the $\mathcal{F}$ and G statistics are the coarse- or fine-heterodyned data generated for the time domain Bayesian method described above. This method was first applied to search for the Vela pulsar in VSR2 data [497] and was recently applied to search for several known pulsars in VSR2/VSR4/S6 data sets: J0534+2200 (the Crab), J0537-6910, J0835-4510 (Vela), J1813-1246, J1833-1034, J1952+3252 and J2022+3842 [489].

Signal Fourier 5 components method. This independent targeted search method consists of three steps:

- extract a small (say 0.1 Hz) band from a Short FFT Database (built from 1024 s data segments, after a time-domain cleaning);
- transform to the time domain and apply Doppler, spin-down and Einstein delay corrections through a re-sampling procedure;
- down-sample the corrected data to a much smaller rate (e.g., 1 Hz).

After an additional cleaning step to remove outliers present in the small analyzed band, data- and signal-template 5-vectors are computed and used to calculate a detection statistic, as described in [501, 497]. The corresponding p-value is then computed in order to establish how compatible the data are with pure noise. If no significant evidence of a signal is obtained, an upper limit on the signal amplitude is determined. This method has been applied, together with the other two coherent pipelines, for the search of CW signals from the Vela pulsar in the VSR2 data [497] beating the spin-down limit by a factor of about 1.6. It was recently applied to the analysis of VSR4/S6 data, obtaining upper limits for Vela (VSR4), Crab (VSR2/VSR4/S6) and a few other low-frequency pulsars [489].

3.16.4 Result Validation Plan

- Results will be validated via our standardized set of software signal injections into a parallel stream of real data from the LIGO and Virgo interferometers. Two sets of 3000 injections each of isolated stars over the sky and frequency band have been created. This is the same mock data challenge being used in the all-sky continuous-wave searches, and is described in more detail in that observing plan. Injections have been made at the ephemerides of our target known pulsars, and the recovered parameters for these from the three pipelines will be critically compared. There are significant differences in the statistical approach of the three pipelines, resulting in slightly different performances dependent on the noise environment around the target frequencies. However, cross-validation between pipelines and injections has produced robust results in initial targeted pulsar results from LIGO and Virgo and the same successful policy will be employed for advanced detector data.

3.16.5 Publication Plan

Following the observing scenarios document [484], we assume here a 3-month run in 2015, a 6-month run in 2016–2017, and a 9-month run in 2017–2018. In each run we assume an intrinsic strain sensitivity significantly improved over previous runs. Hence we envision at least one targeted publication from each run and perhaps more if specific sources pass the spin-down sensitivity point. We expect the publication rate for targeted searches to be more rapid in the ADE than for the S5–S6/VSR1–4 era. We have established reviewed, tagged pipelines, based on those used for the final S6/VSR2–4 searches. Although refinements will likely be explored in the ADE, there will be default pipelines for which the reviews can focus on analysis results, not source code. The expected pace of the data runs and the sensitivity improvements make it attractive to put a premium on rapid analysis and review.

Targeted searches are constrained by the availability of radio and X-ray pulsar ephemerides. It is good-practice to get timing solutions for our targets that cover the entire science run to ensure that there are observations at the end that confirm the pulsar has not undergone a glitch over the period. As a result, we must wait for our EM partners to deliver final ephemerides before the analysis is completed, and past experience has shown this to be the rate-limiting step in our multi-target publications (note that the Crab Pulsar ephemeris is kept up-to-date online, and there are no plans by the University of Manchester Jodrell Bank to stop this service).

Specifically, for each of the first three runs, we plan the following targeted publications:

1. A synoptic publication, updating upper limits on the entire target list.
2. When appropriate, one or possibly two publications highlighting a milestone passed on particular sources.

Although we assume here that published results from the first three runs will be upper limits on CW sources, we may well detect a signal. In that event we would, of course, publish one or more papers on the detected source and would collaborate with radio, optical, X-ray and γ-ray astronomers on deriving astrophysical insight from joint observations. The relationships we have developed already to use detailed unpublished ephemeris data from these communities in targeted searches should assist in joint follow-up observations.

If detector sensitivity does not improve significantly from one year to the next, *e.g.*, the 2016 run is comparable to that in 2015, then the searches will nonetheless be carried out, but the journal (or perhaps conference proceedings) to which results are submitted for publication may be adjusted accordingly.

3.16.6 Technical requirements and development plan

Pipeline	**Published observational results to date**	**Search& review status**	
Time Domain Bayesian	*PRD* 69 (2004) 082004 (S1) *PRL* 94 (2005) 181103 (S2) *PRD* 76 (2007) 042001 (S3/S4) *ApJL* 683 (2008) 45 (S5) *ApJ* 713 (2010) 671 (S5) *ApJ* 737 (2011) 93 (VSR2) *ApJ* 785 (2014) 119 (S6/VSR2/VSR4)	Review complete	
Time Domain Matched Filter	*ApJ* 737 (2011) 93 (VSR2) *ApJ* 785 (2014) 119 (S6/VSR2/VSR4)	Review complete	
Signal Fourier 5-component	*ApJ* 737 (2011) 93 (VSR2) *PRD* 91 (2015) 022004 (S6/VSR2/VSR4) *ApJ* 785 (2014) 119 (S6/VSR2/VSR4)	Review complete	

Time domain Bayesian Pipeline

This pipeline has been intensively reviewed and used for every science run of LIGO and Virgo. Additions and refinements to the pipeline are reviewed as they come into use, but are generally improvements in efficiency, and a fully-reviewed pipeline is always available for production runs. Most recently the pipeline was used for S6/VSR2/VSR4 analyses and the majority of the pipeline is identical to the one used and reviewed for S5. The only changes were to increase automation of the post-processing scripts and to incorporate priors from searches in older data. These changes have now been reviewed. A new method that uses spectral interpolation of Short Fourier Transforms (SFTs) to create a downsampled timeseries is currently under development, and a reviewed pipeline for ADE, thoroughly tested against the current heterodyne pipeline and software/hardware injections, is expected by late 2015.

The new parameter estimation code has already been validated against the current code using both fake data and simulated signals [502, 503]. The central algorithms for posterior sampling have also been reviewed as part of the CBC parameter estimation code review. However, the code calling these sampler functions and providing the specific pulsar model, prior initialisation and data reading functions, requires review. This code refinement is expected to be in place for the ADE.

Time domain matched-filter method using the $\mathcal{F}$ and G statistics

The pipeline has been extensively tested and reviewed during its application to the search of the VSR2 data for the Vela pulsar. The current version of the pipeline is a simple extension that is able to read any coarse and fine heterodynes produced by the time domain Bayesian method and can use any pulsar ephemeris file. Also the Monte Carlo simulations to establish the upper limits are automated.

Signal Fourier 5 components method

The single-detector pipeline has been extensively reviewed and tested in the past. Recently the method was extended to allow a coherent analysis of different datasets, coming from the same or different detectors [504]. Moreover, a new method for computing upper limits in the frequentist framework has been developed, which overcomes some problems of the standard frequentist methods [505], and a methodological paper is in preparation. The review of these method and software updates was completed in summer 2013. We are also working on the method extension to narrow-band searches, allowing a small mismatch between the EM and GW frequency and frequency derivatives.

3.17 Directed Searches for Gravitational Waves from Isolated Neutron Stars

3.17.1 Abstract

Gravitational waves (GWs) are small ripples in the geometry of space-time, propagating at the speed of light, predicted to exist by Einstein's General Theory of Relativity. Rapidly spinning neutron stars in our galaxy may emit such waves if they are not perfectly symmetric around their spin axis. Their deviation from such axisymmetry is usually characterized by a parameter called the *ellipticity* of the neutron star which, for example, might be as large as 10^{-6} for a broad, 1-cm-high *bulge* across the surface of the 20-km-diameter star.

Searches have been carried out in data from the LIGO and Virgo gravitational wave detectors for signals from the centers of supernova remnants, where isolated neutron stars are thought to reside. No GW signals have been detected from these stars, leading to upper limits on the strengths of possible gravitational radiation emitted by these stars. These limits improve substantially upon searches carried out in data from the initial LIGO detector. We describe here the plans to extend directed searches for gravitational waves from isolated neutron stars to the next generation of ground based detectors: Advanced LIGO and Advanced Virgo.

3.17.2 Scientific Justification

Rapidly rotating neutron stars are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band. These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [451, 452, 453], magnetic deformations [454, 455], unstable r-mode oscillations [456, 451, 457], and free precession [458]. A review of these emission mechanisms can be found in [459].

Here, we focus on the directed search for gravitational waves from isolated neutron stars, where we know a potential source location to high accuracy, but have little prior information on the rotational frequency of the star. The most promising such source is probably the compact central object at the center of the Cassiopeia A (Cas A) supernova remnant (SNR). There is excellent reason to believe a neutron star was formed there just over 300 years ago in a supernova event, a star that is rapidly cooling and may have significant residual quadrupole asymmetry. In addition, there are of O(10) other SNRs in the galaxy young enough and near enough to warrant a search. A list of SNRs for which we have searched in S6 data is given in Table 5. Aside from Cas A, perhaps the most interesting source in the table is Vela Jr. for which two entries appear, corresponding to two distinct possible interpretations of astrophysical observations to date, one of which places the star quite close and quite young (so much so that it's hard to understand how it could have been missed by medieval astronomers). The weight of evidence, however, argues for the more distant and older interpretation, which is nonetheless comparable in its age-based strain limit to Cas A. Because of the large uncertainties in this source's properties and its likely greater age than the very young Cas A, it is given a lower priority in current search planning. Another intrinsically interesting direction is the galactic center, where enhanced star formation could create young neutron stars, but where dust could obscure the present-day SNR and could have obscured the optical signal of a supernova from earlier astronomical observers. Other interesting directions where undiscovered isolated neutron stars might reside are at the cores of globular clusters. Since globular clusters host very old stars, detected milli-second pulsars dominate over young pulars. In addition, it is possible that a neutron star was created in Supernova 1987A, making the center of that supernova remnant an intriguing direction in which to search, despite the large distance (51 kpc) to it.

Our pipelines exploit the known sky location of the source to search more deeply and for younger objects than is possible for all-sky searches. The increased depth comes from increasing the coherence time of each

Table 5: List of supernova remnant objects for which an S6 directed search has been carried out.

SNR (G name)	Other name	RA+dec (J2000)	D (kpc)	a (kyr)	h_0^{age}
1.9+0.3		174846.9−271016	8.5	0.1	8.4×10^{-25}
18.9−1.1		182913.1−125113	2	4.4	5.4×10^{-25}
93.3+6.9	DA 530	205214.0+551722	1.7	5	6.0×10^{-25}
111.7−2.1	Cas A	232327.9+584842	3.4	0.3	1.3×10^{-24}
189.1+3.0	IC 443	061705.3+222127	1.5	3	8.7×10^{-25}
266.2−1.2	Vela Jr.	085201.4−461753	0.2	0.69	1.4×10^{-23}
266.2−1.2	Vela Jr.	085201.4−461753	0.75	4.3	1.5×10^{-24}
291.0−0.1	MSH 11−62	111148.6−603926	3.5	1.2	5.9×10^{-25}
347.3−0.5		171328.3−394953	0.9	1.6	2.0×10^{-24}
350.1−0.3		172054.5−372652	4.5	0.6	6.5×10^{-25}

Table 6: Values of distance D and age a are at the optimistic (nearby and young) end of ranges given in the literature, except for the second search for Vela Jr.

segment searched, and the probing of young ages comes from searching over higher frequency derivatives, normally neglected beyond the first spin-down parameter in all-sky searches.

The conventional spindown limit for known pulsars, based on measured spin and spindown, does not directly apply to a source with an unknown frequency, but if the source's approximate distance D and age τ are known, as is the case for many SNRs, one can derive a frequency-independent age-based limit under the assumption that the star has been a gravitar since birth (spindown dominated by gravitational wave emission)[506]

$$h_0^{\mathrm{age}} = \frac{1}{D}\sqrt{\frac{5GI_{zz}}{8c^3\tau}} \tag{3}$$

Here I_{zz} is the star's moment of inertia (as estimated by theory as these are yet to be directly observed) and could be higher than the fiducial value by a factor of up to three. For nominal Cas A parameters, this can be conveniently rewritten:

$$h_0^{age} = 1.2 \times 10^{-24}\left(\frac{3.4kpc}{D}\right)\sqrt{\left(\frac{I_{zz}}{10^{45}\mathrm{gcm}^2}\right)\left(\frac{300\mathrm{years}}{\tau}\right)} \tag{4}$$

In designing a search, it is customary [506] to determine the band over which a search of fixed computational cost can beat this age-based limit. For the S5 Cas A search, this band was 100-300 Hz, while for the nearly complete S6 coherent search the band widened to 91-573 Hz, and among other supernova remnants coherently searched in S6 data, band widths varied from ∼150 Hz to ∼2000 Hz. The ongoing Einstein@Home semi-coherent search for Cas A is searching up to 1000 Hz. The boundaries of the band to be searched are defined by the intersections of the age-based strain limit above and the expected strain sensitivity of the search. All of the nominal values of h_0^{age} in Table 6 can be beaten with advanced detector data over at least some frequency band.

Assuming that pulsations are indeed observed in one or more electromagnetic bands in follow-up, there is immediate astrophysical payoff. If the gravitational wave frequency is twice that of the rotation frequency, with no GW emission observed at the rotation frequency, then a rotating triaxial ellipsoid model is favored. On the other hand, if the two frequencies agree, then precession probably plays an important role. If the GW frequency is approximately 4/3 the rotation frequency, then r-modes are strongly favored, yielding direct information on the interior fluid motion of a neutron star. In the conventional scenario of $f_{\rm GW} = 2f_{\rm Rot}$, the phase of the GW emission relative to the phase of the electromagnetic pulse(s) gives insight into the relationship between the stellar quadrupole asymmetry and the magnetic poles. Similarly, the inclination angle and transverse polarization angle of the star's spin axis can be inferred from the GW emission, and if deep electromagnetic observations detect a pulsar wind nebula around the star, it may be possible to compare the GW- and EM-inferred angles. More generally, one can envision a sequence of beyond-discovery collaboration publications, as we and electromagnetic astronomers continue studying a source that (we hope) remains detectable throughout the ADE.

Another interesting follow-up possibility after a discovery is to exploit advanced interferometer narrow-banding to improve parameter estimation. Whether the potentially improved signal-to-noise ratio from narrow-banding is worth the cost in sensitivity in other bands and in interferometer downtime is difficult to judge in advance, however. Unless all interferometers are narrow-banded together, the signal-to-noise ratio gain from narrow-banding over broadband operation may be modest, and given current astrophysical knowledge, there is no compelling reason to choose narrow-banding in a directed isolated-star search. We can't even say with confidence that detection of an unknown high-frequency millisecond star is more or less likely than detection of a young, low-frequency star. (Even for low-mass X-ray binaries, such as Scorpius X-1, which were often cited in the past as targets for narrow-banding, it is not at all clear that narrow-banding is truly sensible.)

3.17.3 Search Description

At present two mature approaches are available for directed searches. The first, used for the S5 Cas A search [507] and for the S6 SNR search, is based on computing the $\mathcal{F}$-Statistic[478] for a single coherence time of O(10 days) and includes an explicit search over the 2nd derivative of the GW frequency. A variation of this first pipeline, based on a more efficient, barycenter-resampled $\mathcal{F}$-Statistic [508] has also been used in an ongoing pilot search for a source in globular cluster NGC 6544. Henceforth this approach based on a single coherent measurement (with or without barycenter-resampling) will be called the *coherent* method.

The second mature approach is based on stacking $\mathcal{F}$-Statistic values semi-coherently from many separate segments of data. This approach (henceforth called *semi-coherent*) was pioneered in the S5 galactic center search [509] but without an explicit search over $\ddot{f}$, and is now being used in an Einstein@Home S6 SNR search that does search over a range of $\ddot{f}$ values consistent with SNR age.

In general, one expects the semi-coherent search to yield better strain sensitivity than the coherent search, even for fixed computation cost, for a long data run and uniform sensitivity throughout the run. Increasing computational resources through the use of Einstein@Home adds further to the potential gain in sensitivity. Figure 8 shows a comparison of directed-search efficiencies for coherent and semi-coherent methods from stage 3 of the ongoing mock data challenge (see below). In the long term it is likely that future directed searches will rely on the semi-coherent method, but for the 3-month O1 run, the intrinsic benefit of the semi-coherent approach is reduced from what it would be for a yearlong run, and if there is substantial improvement in detector performance during the run (as has often been the case), then a coherent search of the most sensitive data subset may give the best sensitivity. For these reasons, and given the long experience of the CW group in coherent directed searches, we expect to continue supporting this approach for the time being.

In the event of an interesting candidate detection, either search can benefit from "zooming in" on interesting sources, using the "Nomad" approach [482] used in the full-S5 Einstein@Home paper [483]. In the case of an isolated neutron star detection, this would allow us to focus promptly on pinning down the source parameters via zooming. There are several important considerations to bear in mind with respect to such follow-ups. First, the signal-to-noise ratio increase from zooming can be an order of magnitude, allowing relatively precise parameter estimation compared to what is found upon first detection. Second, one can

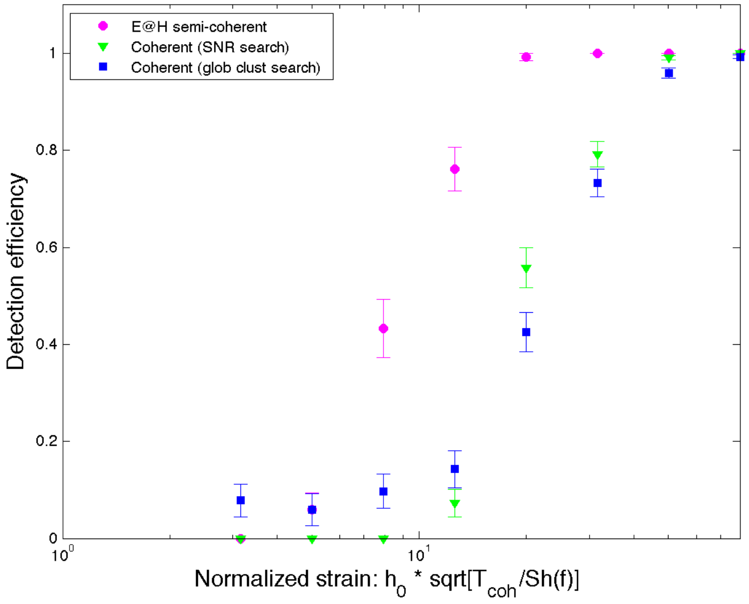

Figure 8: Approximate detection efficiencies of coherent pipelines with (glob cluster) and without (supernova remnant) resampling; and of a semicoherent directed search pipeline, based on mock data challenge studies. The semi-coherent results are taken from a cluster emulation of what can be achieved with Einstein@Home, using a 20 mHz search band and ignoring $\ddot{f}$.

then go back to older, less sensitive data and still find a source with very good signal-to-noise ratio that had been missed previously; moreover, that expanded time baseline itself offers improved parameter precision, at the cost of requiring an expanded search range to account for extrapolation errors due to uncertainties in the initial (discovery) parameters. Finally, although the semi-coherent search in all-sky searches has a sensitivity that improves only as the fourth root of observation time, once a true signal has been pinpointed, then the signal-to-noise ratio and parameter estimation from coherent followups improve as the square root of observation time (assuming stable detector noise during the data run).

Electromagnetic follow-ups play a critical role in assessing the scientific impact of a detection in this search. Before publishing a CW discovery, we wish to consult electromagnetic partners with access to radio, X-ray and γ-ray telescopes on what they can see in the direction of the source. With a year's observation time and good signal-to-noise ratio, we expect to achieve an angular resolution of O(arcsec), which should suffice for excellent electromagnetic follow-up. With an ephemeris provided by GW measurements, these follow-ups can be more sensitive in detecting pulsations for an already-known star or in confirming a new source. We already have partnerships in the radio, X-ray and γ-ray communities for obtaining ephemerides for our targeted pulsar searches.

Neither the coherent nor the semi-coherent pipeline is thought to be well suited to a search for Supernova 1987A because such a young object, if emitting detectable gravitational radiation, is highly likely to be spinning down at such a rate that explicit searching over higher-order spindowns would be necessary. A more robust approach based on cross correlation techniques [510] is under development, but is not expected to be ready for O1 use.

The criteria for prioritizing targets is likely to follow past practice, with highest priority given to Cas A and other young, close SNRs, to the galactic center and to close globular clusters (pending new electromagnetic observations). One concern is searches for young pulsars is glitching that leads to small changes in frequency and its derivatives. These happen infrequently and would be unlikely to affect a short-coherence-time search, but become more of a concern for a semi-coherent search over many months.

3.17.4 Result Validation Plan

As discussed in section 3.17.3, "zooming" in directed-search candidates offers the potential for a large increase in signal-to-noise ratio for a source that follows the signal model of an isolated star with constant spindown over the observation span. With high SNRs confirmed by the three independent targeted-search pipelines described elsewhere, there will be no uncertainty as to whether a detection has been made.

Given the significant differences in achievable sensitivities between coherent and semi-coherent searches, validating upper limits from the semi-coherent search in the absence of detection will rely mainly on careful review.

3.17.5 Publication Plan

Following the observing scenarios document [484], we assume here a 3-month run in 2015, a 6-month run in 2016–2017, and a 9-month run in 2017–2018. In each run we assume an intrinsic strain sensitivity significantly improved over previous runs. Hence we envision at least one directed publication from each run and perhaps more, separated by source type (SNR, galactic center, other) and with a possible quick-look publication based on coherent searches (within $\sim$3-6 months of run completion), followed later by a deep-search paper using the semi-coherent pipeline. Whether or not to publish the latter paper will depend on the improvement in depth and breadth of the search, and on the timeliness of the results with respect to new data sets. We expect the publication rate for directed searches to be more rapid in the ADE than for the S5–S6/VSR1–4 era. We will have established reviewed, tagged pipelines, based on those used for the final S6/VSR2–4 searches. Although refinements will likely be explored in the ADE, there will be default pipelines for which the reviews can focus on analysis results, not source code. The expected pace of the data runs and the sensitivity improvements make it attractive to put a premium on rapid analysis and review.

If detector sensitivity does not improve significantly from one year to the next, *e.g.*,the 2016 run is comparable to that in 2015, then the searches will nonetheless be carried out, but the journal (or perhaps conference proceedings) to which results are submitted for publication may be adjusted accordingly.

A discovery publication might include co-authors from those communities who carry out follow-ups at our request (whether or not they confirm the source), a blanket request we envision going out through our existing liaisons in those communities[17]. Inclusion of these authors in the discovery paper (as opposed to later papers or companion papers), would depend on the timeliness and astrophysical importance of their results. We fully expect results from some observers that would be rapid and enhance the astrophysical impact of the discovery paper.

---

[17] We envision a very lightweight MOU arrangement with electromagnetic observers. In the event of a GW discovery, observers would be offered the opportunity to sign an *ad hoc* MOU setting out publications ground rules, in order to obtain the GW source coordinates and ephemeris. While leaks of such parameters to MOU non-signers are likely, the pulsar community has established an etiquette for protecting priority on discoveries that makes such leaks benign, we believe.

3.17.6 Technical requirements and development plan

Pipeline	Published observational results to date	Search& review status
Coherent	*APJ* 722 (2010) 1504 (S5)	Review complete in Q4 2014 Review of resampling enhancement complete in Q1 2015
Semi-coherent	*PRD* 88 (2013) 102022 (S5)	Review completed 2013 Review of $\ddot{f}$ incorporation complete in Q3 2015

Coherent Pipeline
The first iteration of this pipeline was originally reviewed for the S5 Cas A search [507]. An updated version of the pipeline, better suited to multiple sources, has received extensive review for the S6 supernova remnant search, a process coming to a close in early Q4 2014. A barycentered-resampled version of the pipeline is being used for the S6 globular cluster search and will begin review in Q4 2014.

Semi-Coherent Pipeline
The first iteration of this pipeline was originally reviewed for the S5 galactic center search [509]. An updated version of the pipeline, including a search over non-zero $\ddot{f}$ began review in Q2 2015 in the context of the Einstein@Home S6 all-sky "bucket search" which shares a code base with this one. Completion of the directed-search part of the code is expected in Q4 2015.

Mock Data Challenge
The performances of various isolated-star pipelines are being evaluated via a standardized set of software signal injections into a parallel stream of real data from the LIGO and Virgo interferometers (S6 and VSR2/4). Two sets of 3000 injections each of isolated stars over the sky and frequency band have been created, and a multi-stage mock data challenge is under way, with approximately six months allotted to each stage. This is the same mock data challenge being used in the all-sky continuous-wave searches, and is described in more detail in that observing plan. Based on results from stage 3 of the MDC, the potential of improved sensitivity with an Einstein@Home-hosted semi-coherent search for a data run of length comparable to the S6 run was tentatively confirmed. Note, however, that unlike the coherent pipelines, which searched over 1-Hz bands and over both $\dot{f}$ and $\ddot{f}$, the semi-coherent search was carried out on the CIT cluster in only a 50-mHz band and with no search over $\ddot{f}$, for practicality in this cluster emulation of an Einstein@Home search. Further study is needed to evaluate and optimize the performance of the semi-coherent search in the event of an LVC-cluster-hosted search.

3.18 Directed Searches for Scorpius X-1 and Other Known Binary Stars

3.18.1 Abstract

Gravitational waves (GWs) are small ripples in the geometry of space-time, predicted to exist by Einstein's Theory of General Relativity, and propagating at the speed of light. Rapidly spinning neutron stars in our galaxy may emit such waves if they are not perfectly symmetric around their spin axis. Scorpius X-1 is a binary system believed to include a neutron star experiencing a high accretion rate from its lighter companion star. This accretion of material could lead to a sustained deviation from pure axisymmetry of the star's mass distribution or lead to excitations of oscillatory vibrations, either of which can cause gravitational wave emission. A comprehensive search has been carried out in data from the LIGO and Virgo gravitational wave interferometers for such radiation from Scorpius X-1. No gravitational-wave signal was detected over a broad frequency range, allowing scientists to constrain the degree to which accretion distorts the shape of the Sco X-1 neutron star or excites oscillations. New data to be taken in the coming year after improvements to the LIGO and Virgo detectors is expected to be still more sensitive, allowing scientists to probe Scorpius X-1 more deeply, in the hope of detecting continuous gravitational waves for the first time.

3.18.2 Scientific Justification

Scorpius X-1 (Sco X-1) offers one of our most intriguing potential sources of continuous gravitational waves. A low mass X-ray binary (LMXB) system, it emits X-rays copiously, indicating a high accretion rate from its companion star. Such accretion could lead to a sustained non-axisymmetry large enough to produce detectable gravitational radiation, or it could lead to excitation of r-modes [511, 512, 513, 514, 457, 451, 456, 515]. The X-ray luminosity from the accretion is a measure of mass accumulation at the surface. As the mass rains down on the surface it can add angular momentum to the star, which in equilibrium may be radiated away in gravitational waves. Hence one can derive a torque-balance limit [511, 512, 451]:

$$h_{\mathrm{torque}} \quad \sim \quad (5\times10^{-27})\sqrt{\left(\frac{600\ \mathrm{Hz}}{f_{\mathrm{GW}}}\right)\left(\frac{\mathcal{F}_{\mathrm{x}}}{10^{-8}\ \mathrm{erg/cm^2/s}}\right)} \tag{5}$$

where $\mathcal{F}_{\mathrm{x}}$ is the observed energy flux at the Earth of X-rays from accretion. Note that this limit is independent of the distance to the star. Sco X-1 is the brightest steady-state X-ray source in the sky (outside of the Sun) with $\mathcal{F}\approx 4\times10^{-7}$ erg/cm^2/s, yielding the nominal torque-balance strain relation:

$$h_{\mathrm{torque}}^{\mathrm{ScoX1}} \quad \sim \quad (3\times10^{-26})\sqrt{\left(\frac{600\ \mathrm{Hz}}{f_{\mathrm{GW}}}\right)} \tag{6}$$

The notion of gravitational wave torque equilibrium is potentially important, given that the maximum observed rotation frequency of neutron stars in LMXBs is substantially lower than one might expect from calculations of neutron star breakup rotation speeds ($\sim$1400 Hz) [516]. It has been suggested [517] that there is a "speed limit" governed by gravitational wave emission that governs the maximum rotation rate of an accreting star. In principle, the distribution of frequencies could have a quite sharp upper frequency cutoff, since the angular momentum emission is proportional to the 5th power of the frequency. For example, for an equilibrium frequency corresponding to a particular accretion rate, doubling the accretion rate would increase the equilibrium frequency by only about 15%.

A number of mechanisms have been proposed by which the accretion leads to gravitational wave emission. The simplest is localized accumulation of matter, *e.g.*, at the magnetic poles (assumed offset from the rotation axis), leading to a non-axisymmetry [451]. One must remember, however, that matter can and will diffuse into the crust under the star's enormous gravitational field. This diffusion of charged matter can be slowed by the also-enormous magnetic fields in the crust, but detailed calculations [518] indicate the slowing is not dramatic. Another proposed mechanism is excitation of r-modes in the fluid interior of the star [513, 451, 514, 456], with both steady-state emission and cyclic spinup-spindown possible [515, 519].

While Sco X-1 is probably the most promising source of its type, given its bright X-ray emission, large uncertainties in neutron star physics and in the accretion process argue for searching other X-ray binary systems for continuous gravitational waves [520]. We envision a program in which the search for Sco X-1 is given highest priority (CW group commitment to deliver publishable results for each observing run), but in which attention is given to other X-ray binary systems, including accreting millisecond pulsars, burst oscillation sources and stars with detected quasi-periodic oscillations, as time and computational resources permit. Any X-ray binary system observed to burst during an ADE observing run would be considered for a directed search. Other potential sources are Fermi satellite gamma-ray sources at high galactic latitudes (suggesting nearness), sources with spectral and temporal properties consistent with pulsars, but for which no pulsations have been detected. For such sources, with neither source frequency nor orbital parameters known, not all search pipelines described below can be sensibly applied.

Three other binary systems are especially intriguing. Two (XTE J1751-305 and 4U 1636-536) are thought to be LMXBs with neutron stars accreting matter and for which a sharp spectral line has been observed in X-rays after demodulating for the known spin frequency of the star. These lines suggest a non-radial mode of oscillation which, in principle, could indicate gravitationally detectable r-modes. The inconsistency of the observed spindowns for these sources with ordinary r-mode emission, however, suggests that a different type of oscillation is being observed [521] or that the r-modes are restricted to the neutron star crust and hence gravitationally much weaker than core r-modes [522]. Were the X-ray oscillation observed in J1751 in 2002 to occur again during a full-sensitivity ADE run and to correspond to a core r-mode, it would very likely be detectable. How much weaker in gravitational waves the signal would be if due to a crustal r-mode is hard to know. For these two sources, the absence of precise ephemerides make truly targeted searches difficult, but searching only a handful of discrete, extremely narrow frequency bands makes these searches computationally cheap.

The third additional binary system of special interest, perhaps surprisingly, is Cygnus X-1 which is an HMXB thought to host a black hole of 14.8 solar masses. An interesting scenario receiving some attention in the theoretical community is the formation of a Bose-Einstein condensate (BEC) of string axions in the vicinity of a black hole [523, 524, 525]. In this scenario the axion BEC undergoes a continual emission of gravitational waves at a frequency (in the BH rest frame) determined by the axion mass and couplings. Detecting such emission in a binary system would require the same corrections for orbital Doppler modulations as for emission from a neutron star.

While it may be difficult to place much confidence in detectable gravitational-wave emission from any of these three intriguing sources, it is desirable to carry out searches for them, using our best available pipelines. The narrowband searches for J1751 and 4U 1636 are computationally quite cheap and should simply be carried out. The search strategy for Cygnus X-1, on the other hand, is not yet well understood. A recent meeting between the CW search group and the proponents [525] of the Cygnus X-1 signal indicated considerable uncertainty concerning signal frequency modulations (potentially a few percent in modulation depth) due to BEC self-interactions. Even the relatively coarse resolution of the Radiometer pipeline may be too fine for such a source. The Cygnus X-1 proponents have been encouraged to examine this issue more quantitatively.

If precise ephemeris information for Sco X-1 or other binary sources were available, *targeted* searches exploiting that information could use optimal detection statistics to exploit the full intrinsic sensitivity of the interferometers, as has been done for binary millisecond pulsars [526]. Here we consider instead sources for which there is little, if any, prior knowledge of the GW source frequency. In addition, frequency evolution is presumably governed by accretion processes subject to fluctuations; hence at least some searches must be robust against phase wandering.

In the event of GW detection with electromagnetic follow-up, successful EM detection would bring an immediate astrophysical payoff. If the gravitational wave frequency is twice that of the rotation frequency, with no GW emission observed at the rotation frequency, then a rotating triaxial ellipsoid model is favored. On the other hand, if the two frequencies agree, then precession probably plays an important role. If the GW frequency is approximately 4/3 the rotation frequency, then r-modes are strongly favored, yielding direct information on the interior fluid motion of a neutron star. In the conventional scenario of $f_{\mathrm{GW}} = 2f_{\mathrm{Rot}}$, the phase of the GW emission relative to the phase of the electromagnetic pulse(s) gives insight into relation of the stellar quadrupole asymmetry to the magnetic poles. Similarly, there is some *a priori* information available for the inclination angle and transverse polarization angle of Sco X-1's spin axis, based on Sco X-1 jet observations, allowing comparison with the corresponding inferred GW values.

3.18.3 Search Description

CW Pipelines used in targeted searches for known pulsars and in directed searches for isolated neutron stars of unknown frequency cannot be used without substantial modification in a search for a neutron star in a binary system because of the frequency modulations due to the star's orbital motion. Searches must allow for those modulations, in addition to allowing for likely phase wandering over long time scales of GW signals from accreting systems. Given the large uncertainties of the astrophysics that govern the signal's time evolution, it seems prudent to apply a diverse suite of pipelines in this search. Potential considerations include broadband coverage, sensitivity for short data stretches, e.g., during detected X-ray outbursts, sensitivity for long data stretches (steady sources), robustness against phase wandering, ability to exploit orbital parameter knowledge, ability to exploit approximate phase evolution for accreting millisecond X-ray pulsars (AMXPs) and computational cost. Highest priority will be given to deploying several pipelines to target Sco X-1 in a search up to 2 kHz, with lower priority given to deploying one or more pipelines to target other promising X-ray binaries, including J1751-305, 4U1636 and Cygnus X-1 (BEC axion scenario). An upper limit on the search band of 2 kHz provides some safety margin above the 1.4 kHz defined by twice the rotation frequency of the fastest known pulsar. There is good reason to expect the accreting systems of most interest to GW searches to sit near the upper range of pulsar spin frequencies, and neutron stars should, in principle, be able to spin as fast as 1 kHz, for a variety of assumed equations of state [516].

Seven existing or potential pipelines in all are under consideration for these searches (one of which is a Stochastic Search Group pipeline). Five of these pipelines are mature enough to have participated fully in the first Sco X-1 mock data challenge [527]: CrossCorr, Sideband, TwoSpect, Polynomial and Radiometer (Stochastic Search Group pipeline). CrossCorr has the best strain sensitivity in this year-long search for which orbital period and phase are known well, while TwoSpect and Radiometer have strain sensitivities several times worse and Sideband has still worse sensitivity when restricted to a 10-day observation span, a limit motivated for a fully coherent search by expected source phase wandering from accretion fluctuations [528]. Based on these MDC results, Polynomial appears to offer no significant advantage for a directed binary search (but is being pursued for all-sky binary searches). All of the pipelines described here are expected to be robust against plausible fluctuations. CrossCorr, Sideband, TwoSpect and Radiometer; they also have modest computational cost and seem appropriate to include in the suite described below.

Two other pipelines were not mature enough to participate fully in the first MDC stage, but are undergoing development and are expected to participate fully in later stages: PowerFlux (developed for isolated-star searches, but adaptable to binary searches) and an Einstein@Home-based F-Statistic search. Further development of PowerFlux for this search is manpower-limited, with prospects uncertain. An Einstein@Home-based $\mathcal{F}$-Statistic search is unlikely to be ready before December 2015.

Below we describe the four search pipelines that are currently mature and likely to be run on upcoming O1 data: CrossCorr, Sideband, Radiometer and TwoSpect.

The Cross-Correlation (CrossCorr) pipeline[480, 529] was developed as an improvement to the radiometer search method, tailored to periodic gravitational-wave sources rather than unmodelled stochastic backgrounds. By using the periodic signal model, it is able to look for correlations not just between data from different detectors at the same times, but also between data taken at different times, from the same or different detectors. By restricting the terms in the cross-correlation to include only SFTs (Short Fourier Transforms) separated in time by less than some maximum lag time T_{lag}, the pipeline can be tuned to trade off sensitivity versus computing cost by adjusting the time lag parameter, as discussed in section 3.18.3.

The expected cross-correlation between two SFTs, which is used to coherently combine all of the included SFT pairs, depends on the signal parameters, which means the method needs to search over a grid of parameter-space points in the phase-modulation parameters. For Sco X-1, where the period is well enough known that it need not be searched over, the parameter space consists of spin frequency, projected semimajor axis, and time of ascension. When T_{lag} is small compared to the orbital period of 19 hours, the density of points needed in each dimension is proportional to T_{lag}. Since the number of SFT pairs for a fixed observation time also scales with T_{lag}, the computing cost should scale like T_{lag}^4. On the other hand, the h_0 to which the search is sensitive scales like $(T_{\text{lag}}T_{\text{obs}})^{-1/4}$. Note that the sensitivity of the search is not affected by the choice of SFT duration T_{sft} (shorter SFTs means more SFT pairs, each of which contributes less to the overall sensitivity). However, the number of SFT pairs for fixed T_{lag} and T_{obs} will scale like T_{sft}^{-2}, so we expect the computing cost to depend similarly on the choice of T_{sft}, which makes it advisable to use the longest SFTs we can get away with. The SFT duration is limited by the validity of the linear phase model, which assumes that the signal can be characterized by a single frequency during each SFT. An empirical Monte Carlo measurement has shown that for $f \lesssim 400\,\text{Hz}$, the sensitivity degrades when T_{sft} exceeds 600 seconds. The computing cost will also increase with frequency because the density of templates (at a given mismatch) in the two orbital parameter directions is proportional to frequency.

The Sideband pipeline [530, 528] is based on the venerable $\mathcal{F}$-Statistic [478] applied over coherence times long compared to the orbital period of the source (18.9 hr for Sco X-1), where Doppler demodulation is carried out to correct for the Earth's motion relative to the Solar System Barycenter, but no demodulation is carried out for the source's orbital motion. The remaining time-dependent frequency modulation in the detector frame allows decomposing the signal into an infinite sum of frequency modulated sidebands. Under the conditions that the observation time is $\gtrsim$ 3 orbital periods and that there is negligible drift in the intrinsic spin frequency of the source (i.e $\dot{\nu} \lesssim T^{-2}$ where T is the observation time) this sum is truncated leaving $M \sim 4\pi f_{\text{gw}} a \sin i/c$ frequency resolvable sidebands where f_{gw} is the intrinsic GW frequency and $a \sin i/c$ is the orbital semi-major axis projected along the line of sight and normalised by the speed of light. Each of the sidebands is uniformly separated from the next by $1/P$ where P is the orbital period, and any orbital eccentricity acts only to redistribute power amongst existing sidebands.

By computing the $\mathcal{F}$-statistic for a given sky position and for a long enough observation time, a signal of adequate amplitude can be extracted by incoherently summing together the $\mathcal{F}$-statistic at each sideband frequency [531, 532]. This is equivalent to convolving the detection statistic frequency series with a "comb" of unit amplitude spikes separated in frequency by the inverse of the orbital period. The incoherent summing to create a "C-Statistic" makes this a non-optimal strategy, but one that can have greater sensitivity to GW signals than a matched–filter approach because its observation length is not computationally limited. When using this approach, the parameter space resolution (and hence the number of search templates) is significantly reduced. It should also be noted that the sensitivity of this search to GWs scales with $T^{-1/2}$, as with a coherent search (and unlike other incoherent searches); the sensitivity, however, also scales as $M^{-1/4}$ (M is the number of sidebands) and hence high frequency, large orbit sources will be harder to detect with this method.

Of the LMXBs it is those of unknown spin frequency, but known sky location and orbital period, to which this search is most suited. The remaining orbital parameters, semi-major axis, time of passage through the ascending node ("time of ascension"), eccentricity, *etc.*, are generally quite poorly known. This scenario suits this search, as the sensitivity is relatively insensitive to all orbital parameters except for the orbital period.

The search code and associated pipeline are complete, and a methods paper detailing the search was published in early 2014 [530]. Search results obtained from a ten-day stretch of S5 data have been reviewed and published [528]. A possible future enhancement to this pipeline is double-demodulation, in which an approximate orbital demodulation is performed in computing the $\mathcal{F}$-Statistic, thereby reducing the number of peaks over which to sum in the creation of the C-Statistic. Semi-coherently combining C-Statistic values from multiple observing times during a data run is another possible enhancement.

The Radiometer pipeline[533] is one of two pipelines used in the Stochastic Group's directed search for an anisotropic gravitational-wave background, described in detail in the corresponding search plan (section **??**). Briefly, the method computes cross-correlations among pairs of detectors after correcting for the expected GW time delay for the source direction and time of observation. The full observation span of a data set is parsed into 60-s, Hann-windowed, 50%-overlapping segments, which are coarse-grained to achieve 0.25Hz resolution. No Doppler correction is applied to the source frequency, given the relatively coarse frequency resolution of the search, although at higher frequencies there can be substantial signal leakage [527] across frequency bins.

The present search program is quite mature, having been used previously for S4 and S5 searches [533, 534], and will be used for a forthcoming paper describing a search in S6/VSR2-4 data. A modest pipeline refinement to reduce sensitivity loss from leakage across bins at high frequencies is under development and will require additional review.

The TwoSpect pipeline[535, 536] was originally developed for an all-sky search for continuous waves from unknown binary systems, but has been more recently adapted for directed searches. The TwoSpect method [535] relies on computing two successive power spectra of the calibrated strain data channel, hence the name TwoSpect. First, the program computes a power spectrum of the time series data, where the coherence time for the first power spectrum depends on the region of parameter space to be covered. For shorter-period binary systems, a shorter coherence time for each SFT is used. These choices ensure the signal remains in one bin during most of each SFT interval. The SFTs are then demodulated based on the sky location, correcting for the Earth's daily and annual motions. The SFTs are noise- and antenna-pattern-weighted in the same manner as for the PowerFlux algorithm. The initial power spectra are mean-subtracted within search bands to ensure that the powers computed in the second-stage spectra are distributed as a χ^2 distribution with two degrees of freedom. The second spectra are taken over a long observation time, e.g., 1 year, for each bin in the first set of spectra. The resulting frequency-by-frequency plot is matched against templates based on the expectation for a nominal CW signal from a binary system. The current pipeline has no explicit search over polarization and uses a polarization weighting for a circularly polarized source.

An all-sky binary search (in S6 and VSR2-3 data) over a $\sim$500-Hz band was recently published [536], which included a low-frequency search for Sco X-1. A dedicated S6 search for Sco X-1 up to $\sim$2 kHz using the directed pipeline is under way. Some additional, modest review of the pipeline refinements for directed searches will be needed. In the longer term, additional pipeline enhancements are under way or planned: coherent summing of SFTs across detectors, source polarization searching and orbital phase exploitation for known binaries like Sco X-1. These enhancements are expected to improve strain sensitivity significantly, but will require additional review. It is not yet clear how many of these enhancements will be ready and reviewed by the start of O1.

Comparison of pipelines

As described below (section 3.18.6), the first stage of a **mock data challenge** has evaluated the pipeline performances (detection and upper limits strain sensitivity, parameter estimation), with a methods paper recently submitted for publication [527]. Figure 9 shows the detection efficiency of the five pipelines to complete the full MDC analysis to date, when tested on a year-long simulated sample of Gaussian white-noise data for H1, L1 and V1, with data segmentations taken from prior LIGO science running. (The

Sideband search used only a 10-day subset of the data, under the assumption that a longer search would require Sco X-1's frequency evolution to be more stable than is astrophysically plausible.) Based on this efficiency comparison, it is likely that published upper limits (assuming no detection) from the O1 run will come from the CrossCorr pipeline, with backup validation provided by TwoSpect and Radiometer, albeit with significantly worse sensitivity. It is not clear that results from the Sideband pipeline, in its current form and with a continued reliance on a coherence time no greater than 10 days, can contribute usefully. Similarly, Polynomial (which has been optimized for all-sky binary searches) will likely not contribute.

It should be noted, however, that the pipelines have different tolerances for astrophysical uncertainties. Specifically, Sideband assumes frequency stability of O(μHz) for O(10 days); CrossCorr's tolerance for frequency wandering depends on the coherence time, but O(10^5 s) lag times are compatible with frequency wandering of O(mHz) or less for the observation time span, which is also the bound on frequency wandering tolerated by TwoSpect; and Radiometer can tolerate O(500 mHz) frequency wandering. The next stage of the ongoing MDC will address this issue, along with the sensitivities of the pipelines in non-Gaussian (recolored S6) data.

Search strategy

There is little chance of observing Sco X-1 in O1 data, given the run's expected sensitivity and observation span of three months. The longer-term prospects, however, are brighter, albeit not at all assured. The mock data challenge results to date indicate that at least one existing pipeline (CrossCorr) will be able to reach h_0 sensitivities of O(5×10^{-26}) in the bucket at full aLIGO design sensitivity with one year of joint H1-L1-V1 data. From equation 5, this reaches the torque-balance limit for GW signal frequencies below 200 Hz. As discussed above, however, signals may well lie at much higher frequencies. On the other hand, the derivation of the torque-balance limit assumes the relevant radius for accretion is that of the neutron star. Given the poorly understood magneto-hydrodynamics of neutron star-accretion disk interactions, the Alfven radius may be more relevant, which could lead to higher angular momentum transfer to the star and hence a greater GW emission in equilibrium. That said, the astrophysical uncertainty can go in the other direction too, with accretion disk models that maintain torque balance without significant GW radiation[537]. Hence we have reason for tempered optimism in a Sco X-1 search at full ADE sensitivity.

Although the large number of pipelines we have in mind to use in searching for Sco X-1 may seem excessive, there are good reasons to field a broad suite of approaches, even more than is the case for other CW searches. In the case of directed searches for binary systems, considerations for running pipelines include

- Best sensitivity for long observation times
- Best sensitivity for short observation times or long observation times of highly varying sensitivity
- Best sensitivity in the case of disparate detector performances
- Best robustness against signal deviations from assumed phase model, *e.g.,* due to accretion fluctuations
- Fastest pipeline for quick looks at data
- Best robustness against non-Gaussian artifacts
- Ability to probe deeply in strain in the event of hierarchical detection ("zooming")

We expect only a few (1-3) pipelines to excel according to these criteria, allowing an eventual winnowing of pipelines, based on realistic MDC studies and based on running competing pipelines on test bands of ADE data, starting with O1. At this point, CrossCorr is likely to prevail on the long-term, most critical criterion, namely best sensitivity for long observation times, and given its structure with a tunable time lag, it is also well positioned to prevail on the important criterion of zooming capability. Pipeline teams will be asked to produce those test results before receiving allocations of substantial computing resources. The two most important criteria are best sensitivity achievable in real data and best robustness against source modeling uncertainty. Note that the use of multiple interferometers differs among the pipelines. Sideband

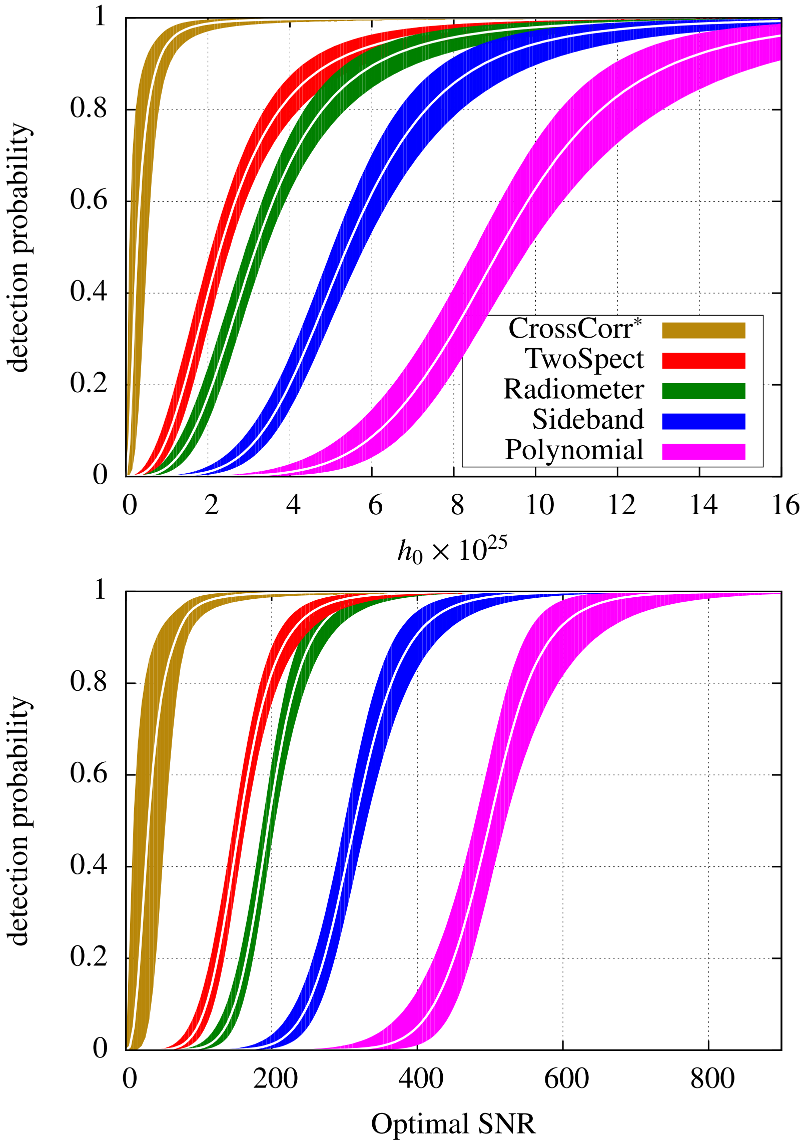

Figure 9: Detection efficiencies *vs* strain amplitude h_0 and *vs* optimal achievable SNR of five participating directed binary search pipelines in the first stage of the Sco X-1 mock data challenge.

and TwoSpect can be run in single-IFO or multi-IFO modes, allowing for coincidence detection or for single- vs multi-IFO signal-to-noise ratio consistency tests, e.g., joint detection confidence should be higher than any single-IFO confidence. In contrast, the Radiometer pipeline can only be run on an network of at least two detectors. CrossCorr can be operated with data from any number of detectors, including one, but a single-detector analysis will involve considerably fewer possible SFT pairs, and therefore be less sensitive, unless the capability is added to include auto-correlation terms. However, consistency checks for Radiometer and CrossCorr are also possible for analyses involving different detector pairs.

Electromagnetic follow-ups would play an important role in assessing the scientific impact of a detection in this search. Before publishing a CW discovery, we wish to consult X-ray astronomers to determine if a pulsation signal consistent with GW observations can be observed. A discovery publication might include co-authors from the X-ray community who carry out follow-ups at our request (whether or not they confirm the source), a blanket request we envision going out through our existing liaisons in the X-ray community[18]. Inclusion of these authors in the discovery paper (as opposed to later papers or companion papers), would depend on the timeliness and astrophysical importance of their results. We fully expect results from some observers that would be rapid and enhance the astrophysical impact of the discovery paper.

Another interesting follow-up possibility after a discovery is to exploit advanced interferometer narrow-banding to improve parameter estimation. Whether the potentially improved SNR from narrow-banding is worth the cost in sensitivity in other bands and in interferometer downtime is difficult to judge in advance, however. Unless all interferometers are narrow-banded together, the SNR gain from narrow-banding over broadband operation may be modest, and given current astrophysical knowledge, there is no compelling reason to choose narrow-banding in even a directed binary search.

3.18.4 Result Validation Plan

As discussed in section 3.18.3, zooming in on directed-search candidates offers the potential for a significant increase in signal-to-noise ratio. With high SNRs confirmed by the three independent targeted-search pipelines described elsewhere, there will likely be no uncertainty as to whether a detection has been made.

Given the significant differences in achievable sensitivities among the pipelines, however, validating upper limits from the most sensitive pipeline in the absence of detection will rely mainly on careful review.

3.18.5 Publication Plan

Following the observing scenarios document [484], we assume here a 3-month run in 2015, a 6-month run in 2016-2017, and a 9-month run in 2017-2018. In each run we assume an intrinsic strain sensitivity significantly improved over previous runs. Hence we envision at least one directed-binary search publication from each run with the first such publication featuring a Sco X-1 search. Subsequent publications for a given run could include additional X-ray binary targets or perhaps a deeper or broader search for Sco X-1 - if justified by substantially improved sensitivity or parameter space coverage. Steadily improving detector sensitivity will argue for expedited analysis of each data set, including a pragmatic approach to detector artifacts and unpromising outliers.

Although we implicitly assume here that published results from the first three runs will be upper limits on CW sources, we may well detect a signal from the directed binary searches. In that event we would, of course, publish one or more papers on the detected source and would collaborate with radio, optical, X-ray and γ-ray astronomers on deriving astrophysical insight from joint observations, as discussed in §3.18.3.

---

[18] We envision a very lightweight MOU arrangement with electromagnetic observers. In the event of a GW discovery, observers would be offered the opportunity to sign an MOU setting out publications ground rules, in order to obtain the GW ephemeris. While leaks of such parameters to MOU non-signers are likely, the pulsar community has established an etiquette for protecting priority on discoveries that makes such leaks benign, we believe.

Pipeline	**Published observational results to date**	**Search & review status status**
CrossCorr	(Only mock data analyzed so far)	Review to begin Q3 2015
Sideband	*PRD* 91 (2015) 062008 (S5)	Reviewed
Radiometer	*PRD* 76 (2007) 082003 (S4) *PRL* 107 (2011) 0271102 (S5)	S6 / VSR2-4 search under way; mods to be reviewed Q2 2015
TwoSpect	*PRD* 90 (2014) 062010 (S6/VSR2) (all-sky binary pipeline)	S6 Sco X-1 search under way; mods to be reviewed Q3 2015

Table 7: Summary of directed binary pipelines, including published observational results from previous data runs, and the current search and review status.

If detector sensitivity does not improve significantly from one year to the next, *e.g.*, 2016 running is comparable to that in 2015, then the searches will nonetheless be carried out, but the journal (or perhaps conference proceedings) to which results are submitted for publication may be adjusted accordingly.

3.18.6 Technical requirements and development plan

Pipelines status and review

CrossCorr has been implemented as lalapps search code within the last year and will soon undergo review, in preparation for O1.

Sideband has been used in a recently published search of LIGO S5 data [528].

Radiometer has been used in published searches of LIGO S4 data [533] and S5 data [534], with an S6 / VSR2-4 search nearing completion in spring 2015. Modest enhancements have been recently implemented to correct for signal leakage across bins at high frequencies, which require additional review.

TwoSpect has been used in its original, all-sky implementation in a published search of LIGO S6 and Virgo VSR2 data [536]. Modest review will be needed for recent minor modifications for use in a directed search, tested in the Sco X-1 MDC [527] and being used in a LIGO S6 search. More substantial review will be needed for additional ongoing enhancements to support coherent detector combination, explicit polarization searching and orbital phase exploitation that improve sensitivity to known binary systems.

Mock Data Challenge The performances of the various pipelines used in this search are being systematically evaluated via a multi-stage mock data challenge. The first stage was based on injection of 100 Sco X-1 signals of various frequencies ranging up to 1500 Hz into one year of Gaussian white noise of amplitude spectral noise density 4×10^{-24} $\mathrm{Hz}^{-\frac{1}{2}}$, where realistic data segmentation was simulated using segmentation from past LIGO science running. An open set of 50 of the injections was used for tuning the parameters of the pipeline that determine detection efficiency and upper limit estimation. Then the pipelines were tested on a closed (blind) set of 50 injections. The resulting performances of the pipelines are summarized as detection efficiency curves in figure 9. A methodological paper describing the results of the first stage of the MDC was recently submitted for publication [527].

The next stage of the MDC, planned for summer 2015, will include evaluation of pipeline performance in the presence of significant phase wandering of the source and in the presence of non-Gaussian noise.

3.19 All-sky Searches for Spinning Neutron Stars in Binary Systems

3.19.1 Abstract

Gravitational waves (GWs) are small ripples in the geometry of space-time, predicted to exist by Einstein's Theory of General Relativity, and propagating at the speed of light. Rapidly spinning neutron stars in our galaxy may emit such waves if they are not perfectly symmetric around their spin axis. Their deviation from such axisymmetry is usually characterized by a parameter called *ellipticity* which, for example, might be as large as 10^{-6} for a broad, 1-cm-high *bulge* across the surface of the 20-km-diameter star. A comprehensive search has been carried out in data from the LIGO and Virgo gravitational wave interferometers for signals from spinning neutron stars in binary systems. No such signal was detected, allowing scientists to rule out rapidly spinning neutron stars with ellipticities greater than 10^{-5} anywhere within 2000 light years of the Earth. Those limits are more stringent (10^{-6}) for neutron stars within 200 light years. We describe here the plans to extend the all-sky search for binary spinning neutron stars to the next generation of ground-based gravitational wave detectors. New data, to be taken in the coming year after improvements to the LIGO and Virgo detectors, is expected to be still more sensitive, allowing scientists to probe deeper into the galaxy and down to smaller ellipticities, in the hope of detecting these expected continuous waves for the first time.

3.19.2 Scientific Justification

Rapidly rotating neutron stars (NS) are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band. These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [451, 452, 453], magnetic deformations [454, 455], unstable r-mode oscillations [451, 456, 457], and free precession [458]. A review of these emission mechanisms can be found in [459]. Here, we focus on the all-sky search for unknown neutron stars in binary systems. The number of undiscovered, electromagnetically quiet neutron stars within 5 kpc can be estimated to be $O(10^6 - 10^7)$ from the neutron star birth rate [460], although it is likely that only a tiny fraction would both be rotating fast enough to be accessible to LIGO [461] and remain bound to the galaxy over the age of the galaxy [462]. Only $\sim$ 2000 radio or x-ray galactic pulsars have been discovered so far [463].

Neutron stars in binary systems are of particular interest because of the phenomenon of "recycling" in which a companion star accretes matter onto the neutron star, imparting angular momentum to it and speeding it up. Such accretion is observed, for example, in low mass X-ray binary systems, such as Scorpius X-1, and most observed millisecond pulsars observed in radio, X-rays and gamma rays, reside in or once resided in systems where the accretion has stopped, but where the neutron stars retain a high angular velocity [461]. The fraction of known millisecond pulsars ($f_{\rm rot} > 100$ Hz) that are binary is more than half, and the fraction of pulsars with $f_{\rm rot} > 400$ Hz that are binary is more than 3/4. The fraction of all known binary pulsars that are millisecond pulsars is $\sim$70% [463].

Unfortunately, neutron stars in unknown binary systems also present extreme challenges for continuous waves searches because the unknown orbital characteristics produce unknown modulations of the source frequency in the Solar System Barycenter (SSB), in addition to calculable modulations due to the Earth's motion with respect to the SSB. As is well known, even the calculable modulations for an assumed source frequency make an all-sky search for unknown isolated stars a formidable computational challenge. Adding unknown modulations makes the problem all the more difficult.

The recent publication [536], however, of the first all-sky binary CW upper limits (based on S6 and VSR2-3 data) indicates this challenge can be met, and there is good reason to believe that the strain sensitivity of future searches will improve significantly with further algorithm development.

Assuming that pulsations are indeed observed in one or more electromagnetic bands in follow-up, there is immediate astrophysical payoff. If the gravitational wave frequency is twice that of the rotation frequency, with no GW emission observed at the rotation frequency, then a rotating triaxial ellipsoid model is favored. On the other hand, if the two frequencies agree, then precession probably plays an important role. If the GW frequency is approximately 4/3 the rotation frequency, then r-modes are strongly favored, yielding direct information on the interior fluid motion of a neutron star. In the conventional scenario of $f_{\rm GW} = 2f_{\rm Rot}$, the phase of the GW emission relative to the phase of the electromagnetic pulse(s) gives insight into relation of the stellar quadrupole asymmetry to the magnetic poles. Similarly, the inclination angle and transverse polarization angle of the star's spin axis can be inferred from the GW emission, and if deep electromagnetic observations detect a pulsar wind nebula around the star, it may be possible to compare the GW- and EM-inferred angles. More generally, one can envision a sequence of beyond-discovery collaboration publications, as we and electromagnetic astronomers continue studying a source that (we hope) remains detectable throughout the ADE.

Another interesting follow-up possibility after a discovery is to exploit advanced interferometer narrow-banding to improve parameter estimation. Whether the potentially improved SNR from narrow-banding is worth the cost in sensitivity in other bands and in interferometer downtime is difficult to judge in advance, however. Unless all interferometers are narrow-banded together, the SNR gain from narrow-banding over broadband operation may be modest, and given current astrophysical knowledge, there is no compelling reason to choose narrow-banding in an all-sky binary-star search.

3.19.3 Search Description

As in all-sky searches for isolated neutron stars, it is prudent to establish and maintain more than one search pipeline, preferably with somewhat different approaches, given our relative ignorance of the underlying astrophysics of neutron stars. In this case, the unknown orbital dynamics of the system only increases the uncertainty as to which regions of parameter space are most promising. We aim to search for GW signal frequencies as least as high as 1000 Hz and to cover orbital periods from as short as a few hours to as long as several months.

Below we describe two existing pipelines, one mature (TwoSpect) with published results (but with significant enhancements planned) and the other (Polynomial) now approaching maturity.

The TwoSpect pipeline[536, 535] relies on computing two successive power spectra of the calibrated strain data channel, hence the name TwoSpect. First, the program computes a power spectrum of the time series data, where the coherence time for the first power spectrum depends on the region of parameter space to be covered. For shorter-period binary systems, a shorter coherence time for each SFT (Short Fourier Transform) is used. These choices ensure the signal remains in one bin during most of each SFT interval. The SFTs are then demodulated based on the sky location, correcting for the Earth's daily and annual motions. The SFTs are noise- and antenna-pattern-weighted in the same manner as for the PowerFlux algorithm [465]. The initial power spectra are mean-subtracted within search bands to ensure that the powers computed in the second-stage spectra are distributed as a χ^2 distribution with two degrees of freedom. From these demodulated spectra, a spectrogram is created, and a second spectrum is taken over each of its rows (along the time axis), to create a plane of strain power squared vs modulation frequency (horizontal) and nominal source frequency (vertical). The pixels in this plane are matched against templates based on the expectation for a nominal CW signal from a binary system. The current pipeline has no explicit search over polarization and uses a polarization weighting optimized for a circularly polarized source, a weighting which is also optimized for performance on an average over other signal polarizations.

An all-sky binary search (in S6 and VSR2-3 data) over a 500-Hz band was recently published [536], which also included a low-frequency search for Scorpius X-1. A dedicated S6 search for Sco X-1 up to $\sim$2 kHz using the directed pipeline is under way. In addition, a more ambitious effort to enhance TwoSpect's sensitivity in both the all-sky and directed modes is under way. Specific improvements under development or planned include 1) coherent summing of SFTs from multiple interferometers prior to taking the second Fourier transforms; 2) weighting for arbitrary elliptical polarization; and 3) orbital phase determination from matched filtering (or exploitation for directed searches).

The present all-sky sensitivity of TwoSpect is severely limited by its first stage of a hierarchical search, which uses a relatively insensitive but computationally efficient incoherent harmonic sum (IHS) in the frequency-frequency plane. This limitation is easily eliminated in the directed search because of the vastly smaller parameter space volume to search. In the all-sky TwoSpect search, one must search over sky location, nominal source frequency, orbital period and orbital modulation depth (circular orbit approximation), while in the search for Sco X-1, there is only the search over frequency and a highly restricted range of orbital modulation depth. A study is under way to assess whether or not shortening the SFT coherence time is sufficient to allow elimination of the IHS stage in the all-sky search, while improving sensitivity on the whole.

The **Polynomial** [538] pipeline is based on searching systematically for signals with up to 2nd order time derivatives (3rd order in phase) in short stretches of data and then correlating different stretches of data for which consistent signal parameters are detected. The method is best suited to short-period binary systems for which frequency evolution is too rapid for more conventional search methods.

The orbital movement of the neutron star around the center of gravity of the binary system may induce large and rapidly changing frequency modulations of the gravitational wave. The frequency $f_{\rm ssb}$ detected in the solar system barycenter may be modeled as

$$f_{\rm ssb} = f_{\rm gw}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right) \tag{7}$$

with $f_{\rm gw}$ the frequency of the gravitational wave in the neutron-star rest frame, γ the Lorentz contraction factor, $\vec{v}$ the velocity of the neutron star with respect to the solar system barycenter, and $\vec{n}$ a unit vector in the direction of the source. Similarly, the change in frequency per unit time may be modeled by

$$\frac{df_{\rm ssb}}{dt} = f_{\rm gw}\gamma\left(1 - \frac{d\vec{v}}{dt}\cdot\frac{\vec{n}}{c}\right) + \frac{df_{\rm gw}}{dt}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right). \tag{8}$$

Assuming that the movement of the neutron star can be described adequately by Keplerian orbits, the phase of the gravitational wave depends on 6 extra parameters (e.g., the phase in the orbit, the orbital period, the mass of the accompanying star, the eccentricity, and the angles of the major and minor axes with respect to $\vec{n}$). For short orbital periods, the derivative of the detected frequency df/dt will be completely dominated by the Doppler shift. As an extreme example, for a neutron star orbiting an object with the same mass in a circular orbit with a period of 5000 s, $df_{\rm ssb}/dt$ may be as large as $0.002 \times f_{\rm gw}/s$.

An extension of coherent search methods to use additional parameters to describe the orbital motion of the neutron star is computationally infeasible for an all-sky binary search (for coherence times in the order of 1 h, the extra number of parameter values needed to cover all likely Keplerian orbits exceed a factor of 10^9). Even a conventional hierarchical search method used for isolated stars, such as the Stack-Slide or Hough transform methods as discussed in Ref. [467] would also incur exorbitant computational costs in all-sky binary search, to accommodate the possible modulation variations.

As an alternative, a set of filters that describe the phase of the gravitational wave as a third-order polynomial in time (and hence the frequency as a second-order polynomial in time) is used. The presence of the gravitational wave may be detected by looking for the correlation of the data with these filters. The polynomial shape of the filters facilitates the analysis (a large reduction in filter parameters is obtained by relying on the fact that translating the polynomial filter in time or in frequency will give another polynomial filter in the same parameter set) and renders a complete scan over years of data computationally feasible.

The filters should be coherent over the time that they are applied, implying that third-order derivatives of the frequency of the gravitational signal should be small. If a correlation between a filter and the data exceeds a threshold and constitutes a hit, then for the hit the frequency is known as a function of time. Therefore, hits between data stretches can be correlated. Development of the algorithm and search strategy are under way; it is believed that the approach offers a means to search a parameter range that is not currently covered by other analysis techniques, i.e., waves with frequency derivatives df/dt up to 2 mHz/s and d^2f/dt^2 up to 10^{-6} Hz/s^2.

The Polynomial pipeline has been implemented and tested on simulated data with white noise (see S. van der Putten's Ph.D. thesis [538]) and has been further tested in the Scorpius X-1 mock data challenge [527]. Review of the pipeline is planned for the calendar year 2015.

Another new method, similar in approach to TwoSpect, but likely to be less computationally intensive, has been proposed [539], in which relatively short-coherence-time SFTs are used, short enough that Doppler effects due to the Earth's motion can be neglected, to construct spectrograms over the course of a data run. Autocorrelations along the time dimension can then pick out turning-point frequencies of orbital modulation, without explicit Earth-motion Doppler demodulation and without use of binary templates. It is not yet clear, however, when a pipeline based on this method will be available for evaluation.

A very recent and promising development is the implementation of sidereal folding of data in the Stochastic Group's all-sky Radiometer search. This enhancement has the potential to permit rapid all-sky, frequency-binned searches for background gravitational radiation, including from neutron stars in binary systems. First indications are that the method could be competitive with the pipelines described here and computationally much faster. It is not yet clear if a full pipeline can be ready and reviewed for O1.

As for all-sky isolated-star CW searches, "zooming in" on interesting signals should be possible via lengthening the relevant coherence time, once a hint of a signal is found, allowing the parameter space volume to shrink by orders of magnitude. If, for example, TwoSpect detected a signal using very short SFTs, lengthening the SFTs would likely improve signal-to-noise ratio and parameter estimation accuracy (except for the shortest-period binaries). Other directed binary searches, such as the Cross-Correlation pipeline [480, 529] could also be put to good use in zooming. Ultimately, one could aim for a fully coherent follow-up for non-accreting systems and perhaps quasi-coherent follow-up for accreting systems. Future stages of the Scorpius X-1 mock data challenge should yield a quantitative understanding of what is feasible, but based on what has been achieved in isolated-star follow-ups, it seems likely that signal-to-noise ratio could improve by more than an order of magnitude via zooming, ensuring a high detection confidence. In isolated-star follow-ups, zooming can lead to signal-to-noise ratios that grow more rapidly with coherence time than the corresponding fixed-false-alarm-rate threshold does from increased trials factors [468, 469, 482].

Electromagnetic follow-ups play a critical role in assessing the scientific impact of a detection in this search. Before publishing a CW discovery, we wish to consult electromagnetic partners with access to radio, X-ray and γ-ray telescopes on what they can see in the direction of the source. With a year's observation time and good SNR, we expect to achieve an angular resolution of O(arcsec), which should suffice for excellent electromagnetic follow-up. We already have partnerships in the radio, X-ray and γ-ray communities for obtaining ephemerides for our targeted pulsar searches. Confirmation in follow-up of a previously unknown source in radio, optical, X-ray or gamma-ray with pulsations consistent with the observed GW signal would add to our own detection confidence and permit the multi-messenger studies discussed in section 3.19.2. The added confidence that would come from confirming a new source not displaying pulsations is harder to assess at this point, and the absence of any EM confirmation would likely increase the effective signal-to-noise threshold at which we would feel confident in declaring a discovery. We would expect partners to be able to acquire telescope time with high priority, *e.g.,* through target of opportunity programs, if following up a credible GW signal; hence results should be timely from at least some EM groups, permitting inclusion in a discovery paper with O(weeks) delay.

Pipeline	**Published observational results to date**	**Search & review status status**
TwoSpect	*PRD* 90 (2014) 062010 (S6/VSR2-3)	S6/VSR2-3 search review completed
Polynomial		Review begins 2015

Table 8: Summary of existing CW all-sky binary pipelines, including published observational results from previous data runs and the current search and review status.

3.19.4 Publication Plan

Following the observing scenarios document [484], we assume here a 3-month run in 2015, a 6-month run in 2016-2017, and a 9-month run in 2017-2018. In each run we assume an intrinsic strain sensitivity significantly improved over previous runs. We envision one all-sky binary publication from each run by default. Subsequent publications using significantly more computing time would be considered if justified by a significant improvement in strain sensitivity or in parameter space coverage. Steadily improving detector sensitivity will argue for expedited analysis of each data set, including a pragmatic approach to detector artifacts and unpromising outliers that can be followed up in more sensitive data.

Although we implicitly assume here that published results from the first three runs will be upper limits on CW sources, we may well detect a signal from the all-sky searches. In that event we would, of course, publish one or more papers on the detected source and would collaborate with radio, optical, X-ray and γ-ray astronomers on deriving astrophysical insight from joint observations, as discussed in §3.19.3. A discovery publication might include co-authors from those communities who carry out follow-ups at our request (whether or not they confirm the source), a blanket request we envision going out through our existing liaisons in those communities[19]. Inclusion of these authors in the discovery paper (as opposed to later papers or companion papers), would depend on the timeliness and astrophysical importance of their results. We fully expect results from some observers that would be rapid and enhance the astrophysical impact of the discovery paper. The target journal for such a discovery paper might depend on whether or not other types of GW detections have been made by that time, and the decision would presumably be made at the executive/steering committee level.

If detector sensitivity does not improve significantly from one year to the next, *e.g.*, 2016 running is comparable to that in 2015, then the searches will nonetheless be carried out, but the journal (or perhaps conference proceedings) to which results are submitted for publication may be adjusted accordingly.

3.19.5 Technical requirements and development plan

Pipelines status and review

TwoSpect has been used in a published search of LIGO S6 and Virgo VSR2-3 data [536]. Hence a reviewed and tagged pipeline is ready for use in O1. Nonetheless, futher enhancements are under way, to improve strain sensitivity and parameter estimation. It is likely that at least one enhancement (coherent

---

[19] We envision a very lightweight MOU arrangement with electromagnetic observers. In the event of a GW discovery, observers would be offered the opportunity to sign an MOU setting out publications ground rules, in order to obtain the GW source coordinates and ephemeris. While leaks of such parameters to MOU non-signers are likely, the pulsar community has established an etiquette for protecting priority on discoveries that makes such leaks benign, we believe.

multi-interferometer SFT summing) will be ready and reviewed for use in O1, along with minor refinements implemented for the first stage of the Sco X-1 mock data challenge [527].

Polynomial too has participated in the Sco X-1 mock data challenge and is preparing for a review to begin in 2015.

3.20 Search for an Isotropic Stochastic Gravitational Wave Background

3.20.1 Abstract

Gravitational waves, predicted by Einstein's general theory of relativity, are ripples in the fabric of spacetime, which propagate at the speed of light. A stochastic gravitational-wave signal is formed from the superposition of many events or processes that are too weak and too numerous to be resolved individually, and which therefore combine to produce a stochastic background. A stochastic background can arise from cosmological sources such as inflation, cosmic strings, and pre-Big-Bang models. Alternatively, it can arise from astrophysical sources such as compact binary coalescences, supernovae, and neutron stars. The prime objective of the Stochastic Group is to measure the stochastic background. A comprehensive search has been carried out using data from initial LIGO and Virgo. No signal was detected, but our results constrain the energy density of the stochastic background to be $\Omega_0 < 5.6 \times 10^{-6}$ at 95% confidence [540]. Advanced detectors are expected to have about 10-times better strain sensitivity than the initial detectors, as well as to extend the sensitive band from 40 Hz down to 10 Hz. Furthermore, the number of detectors operating in a worldwide network is expected to increase, eventually including sites at LIGO-Hanford, LIGO-Livingston, Virgo, GEO-HF (at high frequencies), KAGRA (Japan), and potentially LIGO India. The significant strain sensitivity improvements and wider bandwidth will enable real breakthroughs in the searches for the stochastic background, with a potential sensitivity of $\Omega_0 < 6 \times 10^{-10}$. The detection of a cosmological background would be a landmark discovery of enormous importance to the larger physics and astronomy community. The detection of an astrophysical background is not unlikely and would also be of great interest.

3.20.2 Scientific Justification

The stochastic isotropic search targets stochastic gravitational-wave backgrounds. A stochastic background arises from the superposition of an ensemble of gravitational-wave sources each of which is too weak to be detected individually. There are two broad categories of stochastic backgrounds. In cosmological scenarios, gravitational-wave backgrounds can be created in the early universe from a variety of processes such as the amplification of vacuum fluctuations following inflation [541], phase transitions in the early universe [542, 543], cosmic strings [544, 545, 546, 547], and pre-Big Bang models [548, 549]. In astrophysical scenarios, the gravitational-wave signals from more ordinary objects, such as neutron star binaries, combine to produce a stochastic signal (see [550] for a review). Astrophysical backgrounds may arise from core collapses to neutron stars or black holes [551, 552, 553, 554], rotating neutron stars [555, 556] including magnetars [557, 558, 559, 560], phase transition [561, 562] or initial instabilities in young neutron stars [563, 564, 565, 564], compact binary mergers [566, 567, 568, 569, 570, 571] and compact objects around supermassive black holes [572, 573]. In the context of detecting a cosmological background, astrophysical sources are sometimes referred to as "foregrounds", as they may mask the cosmological contribution. One hope is that the astrophysical background has a different statistical signature and can be identified and removed. For instance, astrophysical sources may not be numerous enough to create a Gaussian "stochastic" background in the sense that the sources do not overlap in the time-frequency domain [568]. For our purposes here, we refer to both astrophysical and cosmological signals as stochastic backgrounds. Together, they constitute the target for the isotropic stochastic search.

The detection of a cosmological background would yield arguably the most exciting science possible with gravitational wave astronomy. There would be enormous interest from the wider physics and astronomy community, and the implications of the discovery would likely be far-reaching. By measuring gravitational waves from the early universe, we could infer the existence of previously conjectural objects such as cosmic strings or, even more exciting, probe physics at energy scales inaccessible through other means. LIGO/Virgo will probe energy scales of $\sim 10^9$–10^{10} GeV in the early universe [574]. The detection of an astrophysical background would also be an important discovery, though, not with the revolutionizing implications of cosmological background detection. Differentiating a cosmologically produced background from one that might be astrophysically produced will not be easy, but parameter estimation techniques could provide the ability to do so [575]; the statistical nature of the signal will also provide information [568]. By detecting the stochastic background from binary neutron stars, for example, we can learn about the properties of a large ensemble of binaries. The stochastic measurement would complement results obtained from the detection of individual binaries, and since the stochastic signal is dominated by very distant objects (redshift $z \approx$ 1–2), the stochastic search probes a different population of binaries than the nearby ones that can be detected individually. While astrophysical backgrounds are interesting in their own right, it is also important to develop an observationally informed understanding of them in order to facilitate a confident detection of a cosmological background.

The theoretical uncertainties for the amplitude of cosmological backgrounds are significant, and observational measurements by LIGO and other experiments/observatories play a crucial role in guiding theoretical models. The amplification of vacuum fluctuations is the most well-known cosmological model, and it is widely regarded as plausible. Unfortunately for gravitational-wave astronomers, the expected amplitude of this canonical model is $\Omega_{\rm gw}(f) \approx 10^{-15}$, which is significantly below what will be achieved by second-generation detectors. However, the canonical model is merely the simplest one to write down, and we should keep in mind that very little is actually known about inflation and the very early universe in general; reality may very well be different. Cosmologists have proposed models, which can produce inflation-era gravitational waves that are detectable with advanced detectors [576, 577]. Similarly, there are regions of parameter space where we may be able to detect cosmic strings [578, 579]. As for astrophysical backgrounds, it is not, perhaps, widely appreciated, but advanced detectors have a good chance of detecting a stochastic background from unresolvable compact binaries [580]. In the event of a non-detection, the upper limits obtained from a stochastic search will be of astrophysical interest. Another point worth considering is the possibility of looking at contributions to stochastic background coming from gravitational waves with nonstandard polarization [581]. While general relativity allows only for two kind of tensorial polarizations, a generic prediction of extended gravitational theories, such as scalar-tensor ones [582, 583], $f(R)$ gravity [584, 585], bimetric [586] and massive [587] gravity theories, is the presence of other physical degrees of freedom, with polarizations of a scalar or vectorial kind. Advanced detectors will probe previously unexplored parameter space in the plane of frequency vs energy density plane. The constraints on cosmological and astrophysical models will probe new and interesting parts of parameter space that are not constrained by other observations.

The isotropic search is the flagship search of the Stochastic Group, and this search will build on a series of LIGO/Virgo papers [588, 589, 578, 540, 590]. Advanced LIGO operating at design sensitivity is expected to surpass initial LIGO sensitivity by four orders of magnitude in energy density for a low frequency (less than 150 Hz) search; this is due to the increase in detector sensitivity by a factor of 10 and the reduction of the lower limit of the search from 40 Hz to 10 Hz. In addition, a factor of 10 increase in advanced Virgo's sensitivity will results in a reduction of the stochastic background energy density upper limit by at least a factor of 100 in the 1 kHz regime.

The final initial detector isotropic searches [540, 590] have been reviewed and published. A search for extended polarizations is also ongoing, and is expected to be ready for a first review at the end of 2015. Searches with advanced detectors are expected to benefit from dramatically improved sensitivity.

3.20.3 Search Description

The primary goal of the isotropic search is to estimate the energy density of the stochastic background:

$$\Omega_{\rm GW}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\rm GW}}{d\ln f} \, , \tag{9}$$

where $\rho_{\rm GW}$ is the energy density of gravitational waves, ρ_c is the critical density of the universe, and f is the frequency. This is accomplished through a well-established cross-correlation procedure, documented in [591, 592], which has served as the basis for all previous LIGO/Virgo stochastic searches, e.g., [588, 589, 578, 540, 590]. The expected signal-to-noise ratio for a stochastic search, using two detectors I and J, is given by

$$\mathrm{SNR} = \frac{3H_0^2}{10\pi^2}\left(2T\int_0^\infty df\,\gamma_{IJ}^2(f)\frac{\Omega_{GW}^2}{f^6 P_I(f)P_J(f)}\right)^{1/2} \tag{10}$$

where H_0 is the present value of the Hubble expansion rate, $\gamma_{IJ}(f)$ is the overlap reduction function (see [592]), and P_I is the strain power spectral density of detector I. Due to the factor of f^6 (and also the $\gamma(f)$), the integral is dominated by the low frequencies $\lesssim 200\,\mathrm{Hz}$.

In order to handle gaps in the data, non-stationarity, and for purposes of computational feasibility, the data for an interferometer pair are divided into many segments of equal duration (typically $60\,\mathrm{s}$), and estimators $\hat{\Omega}$ and σ_Ω are calculated for each segment. The loss in duty-cycle due to the finite segment size is small $\lesssim 1\%$. The LIGO data are resampled from $16384\,\mathrm{Hz}$ to $4096\,\mathrm{Hz}$ ($20000\,\mathrm{Hz}$ to $4000\,\mathrm{Hz}$ for Virgo) and high-passed filtered with a $9\,\mathrm{Hz}$, $n = 16$ Butterworth filter[20]. They are also Hann-windowed to avoid spectral leakage from strong lines present in the data. Since Hann-windowing effectively reduces the interval length by 50%, the data segments are overlapped by 50% to recover the original signal-to-noise ratio.The effects of windowing are discussed in [593].

The power spectral densities for each segment (needed for the calculation of $\hat{\Omega}$ and σ_Ω) are calculated using the two neighboring segments. (This approach avoids a bias that would otherwise exist due to a non-zero covariance between the cross-power and the power spectra estimated from the same data.) Furthermore, by comparing σ_Ω calculated using the neighboring segments with σ'_Ω calculated using each segment by itself (not its neighbors), we identify segments containing large noise transients and reject them from the analysis. In addition to this stationarity cut, we impose some data-quality flags (such as $30\,\mathrm{s}$ before lock-loss), a large-σ cut (rejecting particularly noisy time-periods), and we ignore frequency bins which are determined to contain instrumental correlations. The segments that pass all the data-quality cuts are combined using a weighted average to yield final estimators.

The stochastic pipeline estimates $\Omega_{\rm gw}(f)$ given some assumed power law $\Omega_{\rm gw}(f) \propto f^\alpha$. For cosmological sources, we emphasize $\alpha = 0$ limits, while $\alpha = 2/3$ is appropriate for the signal from binaries. We will report results for arbitrary spectral indices [578, 575]. Also, we will use the results to place constraints on specific models such as cosmic strings [578] and astrophysical backgrounds from binary coalescence [575]. In the event of a detection, we will employ consistency checks, spectral fitting, and tests of isotropy in order to speculate about the origin of the signal; see [575].

The search for an isotropic stochastic gravitational-wave background in a single detector is confounded by noise, and it is therefore necessary to cross-correlate the output of pairs of detectors. The detection statistic that maximizes the signal-to-noise ratio depends on the exact shape of the stochastic gravitational-wave background spectral energy density and it is shown to be optimal in the case the signal is stationary, isotropic and Gaussian. Recent mock data challenges have also proven this method is still nearly optimal for non-Gaussian astrophysical backgrounds [594, 595]. The stochastic group has developed a single data analysis pipeline, based on the optimal method, to search for an isotropic stochastic gravitational-wave background. In the event of an observed signal, the stochastic group will employ a detection checklist[21]. In the event of an observed signal the a spherical-harmonic decomposition analysis will be conducted to see the distribution of the background across the sky [534]. By summing the spherical-harmonics the isotropic signal will be recovered. While the spherical-harmonic decomposition analysis is not a completely independent second pipeline, it will provide important validation of the detection.

---

[20] http://tinyurl.com/kgbffx7

[21] https://dcc.ligo.org/LIGO-T1500251

The search for nonstandard polarizations uses the same pipeline of the standard isotropic search, with minimal modifications which take into account the different coherence structure of the expected signal, parameterized by the overlap reduction function. Only the postprocessing phase is different in a significant way, as it is designed to disentangle different degrees of freedom in order to obtain detailed upper limits and constraints on extended gravitational theories. It should be noted that the possibility of separating in a model independent way different contributions is greatly improved by using a network with more than one pair of detectors.

3.20.4 Publication Plan

By the end of the Advanced LIGO and Advanced Virgo observations, we expect to improve on our initial low frequency (less than 150 Hz) sensitivity to $\Omega_{\rm gw}$ by four orders of magnitude. In the event of an unambiguous detection, we plan to publish as quickly as we can vet the detection and produce a paper describing the implications of the discovery. If no signal is present, we plan to publish after we have achieved at least a factor of two improvement in $\Omega_{\rm gw}$ sensitivity *and* we have reached a convenient milestone in data collection, e.g., the end of a science run. (One science run with a factor of eight improvement in $\Omega_{\rm gw}$ will yield one paper, not three.) It should be noted that the upper limits to be set on the energy density of the stochastic background should evolve rapidly during the commissioning phase for advanced LIGO. Based on the predicted typical sensitivities and observational runs given in [596], we can expect to improve our S5 upper limit by a factor of 10 (to $\Omega_{\rm gw} \sim 6.9 \times 10^{-7}$) with 70 days of data in the early commissioning era. With 6 months of data in the mid era we should set an upper limit of $\Omega_{\rm gw} \sim 2.4 \times 10^{-8}$. In the late era 9 months of data should allow for an upper limit of $\Omega_{\rm gw} \sim 5 \times 10^{-9}$. Once advanced LIGO hits its target sensitivity one year of data will allow for an upper limit of $\Omega_{\rm gw} \sim 1 \times 10^{-9}$, while 3 years of data will give a limit of $\Omega_{\rm gw} \sim 6 \times 10^{-10}$. As these numbers show, rapid progress will be made at every commissioning stage, thereby justifying publishing upper limit results at the end of each observational run during the commissioning era. At the present time no upper limits for backgrounds with nonstandard polarizations have been published by the collaboration. We expect that a separate observational paper based on the exploratory study in progress will be published within a year from now. Results are expected to improve with Advanced LIGO and Advanced Virgo by a factor analogous to the standard search's one, and we plan to publish a standalone paper based on these also.

A *detection paper* will establish that a stochastic signal is present. It will discuss the statistical significance of the signal and the data-quality and sanity-check steps we have taken to ensure that the signal is real. The paper will include spectral fits and tests of isotropy. We will comment on possible origins of the signal. In the event of a *marginal detection*, we advocate collecting more data since the stochastic signal-to-noise ratio grows like the square root of observation time. If this is not possible, we advocate publishing with the data available and stressing that the apparent signal is marginal. In the event of a *null result*, we will present upper limits on energy density as well as constraints on models such as cosmic string backgrounds and compact binary backgrounds.

The stochastic pipeline is well established. It runs quickly and reliably. The isotropic results are straightforward to interpret. We have many years of experience running this analysis. Thus, we plan to analyze data as it is collected. If no signal is apparent, we plan to produce paper drafts to submit for review within $\approx 2\,\mathrm{months}$ after meeting the requirements for a new paper (see above). In the event of a detection, we will carry out additional tests to make sure the detection is real. We will also characterize the isotropy and spectral shape of the signal. We expect this extra work to take an additional two to four months. Due to the potentially far-reaching implications, we recommend publishing a detection paper in a high-visibility journal. Upper-limit papers can be submitted to a variety of journals. We leave for future discussion the question of what constitutes an important milestone as this will depend in part on the state of the literature (theoretical and observational) at the time of publication.

There is no overlap with any current search plans.

3.20.5 Technical Requirements and Development Plan

The stochastic pipeline is Matlab-based code that lives in the matapps repository[22]. The main function is stochastic.m.

The stochastic pipeline can be traced back to S1 [593]. While the pipeline has evolved since then, e.g., to carry out the stochastic directional searches, the basic isotropic code has been remarkably stable over many years and publications (see, e.g., [588, 589, 578, 540, 590]). The code is reviewed and ready to be used for advanced-detector analyses. No development is required. We currently have the filters needed to reduce the low frequency cut off to 10 Hz; the addition of data from more interferometers is trivial

The only real difference with the advanced detector pipeline versus the initial detector pipeline pertains the the low frequency cut off difference, 10 Hz for advanced detectors and 40 Hz for the initial detectors. The group has already successfully designed and tested (in the mock data challenge) a new high pass filter. Involving more detector pairs in the advanced detector era will have a computational cost, as the number of CPU is proportional to the number detector pairs; however, this will not change the fact that the computational cost for the stochastic isotropic search is very low.

A mock science data and science challenge is an on-going exercise of the stochastic group. The purpose of a recently completed mock data and science challenge [595] was to prepare the data analysis and science interpretation for the future generation of gravitational-wave experiments by advanced LIGO - advanced Virgo in the search for a stochastic gravitational wave background signal from an astrophysical origin. The results from this mock data and science challenge have shown that advanced LIGO and advanced Virgo are ready and able to make make a detection of the stochastic background within a few years of operations of the advanced detectors, given a high enough rate of compact binary coalescing events. The continuation of mock science data and science challenges will be an important part of the verification process for LIGO and Virgo when a stochastic gravitational-wave background is eventually observed.

The stochastic pipeline is ready to analyze advanced detector data. One innovation that we plan to implement for the advanced detector era is to create a stochastic monitoring webpage: stochmon[23]. Stochmon includes standard result plots as well as diagnostic plots such as coherence spectra; it will provide detector characterization information that is important for the stochastic search, as well as the status of the correlation between pairs of interferometer data. The nominal goal is for stochmon to update on an hourly basis. Once operational, stochmon will be monitored by members of the stochastic group as part of “stochastic monitoring shifts.” We will use stochmon to track the sensitivity of the stochastic search and to identify as quickly as possible detector artifacts that may affect the analysis. Currently, we are working to run a prototype version of stochmon, which we plan to test with ER7 data. We hope to have it fully functioning and vetted (by showing it produces results comparable to stochastic.m) in the coming several months. We note that while we expect stochmon to facilitate early study of stochastic data and to provide helpful detector characterization information, it is not essential to carrying out the isotropic analysis. A mock data challenge has been conducted in order to validate the different aspects of the search [595]; no cause for concern has been displayed by this exercise. A stochastic detection checklist exists[24].

---

[22] http://tinyurl.com/kpmyruv

[23] https://svn.ligo.caltech.edu/svn/sgwb/trunk/stochmon

[24] https://dcc.ligo.org/LIGO-T1500251

3.20.6 Resources

Code infrastructure. The stochastic isotropic search utilizes the Matlab-based stochastic pipeline, which lives in matapps[25]. The code is reviewed and ready to be used for advanced-detector analyses. Very minimal additions have been introduced to allow for the search of backgrounds with nonstandard polarizations, and will need nominally a review which we expect to be quite simple. These additions do not interfere with the standard pipeline. No development is required.

Person-power. The isotropic search requires the FTE of one postdoc or 1–2 graduate students plus a senior mentor. Additional members of the group, such as the chairs and the advisors of the the students/postdocs, are also actively involved in the analysis as well as data-quality efforts and validation. We do not expect the FTE requirements to grow considerably in the advanced detector era.

Detector characterization for stochastic searches. The long integration time used in the stochastic analysis can reveal noise lines. Common noise lines at two sites can occur due to similar equipment in the laboratory, or common data acquisition systems. A strong line in one interferometer, along with a large random fluctuation in the other, can also produce an apparent narrowband signal in the stochastic search pipeline. In the advanced detector era, the coherence between pairs detectors' output will be calculated in near real time; (see the discussion of stochmon above). In this way noise lines that would affect the stochastic search (and by extension, also the CW search) can be identified during the science runs, and possibly addressed at the sites. We will work with the detector characterization and continuous wave groups to identify and find the sources of noise lines using all available tools.

The calibration requirements for the stochastic group are described in [368]. They are comparable to what has been requested by other search groups for the advanced detector era.

We continue to investigate noise correlations between the LIGO Hanford and Livingston sites, as well as with Virgo. We have previously observed correlated magnetic fields in magnetometer channels at widely separated detectors [597]. The primary sources of these correlated fields are Schumann resonances [597]. Having documented this potential source of correlated noise, we will be careful to determine if correlated noise is affecting the stochastic search. Investigations are in progress to determine how well the magnetometers at LIGO and Virgo sites can measure the correlated Schumann resonance noise, and whether more antennas dedicated to Schumann resonance observations are needed. Noise subtraction techniques, especially with respect to the correlated electromagnetic noise, are being studied. If the correlated Schumann resonance magnetic fields are limiting the stochastic search, then it will be necessary to directly measure the Schumann resonance magnetic fields at each interferometer site and implement noise subtraction techniques. Data from radio frequency (RF) monitors at the sites will also be studied. Other physical and environment monitors will also be carefully observed.

Calibration. The calibration requirements for aLIGO have been detailed in LIGO-T1200562. The phase error is requested to be less than 9 degrees. The relative timing error is requested to be less than 24 μs. The amplitude error is requested to be less than 9%.

Review. The stochastic review committee has reviewed numerous upper limit papers. The code is unchanged, so code review will be trivial. Scientific conclusions about constraints placed on models because of the upper limits are always carefully reviewed, but the timescale is not prohibitive (note the pace of the S5 Nature paper [578] result). The review committee also carefully reviews and critiques the paper. There was a recent review associated with the published S6 stochastic isotropic analysis [540]. A detection result will certainly be more complicated, but there are numerous sanity checks available. The estimate is that a detection result will take an extra 2 to 4 months to review.

---

[25] http://tinyurl.com/kpmyruv

3.21 Directional Search for Persistent Gravitational Waves

3.21.1 Abstract

The stochastic directional search has several important goals. First, it provides a crucial follow-up for the isotropic search by characterizing the anisotropy present in a stochastic detection. Second, it facilitates the detection of highly anisotropic stochastic sources (e.g., clustered in the Galactic plane) that might be missed by the isotropic search. Third, it provides a robust and sensitive search for narrowband point signals from interesting sources such as nearby low-mass X-ray binaries like Sco X-1. While the isotropic background search focuses on the frequency content of the background, the directional search provides additional information on the anisotropy of the SGWB, hence providing a powerful tool to distinguish between different SGWB models. Searches for anisotropic SGWB have been carried out using data from initial LIGO and Virgo. No signal was detected, but our results constrain the GW strain power at 90% CL with typical values $2 - 20 \times 10^{-50}$ $\mathrm{strain}^2\ \mathrm{Hz}^{-1}$ and $5 - 35 \times 10^{-49} \mathrm{strain}^2\ \mathrm{Hz}^{-1}\ \mathrm{sr}^{-1}$ for pointlike and extended sources respectively. In the same publication, we also searched for persistent narrowband signals from the Galactic Center, SN1987A, and Sco X-1. No signals were detected, but we placed upper limits on strain as a function of frequency. The limits on Sco X-1 and SN1987A are the most constraining to date while the Galactic Center results include parameter space not probed by other searches. The improved strain sensitivity and wider observation band of advanced detectors will facilitate dramatic improvements in the stochastic directional search, potentially facilitating the detection of an anisotropic background and/or a nearby point source. The fulfills an important role in support of the isotropic search.

3.21.2 Scientific Justification

The directional search has multiple targets. It provides a follow-up to characterize stochastic signals detected by the isotropic search, it targets highly anisotropic stochastic signals that might be missed by the isotropic search, and, when run in narrowband mode, it is used to search for persistent point sources, e.g., low-mass X-ray binaries such as Sco X-1. During O1, we will extend the radiometer search to target *unknown*, narrowband point sources such as rotating neutron stars in binary systems using a new folded data scheme [598, 599], described below.

Stochastic gravitational wave backgrounds (SGWB) can be either cosmological or astrophysical in origin. In cosmological scenarios, gravitational-wave backgrounds can be created in the early universe from a variety of processes such as the amplification of vacuum fluctuations following inflation [541], mechanisms that terminate inflation and may give contributions at high frequencies [576, 577], phase transitions in the early universe [542, 543], cosmic strings [544, 545, 546, 547], and pre-Big Bang models [548, 549]. While in most models the SGWB is predicted to be isotropic, there are mechanisms that could introduce anisotropy [547, 600]. Astrophysical backgrounds may arise from binary mergers [580, 601, 566], core-collapse supernovae [602, 603], neutron-star excitations [604, 561], persistent emission from neutron stars [605, 606], and compact objects around supermassive black holes [572, 573]. Depending on the rate and redshift distribution of these objects, the corresponding SGWB could be anisotropic or isotropic. For example, the superposition of all neutron stars in the Milky Way would lead to an extended (but anisotropic), broadband SGWB. Such an anisotropic signal may appear with greater statistical significance in the anisotropic search than in the isotropic search. The search provides information on the angular content of the SGWB in the form of a map of the gravitational-wave sky, and is therefore a powerful tool for distinguishing among different possible sources of the SGWB. The anisotropic search is a critical follow-up in the isotropic stochastic search (proposed elsewhere).

Previously, the narrowband search has been associated with particular direction on the sky such as the Galactic Center, Sco-X1, and Supernova 1987A. Sco X-1, a low-mass X-ray binary, and the brightest source of X-rays aside from the sun, exemplifies the kind of source we target with the narrowband radiometer.

Sco X-1 is known to contain a neutron star with unknown period, which is likely to have been spun up through accretion torque. Since the phase evolution of the neutron star signal is unknown, and since it is modulated in a complicated way by its binary motion (and possibly spin wandering), cross correlation provides a powerful and robust tool for detecting this persistent, but difficult-to-model, source. The robustness of the Sco X-1 radiometer search makes it a highly complementary to CW searches targeting the same source.

Recent work [599] demonstrates that data compressed using sidereal folding [598] can be used to facilitate an extremely efficient narrowband search looking in all directions and at all frequencies. The all-sky, all-frequency extension to the radiometer will target unknown neutron stars in binary systems as we well as all other narrowband searches that do not conform to a canonical CW template. In this way, the stochastic radiometer provides a sensitive tool for discovering *any* persistent point source, which does not conform to the assumptions made by template-based searches.

The detection of a cosmological background would yield arguably the most exciting science possible with gravitational-wave astronomy. There would be enormous interest from the wider physics and astronomy community, and the implications of the discovery would likely be far-reaching. By measuring gravitational waves from the early universe, we could infer the existence of previously conjectural objects such as cosmic strings or, even more exciting, probe physics at energy scales inaccessible through other means. The anisotropic stochastic search will be critical in establishing the origin of a signal detected in an isotropic search.

A detection of an astrophysical background would also be an important discovery. In this case, the anisotropic search could provide information about a variety of astrophysical objects. For example, detection of the stochastic background from compact binary coalescences is plausible with advanced detectors. Measurement of the angular content of this background (which is expected to be nearly isotropic) could provide information about the evolution of matter, and therefore inform population synthesis models. Similarly, measurement of the GW background generated by neutron stars within the Milky Way will allow us to estimate the number of neutron stars in the galaxy as well as to constrain the average ellipticity of neutron stars; both of these are important for understanding the evolution of matter and the state of the matter in neutron stars.

The search could also yield the discovery of a persistent gravitational-wave point source. The detection of gravitational waves from a known astrophysical object such as Sco X-1 would measure interesting astrophysical observables such as the neutron star's rotational period, which could facilitate a better understanding of the inner workings of low-mass X-ray binary systems. Gravitational wave detections from the Galactic Center and/or SN1987A would be similarly spectacular discoveries. Finally, it is worth noting that it may be possible to detect an unexpected point source.

The advanced detectors are expected to reach the sensitivity to normalized energy density of $\Omega_{\rm GW} \sim 10^{-9}$ or better. Many models of stochastic background predict amplitudes that would be within reach of this sensitivity. This includes some inflationary models where the physics of the late stages of inflation generates a boost in GW production that may be detectable in the LIGO/Virgo band [576, 577]. Similarly, a large fraction of the parameter space in cosmic (super)string models, based on GWs produced by cusps and kinks in the string loops, will be within reach of the advanced detector sensitivity [578, 607]. Among the astrophysical models, the GW background due to compact binary coalescences (dominated by binaries at redshifts of 1-2) is probably the most likely to be accessible to advanced detectors [580], although parts of the parameter space in magnetar- or pulsar-based models will also be within reach [608].

There is therefore a very real prospect for detection of stochastic GW background with advanced detectors, and the anisotropic search will play a critical role in understanding the angular structure of the detected background, and therefore in distinguishing between different models/sources of the background. Even the null result will be very interesting as it would rule out significant parts of the parameter space in these models (not constrained by other observations), constraining the relevant physics and providing guidance for future development of these models.

In the case of individual astrophysical sources, the uncertainties in the expected signal amplitude are significantly larger, stemming from the complex nature of the accretion processes or unknown ellipticities of neutron stars. Even in this case, though, we expect to reach interesting sensitivities, for example reaching ellipticities of order 10^{-7}

The anisotropic stochastic search will build on the past two completed anisotropic searches using LIGO/Virgo data [533, 534]. Advanced LIGO operating at design sensitivity is expected to lead to improvements in the sensitivity of the anisotropic search by 2-3 orders of magnitude, and to allow (frequency-dependent) improvements in the angular resolution of the measured background.

The remaining initial detector anisotropic search using the S6/VSR2/VSR3 data is designed to improve on the sensitivity to point sources by implementing a more intelligent algorithm for frequency binning. This search is expected to have $\approx 2\times$ better strain sensitivity to Sco X-1 than previous radiometer measurements [534], and is expected to be completed by the end of summer 2014. The new algorithm will be implemented in the anisotropic searches with the Advanced LIGO/Virgo data proposed here (which will, of course benefit from dramatically improved strain sensitivity).

3.21.3 Search Description

Similarly to the isotropic search, the anisotropic SGWB search estimates the energy density of the stochastic background, but keeps the directional information [534]:

$$\Omega_{\rm GW}(f) \equiv \frac{1}{\rho_c}\,\frac{d\rho_{\rm GW}}{d\ln f} = \frac{2\pi^2 f^3}{3H_0^2}\int d\hat{\Omega}\; H(f)\; P(\hat{\Omega}) \tag{11}$$

where $\rho_{\rm GW}$ is the energy density of gravitational waves, ρ_c is the critical density of the universe, f is the frequency, H_0 is the Hubble parameter and $\hat{\Omega}$ is sky location. The frequency spectrum is typically assumed to be a power law in the frequency band of GW detectors: $H(f) = (f/f_0)^{\beta}$. For a given value of the power index β (typically $\beta = -3$ for cosmological models and $\beta = 0$ for astrophysical models), the objective of the search is to estimate $P(\hat{\Omega})$. Two approaches are pursued. In the radiometer algorithm, we assume the signals is characterized by a point source

$$P(\hat{\Omega}) = \eta(\hat{\Omega}_0)\delta^2(\hat{\Omega},\hat{\Omega}_0), \tag{12}$$

and in the spherical harmonic decomposition (SHD) algorithm we assume that the signal can be written as a superposition of spherical harmonics

$$P(\hat{\Omega}) = \sum_{lm} P_{lm} Y_{lm}(\hat{\Omega}). \tag{13}$$

In either basis, we can compute the "dirty map" X_ν and the corresponding Fisher matrix $\Gamma_{\mu\nu}$:

$$X_\nu = \sum_{ft} \gamma_\nu^*(f,t)\frac{H(f)}{P_1(f,t)P_2(f,t)}C(f,t) \tag{14}$$

$$\Gamma_{\mu\nu} = \sum_{ft} \gamma_\mu^*(f,t)\frac{H^2(f)}{P_1(f,t)P_2(f,t)}\gamma_\nu(f,t) \tag{15}$$

where the indices μ, ν run over the lm's in the SHD algorithm and over the pixels on the sky in the radiometer algorithm, $C(f,t)$ is the cross spectral density for two GW detectors (evaluated at frequencies f and at the times t), and $P_i(f,t)$ $(i = 1, 2)$ are the power spectral densities for the two detectors. The functions

$\gamma_\mu(f,t)$ capture the angular decomposition of the overlap reduction function, and can be computed in either basis [534]. Likelihood maximization then leads to the estimators of the angular content in the two bases:

$$\hat{\eta}_{\hat{\Omega}} = \left(\Gamma_{\hat{\Omega}\hat{\Omega}}\right)^{-1} X_{\hat{\Omega}} \tag{16}$$

$$\hat{P}_{lm} = \sum_{l'm'} (\Gamma^{-1})_{lm,l'm'} X_{l'm'} \tag{17}$$

Analyses using folded data proceed in much the same way except that the data are pre-processed to produce on sideral day's worht of data for the entire observing run; see [598, 599] for additional details.

In order to handle gaps in the data, non-stationarity, and for purposes of computational feasibility, the data for an interferometer pair are divided into many segments of equal duration (typically $60\,\mathrm{s}$), and the above estimators are calculated for each segment. The data are resampled from $16384\,\mathrm{Hz}$ to $4096\,\mathrm{Hz}$ and high-passed filtered with a $9\,\mathrm{Hz}$, $n = 16$ Butterworth filter[26]. They are also Hann-windowed to avoid spectral leakage from strong lines present in the data. Since Hann-windowing effectively reduces the interval length by 50%, the data segments are overlapped by 50% to recover the original signal-to-noise ratio.

The power spectral densities for each segment are calculated using the two neighboring segments. (This approach avoids a bias that would otherwise exist due to a non-zero covariance between the cross-power and the power spectra estimated from the same data.) Furthermore, by comparing the variance calculated using the neighboring segments with the variance calculated using each segment by itself (not its neighbors), we identify segments containing large noise transients and reject them from the analysis. In addition to this stationarity cut, we impose some data-quality flags (such as $30\,\mathrm{s}$ before lock-loss), a large-σ cut (rejecting particularly noisy time-periods), and we ignore frequency bins which are determined to contain instrumental correlations. The segments that pass all the data-quality cuts ($\gtrsim 95\%$) are combined using a weighted average to yield final estimators.

The stochastic group has developed a single data analysis pipeline to search for anisotropies in the stochastic background and for excess coherence from point sources. However, the spherical harmonic decomposition algorithm is compared with the radiometer algorithm algorithm as a cross-check. In the event of an observed signal, the stochastic group will employ a detection checklist[27]. In the event of an observed anisotropic stochastic signal, the result will be checked for consistency with the isotropic analysis. In the event of a point source detection, the result will be checked with the appropriate CW pipelines, which may be able to provide confirmation if, for example, the signal is from a neutron star in a binary system.

3.21.4 Publication Plan

At the end of an extended (~ 1 year) science run at the expected Advanced LIGO sensitivity we expect to improve on our initial LIGO/Virgo sensitivity by 2-3 orders of magnitude. In the event of an unambiguous detection, we plan to publish as quickly as we can vet the detection and produce a paper describing the implications of the discovery. Detection claims will be vetted through standard cross-checks to make sure that the signal behaves as expected in frequency/time and that it cannot be accounted for due to data quality artifacts such as correlated noise. If no signal is present, we plan to publish after we have achieved at least a factor of two improvement in sensitivity for P_{lm}'s and $\eta_{\hat{\Omega}}$ *and* we have reached a convenient milestone in data collection, e.g., the end of a science run. (One science run with a factor of eight improvement in sensitivity will yield one paper, not three.)

A *detection paper* will establish that a stochastic signal is present and it will constrain its angular distribution. It will discuss the statistical significance of the measured angular distribution and the data-quality

---

[26] http://tinyurl.com/kgbffx7

[27] https://dcc.ligo.org/LIGO-T1500251

and sanity-check steps we have taken to ensure that the signal is real. Other cross checks will include comparisons of the SHD and radiometer results, comparisons of the statistics in different parts of the sky (e.g. top vs bottom hemisphere), tests of the time variability in the measured angular distribution etc. We will also comment on possible origins of the signal, whether it is an object in a specific direction in the sky (appropriate if detecting a point source) or consistency with isotropy (appropriate when a signal is detected, but no angular structure is observed).

In the event of a *marginal detection*, we advocate collecting more data since the stochastic signal-to-noise ratio grows like the square root of observation time. If this is not possible, we advocate publishing with the data available and stressing that the apparent signal is marginal.

In the event of a *null result*, we will present upper limit maps on the energy density, similarly to what was done with initial detector data [533, 534].

The anisotropic stochastic pipeline is well established. It runs quickly and reliably, and the statistics of the relevant estimators is well understood, allowing straightforward interpretation of the results. We have many years of experience running this analysis. Thus, we plan to analyze data as it is collected.

If no signal is apparent, we plan to produce paper drafts to submit for review within $\approx 4\,\mathrm{months}$ after meeting the requirements for a new paper (see above). In the event of a detection, we will carry out additional tests to make sure the detection is real. We expect this extra work to take an additional two to four months. Due to the potentially far-reaching implications, we recommend publishing a detection paper in a high-visibility journal .

Upper-limit papers can be submitted to a variety of journals. We leave for future discussion the question of what constitutes an important milestone as this will depend in part on the state of the literature (theoretical and observational) at the time of publication.

The narrowband radiometer search overlaps with multiple efforts in the CW group. The radiometer search team participated in a CW-led mock data challenege to study the sensitivity of different pipelines to gravitational waves from Sco X-1 [527]. We will continue to coordinate this search with related CW searches.

3.21.5 Technical Requirements and Development Plan

The stochastic pipeline is Matlab-based code that lives in the matapps repository[28]. The main function is stochastic.m. The stochastic radiometer pipeline can be traced back to S4 [533], with the SHD algorithm added in the S5 analysis [534]. The pipeline has not evolved much since then, although we are planning relatively minor modifications in the analysis of S6 data, so as to include a more intelligent frequency binning algorithm, which will result is improvements in sensitivity to some of the point sources. Hence, a large majority of the code that will be used in searches with advanced detector data has already been reviewed, and has been stable for several years.

As noted above, minor development of the anisotropic stochastic search code is expected, in order to improve the frequency binning (and therefore improve the sensitivity to some of the point sources). This upgrade is already well understood, and we are in the process of implementing it as a part of the analysis of the S6/VSR2/VSR3 data. We expect this modification to be completed before the summer 2014.

The anisotropic stochastic pipeline will be ready to analyze advanced detector data by the summer 2014. New filters for the lower frequency cutoff (10 Hz) have been developed.

3.21.6 Resources

Code infrastructure. The stochastic anisotropic search utilizes the Matlab-based stochastic pipeline, which lives in matapps. The code is reviewed and ready to be used for advanced-detector analyses. Mi-

---

[28] http://tinyurl.com/kpmyruv

nor modifications are expected before the summer 2014 (see above).

Person-power. The anisotropic search requires the FTE of one postdoc and one graduate students plus a senior mentor. Additional members of the group, such as the chairs and the advisors of the the students/postdocs, are also actively involved in the analysis as well as data-quality efforts and validation. We do not expect the FTE requirements to grow considerably in the advanced detector era.

Detector characterization for stochastic searches. The long integration time used in the stochastic analysis can reveal noise lines. Common noise lines at two sites can occur due to similar equipment in the laboratory, or common data acquisition systems. A strong line in one interferometer, along with a large random fluctuation in the other, can also produce an apparent narrowband signal in the stochastic search pipeline. In the advanced detector era, the coherence between pairs detectors' output will be calculated in near real time; (see the discussion of stochmon in the search plan for isotropic stochastic search). In this way noise lines that would affect the stochastic search (and by extension, also the CW search) can be identified during the science runs, and possibly addressed at the sites. We will work with the detector characterization and continuous wave groups to identify and find the sources of noise lines using all available tools.
We continue to investigate noise correlations between the LIGO Hanford and Livingston sites, as well as with Virgo. We have previously observed correlated magnetic fields in magnetometer channels at widely separated detectors [597]. The primary sources of these correlated fields are Schumann resonances [597]. Having documented this potential source of correlated noise, we will be careful to determine if correlated noise is affecting the stochastic search. Investigations are in progress to determine on how well the magnetometers at LIGO and Virgo sites can measure the correlated Schumann resonance noise, and whether more antennas dedicated to Schumann resonance observations are needed. Noise subtraction techniques, especially with respect to the correlated electromagnetic noise, are being studied. If the correlated Schumann resonance magnetic fields are limiting the stochastic search, then it will be necessary to directly measure the Schumann resonance magnetic fields at each interferometer site and implement noise subtraction techniques. (A preliminary study of the necessary resources has been made [609].) Data from radio frequency (RF) monitors at the sites will also be studied. Other physical and environment monitors will also be carefully observed. While optimal re-weighting techniques can be used to help ameliorate the effect of correlated magnetic field noise for searches for broadband signals, they will provide no help for cross-correlation searches for narrowband signals. Correlated magnetic field noise could be a significant problem for coherent directional searches below 25 Hz with advanced LIGO and advanced Virgo. Diminishing the magnetic coupling to the interferometer masses will help to address this problem, as would low noise detection of the Schumann resonances and the subsequent implementation of noise subtraction (Wiener filtering) methods.

Calibration. The calibration requirements for aLIGO have been detailed in LIGO-T1200562. The phase error is requested to be less than 9 degrees. The relative timing error is requested to be less than 24 μs. The amplitude error is requested to be less than 9%.

Review. The stochastic review committee has reviewed numerous upper limit papers. The code will not change much, so code review will be trivial. Scientific conclusions about constraints placed on models because of the upper limits are always carefully reviewed, but the timescale is not prohibitive (note the pace of the S5 Nature paper [578] result). The review committee also carefully reviews and critiques the paper. A detection result will certainly be more complicated, but there are numerous sanity checks available. The estimate is that a detection result will take an extra 2 to 4 months to review. The S6 radiometer analysis is nearing completion and the review is expected to begin during the northern hemisphere summer. The group has requested a review of the data folding procedure [598] used for all-sky, all-frequency extension [599].

3.22 Contributions of long transient gravitational waves to the stochastic background search

3.22.1 Abstract

The search for an isotropic stochastic background includes contributions from different gravitational-wave frequencies, from different directions in the sky, and from different data segments. Each of these parameters—frequency, direction, and time—can be thought of as variable, over which the isotropic search integrates. By carrying out all but one integral, it is possible to gain insights into the nature of the stochastic signal while potentially revealing hidden signals diluted through integration. For example, the spherical harmonic decomposition search shows how the stochastic background varies across the sky. Analogously, the directed radiometer search shows how the stochastic signal varies with frequency at a given point on the sky. Here, we describe an analysis designed to explore time-dependent features of the stochastic background. The search characterizes the temporal behavior of an apparent stochastic signal while potentially identifying very long-lived transient signals ($\gtrsim 10\,\mathrm{hr}$) that might be otherwise overlooked.

3.22.2 Scientific Justification

The scientific rationale for this search is based on the following observation: one or more long-lived transient signals (or coherent long-duration noise) can produce an apparent signal in either isotropic or directional stochastic searches, while simultaneously evading detection in searches for short-duration transients. Thus, a dedicated search is necessary to understand the origin of apparent stochastic signals. Consider, for example, a $11\,\mathrm{hr}$-long signal observed during a $90\,\mathrm{day}$ science run. A $11\,\mathrm{hr}$-long signal, with a stochastic signal-to-noise ratio [591] of SNR $= 40$, would produce an SNR $= 3$ signal through dilution [610]. The same signal would produce only a SNR $= 0.20$ signature in a short $1\,\mathrm{s}$-long data segment. This back-of-the-envelope argument illustrates how long signals can influence stochastic measurements while producing little trace of their existence on short time scales.

The astrophysics of very-long transients, spanning hours to weeks, was first explored in [611][29]. The authors of [611] review a number of somewhat speculative scenarios associated with neutron stars including gravitational-wave emission lasting days to months from non-axisymmetric Ekman flow following a glitch [612, 613, 614], Alfvén oscillations from giant magnetic flares (also lasting days to months) [615, 616], emission from free precession (with a damping time possibly lasting from weeks to years) [617, 618, 619], magnetic instabilities in newborn neutron stars (lasting days) [620], and gravitational-waves from r-modes [621, 622]. Generic rotational instabilities in newborn neutron stars, potentially powered by fallback accretion [340, 623], may persist on a timescale of hours [338]. Somewhat more speculatively, we note that observations of intermittent pulsars, which become quiescent on timescales of days (e.g., [624]), suggest that neutron star dynamics vary on the timescales considered here, motivating exploration of this region of parameter space. Similarly, variability in accretion on these timescales may affect gravitational-wave emission from accretion-supported mountains [625]. Finally, it is worthwhile to be prepared for a surprise: a very long lived transient signal from an unexpected source. Recent work proposing gravitational-wave emission from gravitationally bound axion clouds [626], potentially starting and stopping on the timescale of a few years, serves to illustrate this possibility. The method we propose below is also applicable to quasi-infinite signals turning on or off during an observing run as well as repeating sources with long-lasting emission periods.

Most previous work on long-duration transients has focused on signals in the regime of $\approx$10–$3000\,\mathrm{s}$. The reality, let alone the detectability of such very long signals is admittedly speculative and highly uncertain. A back-of-the-envelope calculation [610] suggests that a very long transient associated with Scorpius X-1 would produce a signal about ten times too weak in strain to detect assuming a hypothesized accretion-torque

---

[29] This paragraph is from [610].

balance scenario [625]. Detection would require a source that is either significantly closer or significantly brighterin gravitational waves. That said, it is important to note that detecting *any* gravitational waves from neutron stars such as Scorpius X-1 is a challenging proposition for a variety of search techniques [527].

However, it is imperative for the stochastic search to confirm whether or not signal of long duration are contributing to a signal from a search for an isotropic stochastic background of gravitational waves. This is part of an *eyes wide open* search strategy. We also note that *non-astrophysical* signals may persist on these time scales as well, e.g., from correlated noise due to time-dependent magnetic phenomena [597, 609]. Finally, it is worthwhile to be prepared for asurprise: a very long lived transient signal not predicted.

Taken together, the stochastic searches tell us: 1) if the LIGO-Virgo detectors exhibit higher than expected correlation, 2) the contribution to this excess from a persistent stochastic background, 3) the contribution from correlated noise, 4) the contribution from persistent sources that are resolvable in sky location, 5) the contribution from non-persistent (transient) sources. Analyzing the time-dependence of cross-correlated data (along with the frequency and directional information) is crucial for understanding both the detector and the signal. It is essential that we understand all of these degrees of freedom in order to confidently make a stochastic detection and to be sure that the signal is a stochastic background and not something else (astrophysical in origin or not).

Relationship to other searches. This search is designed to provide information on the time-dependence of an apparent stochastic signal while potentially revealing very long-lived transient signals. Like the directional search and the narrowband radiometer, this is a core activity of the Stochastic Working Group, which allows us to parse the stochastic signal in several different ways. This search complements dedicated burst searches, which target signals of shorter duration (less than an hour). There is also an overlap with the CW group, which is currently developing a search for glitching pulsars and type I bursts in accreting millisecond pulsars, which could result in transient signals on very long durations (up to weeks). However, this method is complementary to CW techniques as we allow for significant variation in the frequency evolution of the gravitational-wave signal, which allows us to probe a completely different parameter space compared to continuous wave searches, which assume very limited changes in the neutron star spin: $\lesssim 6 \times 10^{-9}\,\mathrm{Hz\,s^{-1}}$ versus $\lesssim 300\,\mathrm{Hz\,s^{-1}}$.

Search Description

This search plan is designed to uncover transient contributions to an apparent stochastic signal while potentially revealing very long-lived signals that might otherwise be missed. To that end, we will constrain the energy density Ω_{gw} [591] due to transient phenomena. We will determine the relative contributions in our data from noise, from a persistent stochastic background, and from astrophysical transient phenomena.

The baseline plan is to carry out these searches using STAMP [337, 352, 347, 341, 342]. STAMP is a Matlab-based code package that resides in the Matapps svn repository [353]. STAMP works by cross-correlating data from two detectors to produce cross-power spectrograms [337]. Gravitational-wave signals appears as tracks of brighter-than-usual spectrogram pixels. STAMP employs a user-specified clustering algorithm (there are a few options [337, 627, 341, 342, 628]) in order to identify statistically significant clusters of pixels. One of these clustering algorithms, which we shall refer to below, is called Stochtrack [341]. Stochtrack works by calculating the coherence between two detectors on randomly drawn curves on a spectrogram to identify tracks left by gravitational-wave signals.

Much of the infrastructure was originally adapted from Stochastic code, though, the pipeline is now used by members of the Bursts, Detchar, and Stochastic Groups. The pipeline was reviewed for an initial LIGO analysis studying long-lived transients coincident with long gamma-ray bursts, which was ultimately published in Physical Review D [347].

A project to carry out an all-sky long duration signal search with initial LIGO data is complete [354]; the review of this search was initiated in July 2014. The STAMP all-sky team has made great strides to demonstrate an all-sky search for long-lived transient signals including: managing the computational cost of the search for signals of duration down to 10 sec, understanding the background and minimizing the effect of noise transients to a level approaching Gaussian noise, and investigating the applicability of the search to eccentric binary waveforms [354]. The S5 and S6 datasets have been analyzed. A paper draft describing the search is available on the DCC [355].

The code for a new seedless clustering algorithm [341, 342] will require review, but the new code makes up a small addition to the previously reviewed code package: two functions consisting of <500 lines of new code along with a handful of minor changes to existing code (estimated FTE=0.5 over the span of a month). The seedless clustering code is highly parallel, and recent work [342] demonstrates that GPUs and multi-core CPUs facilitate dramatic speed-ups.

We will analyze data on timescales of $\approx 10\,\mathrm{hr}$–$1\,\mathrm{month}$ in order to determine if there are individual long-lived transient signals contributing to the isotropic or directional stochastic measurements. We will run STAMP in all-sky mode on all O1 data used in the stochastic search. In order to analyze these very long signals, we will add an extra stage of pre-processing in which the data are compressed through time-averaging as described in [610]. This new pre-processing step will require minimal new infrastructure/review. This extension will also be made available to STAMP-PEM, a detector characterization application of the STAMP algorithm. STAMP-PEM will then be used to search for instrumentally or environmentally caused long transients in the O1 data, for example due to magnetic field effects, on the time scales longer than $\approx 10\,\mathrm{hr}$. This effort will have a direct impact on the long-duration studies discussed here, as well as on understanding the origin of apparent stochastic signals. The compressed data will be analyzed using seedless clustering algorithm described in [342].

3.22.3 Publication plan

We foresee several possible scenarios:

- If a stochastic signal is observed (either isotropic or directional) and no evidence is found for long-duration transient signals, then the results of this search will be used to characterize the observed stochastic signal as persistent. The discovery of a stochastic background should be submitted to a high-impact journal.
- If a stochastic signal is observed (either isotropic or directional) and we find evidence for one or more possible very long transient signals, we will estimate the energy density contribution of these transient signals to the overall stochastic signal, as well as the residual persistent stochastic signal. In this very positive scenario, we anticipate submitting two papers to high-impact journals, one on the stochastic background and one on the very long transient(s). In the event of a very long transient identification, we will coordinate with the Burst and CW groups to interpret the event.
- If the O1 data are consistent with noise, we will place constraints on the energy density and the rate of the long-lived transient phenomena. We recommend publishing a paper with limits if/when the constraints on energy density improve by a factor of at least ten *and* there is a convenient break in data-taking (the end of a run).

3.22.4 Technical requirements and development plan

Pipeline. The baseline plan is to carry out the analysis with STAMP [337, 352, 347]. STAMP (and Stochtrack) are Matlab-based code. They are part of Matapps [353]. Most of the STAMP code was reviewed for the published initial LIGO targeted analysis [347], and some recent additions relevant for all-sky searches are currently being reviewed as a part of the S5/S6 STAMP all-sky review. The new Stochtrack clustering algorithm [341, 342] will require review, but the job is modest: two (very well-documented) functions with $<$500 lines of code. The new clustering algorithms and their performance are documented in publications [341, 342]. There are a handful of changes to the existing code suite, but they are minor. Future changes to facilitate the analysis of longer-duration signals, lasting from hours to months, are also expected to require only small additions to the processing code.

Data quality. We do not anticipate that the long-transient search will require significant detector characterization work beyond what is already planned for other searches. During initial LIGO, STAMP searches were shown to produce well-behaved noise distributions (similar to stationary Gaussian noise) through the application of notches for instrumental lines, basic data quality flags, and a glitch identification flag [352, 347]. This is attributable to the long time scales probed by the search as well as the fact that STAMP uses cross-correlation [352]. Subsequent investigations have applied seedless clustering techniques to recolored initial LIGO noise leading to comparable results [341, 342]. There are no special requirements for calibration beyond what has already been requested by the Stochastic Group [368].

Manpower. The analysis will require the full-time attention of at least one graduate student or postdoc. The attention of a mentoring advisor is also required. Continued code maintenance and support should be strong and steady. There is an active community of STAMP users from multiple analysis groups including $\gtrsim$ 7 regular/expert users with a wide base of experience.

Early testing. The upgraded seedless clustering code has been tested and benchmarked with both Monte Carlo and recolored noise [341, 342]. The first real opportunity to learn significantly new information about the search will come when actual Advanced LIGO strain data is available. In order to catch problems early and facilitate a smooth analysis, we plan to start analyzing data immediately as it becomes available with a goal of estimating the background and sensitivity on an approximately biweekly basis. (This was not possible during initial LIGO/Virgo since the pipeline was still undergoing significant development during the final science runs.) We will coordinate with potentially overlapping efforts to maintain a mock dataset with long-lived injections in order to compare the domain of utility of different searches/algorithms.

Broader impact. The STAMP code package has produced spin-off technology that has proven useful for detector characterization [629, 630] and follow-up/visualization of CBC triggers [628]. We expect continued development and maintenance of STAMP will be broadly useful for Stochastic Group activities and the wider LSC/Virgo community.

4 Characterization of the Detectors and Their Data

4.1 LSC Detector Characterization

4.1.1 Introduction

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is an NSF-funded project with the mission of directly detecting gravitational waves from astrophysical sources. LIGO has now completed the installation and initial commissioning of a major upgrade, called Advanced LIGO, that will increase the strain sensitivity of each interferometer by more than a factor of 10 and the volume of the universe observable by gravitational waves by a factor of 1000. It is currently undertaking its first Observing run, having met the target sensitivity and duty cycle. Advanced LIGO, in partnership with Virgo, is expected to make the first gravitational-wave detections and begin an era of gravitational-wave astronomy.

The LSC detector characterization group [631] directly supports LIGO's mission because a thorough characterization of the detectors is required to confidently detect gravitational waves. The collaboration's ability to make detections and the level at which upper limits for gravitational-wave emission are set depend critically on detector performance characteristics, such as the overall level of the noise-limited detector spectrum, the probability distribution of transients in the detector output, the degree to which the noise components are stationary, and lines and features that are present in the data. Detector characterization is also an important aid to the commissioning process. Characterization efforts identify issues and provide clues to commissioners, who use these to improve the instruments. These two rather different elements – advising commissioning, and providing the information for astrophysical analyses to make the best use of the data – are synergistic products of the Detector Characterization group.

Detector characterization is carried out by a variety of scientists focused on different problems. Commissioners operate at the cutting edge of detector characterization, modifying the detector systems to increase their performance in terms of noise, lines, and robustness. Their investigations may focus on interferometer-based detector characterization, such as investigation of noise sources, lines and features, and environmental disturbances. Members of the analysis groups also make important contributions to detector characterization. They often direct their efforts toward impediments to astrophysical searches, such as coherent or accidentally coincident glitches that pollute compact binary and burst searches, features that could blind searches for periodic or stochastic background sources, or wandering line features that could mimic a pulsar. The core Detector Characterization activity includes elements of these 'extremes' and also has a vital middle ground of analysis and tool development which then enables the instrument commissioners and the data analysts.

During intense commissioning, it is difficult to evaluate the long-term performance of the instruments. Science and engineering runs serve as testing grounds for interferometer stability and for rapid communication between commissioning, detector characterization, and data analysis groups. As experience has accumulated, tools to evaluate the search backgrounds and instrument stability have been and continue to be developed by the Detector Characterization group, and the latency of diagnostic feedback about the noise and transient behavior of the instrument has decreased greatly.

However, even after years of commissioning and detector characterization the data recorded for scientific analysis contains unforeseen artefacts that decrease the sensitivity of or even blind some searches if left unchecked. For that reason, the detector characterization group has a strong effort to identify and remove non-astrophysical artifacts from the recorded data. For transient searches this is done using data quality flags and vetoes. For periodic and stochastic searches, times and/or specific frequency ranges are identified and removed from the analyses. These efforts have led to improved upper limits in the searches performed to date.

As new artefacts are found, new characterization methods are developed. If the artefacts persist, the group works to automate the relevant methods for more rapid detection of the artefacts and to pass information about them on to the search pipelines. For initial LIGO, the online monitoring systems included the Data Monitoring Tool (DMT)[632] with a number of targeted monitors, the controls system software (EPICS)[633], and search-oriented monitors such as the trigger generators Omega [634] and KleineWelle [635] for the burst search, and daily iHope [636] for CBC search, as well as a variety of customized tools written in e.g., python, C++, and Matlab. It also included a human element, namely the attentiveness and active data exploration by interferometer operators and scientific monitors (scimons). For the advanced detector era, we have built upon this experience, automating and adding straightforward improvements to what worked, and developing new strategies to address issues that were encountered.

The LSC Detector Characterization community has a broad membership, including full-time commissioners, full-time data analysts, and many in between. Every search working group has active DetChar members, and information flows in both directions, to mutual benefit. The DetChar working group has a unique scope of providing characterization tools and monitors and on providing characterization (most directly data quality flags and vetoes, and diagnostic information) of interferometer data in science runs for astrophysical analysis. We work closely with search and instrumental groups to identify and solve issues with Advanced LIGO data, and document the outcomes in observatory log books, data analysis wiki pages, technical documents, and short author list publications on limited sets of interferometer data.

In the following subsections, we describe the foci for detector characterization for our collaboration to achieve its scientific goals during the advanced detector era (4.1.2), the LIGO-specific priorities for detector characterization (4.1.3), status and plans for data run support (4.1.4), commissioning support (4.1.5), and software infrastructure (4.1.6), as well as the activities and priorities of the different working groups, noise transients (4.1.7), and spectral features (4.1.8).

4.1.2 Overview Requirements to Achieve aLIGO Science

Advanced LIGO began scientific data collection in Fall 2015 and will perform observations of increasing sensitivity and duration in the years that follow [124]. The overarching goal of the LIGO detector characterization group over the next few years is to ensure that Advanced LIGO data are of high enough quality and accurately calibrated to maximize the scientific information that can be extracted from the first Advanced LIGO observations – especially gravitational-wave detections. The following requirements for detector characterization in the advanced detector era are driven by the science requirements of the search groups described in further sections of this whitepaper. Here we present the requirements but do not detail how they will be addressed until Subsection 4.1.3.

1. **Provide accurate state information for the gravitational-wave channels.** It is imperative for confident detections and deep searches that the state of the detectors be accurately recorded (e.g. to signal what data should be analyzed and the expected quality of those data).

2. **Remove egregious data quality problems in the detectors.** Loud transients or dramatic non-stationarity could complicate or blind a potential detection, in particular where an agnostic ‘burst’ signal is sought. To maximize detections, BNS searches require that egregious data quality problems be identified and removed in the detectors.

3. **Provide a well-understood and documented physical environmental monitoring system.** Among the most important requirements for both improving the detectors and following up potential signals is a well-monitored and understood physical environment at the detector sites. Only a suite of well understood environmental monitors will allow us to say with confidence that potential coincident signals did not arise from anthopogenic, terrestrial, atmospheric, etc, external effects.

4. **Remove the majority of short duration transients in the detector outputs.** LIGO data show a non-negligible rate of non-Gaussian noise transients (glitches). These act to increase the background in

(and therefore decrease the sensitivity of) searches for shorter duration, less well modeled waveforms such as burst sources, and higher mass CBC systems (BHBH and NSBH). In addition, noise transients may effect parameter estimation for BNS detections. While Advanced LIGO has a smaller glitch rate than Initial LIGO, these searches will profit from a further reduction. The best way to achieve this is through mitigating noise sources in the detector, which can be achieved through early characterization and closely working with commissioners. However, improved data quality products such as flags and vetoes will also be required.

5. **Remove lines or spectral features that limit continuous-wave and stochastic background searches.** For continuous waves and stochastic searches the detector characterization group should identify and help remove lines or spectral features, prioritizing those that are coherent in the instruments and/or occur at frequencies targeted by the searches. Those lines which cannot be removed via instrument changes must be adequately characterized to be 'vetoed' from the CW and Stochastic searches.

6. **Provide high data quality with low latency.** Providing this information is necessary to carry out sensitive low-latency searches, including searches with electromagnetic followup. High data quality includes accurate calibration, timing, and information about the interferometer state, as well as the automatic removal of as many data artefacts as possible.

4.1.3 Priorities for LIGO Detector Characterization

In this section we set priorities for detector characterization during the upcoming year by choosing activities that will ensure that those requirements are met for the first and second aLIGO observing runs, in late 2015 and 2016, respectively.

1. **Characterize the Advanced LIGO subsystems in situ.** [*Supports requirements 2,3,5 above.*] We will continue to investigate the data quality of Advanced LIGO subsystems via their own diagnostics and the diagnostics available from the complete instrument. Each subsystem has a detector characterization “lead" responsible for coordinating investigations for that subsystem. Investigations will include,

   - Documenting information about auxiliary channels, such as their name, meaning, sample rate, dates of recording, physical location, and calibration function and units.
   - Checking the fidelity of the recorded signals. For example, that each channel is recorded above ADC noise over its useful frequency range, and that the signals do not saturate during nominal operation.
   - Recording accurate and authoritative information about the state of each subsystem and of the entire interferometer.
   - Identifying artefacts (glitches, lines, features) in the key channels for each subsystem, and helping to find their exact origins and reduce them.
   - Contributing to the “noise budget" (a tool for understanding of the various noise contributions) for each subsystem (primarily development done within the LIGO Lab).
   - Performing deeper and more specific investigations on the subsystems such as those listed in Section 4.1.5.
   - Determine the coupling of environmental influences to the aLIGO interferometers by direct measurement, so called “PEM injections".

2. **Characterize and refine the LIGO Physical Environmental Monitoring Systems for Advanced LIGO** [*Primarily supports requirement 2, but supports all goals above.*] The LIGO PEM system must be actively characterized, refined, and feedback must be given to the LIGO Laboratory with respect to the operation, maintenance and upgrades of the system. The detector characterization group is uniquely qualified to provide this type of feedback because of the tools and investigations they have developed to scan large channel spaces and uncover correlations between environmental channels and artefacts in the detector outputs.

3. **Participate actively in the Advanced LIGO Observing runs.** [*Supports all requirements above.*] Having used the series of engineering runs to test important software infrastructure, establish procedures for software release/maintenance during aLIGO, perform detector characterization early using real subsystem data, and measure progress of the analysis groups toward key science goals, we support the Observing runs with these seasoned tools and techniques. In the detector characterization group we will work toward having key investigations maintain pace with observation and critical software (calibration, timing, state, data quality monitoring, etc.) operating smoothly. We expect observation runs to provide excellent opportunities to observe the longer term stability of the interferometer and its subsystems than is often possible during heavy commissioning and to point to inter-run commissioning priorities.

4. **Develop improved methods to uncover the causes of and veto noise transients.** [*Supports requirement 2,4,5 above.*] We have had success using burst and CBC search algorithms [635] to parameterize glitches in the detector outputs and a large number of auxiliary channels and then using automated tools such as UPV [637] and hveto [638] to generate "veto" segments based on statistical correlation. To achieve requirement 5 above we will need to improve upon the performance of these algorithms. Promising avenues of research that must be followed are:

    - Investigations of the utility of other physical inputs (than glitch parameters) as an indicator of glitches in the detector output. For example mean values or RMS of slow auxiliary channels (e.g. alignment).
    - Extending veto techniques by straightforward refinement, or using methods such as multivariate classifiers, bilinear coupling indicators, etc.
    - Data mining techniques that identify connections between times subject to glitches and the values of a wide array of control and monitoring signals. This will allow the exploration of the possibility that saturation of error signals in control systems causes extra sensitivity to environmental disturbances, as well as other mechanisms that can cause time-varying couplings between control channels and the gravitational wave output.

    Further discussion of this goal is in Section 4.1.7.

5. **Continue to refine tools during Observation runs** The approaches to efficiently and creatively identifying problems in the data are under constant improvement, and constant modification to track instrument evolution. In particular, data quality flags and veto segments should be shown to have a significant beneficial effect on data sets containing known artefacts.

6. **Provide data quality support to search groups for the first observing run analyses.** The detector characterization group is providing support for data quality issues for all gravitational-wave analyses performed on the Advanced LIGO engineering runs and observing runs.

7. **Continue to document the detector characterization work that has had an impact on the LIGO detectors and searches.** This includes contributing to S6/VSR2,3 analysis papers, writing an overview

detector characterization paper for LIGO S6, documenting the PEM system, and completing papers describing methods had an impact on iLIGO or early aLIGO data.

8. **Characterize and reduce low-level correlated noise.** [*Supports requirement 5 above.*] This includes studies to answer the question, is there any correlated noise hidden below the uncorrelated noise curve that could affect stochastic searches? An important example of such a noise source are Schumann resonances (described in more detail in Section **??**), electromagnetic resonances of the cavity formed between the Earth's surface and the ionosphere that are excited globally by lightning strikes. More broadly, the survey and evaluation of potential correlations at specific times will be crucial to vet any gravitational-wave candidates.

4.1.4 Observing Run Support

For the Advanced Detector epoch, the LSC has put in place a LIGO Fellows Program LIGO-M1400310-v8. Two key goals for this program are, i) to help increase the scientific output of the LIGO detectors, characterized by the product of searched volume and time, particularly by maximizing the amount of observing time available to searches by running as often as possible (as other astronomical observatories do) and ii) to form a bridge between the LIGO sites and the broader LSC to maximize the astrophysical potential of the LIGO detectors. The Fellows will act as an element of Detector Characterization during the Observing runs, and complement the more seasoned Data Quality Shift staff (see https://dcc.ligo.org/LIGO-L1500110/public for the Data Quality Shift Policy) who can draw on a range of tools and colleagues to give both real-time support and to identifying major contributors to the astrophysical search backgrounds triggered by data during observing runs.

The Data Quality Shifts are taking on a larger role, and one that is more closely integrated with the vetting procedure for low-latency candidate events. DetChar staff with DQ training are 'on call' 24/7 to help verify that candidate events are valid and should be sent to EM partners, and that significant events have appropriate real-time followup to ensure the most complete environmental and instrumental data around the time of the event. Event-followup DQ shifters must be able to depend on a robust team to follow up on leads around events.

Besides human resources, working together with the search groups and the commissioning group to identify and display key figures of merit and developing and tailoring daily summary pages of the observatory data will also be key support to data runs - helping to quickly identify issues and inform data run decisions.

Another important aspect of data run support is injection of simulated astrophysical signals into the detector hardware to validate data analysis pipelines with high confidence. LIGO Laboratory and LSC scientists have provided the human resources to set up the injection infrastructure and carry out the injections. We will develop a version of the injection software which is compatible with the Advanced LIGO Guardian sequencing programs and which makes the best use of the actuation systems and their characteristics.

4.1.5 Commissioning Support

Heavy commissioning of Advanced LIGO is expected to continue over the next several years. This commissioning is the primary way for the completed Advanced LIGO detectors to improve in astrophysical sensitivity and robustness over time. The detector characterization group will continue to work ever closely with the commissioning team over the next few years. Below are two main areas for joint work between characterization and commissioning groups.

Detector Commissioning and Characterization Projects The detector characterization group works closely with the commissioning and detector groups and communicates regularly about data quality and interferometer robustness issues via email and the detector electronic logs. The two groups periodically hold joint teleconferences to discuss data quality issues that have been identified during engineering and observing runs, and plans to follow them up.

Subsystem characterization Characterization of the Advanced LIGO subsystems is Priority 1 from Section 4.1.3 because we think this will help us identify issues (such as with glitches, noise lines and robustness) early so that they can be fixed and will lead to a deeper understanding of the detector systems among the detector characterization group leading to better characterization and cleaner data. Strengthening the relationship between the detector characterization and commissioning teams is very important for these goals.

To foster this, we have set up an organizational scheme based around a “Subsystem Matrix” [639] that lists each subsystem (for example, the Pre-Stabilized Laser, Seismic Isolation, Data Acquisition System) as horizontal rows, with vertical columns for the projects/tasks that should be accomplished for each (for example, documenting the meaning of channels, monitoring their transient behavior, and checking signal fidelity). One person, a lead/liaison is assigned to each subsytem, as the primary person of contact from detector characterization, responsible for communicating with the commissioning experts for that subystem and organizing and reporting on the characterization work. In addition, there are lead/liaisons for the projects, who help define and are the primary contact for the work that should be done to complete that project.

4.1.6 Software Infrastructure

Over the years, many tools have been developed for on- and off-line monitoring of detector status and data quality. Many of these software tools (EPICS[633], DTT[640] and DataViewer) are used interactively in the Observatories’ control rooms by operators, commissioners and scientific monitors, and have proved to be essential to operations and commissioning. These tools were developed by the LIGO Laboratory, and we expect these tools to be maintained, and improved when appropriate, for Advanced LIGO operations.

The Data Monitoring Tools system, or DMT[640], is used as a background-process environment for continuous monitoring. The DMT system provides many critical functions, including the online production of calibrated strain files that are were used for low-latency analyses, online production of data quality information that was used for selecting appropriate data to run on and vetoing noise transients in those analyses, and continuous graphical monitoring of the data and the environment displayed in the control room. Although programs used by the DMT system were written by scientists in many different institutions, the maintenance of the infrastructure is performed by the LIGO Laboratory.

Data quality monitoring involves reviewing the results produced by monitors such as those run in DMT, veto selection programs such as hveto [638] and UPV [637], noise transient monitoring, coherence and line monitoring, results of data analysis such as search background, and other scripts running in the LIGO Data Grid clusters at the Observatories. This review is done continuously during Observing runs by fellows and DQ (Data Quality) shift-takers at a basic level, and periodically at a deeper level by members of the Glitch and Spectral features subgroups. These investigations result in the diagnosis, identification, and sometimes fixing of artefacts in the gravitational wave channel, which reduce the background for the searches of astrophysical signals in the data. For Advanced LIGO we have brought all of these types of monitors together on a single easily digestible page (with many sub-pages) for each detector, called a LIGO summary page (with GEO and Virgo pages available from the same system).

The DetChar group compiles information from its investigations to create a repository of data quality (DQ) information for each run. The information is used to form time intervals which are flagged with data quality flags, which are incorporated into a database (which also has the information on which times are in nominal science mode). The database can be queried by the astrophysical search programs, as well as by members of the group diagnosing problems in the data. We now have a reliable system where flags are introduced online by well tested DMT monitoring programs, and offline by the DetChar group. For aLIGO have added monitoring of critical state information of the interferometers and their subsystems to the front-end systems. This will enable the production of low-latency and authoritative state information that can be used automatically by searches, and downstream monitors.

For detector characterization it is important that collaboration members have easy and reliable access to the LIGO data, including the gravitational wave channels and the many auxiliary channels (and their archived trends). For most of the initial LIGO era access to LIGO data from outside the observatory control rooms required signficant effort - and this was an impediment to engaging more collaboration members in detector characterization. This situation was greatly improved in 2010 with the LIGO Laboratory's development of a secure network data server, NDS2. This system now reliably serves raw, archived, and trend data, and is robust enough for use by the entire collaboration. However, because Advanced LIGO is significantly more complex than initial LIGO, it will have many more channels, and in the era leading to first detections demand for served data will be greater. It is critical for detector characterization work in aLIGO that NDS2 be supported to reliably serve all raw, archived and trend data available on frames to a large number of users.

For Advanced LIGO we also require data viewing and signal processing tools to read the data served by NDS2 and make a variety of plots or results ranging from quick looks at timeseries and spectra to more complex analyses. The tools currently under active use and development are the Matlab-based graphical user interface LIGO Data Viewer, ligoDV [641], a script-based Matlab interface, mDV [642], and a python interface to the NDS server, GWpy, sand ligoDV-web (`ldvw.ligo.caltech.edu`), was developed. This service has made it possible to access LIGO data through a web browser on your desktop, laptop, tablet or smartphone, and only requires users to have a valid ligo.org username and password for authentication.

Software priorities This section describes priorities for detector characterization software work over the next few years. These activities will be coordinated with the Software Working Group.

1. Implement and test glitch parametrization software that can run online and continuously on the detector output and auxiliary channels and generates output that is improved (in sensitivity, particularly at low frequencies, SNR and frequency accuracy) with respect to the triggers that were produced in S6. Prepare to process hundreds of fast channels per detector from 1Hz to 6kHz. In order to make this information easily accessible by other characterization tools, we adopt the common trigger handling format defined in T1300468 [643] for the upcoming observation runs.

2. Develop and implement Online Detector Characterization (ODC) channels to be deployed in the aLIGO front-end systems that will monitor key aspects of the interferometers and their subsystems and provide critical and authoritative information about the interferometer state.

3. Continue development of Channel Information System (CIS) `cis.ligo.org` containing channel names, sample frequencies, editable descriptions, links to appropriate subsystem models, and other information.

4. Automate and improve upon current data quality flag and veto segment performance validation tools. For Advanced LIGO these should be capable of running daily (and on longer timescales) for all data quality and vetoes and report individual and cumulative efficiency, deadtime, used percentage and safety with respect to hardware signal injections.

5. Improve the current dead channel monitor with a lower false alarm, integrated reporting, and more direct ties to the segment database. The LIGO detector uses thousands of auxiliary channels to validate instrumental behavior and to reveal environmental or instrumental disturbances coupled to the gravitational-wave (GW) strain channel. This information is invaluable for identifying excess detector

noise and to help reduce false candidate events in gravitational-wave searches. However, the associated sensors can become faulty or disconnected. Hence, commissioners require having a diagnostic tool for monitoring auxiliary channels. The utilities should include locating a malfunctioning channel, graphic information of channel's time series and spectral data, and spectral change. In addition, since the GW strain channel can be affected by various band-limited environmental disturbances of non-astrophysical origin, a monitoring tool providing band-wise information is required. The detector characterization group will fully develop and test such a tool during the next year.

6. Produce software to monitor the subsystems of Advanced LIGO that form the foundation for data quality flags in the first runs.
7. Continue development of the LIGO segment database to increase input and output speed, robustness and to improved user interface tools.
8. Develop a new trigger database appropriate for the storage of short-duration veto information. This should be able to store parameters such as central time, duration, central frequency and SNR.
9. Continue development of NDS2, and data access/viewer/processing tools such as ligoDV, ligoDV-Web, pynds, to ensure easy and reliable access to LIGO data and standard signal processing techniques for detector characterization.
10. Continue refinement of veto production algorithms and test these improvements on aLIGO subsystem data.
11. Maintain appropriate reduced data sets for Advanced LIGO to be used for detector characterization and for data analysis. This includes data from engineering runs.

4.1.7 Noise Transients

The goals of the data quality subgroup with respect to noise transients are:

- To identify the times of brief transients in the data taken during engineering and science runs that will affect the astrophysical searches.
- To investigate the causes of these transients using information from auxiliary instrumental and environmental channels and other information such as logbook entries.
- To work with commissioners and experimentalists to confirm the suspected causes of transients and attempt to mitigate them by changes to the instrument.
- To produce data quality flags or other information that can be used by the astrophysical searches to reduce the effect of any transients that are impossible or impractical to mitigate.
- To provide information to experimentalists and builders of future detectors to achieve interferometer noise that is stationary and Gaussian.

The priorities for the coming year are:

- Participate in engineering and observing runs, including staffing Data Quality shifts during times designated for intensive characterization of the instrument (cf. https://dcc.ligo.org/LIGO-L1500110)

- Automate production of graphical visualization of the data products needed for evaluating data quality and identifying transients. Work with commissioners to configure these plots so that the most important channels and information are emphasized. The automated plots should be useful to commissioners and also the primary tools used for glitch shifts.
- Develop standard filter tools for producing audio from instrument data, to use human audio signal processing to aid in understanding noise sources and environmental coupling; develop a library of sounds associated with phenomena to aid in future studies
- Work to provide low-latency tools for processing and identifying instrument and data defects, to both support the low-latency pipelines and to aid in the DetChar event vetting responsibility
- Devise improved ways to diagnose problems arising from data acquisition, data sampling and/or imperfect timing in digital control systems.
- Tune and improve currently existing code, and develop new approaches, for finding and diagnosing data quality problems.

Going forward, we would like use the daily summary pages to provide the automated plots which are used for the Data Quality shifts to be undertaken during engineering and Observing runs.

The first and foremost goal for the advanced detector era is to enable the astrophysical searches to make confident detections. This requires a deep cleaning of the background by understanding nearly all of the glitches that are of concern to the astrophysical searches, and either mitigating them or creating data quality flags that identify them with good accuracy. There are a number of sub-goals that will facilitate requires better ways to analyze the auxiliary channels that provide information about the state of the instrument and the environment, since it is this information that predicts the occurrence of glitches.

4.1.8 Spectral Features

Another working group of the LSC Detector Characterization group is charged with investigating spectral features of the gravitational wave spectral noise density, which is especially important for the searches of gravitational waves from rotating stars and stochastic background. Many of the spectral features are due to environmental disturbances, including seismic activity, high wind, acoustic noise, and electromagnetic interference. Some sources are natural, but many are also anthropogenic, including sources from observatory infrastructure, e.g., nearby motors and HVAC systems. A wide variety of environmental channels have been commissioned and are monitored in initial, enhanced and advanced LIGO, but unusual artefacts typically require detailed on-site investigations and eventually mitigation, work carried out by scientists from the observatories and from LSC institutions, as part of commissioning. Acoustic mitigation has played an especially critical role in lowering interferometer noise floors[23]. Significant increase in gravitational wave channel noise is seen during high seismic noise times, due to not very well understood upconversion of the noise into the gravitational wave band (40Hz-6kHz). Environmental disturbances may also, of course, be manifested through linear couplings to the interferometer as direct glitches or lines, for sources with characteristic frequencies in the LIGO band of sensitivity.

The list of high priority activities related to characterizing spectral features in the upcoming year are:

- Continue to analyze and investigate noise lines affecting data quality in strain data, including summarizing frequency line issues.
- *Noise budget for subsystems*: Measure the environment about advanced LIGO subsystems to identify periodic signals so as to develop a catalog of potential noise lines that could enter these sub-systems. Conduct noise injection tests to measure the transfer function of different environmental noise sources.

- *List of lines and line monitors in subsystems*: Apply the existing noise line finding tools in order to characterize the noise environment of advanced LIGO. Use seismometers, accelerometers, microphones, magnetometers, voltage line monitors and other devices to map out noise, and how it couples into advanced LIGO subsystems. Use existing line finding tools, such as Fscan (a pulsar search code, applied to auxiliary channels), coherence (which calculates the coherence between the gravity wave channel and auxiliary channels), and NoEMI (Noise Event Miner, developed at Virgo).
- *Investigate coherence of environmental channels with the different subsystems*: Use the coherence tool to monitor the coherence between various signals. The Stochastic Transient Analysis Multi-detector Pipeline (STAMP) also allows for the long-term monitoring of the coherence between different channel pairs. These tools will be used to monitor noise signals in subsystems, producing an executive summary for each system. There will also be a need to study non-linear frequency up-conversion of noise; STAMP, as well as bicoherence code, will be used to study up-conversion of noise in subsystems.

5 Calibration

5.1 LIGO calibration

For the LIGO interferometers, *calibration* involves converting data streams from channels that monitor the feedback control loop that maintains the differential arm length into a derived time series that represents the inferred differential arm length variations, h(t), which is normalized to the average arm length, approximately 4000m. h(t) is referred to as *interferometer strain* or just *strain*. The analog and digital filters used in h(t) production are first produced in the frequency domain by the calibration and commissioning team.

Calibration of the LIGO interferometers is a task critical to the success of the data analysis algorithms, and the confidence associated with their results. As such, the LSC created in its bylaws a Calibration Committee, separate from the Detector Characterization group, although there are still many common members and activities. The goal of the Calibration Committee is to provide calibrated h(t) with sufficiently small uncertainities in amplitude, phase, and timing. The current tentative goal is to have maximum calibration errors of roughly 10 percent in amplitude, a few degrees in phase, and about 10 microsecond in timing.

Calibration of a detector is a complex task that involves instrumental hardware measurements, detector modeling, computer programs, and extensive validation and review. The time domain calibrated data is the main data product, and its generation is sufficiently complex that it needs a dedicated team for calibration and another one for review. The Calibration Committee is therefore co-chaired by a time-domain chair and an experimental chair, and includes LIGO Laboratory and other LSC scientists. It works along with a dedicated Calibration Review Committee which provides advice and vetting of this work. The Calibration Committee's results are posted and documented on a web page[644] available to the LSC, and as with previous science runs, will continue to be recorded in the electronic logs, software repositories, and LIGO documents[645].

The scope of the calibration itself was expanded during and after the S5 run to include the timing of the LIGO data. If the interferometer model used for calibration is incorrect, it could skew the timing of LIGO data even if the clock system is working perfectly. See 5.1.1.

Estimation and reduction of the errors in the calibration data products will be a major effort in aLIGO. Towards that end multiple methods of calibration will be used, including a method using auxiliary laser pressure actuation ("photon calibrator")[**?**] and a method using interferometer laser frequency modulation[**?**], both of which were used in initial LIGO science runs. The aLIGO photon calibrator subsystem has been installed and has been used successfully in understanding the calibration of the seventh engineering run, and will continue to be an important part of calibration going into the first observing run.

Production and analysis of the time-dependent calibration coefficients is an essential tool for calibration validations. They can be used to estimate the systematic and statistical uncertainties of the calibration as well as the time offset changes. These studies will be continued in the future. Development of on-line tools to monitor the time-dependent calibration factors, and more generally h(t) data quality, is essential.

The Calibration Committee's membership has been augmented in recent years by graduate students and scientists alike from several LSC institutions. Each site will have a dedicated LIGO lab person responsible for the calibration, but the Calibration Committee expects additional manpower of about 3 people per site, on time scales of 6-8 weeks per year, to be necessary to get calibration out and vetted in a timely manner around science runs. This manpower would be in addition to those working on the calibration software pipelines and those maintaining close communication with various aLIGO subsystem groups. This work provides students valuable instrumental training. It would be highly desirable to sustain this broad participation.

In anticipation of the aLIGO science runs we will be creating and maintaining communication channels between aLIGO and other projects' calibration teams and reviewers. In collaboration with Virgo and GEO, the calibration team will also work on improving h(t) generation techniques, and the development of pre-processed h(t) products such as whitened, cleaned, and coherent data streams. Also important is an exchange of ideas about the review process.

The work of the calibration team is currently focused on preparations for the advanced detector era. New techniques are being commissioned to produce h(t) data with second and sub-second latencies (during S6 the latency was 1 minute). These techniques include moving the generation of h(t) to the front end of the interferometer (CDS) and a gstreamer-based algorithm. In addition, online tools to monitor the quality of the data produced on the fly, and the development of pre-processed h(t) products (e.g. whitened, cleaned, and coherent data streams) are being developed.

The front end calibration effort is intended to develop the necessary code to perform time domain calibration on the CDS computers that directly runs the interferometer. This code is directly embedded in the controls code. This method has the advantage of providing the lowest latency possible as it works directly with the data before it is sent on to be recorded, and can thus be included directly in the recorded frame data. The disadvantage of this method is that has difficultly handling the effects of poles and zeros in the actuation and sensing functions above the nyquist frequency that the front end models run at.

The calibration team is commisioning a low-latency (sub-second latency) gstreamer-based pipeline for time domain calibration in aLIGO. This is a robust pipeline with both frame file and shared memory I/O capabilities, thus allowing for the same pipeline to run for both online and offline calibration. The online infrastructure required for the aLIGO gstreamer-based calibration pipeline was successfully tested during aLIGO's seventh engineering run. The offline infrastructure for the pipeline has also been tested during offline data reproduction after the sixth engineering run. The calibration team is continuing to commission both the infrastructure required for and the inner-workings of the low-latency aLIGO time domain calibration pipeline.

The current plan for the eigth engineering run and first observing run is to combine the front end calibration code outputs with the low-latency gststreamer pipline. The front end code was easy for the calibrators on site, who did the calibration measurements of the interferometer, to keep up to date with changes in the actuation and sensing of the interferometer. The low-latency gst-lal calibration pipeline is able to take data products from that front end code and properly apply poles and zeros above the nyquist frequency, as well as properly handle time delay, and generate the final h(t) stream. In the case where calibration needs to be recalculated after data taking due to errors, the offline gstreamer pipeline is capable of producing from just the raw differential arm control signals the same data products as the front end code. This mode of operation was tested successfully against the front end calibration code during the seventh engineering run.

5.1.1 Timing

Traceable and closely monitored timing performance of the GW detectors is mission critical for reliable interferometer operation, astrophysical data analysis, and discoveries. For example, (a) timing jitter of digitization of the GW signal could directly contribute to the noise level degrading the astrophysical reach of the LIGO interferometers, (b) coincident and coherent observation using the network of GW detectors is only possible if the absolute timing of the data streams agree within a high degree of accuracy, (c) a network of interferometric GW detectors can only recover both the polarization and sky direction information for a detected event if the absolute timing of their data-streams are known, and (d) multimessenger astronomy with external observatories also require traceable and accurate absolute timing.

The Timing Stability Working Group (TSWG) includes scientists from both the LSC and the LIGO Laboratory. The group shall be responsible for (a.) the availability and diagnostics of timing information and signals provided for various subsystems (e.g., LSC, OMC, etc.), (b.) measuring and documenting the timing performance of mission critical digital subsystems such as LSC and OMC DAQs, (c.) in close collaboration with the Calibration team (also see **??**), the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator[646], characterization of analog modules, etc.), and (d.) the documented review and certification of the physical/software implementation and verification of the availability of precise documentation of timing related parts of mission critical subsystems. While it is quite likely that issues with the timing performance of subsystems are discovered by the timing team, it is the responsibility of the subsystems to address the problem; the timing team is responsible only for the certification that the issue was indeed eliminated.

The construction, testing and diagnostics tasks have already provided fertile ground for undergraduate and graduate student research involvement and diversity in the program is strongly encouraged for the future.

The next challenge in timing diagnostic is long term. Several projects will be executed in preparation of the advanced detector era, such as:

- Further develop and test injection techniques to determine accurate timing through direct test mass excitations
- Augment and expand the capabilities of data monitoring tools related to timing and phase calibration
- Enhance the availability of timing diagnostics capabilities provided for various subsystems
- Measure and document the timing performance of mission critical digital subsystems
- Measure and document the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator, characterization of analog modules, etc.)
- Review and certify the physical/software implementation and verify of the availability of precise documentation of timing related parts of mission critical subsystems

References

[1] the LIGO Scientific Collaboration, the Virgo Collaboration: J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, and et al. All-sky search for gravitational-wave bursts in the second joint LIGO-Virgo run. *Physical Review D*, 85(12):122007, June 2012.

[2] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev. D*, 72:122002, 2005.

[3] S. Klimenko et al. A coherent method for detection of gravitational wave bursts. *Class. Quantum Grav.*, 25(11):114029–+, June 2008.

[4] S. Klimenko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin. Localization of gravitational wave sources with networks of advanced detectors. *Phys. Rev. D*, 83:102001, May 2011.

[5] P. J. Sutton, G. Jones, S. Chatterji, P. Kalmus, I. Leonor, S. Poprocki, J. Rollins, A. Searle, L. Stein, M. Tinto, and M. Was. X-Pipeline: an analysis package for autonomous gravitational-wave burst searches. *New Journal of Physics*, 12:053034–+, 2010.

[6] E. Thrane and M. Coughlin. Seedless clustering in all-sky searches for gravitational-wave transients. *ArXiv e-prints*, January 2014. arXiv:1401.8060.

[7] S. Chamberlin et al. Applications of the ExcessPower Data Analysis Pipeline to Gravitational Wave Detection Efforts, 2013.

[8] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21:S1809, 2004.

[9] Tyson B. Littenberg and Neil J. Cornish. Separating gravitational wave signals from instrument artifacts. *Phys. Rev. D*, 82:103007, Nov 2010.

[10] LIGO Scientific Collaboration, Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, and et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *arXiv:1304.0670*, April 2013.

[11] J. Aasi, J. Abadie, B.P. Abbott, R. Abbott, T. Abbott, et al. Constraints on cosmic strings from the LIGO-Virgo gravitational-wave detectors. *Phys.Rev.Lett.*, 112:131101, 2014.

[12] J. Abadie et al. Predictions for the rates of compact binary coalescences observable by ground-based gravitational-wave detectors. *Class. Quantum Grav.*, 27:173001, March 2010.

[13] J. M. Lattimer. Observed Neutron Star Masses, 2015. [Online at http://www.stellarcollapse.org/nsmasses; accessed 17-May-2015].

[14] Bulent Kiziltan, Athanasios Kottas, and Stephen E. Thorsett. The Neutron Star Mass Distribution, 2010.

[15] Ph. Podsiadlowski, N. Langer, A. J. T. Poelarends, S. Rappaport, A. Heger, and E. Pfahl. The effects of binary evolution on the dynamics of core collapse and neutron star kicks. *The Astrophysical Journal*, 612(2):1044, 2004.

[16] E. P. J. van den Heuvel. Double neutron stars: Evidence for two different neutron-star formation mechanisms. *AIP Conference Proceedings*, 924(1):598–606, 2007.

[17] S.E. Woosley, A. Heger, and T.A. Weaver. The evolution and explosion of massive stars. *Rev.Mod.Phys.*, 74:1015–1071, 2002.

[18] M. Kramer and N. Wex. The double pulsar system: a unique laboratory for gravity. *Classical and Quantum Gravity*, 26:073001, 2009.

[19] Jason W.T. Hessels, Scott M. Ransom, Ingrid H. Stairs, Paulo Cesar Carvalho Freire, Victoria M. Kaspi, et al. A radio pulsar spinning at 716-hz. *Science*, 311:1901–1904, 2006.

[20] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Physical Review D*, 86(8):084017, October 2012.

[21] Lawrence E. Kidder. Coalescing binary systems of compact objects to (post)$^{5/2}$-Newtonian order. V. Spin effects. *Physical Review D*, 52(2):821–847, 1995.

[22] Lawrence E. Kidder, Clifford M. Will, and Alan G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. *Physical Review D*, 47(10):R4183–R4187, May 1993.

[23] Alessandra Buonanno, Bala Iyer, Evan Ochsner, Yi Pan, and B.S. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys.Rev.*, D80:084043, 2009.

[24] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 17:2, 2014.

[25] B. D. Metzger and E. Berger. What is the Most Promising Electromagnetic Counterpart of a Neutron Star Binary Merger? *Astrophysical Journal*, 746:48, February 2012.

[26] S. Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New Journal of Physics*, 11:123006, December 2009.

[27] Carl L. Rodriguez, Benjamin Farr, Vivien Raymond, Will M. Farr, Tyson B. Littenberg, et al. Basic Parameter Estimation of Binary Neutron Star Systems by the Advanced LIGO/Virgo Network. *Astrophys.J.*, 784:119, 2014.

[28] F. Beauville et al. Detailed comparison of LIGO and Virgo inspiral pipelines in preparation for a joint search. *Class. Quantum Grav.*, 25:045001, 2008.

[29] Kipp Cannon, Adrian Chapman, Chad Hanna, Drew Keppel, Antony C. Searle, and Alan J. Weinstein. Singular value decomposition applied to compact binary coalescence gravitational-wave signals. *Physical Review D*, 82(4):044025, Aug 2010.

[30] Kipp Cannon et al. Toward early-warning detection of gravitational waves from compact binary coalescence. *The Astrophysical Journal*, 748(2):136, 2012.

[31] Jing Luan, Shaun Hooper, Linqing Wen, and Yanbei Chen. Towards low-latency real-time detection of gravitational waves from compact binary coalescences in the era of advanced detectors. *Phys.Rev.*, D85:102002, 2012.

[32] Shaun Hooper, Shin Kee Chung, Jing Luan, David Blair, Yanbei Chen, et al. Summed Parallel Infinite Impulse Response (SPIIR) Filters For Low-Latency Gravitational Wave Detection. *Phys.Rev.*, D86:024012, 2012.

[33] Yuan Liu, Zhihui Du, Shin Kee Chung, Shaun Hooper, David Blair, et al. GPU-accelerated low-latency real-time searches for gravitational waves from compact binary coalescence. *Class.Quant.Grav.*, 29:235018, 2012.

[34] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. *Physical Review D*, 88:062001, 2013.

[35] J. M. Lattimer. The Nuclear Equation of State and Neutron Star Masses. *Annual Review of Nuclear and Particle Science*, 62:485–515, November 2012.

[36] Jocelyn S. Read, Luca Baiotti, Jolien D. E. Creighton, John L. Friedman, Bruno Giacomazzo, et al. Matter effects on binary neutron star waveforms. *Phys.Rev.*, D88:044042, 2013.

[37] Walter Del Pozzo, Tjonnie G. F. Li, Michalis Agathos, Chris Van Den Broeck, and Salvatore Vitale. Demonstrating the feasibility of probing the neutron star equation of state with second-generation gravitational wave detectors. *Phys. Rev. Lett.*, 111:071101, 2013.

[38] M. Favata. Systematic Parameter Errors in Inspiraling Neutron Star Binaries. *Physical Review Letters*, 112(10):101101, March 2014.

[39] Leslie Wade, Jolien D. E. Creighton, Evan Ochsner, Benjamin D. Lackey, Benjamin F. Farr, et al. Systematic and statistical errors in a bayesian approach to the estimation of the neutron-star equation of state using advanced gravitational wave detectors. *Phys.Rev.*, D89:103012, 2014.

[40] Paul D. Lasky, Brynmor Haskell, Vikram Ravi, Eric J. Howell, and David M. Coward. Nuclear Equation of State from Observations of Short Gamma-Ray Burst Remnants. *Phys.Rev.*, D89:047302, 2014.

[41] M. Tanaka, K. Hotokezaka, K. Kyutoku, S. Wanajo, K. Kiuchi, Y. Sekiguchi, and M. Shibata. Radioactively Powered Emission from Black Hole-Neutron Star Mergers. *Astrophysical Journal*, 780:31, January 2014.

[42] J. M. Weisberg, D. J. Nice, and J. H. Taylor. Timing Measurements of the Relativistic Binary Pulsar PSR B1913+16. *Astrophysical Journal*, 722:1030–1034, October 2010.

[43] M. Kramer, I. H. Stairs, R. N. Manchester, M. A. McLaughlin, A. G. Lyne, R. D. Ferdman, M. Burgay, D. R. Lorimer, A. Possenti, N. D'Amico, J. M. Sarkissian, G. B. Hobbs, J. E. Reynolds, P. C. C. Freire, and F. Camilo. Tests of General Relativity from Timing the Double Pulsar. *Science*, 314:97–102, October 2006.

[44] L. Blanchet and B. S. Sathyaprakash. Signal analysis of gravitational wave tails. *Classical and Quantum Gravity*, 11:2807–2831, November 1994.

[45] L. Blanchet and B. S. Sathyaprakash. Detecting a Tail Effect in Gravitational-Wave Experiments. *Physical Review Letters*, 74:1067–1070, February 1995.

[46] Michalis Agathos, Walter Del Pozzo, Tjonnie G. F. Li, Chris Van Den Broeck, John Veitch, et al. TIGER: A data analysis pipeline for testing the strong-field dynamics of general relativity with gravitational wave signals from coalescing compact binaries. *Phys.Rev.*, D89:082001, 2014.

[47] Stas Babak, Balasubramanian, David Churches, Thomas Cokelaer, and B.S. Sathyaprakash. A template bank to search for gravitational waves from inspiralling compact binaries i: physical models. *Class. Quantum Grav.*, 23:5477–5504, 2006.

[48] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D.E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys.Rev.*, D85:122006, 2012.

[49] Bruce Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Physical Review D*, 71:062001, 2005.

[50] Andres Rodríguez. Reducing false alarms in searches for gravitational waves from coalescing binary systems. Master's thesis, Louisiana State University, 2007.

[51] C. A. K. Robinson, B. S. Sathyaprakash, and Anand S. Sengupta. Geometric algorithm for efficient coincident detection of gravitational waves. *Physical Review D*, 78(6):062002, 2008.

[52] S. Babak, R. Biswas, P. R. Brady, D. A. Brown, K. Cannon, C. D. Capano, J. H. Clayton, T. Cokelaer, J. D. E. Creighton, T. Dent, A. Dietz, S. Fairhurst, N. Fotopoulos, G. González, C. Hanna, I. W. Harry, G. Jones, D. Keppel, D. J. A. McKechan, L. Pekowsky, S. Privitera, C. Robinson, A. C. Rodriguez, B. S. Sathyaprakash, A. S. Sengupta, M. Vallisneri, R. Vaulin, and A. J. Weinstein. Searching for gravitational waves from binary coalescence. *Physical Review D*, 87(2):024033, January 2013.

[53] B. Abbott et al. Analysis of LIGO data for gravitational waves from binary neutron stars. *Physical Review D*, 69:122001, 2004.

[54] B. Abbott et al. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Physical Review D*, 72:082001, 2005.

[55] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D*, 77:062002, 2008.

[56] B. Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. *Physical Review D*, 79:122001, 2009.

[57] B.P. Abbott et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run. *Physical Review D*, 80:047101, 2009.

[58] J. Abadie et al. Search for compact binary coalescences in LIGO and Virgo data from S5 and VSR1. *Physical Review D*, 82:102001, 2010.

[59] J. Abadie et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3. *Phys.Rev.*, D85:082002, 2012.

[60] J. Abadie et al. Evidence for the Direct Detection of Gravitational Waves from a Black Hole Binary Coalescence. Technical Report LIGO-P1000146, LIGO Scientific Collaboration and Virgo Collaboration, 2011.

[61] Kipp Cannon, Chad Hanna, and Drew Keppel. Method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence. *Phys. Rev. D*, 88:024025, Jul 2013.

[62] G. Gonzalez et al. LSC and Virgo policy on releasing gravitational wave triggers to the public in the advanced detectors era. Technical Report LIGO-M1200055-v2, 2012.

[63] F. Raab. Working Timeline to First Observing Run of the Advanced Detector Era. Technical Report LIGO-G1301309, LIGO Scientific Collaboration and Virgo Collaboration, 2014.

[64] The LIGO Scientific Collaboration and the Virgo Collaboration. Commissioning and Observing Scenarios for the Advanced LIGO and Advanced Virgo Gravitational-Wave Observatories. Technical Report LIGO-P1200087, The LIGO Scientific Collaboration and the Virgo Collaboration, 2013. https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=94147.

[65] BNS Spin MDC with PyCBC pipeline. `https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/BNS/offline_pycbc_search/spinnospininv`.

[66] Tito Dal Canton, Andrew P. Lundgren, and Alex B. Nielsen. Impact of precession on aligned-spin searches for neutron-star-black-hole binaries. *Phys.Rev.*, D91(6):062010, 2015.

[67] J. Aasi et al. The LSC-Virgo Policy on Multiple Pipelines. Technical Report LIGO-M1500027-v1, LIGO Scientific Collaboration and Virgo Collaboration, 2015.

[68] BNS MultiPipeline MDC results summary and investigations. `https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/BNS/MDC/MultiPipelineResultsSummary`.

[69] Proposal for combining offline CBC searches. `https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/Face2Face20150316/sources/offlinecombineproposal`.

[70] Proposal for O1 Low Latency MDCs and mass/spin search space. `https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/BNS/MDC/LLMDC`.

[71] Detector characterization group. Plan for Providing Data Quality Information for Advanced LIGO Searches. *LIGO Document T1400002*, 2014.

[72] Y Hu. Estimation of gravitational wave event statistical significance. Technical Report LIGO-1500528-v1, 2015.

[73] W. M. Farr, J. R. Gair, I. Mandel, and C. Cutler. Counting and confusion: Bayesian rate estimation with multiple populations. *Physical Review D*, 91(2):023005, January 2015.

[74] J. Veitch, V. Raymond, B. Farr, W. Farr, P. Graff, et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys.Rev.*, D91(4):042003, 2015.

[75] L. P. Singer. *The needle in the hundred square degree haystack: The hunt for binary neutron star mergers with LIGO and Palomar Transient Factory*. PhD thesis, California Institute of Technology, 2015.

[76] Leo P. Singer, Larry R. Price, Ben Farr, Alex L. Urban, Chris Pankow, et al. The First Two Years of Electromagnetic Follow-Up with Advanced LIGO and Virgo. *Astrophys.J.*, 795(2):105, 2014.

[77] C. P. L. Berry, I. Mandel, H. Middleton, L. P. Singer, A. L. Urban, A. Vecchio, S. Vitale, K. Cannon, B. Farr, W. M. Farr, P. B. Graff, C. Hanna, C.-J. Haster, S. Mohapatra, C. Pankow, L. R. Price, T. Sidery, and J. Veitch. Parameter estimation for binary neutron-star coalescences with realistic noise during the Advanced LIGO era. *ArXiv e-prints*, November 2014.

[78] B Farr. *Extracting Astrophysical Information from the Gravitational Waves of Compact Binary Mergers and Their Electromagnetic Counterparts*. PhD thesis, Northwestern University, 2015.

[79] C. Pankow, P. Brady, E. Ochsner, and R. O'Shaughnessy. A novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences. *ArXiv e-prints*, February 2015.

[80] C.-J. Haster, I. Mandel, and W. M. Farr. Efficient method for measuring the parameters encoded in a gravitational-wave signal. *ArXiv e-prints*, February 2015.

[81] T. B. Littenberg, M. Coughlin, B. Farr, and W. M. Farr. Fortifying the characterization of binary mergers in LIGO data. *ArXiv e-prints*, July 2013.

[82] Tyson B. Littenberg and Neil J. Cornish. Bayesian inference for spectral estimation of gravitational wave detector noise. *Phys.Rev.*, D91(8):084034, 2015.

[83] S. Vitale. Calibration requirements for CBC in the Advanced Detector Era. Technical Report LIGO-G1300788-v2, 2013.

[84] Salvatore Vitale, Walter Del Pozzo, Tjonnie G.F. Li, Chris Van Den Broeck, Ilya Mandel, et al. Effect of calibration errors on Bayesian parameter estimation for gravitational wave signals from inspiral binary systems in the Advanced Detectors era. *Phys.Rev.*, D85:064034, 2012.

[85] Bruce Allen. LIGO calibration accuracy. Technical Report LIGO-T960189-00-Z, LIGO Project, 1996.

[86] K. Belczynski, V. Kalogera, F. A. Rasio, R. E. Taam, A. Zezas, T. Bulik, T. J. Maccarone, and N. Ivanova. Compact Object Modeling with the StarTrack Population Synthesis Code. *ApJS*, 174:223–260, January 2008.

[87] J.M.B. Downing, M.J. Benacquista, M. Giersz, and R. Spurzem. Compact Binaries in Star Clusters I - Black Hole Binaries Inside Globular Clusters. *Mon.Not.Roy.Astron.Soc.*, 407:1946, 2010.

[88] Feryal Ozel, Dimitrios Psaltis, Ramesh Narayan, and Jeffrey E. McClintock. The Black Hole Mass Distribution in the Galaxy. *Astrophys.J.*, 725:1918–1927, 2010.

[89] W. M. Farr, N. Sravan, A. Cantrell, L. Kreidberg, C. D. Bailyn, I. Mandel, and V. Kalogera. The Mass Distribution of Stellar-Mass Black Holes. *ArXiv e-prints*, November 2010.

[90] A. Merloni. Cosmological evolution of supermassive black holes and AGN: a synthesis model for accretion and feedback. Mem. Soc. Astron. Italiana, 79:1310, 2008.

[91] Laura Kreidberg, Charles D. Bailyn, Will M. Farr, and Vassiliki Kalogera. Mass Measurements of Black Holes in X-Ray Transients: Is There a Mass Gap? *Astrophys.J.*, 757:36, 2012.

[92] K. Belczynski, G. Wiktorowicz, C. Fryer, D. Holz, and V. Kalogera. Missing Black Holes Unveil The Supernova Explosion Mechanism. *Astrophys.J.*, 757:91, 2012.

[93] C.L. Fryer, K. Belczynski, G. Wiktorowicz, M. Dominik, V. Kalogera, et al. Compact Remnant Mass Function: Dependence on the Explosion Mechanism and Metallicity. *Astrophys.J.*, 749:91, 2012.

[94] Richard W. O'Shaughnessy, J. Kaplan, V. Kalogera, and K. Belczynski. Bounds on expected black hole spins in inspiraling binaries. *Astrophys.J.*, 632:1035–1041, 2005.

[95] P. B. Demorest, T. Pennucci, S. M. Ransom, M. S. E. Roberts, and J. W. T. Hessels. A two-solar-mass neutron star measured using Shapiro delay. *Nature*, 467:1081–1083, October 2010.

[96] John Antoniadis, Paulo C.C. Freire, Norbert Wex, Thomas M. Tauris, Ryan S. Lynch, et al. A Massive Pulsar in a Compact Relativistic Binary. *Science*, 340:6131, 2013.

[97] Jeffrey M. Silverman and Alexei V. Filippenko. On IC 10 X-1, the Most Massive Known Stellar-Mass Black Hole. *Astrophys.J.*, 678:L17–L20, 2008.

[98] P. A. Crowther, R. Barnard, S. Carpano, J. S. Clark, V. S. Dhillon, and A. M. T. Pollock. NGC 300 X-1 is a Wolf-Rayet/black hole binary. *MNRAS*, pages L11+, January 2010.

[99] Krzysztof Belczynski et al. On The Maximum Mass of Stellar Black Holes. *Astrophysical Journal*, 714:1217–1226, 2010.

[100] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. C. Miller, and M. Walczak. The Formation and Gravitational-wave Detection of Massive Stellar Black Hole Binaries. *Astrophysical Journal*, 789:120, July 2014.

[101] Michal Dominik, Krzysztof Belczynski, Christopher Fryer, Daniel Holz, Emanuele Berti, et al. Double Compact Objects I: The Significance of the Common Envelope on Merger Rates. *Astrophys.J.*, 759:52, 2012.

[102] M. Dominik, E. Berti, R. O'Shaughnessy, I. Mandel, K. Belczynski, et al. Double Compact Objects III: Gravitational Wave Detection Rates. 2014.

[103] J.M. Miller, C.S. Reynolds, A.C. Fabian, G. Miniutti, and L.C. Gallo. Stellar-mass Black Hole Spin Constraints from Disk Reflection and Continuum Modeling. *Astrophys.J.*, 697:900–912, 2009.

[104] R. Shafee, J. E. McClintock, R. Narayan, S. W. Davis, L.-X. Li, and R. A. Remillard. Estimating the Spin of Stellar-Mass Black Holes by Spectral Fitting of the X-Ray Continuum. *Astrophysical Journal Letters*, 636:L113–L116, January 2006.

[105] Jeffrey E. McClintock et al. The Spin of the Near-Extreme Kerr Black Hole GRS 1915+105. *Astrophysical Journal*, 652:518–539, 2006.

[106] J. Liu, J. E. McClintock, R. Narayan, S. W. Davis, and J. A. Orosz. Erratum: "Precise Measurement of the Spin Parameter of the Stellar-mass Black Hole M33 X-7". *Astrophysical Journal Letters*, 719:L109, August 2010.

[107] L. Gou, J. E. McClintock, J. Liu, R. Narayan, J. F. Steiner, R. A. Remillard, J. A. Orosz, S. W. Davis, K. Ebisawa, and E. M. Schlegel. A Determination of the Spin of the Black Hole Primary in LMC X-1. *Astrophysical Journal*, 701:1076–1090, August 2009.

[108] S. W. Davis, C. Done, and O. M. Blaes. Testing Accretion Disk Theory in Black Hole X-Ray Binaries. *Astrophysical Journal*, 647:525–538, August 2006.

[109] L.-X. Li, E. R. Zimmerman, R. Narayan, and J. E. McClintock. Multitemperature Blackbody Spectrum of a Thin Accretion Disk around a Kerr Black Hole: Model Computations and Comparison with Observations. *ApJS*, 157:335–370, April 2005.

[110] Jeffrey E. McClintock, Ramesh Narayan, Shane W. Davis, Lijun Gou, Akshay Kulkarni, et al. Measuring the Spins of Accreting Black Holes. *Class.Quant.Grav.*, 28:114009, 2011.

[111] Jeffrey E. McClintock, Ramesh Narayan, and James F. Steiner. Black Hole Spin via Continuum Fitting and the Role of Spin in Powering Transient Jets. *Space Sci. Rev.*, pages 1–28, July 2013.

[112] James F. Steiner and Jeffrey E. McClintock. Modeling the Jet Kinematics of the Black Hole Microquasar XTE J1550-564: A Constraint on Spin-Orbit Alignment. *Astrophys.J.*, 745:136, 2012.

[113] T. Fragos, M. Tremmel, E. Rantsiou, and K. Belczynski. Black Hole Spin-Orbit Misalignment in Galactic X-ray Binaries. *Astrophysical Journal Letters*, 719:L79–L83, August 2010.

[114] Krzysztof Belczynski, Matthew Benacquista, and Tomasz Bulik. Double Compact Objects as Low-frequency Gravitational Wave Sources. *Astrophys.J.*, 725:816–823, 2010.

[115] V. Kalogera. Spin-orbit misalignment in close binaries with two compact objects. *Astrophysical Journal*, 541:042003, 2000.

[116] R. O'Shaughnessy, C. Kim, V. Kalogera, and K. Belczynski. Constraining Population Synthesis Models via Empirical Binary Compact Object Merger and Supernova Rates. *Astrophysical Journal*, 672:479–488, January 2008.

[117] T. Bulik and K. Belczyński. Constraints on the Binary Evolution from Chirp Mass Measurements. *Astrophysical Journal Letters*, 589:L37–L40, May 2003.

[118] Ilya Mandel. Parameter estimation on gravitational waves from multiple coalescing binaries. *Physical Review D*, D81:084029, 2010.

[119] R. OÕShaughnessy. Data-driven methods to explore a large space of computationally costly compact binary progenitor models. *Phys.Rev.*, D88(8):084061, 2013.

[120] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. *Physical Review D*, 83:122005, 2011.

[121] J. Aasi et al. Search for Gravitational Waves from Binary Black Hole Inspiral, Merger and Ringdown in LIGO-Virgo Data from 2009-2010. *Phys.Rev.*, D87:022002, 2013.

[122] Collin Capano, Ian W. Harry, and Alessandra Buonanno. A search for binary black holes with non-precessing spins. In preparation.

[123] J. Abadie et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys.Rev.*, D81:102001, 2010.

[124] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *arXiv:1304.0670*, 2013. `http://arxiv.org/abs/1304.0670`.

[125] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, D. A. Hemberger, L. E. Kidder, G. Lovelace, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, B. Szilágyi, N. W. Taylor, and A. Zenginoglu. Effective-one-body model for black-hole binaries with generic mass ratios and spins. *Physical Review D*, 89(6):061502, March 2014.

[126] Michael P§rrer. Frequency domain reduced order models for gravitational waves from aligned-spin compact binaries. *Class.Quant.Grav.*, 31(19):195010, 2014.

[127] Mark Hannam, Patricia Schmidt, Alejandro BohŐ, Leila Haegel, Sascha Husa, et al. A simple model of complete precessing black-hole-binary gravitational waveforms. *Phys.Rev.Lett.*, 113:151101, 2014.

[128] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A.J. Weinstein. Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Phys.Rev.*, D89(8):084041, 2014.

[129] Ian W. Harry, Alexander H. Nitz, Duncan A. Brown, Andrew P. Lundgren, Evan Ochsner, and Drew Keppel. Investigating the effect of precession on searches for neutron-star–black-hole binaries with advanced ligo. *Phys. Rev. D*, 89:024010, Jan 2014.

[130] Y. Pan, A. Buonanno, A. Taracchini, L. E. Kidder, A. H. Mroué, H. P. Pfeiffer, M. A. Scheel, and B. Szilágyi. Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism. *Physical Review D*, 89(8):084006, April 2014.

[131] Collin Capano, Yi Pan, and Alessandra Buonanno. Impact of Higher Harmonics in Searching for Gravitational Waves from Non-Spinning Binary Black Holes. *Phys.Rev.*, D89:102003, 2014.

[132] Duncan A. Brown, Prayush Kumar, and Alexander H. Nitz. Template banks to search for low-mass binary black holes in advanced gravitational-wave detectors. *Phys.Rev.*, D87:082004, 2013.

[133] The LIGO Scientific Collaboration, Virgo Collaboration: J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, and et al. First Low-Latency LIGO+Virgo Search for Binary Inspirals and their Electromagnetic Counterparts, 2011.

[134] Kipp Cannon, Chad Hanna, and Drew Keppel. Method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence. *Phys.Rev.*, D88(2):024025, 2013.

[135] Rahul Biswas, Patrick R. Brady, Jolien D. E. Creighton, and Stephen Fairhurst. The Loudest Event Statistic: General Formulation, Properties and Applications. *Class. Quantum Grav.*, 26:175009, 2009.

[136] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. *Phys.Rev.*, D88:062001, 2013.

[137] Scott E. Field, Chad R. Galley, Frank Herrmann, Jan S. Hesthaven, Evan Ochsner, et al. Reduced basis catalogs for gravitational wave templates. *Physical Review Letters*, 106:221102, 2011.

[138] Frank Herrmann, Scott E. Field, Chad R. Galley, Evan Ochsner, and Manuel Tiglio. Towards beating the curse of dimensionality for gravitational waves using Reduced Basis. *Phys.Rev.*, D86:084046, 2012.

[139] Harbir Antil, Scott E. Field, Frank Herrmann, Ricardo H. Nochetto, and Manuel Tiglio. Two-step greedy algorithm for reduced order quadratures, 2012.

[140] Priscilla Canizares, Scott E. Field, Jonathan R. Gair, and Manuel Tiglio. Gravitational wave parameter estimation with compressed likelihood evaluations. *Phys.Rev.*, D87(12):124005, 2013.

[141] Scott E. Field, Chad R. Galley, Jan S. Hesthaven, Jason Kaye, and Manuel Tiglio. Fast prediction and evaluation of gravitational waveforms using surrogate models. *Phys.Rev.*, X4:031006, 2014.

[142] Rory Smith, Chad Hanna, Ilya Mandel, and Alberto Vecchio. Rapidly evaluating the compact binary likelihood function via interpolation. *Phys.Rev.*, D90:044074, 2014.

[143] T. G. F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, J. Veitch, K. Grover, T. Sidery, R. Sturani, and A. Vecchio. Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence. *Physical Review D*, 85(8):082003, April 2012.

[144] Francois Foucart, M. Brett Deaton, Matthew D. Duez, Evan O'Connor, Christian D. Ott, et al. Neutron star-black hole mergers with a nuclear equation of state and neutrino cooling: Dependence in the binary parameters. *Phys.Rev.*, D90(2):024026, 2014.

[145] Benjamin D. Lackey, Koutarou Kyutoku, Masaru Shibata, Patrick R. Brady, and John L. Friedman. Extracting equation of state parameters from black hole-neutron star mergers: aligned-spin black holes and a preliminary waveform model. *Phys.Rev.*, D89(4):043009, 2014.

[146] Mark Hannam, Duncan A. Brown, Stephen Fairhurst, Chris L. Fryer, and Ian W. Harry. When can gravitational-wave observations distinguish between black holes and neutron stars? *Astrophys.J.*, 766:L14, 2013.

[147] Salvatore Vitale, Ryan Lynch, John Veitch, Vivien Raymond, and Riccardo Sturani. Measuring the spin of black holes in binary systems using gravitational waves. *Phys.Rev.Lett.*, 112:251101, 2014.

[148] Katerina Chatziioannou, Neil Cornish, Antoine Klein, and Nicolas Yunes. Spin-Precession: Breaking the Black Hole–Neutron Star Degeneracy, 2014.

[149] R. O'Shaughnessy, B. Farr, E. Ochsner, H.S. Cho, V. Raymond, et al. Parameter Estimation of Gravitational Waves from Precessing BH-NS Inspirals with higher harmonics, 2014.

[150] I. Mandel and R. O'Shaughnessy. Compact binary coalescences in the band of ground-based gravitational-wave detectors. *Classical and Quantum Gravity*, 27(11):114007, June 2010.

[151] Vassiliki Kalogera. Spin orbit misalignment in close binaries with two compact objects. *Astrophys.J.*, 541:319–328, 2000.

[152] Victoria M. Kaspi, M. Bailes, R.N. Manchester, B.W. Stappers, J.S. Sandhu, et al. Discovery of the young, energetic radio pulsar PSR J1105-6107. *Astrophys.J.*, 1997.

[153] George Hobbs, D.R. Lorimer, A.G. Lyne, and M. Kramer. A Statistical study of 233 pulsar proper motions. *Mon.Not.Roy.Astron.Soc.*, 360:974–992, 2005.

[154] Z. Arzoumanian, D.F. Chernoffs, and J.M. Cordes. The Velocity distribution of isolated radio pulsars. *Astrophys.J.*, 568:289–301, 2002.

[155] Brad M.S. Hansen and E. Sterl Phinney. The Pulsar kick velocity distribution. *Mon.Not.Roy.Astron.Soc.*, 291:569, 1997.

[156] Jonathan Grindlay, Simon Portegies Zwart, and Stephen McMillan. Short gamma-ray bursts from binary neutron star mergers in globular clusters. *Nature Physics*, 2:116–119, 2006.

[157] Aleksander Sadowski, Krzysztof Belczynski, Tomasz Bulik, Natalia Ivanova, Frederic A. Rasio, et al. The Total Merger Rate of Compact Object Binaries In The Local Universe. *Astrophys.J.*, 2007.

[158] William H. Lee, Enrico Ramirez-Ruiz, and Glenn van de Ven. Short gamma-ray bursts from dynamically-assembled compact binaries in globular clusters: pathways, rates, hydrodynamics and cosmological setting. *Astrophys.J.*, 720:953–975, 2010.

[159] Matthew J. Benacquista and Jonathan M.B. Downing. Relativistic Binaries in Globular Clusters. *Living Rev.Rel.*, 16:4, 2013.

[160] Tsing-Wai Wong, Francesca Valsecchi, Tassos Fragos, and Vassiliki Kalogera. Understanding Compact Object Formation and Natal Kicks. III. The case of Cygnus X-1. *Astrophys.J.*, 747:111, 2012.

[161] Emmanouela Rantsiou, Shiho Kobayashi, Pablo Laguna, and Frederic Rasio. Mergers of Black Hole - Neutron Star binaries. 1. Methods and First Results. *Astrophys.J.*, 680:1326, 2008.

[162] Koutarou Kyutoku, Hirotada Okawa, Masaru Shibata, and Keisuke Taniguchi. Gravitational waves from spinning black hole-neutron star binaries: dependence on black hole spins and on neutron star equations of state. *Phys.Rev.*, D84:064018, 2011.

[163] Zachariah B. Etienne, Vasileios Paschalidis, and Stuart L. Shapiro. General relativistic simulations of black hole-neutron star mergers: Effects of tilted magnetic fields. *Phys.Rev.*, D86:084026, 2012.

[164] Francois Foucart, Matthew D. Duez, Lawrence E. Kidder, Mark A. Scheel, Bela Szilagyi, et al. Black hole-neutron star mergers for 10 solar mass black holes. *Phys.Rev.*, D85:044015, 2012.

[165] Francois Foucart, M. Brett Deaton, Matthew D. Duez, Lawrence E. Kidder, Ilana MacDonald, et al. Black hole-neutron star mergers at realistic mass ratios: Equation of state and spin orientation effects. *Phys.Rev.*, D87:084006, 2013.

[166] Francois Foucart, Luisa Buchman, Matthew D. Duez, Michael Grudich, Lawrence E. Kidder, et al. First direct comparison of nondisrupting neutron star-black hole and binary black hole merger simulations. *Phys.Rev.*, D88(6):064017, 2013.

[167] Luciano Rezzolla, Bruno Giacomazzo, Luca Baiotti, Jonathan Granot, Chryssa Kouveliotou, et al. The missing link: Merging neutron stars naturally produce jet-like structures and can power short Gamma-Ray Bursts. *Astrophys.J.*, 732:L6, 2011.

[168] Francois Foucart. Black Hole-Neutron Star Mergers: Disk Mass Predictions. *Phys.Rev.*, D86:124007, 2012.

[169] Andrea Maselli, Vitor Cardoso, Valeria Ferrari, Leonardo Gualtieri, and Paolo Pani. Equation-of-state-independent relations in neutron stars. *Phys.Rev.*, D88(2):023007, 2013.

[170] Hideyuki Tagoshi, Chandra Kant Mishra, Archana Pai, and K.G. Arun. Parameter estimation of neutron star-black hole binaries using an advanced gravitational-wave detector network: Effects of the full post-Newtonian waveform. *Phys.Rev.*, D90:024053, 2014.

[171] K.G. Arun, Hideyuki Tagoshi, Chandra Kant Mishra, and Archana Pai. Synergy of short gamma ray burst and gravitational wave observations: Constraining the inclination angle of the binary and possible implications for off-axis GRBs. *Phys.Rev.*, D90:024060, 2014.

[172] E. Berger. The Host Galaxies of Short-Duration Gamma-Ray Bursts: Luminosities, Metallicities, and Star Formation Rates. *Astrophys.J.*, 690:231–237, 2009.

[173] Francesco Pannarale, Luciano Rezzolla, Frank Ohme, and Jocelyn S. Read. Will black hole-neutron star binary inspirals tell us about the neutron star equation of state? *Phys.Rev.*, D84:104017, 2011.

[174] Jocelyn S. Read, Charalampos Markakis, Masaru Shibata, Koji Uryu, Jolien D.E. Creighton, et al. Measuring the neutron star equation of state with gravitational wave observations. *Phys.Rev.*, D79:124033, 2009.

[175] Krzysztof Belczynski, Tomasz Bulik, Ilya Mandel, B.S. Sathyaprakash, Andrzej A. Zdziarski, et al. Cyg X-3: a Galactic double black hole or black hole-neutron star progenitor. *Astrophys.J.*, 764:96, 2013.

[176] J. Abadie et al. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *Classical & Quantum Gravity*, 27:173001, 2010.

[177] Feryal Ozel, Dimitrios Psaltis, Ramesh Narayan, and Antonio Santos Villarreal. On the Mass Distribution and Birth Masses of Neutron Stars. *Astrophys.J.*, 757:55, 2012.

[178] Philipp Podsiadlowski, N. Langer, A.J.T. Poelarends, S. Rappaport, A. Heger, et al. The effects of binary evolution on the dynamics of core collapse and neutron - star kicks. *Astrophys.J.*, 612:1044–1051, 2004.

[179] E.P.J.van den Heuvel. Double Neutron Stars: Evidence For Two Different Neutron-Star Formation Mechanisms. In T. di Salvo, G. L. Israel, L. Piersant, L. Burderi, G. Matt, A. Tornambe, and M. T. Menna, editors, *The Multicolored Landscape of Compact Objects and Their Explosive Origins*, volume 924 of *American Institute of Physics Conference Series*, pages 598–606, August 2007.

[180] Will M. Farr, Niharika Sravan, Andrew Cantrell, Laura Kreidberg, Charles D. Bailyn, et al. The Mass Distribution of Stellar-Mass Black Holes. *Astrophys.J.*, 741:103, 2011.

[181] Andrea Merloni. Cosmological evolution of supermassive black holes and AGN: a synthesis model for accretion and feedback. *Mem. Soc. Astron. Italiana*, 79:1310, 2008.

[182] P.A. Crowther, R. Barnard, S. Carpano, J.S. Clark, V.S. Dhillon, et al. NGC 300 X-1 is a Wolf-Rayet/Black-Hole binary. *MNRAS*, 403:L41–L45, March 2010.

[183] Krzysztof Belczynski, Tomasz Bulik, Chris L. Fryer, Ashley Ruiter, Jorick S. Vink, et al. On The Maximum Mass of Stellar Black Holes. *Astrophys.J.*, 714:1217–1226, 2010.

[184] M. Kramer and N. Wex. The double pulsar system: A unique laboratory for gravity. *Class.Quant.Grav.*, 26:073001, 2009.

[185] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash. A template bank to search for gravitational waves from inspiralling compact binaries. I: Physical models. *Class. Quant. Grav.*, 23:5477–5504, 2006.

[186] Bruce Allen. A chi**2 time-frequency discriminator for gravitational wave detection. *Phys.Rev.*, D71:062001, 2005.

[187] C. A. K. Robinson, B. S. Sathyaprakash, and Anand S. Sengupta. A geometric algorithm for efficient coincident detection of gravitational waves. *Phys. Rev.*, D78:062002, 2008.

[188] S. Babak, R. Biswas, P.R. Brady, D.A. Brown, K. Cannon, et al. Searching for gravitational waves from binary coalescence. *Phys.Rev.*, D87:024033, 2013.

[189] Yi Pan, Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. A physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single-spin binaries. *Physical Review D*, 69:104017, 2004. Erratum-ibid. 74 (2006) 029905(E).

[190] Alessandra Buonanno, Yanbei Chen, Yi Pan, Hideyuki Tagoshi, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects. ii. search implementation for low-mass binaries. *Physical Review D*, 72:084027, 2005.

[191] Chris Van Den Broeck et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Physical Review D*, 80:024009, 2009.

[192] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D*, 77:062002, 2008.

[193] B.P. Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. *Phys.Rev.*, D79:122001, 2009.

[194] J. Abadie et al. Search for Gravitational Waves from Compact Binary Coalescence in LIGO and Virgo Data from S5 and VSR1. *Phys.Rev.*, D82:102001, 2010.

[195] Tito Dal Canton, Alexander H. Nitz, Andrew P. Lundgren, Alex B. Nielsen, Duncan A. Brown, et al. Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors. *Phys.Rev.*, D90(8):082004, 2014.

[196] I.W. Harry, A.H. Nitz, Duncan A. Brown, A. Lundgren, Evan Ochsner, et al. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. *Phys. Rev. D*, 89(2):024010, January 2014.

[197] Duncan A. Brown, Andrew Lundgren, and R. O'Shaughnessy. Nonspinning searches for spinning binaries in ground-based detector data: Amplitude and mismatch predictions in the constant precession cone approximation. *Phys.Rev.*, D86:064020, 2012.

[198] Duncan A. Brown, Ian Harry, Andrew Lundgren, and Alexander H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Phys.Rev.*, D86:084017, 2012.

[199] Diego Fazi. *Development of a Physical-Template Search for Gravitational Waves from Spinning Compact-Object Binaries with LIGO*. PhD thesis, Bologna University, 2009.

[200] Kipp Cannon, Romain Cariou, Adrian Chapman, Mireia Crispin-Ortuzar, Nickolas Fotopoulos, et al. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. *Astrophys.J.*, 748:136, 2012.

[201] J. Abadie et al. First low-latency LIGO+Virgo search for binary inspirals and their electromagnetic counterparts. *Astron. Astrophys.*, 541:A155, May 2012.

[202] Marc van der Sluys et al. Parameter estimation of spinning binary inspirals using Markov-chain Monte Carlo. *Class. Quant. Grav.*, 25:184011, 2008.

[203] M. V. van der Sluys et al. Gravitational-Wave Astronomy with Inspiral Signals of Spinning Compact-Object Binaries, 2007.

[204] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys.Rev.*, D81:062003, 2010.

[205] Alexander H. Nitz, Andrew Lundgren, Duncan A. Brown, Evan Ochsner, Drew Keppel, et al. Accuracy of gravitational waveform models for observing neutron-star–black-hole binaries in Advanced LIGO. *Phys. Rev. D*, 88(12):124039, December 2013.

[206] Andrea Taracchini, Yi Pan, Alessandra Buonanno, Enrico Barausse, Michael Boyle, et al. Prototype effective-one-body model for nonprecessing spinning inspiral-merger-ringdown waveforms. *Phys.Rev.*, D86:024011, 2012.

[207] Andrea Taracchini, Alessandra Buonanno, Yi Pan, Tanja Hinderer, Michael Boyle, et al. Effective-one-body model for black-hole binaries with generic mass ratios and spins. *Phys.Rev.*, D89:061502, 2014.

[208] L. Santamaria, F. Ohme, P. Ajith, B. Bruegmann, N. Dorband, et al. Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries. *Phys.Rev.*, D82:064016, 2010.

[209] L. Blackburn et al. Equinox Event. Technical Report LIGO-G080559, LIGO Scientific Collaboration and Virgo Collaboration, 2008.

[210] Alessandra Buonanno, Yan-bei Chen, Yi Pan, and Michele Vallisneri. A quasi-physical family of gravity-wave templates for precessing binaries of spinning compact objects: Ii. application to double-spin precessing binaries. *Physical Review D*, 70:104003, 2004. Erratum-ibid. 74 (2006) 029902(E).

[211] T.G.F. Li, W. Del Pozzo, S. Vitale, C. Van Den Broeck, M. Agathos, et al. Towards a generic test of the strong field dynamics of general relativity using compact binary coalescence. *Phys.Rev.*, D85:082003, 2012.

[212] B. D. Metzger and E. Berger. What is the Most Promising Electromagnetic Counterpart of a Neutron Star Binary Merger? *Astrophysical Journal*, 746:48, February 2012.

[213] J. Clark, H. Evans, S. Fairhurst, I. W. Harry, E. Macdonald, D. Macleod, P. J. Sutton, and A. R. Williamson. Prospects for joint gravitational wave and short gamma-ray burst observations. 2014.

[214] E. Berger, W. Fong, and R. Chornock. An r-process Kilonova Associated with the Short-hard GRB 130603B. *Astrophys.J.*, 774:L23, 2013.

[215] I. Leonor, P. J. Sutton, R. Frey, G. Jones, S. Marka, et al. Estimating detection rates for the LIGO-Virgo search for gravitational-wave burst counterparts to gamma-ray bursts using inferred local GRB rates. *Class.Quant.Grav.*, 26:204017, 2009.

[216] D. Wanderman and T. Piran. The luminosity function and the rate of Swift's Gamma Ray Bursts. *Mon. Not. Roy. Astron. Soc.*, 406:1944–1958, 2010.

[217] R. Chapman, N. R. Tanvir, R. S. Priddey, and A. J. Levan. How common are long Gamma-Ray Bursts in the Local Universe? *Mon. Not. R. Astron. Soc.*, L21:382, 2007.

[218] S. Desai, E.O. Kahya, and R.P. Woodard. Reduced time delay for gravitational waves with dark matter emulators. *Phys.Rev.*, D77:124041, 2008.

[219] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310, September 1986.

[220] N. Dalal, D. E. Holz, S. A. Hughes, and B. Jain. Short GRB and binary black hole standard sirens as a probe of dark energy. *Phys.Rev.*, D74:063006, 2006.

[221] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, C. Hirata, A. G. Riess, et al. Observational Probes of Cosmic Acceleration. *Phys.Rept.*, 530:87–255, 2013.

[222] Michal Was, Patrick J. Sutton, Gareth Jones, and Isabel Leonor. Performance of an externally triggered gravitational-wave burst search. *Phys.Rev.*, D86:022003, 2012.

[223] F. Pannarale and F. Ohme. Prospects for joint gravitational-wave and electromagnetic observations of neutron-star–black-hole coalescing binaries. *Astrophys.J.*, 791:L7, 2014.

[224] Patrick J. Sutton, Gareth Jones, Shourov Chatterji, Peter Michael Kalmus, Isabel Leonor, et al. X-Pipeline: An Analysis package for autonomous gravitational-wave burst searches. *New J.Phys.*, 12:053034, 2010.

[225] Ian W. Harry and Stephen Fairhurst. A targeted coherent search for gravitational waves from compact binary coalescences. *Phys. Rev.*, D83:084002, 2011.

[226] Andrew Williamson et al. A search for gravitational waves associated with short gamma-ray bursts. *Phys.Rev.*, D, 2014.

[227] IceCube Collaboration. `http://gcn.gsfc.nasa.gov/gcn3/14520.gcn3`.

[228] B. Abbott et al. Implications for the Origin of GRB 070201 from LIGO Observations. *ApJ*, 681:1419–1428, 2008.

[229] J. Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Phys.Rev.*, D88:122004, 2013.

[230] E. Thrane and M. Coughlin. Searching for gravitational-wave transients with a qualitative signal model: seedless clustering strategies. Preprint arXiv:1308.5292, 2013.

[231] M. C. Miller and E. J. M. Colbert. Intermediate-Mass Black Holes. *International Journal of Modern Physics D*, 13:1–64, January 2004.

[232] N. W. C. Leigh, N. Lützgendorf, A. M. Geller, T. J. Maccarone, C. Heinke, and A. Sesana. On the coexistence of stellar-mass and intermediate-mass black holes in globular clusters. *MNRAS*, 444:29–42, October 2014.

[233] J. M. Fregeau, S. L. Larson, M. C. Miller, R. O'Shaughnessy, and F. A. Rasio. Observing IMBH-IMBH Binary Coalescences via Gravitational Radiation. *Astrophysical Journal Letters*, 646:L135–L138, August 2006.

[234] P. Amaro-Seoane and M. Freitag. Intermediate-Mass Black Holes in Colliding Clusters: Implications for Lower Frequency Gravitational-Wave Astronomy. *Astrophysical Journal Letters*, 653:L53–L56, December 2006.

[235] Jonathan R. Gair, Ilya Mandel, M. Coleman Miller, and Marta Volonteri. Exploring intermediate and massive black-hole binaries with the Einstein Telescope. *Gen.Rel.Grav.*, 43:485–518, 2011.

[236] K. Belczynski, A. Buonanno, M. Cantiello, C. L. Fryer, D. E. Holz, I. Mandel, M. C. Miller, and M. Walczak. The Formation and Gravitational-wave Detection of Massive Stellar Black Hole Binaries. *Astrophysical Journal*, 789:120, July 2014.

[237] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, K. Ackley, C. Adams, T. Adams, and et al. Search for gravitational radiation from intermediate mass black hole binaries in data from the second LIGO-Virgo joint science run. *Physical Review D*, 89(12):122003, June 2014.

[238] J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, R. Adhikari, C. Affeldt, and et al. Search for gravitational waves from intermediate mass binary black holes. *Physical Review D*, 85(10):102004, May 2012.

[239] The LIGO Scientific Collaboration, the Virgo Collaboration, J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, and et al. Search for gravitational wave ringdowns from perturbed intermediate mass black holes in LIGO-Virgo data from 2005-2010. *ArXiv e-prints*, March 2014. arXiv:1403.5306.

[240] J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, C. Adams, T. Adams, and et al. Search for gravitational waves from binary black hole inspiral, merger, and ringdown in LIGO-Virgo data from 2009-2010. *Physical Review D*, 87(2):022002, January 2013.

[241] Gstreamer plugins for the LSC Algorithm Library.

[242] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan. Improving the sensitivity of a search for coalescing binary black holes with non-precessing spins in gravitational wave data. *ArXiv e-prints*, October 2013.

[243] Sergey Klimenko and Chris Pankow. Detection statistic for multiple algorithms, networks and data sets. Technical Report LIGO-T1300869, LIGO Project, 2013.

[244] E. A. Huerta and D. A. Brown. Effect of eccentricity on binary neutron star searches in advanced LIGO. *Phys. Rev. D*, 87(12):127501, 2013.

[245] C. Hopman and T. Alexander. The Effect of Mass Segregation on Gravitational Wave Sources near Massive Black Holes. *Astrophys. J.*, 645:L133–L136, July 2006.

[246] Ryan M. O'Leary, Bence Kocsis, and Abraham Loeb. Gravitational waves from scattering of stellar-mass black holes in galactic nuclei. *Mon. Not. Roy. Astron. Soc.*, 395:2127–2146, 2009.

[247] Bence Kocsis and Janna Levin. Repeated bursts from relativistic scattering of compact objects in galactic nuclei. *Phys. Rev. D*, 85:123005, Jun 2012.

[248] A. C. Fabian, J. E. Pringle, and M. J. Rees. Tidal capture formation of binary systems and X-ray sources in globular clusters. *Mon. Not. Roy. Astron. Soc.*, 172:15P–+, August 1975.

[249] D. Pooley, W. H. G. Lewin, S. F. Anderson, H. Baumgardt, A. V. Filippenko, B. M. Gaensler, L. Homer, P. Hut, V. M. Kaspi, J. Makino, B. Margon, S. McMillan, S. Portegies Zwart, M. van der Klis, and F. Verbunt. Dynamical Formation of Close Binary Systems in Globular Clusters. *Astrophys. J.*, 591:L131–L134, July 2003.

[250] W. H. Lee, E. Ramirez-Ruiz, and G. van de Ven. Short Gamma-ray Bursts from Dynamically Assembled Compact Binaries in Globular Clusters: Pathways, Rates, Hydrodynamics, and Cosmological Setting. *Astrophys. J.*, 720:953–975, September 2010.

[251] Linqing Wen. On the Eccentricity Distribution of Coalescing Black Hole Binaries Driven by the Kozai Mechanism in Globular Clusters. *Astrophys. J.*, 598:419–430, 2003.

[252] M. Coleman Miller and Douglas P. Hamilton. Four-body effects in globular cluster black hole coalescence. *Astrophys. J.*, 576(2):894, 2002.

[253] S. J. Aarseth. Mergers and ejections of black holes in globular clusters. *Mon. Not. Roy. Astron. Soc.*, 422:841–848, May 2012.

[254] Fabio Antonini and Hagai B. Perets. Secular evolution of compact binaries near massive black holes: Gravitational wave sources and other exotica. *Astrophys. J.*, 757(1):27, 2012.

[255] T. A. Thompson. Accelerating Compact Object Mergers in Triple Systems with the Kozai Resonance: A Mechanism for ”Prompt” Type Ia Supernovae, Gamma-Ray Bursts, and Other Exotica. *Astrophys. J.*, 741:82, November 2011.

[256] S. Naoz, B. Kocsis, A. Loeb, and N. Yunes. Resonant Post-Newtonian Eccentricity Excitation in Hierarchical Three-body Systems. *Astrophysical Journal*, 773:187, August 2013.

[257] B. Katz and S. Dong. The rate of WD-WD head-on collisions may be as high as the SNe Ia rate. *ArXiv e-prints*, November 2012.

[258] P. C. Peters. Gravitational radiation and the motion of two point masses. *Phys. Rev.*, 136:B1224–B1232, 1964.

[259] Neil J. Cornish and Tyson B. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, 2014.

[260] Kai Sheng Tai, Sean T. McWilliams, and Frans Pretorius. Detecting gravitational waves from highly eccentric compact binaries. *Physical Review D*, 90:103001, 2014.

[261] M. Millhouse, N. Cornish, and T. Littenberg, 2014. In preparation.

[262] E. A. Huerta, P. Kumar, S. T. McWilliams, R. O'Shaughnessy, and N. Yunes. Accurate and efficient waveforms for compact binaries on eccentric orbits, 2014. arXix:1408.3406 [gr-qc].

[263] M. Coughlin, E. Thrane, and N. Christensen. Detecting compact binary coalescences with seedless clustering. *Phys. Rev. D*, 90:083005, 2014.

[264] I. Mandel, D. A. Brown, J. R. Gair, and M. C. Miller. Rates and Characteristics of Intermediate Mass Ratio Inspirals Detectable by Advanced LIGO. *Astrophysical Journal*, 681:1431–1447, July 2008.

[265] D. A. Brown, J. Brink, H. Fang, J. R. Gair, C. Li, G. Lovelace, I. Mandel, and K. S. Thorne. Prospects for Detection of Gravitational Waves from Intermediate-Mass-Ratio Inspirals. *Physical Review Letters*, 99(20):201102–+, November 2007.

[266] C. L. Rodriguez, I. Mandel, and J. R. Gair. Verifying the no-hair property of massive compact objects with intermediate-mass-ratio inspirals in advanced gravitational-wave detectors. *Physical Review D*, 85(6):062002, March 2012.

[267] L. S. Finn and K. S. Thorne. Gravitational waves from a compact star in a circular, inspiral orbit, in the equatorial plane of a massive, spinning black hole, as observed by LISA. *Physical Review D*, 62(12):124021, December 2000.

[268] E. A. Huerta and J. R. Gair. Importance of including small body spin effects in the modelling of extreme and intermediate mass-ratio inspirals. *Physical Review D*, 84(6):064023, September 2011.

[269] E. A. Huerta, J. R. Gair, and D. A. Brown. Importance of including small body spin effects in the modelling of intermediate mass-ratio inspirals. II. Accurate parameter extraction of strong sources using higher-order spin effects. *Physical Review D*, 85(6):064023, March 2012.

[270] E. A. Huerta and J. R. Gair. Intermediate-mass-ratio inspirals in the Einstein Telescope. I. Signal-to-noise ratio calculations. *Physical Review D*, 83(4):044020, February 2011.

[271] J. Veitch, M. Pürrer, and I. Mandel. Measuring intermediate mass black hole binaries with advanced gravitational wave detectors. *ArXiv e-prints*, March 2015.

[272] I. Mandel and J. R. Gair. Can we detect intermediate mass ratio inspirals? *Classical and Quantum Gravity*, 26(9):094036, May 2009.

[273] Alessandra Buonanno et al. Toward faithful templates for non-spinning binary black holes using the effective-one-body approach. *Phys. Rev. D*, 76:104049, 2007.

[274] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, H. P. Pfeiffer, and M. A. Scheel. Inspiral-merger-ringdown multipolar waveforms of nonspinning black-hole binaries using the effective-one-body formalism. *Physical Review D*, 84(12):124052, December 2011.

[275] T. Callister. Improving models of eccentric compact binary inspirals, MPhil Thesis (Cambridge University), LIGO-DCC-P1400238, 2014.

[276] R. J. E. Smith, I. Mandel, and A. Vecchio. Studies of waveform requirements for intermediate mass-ratio coalescence searches with advanced gravitational-wave detectors. *Physical Review D*, 88(4):044010, August 2013.

[277] Y. Aso, Z. Márka, C. Finley, J. Dwyer, K. Kotake, and S. Márka. Search method for coincident events from LIGO and IceCube detectors. *Classical and Quantum Gravity*, 25(11):114039–+, June 2008.

[278] T. Pradier. Coincidences between Gravitational Wave Interferometers and High Energy Neutrino Telescopes. *ArXiv e-prints*, 807, July 2008. arXiv:0807.2562.

[279] V. van Elewyck, S. Ando, Y. Aso, B. Baret, M. Barsuglia, I. Bartos, E. Chassande-Mottin, I. di Palma, J. Dwyer, C. Finley, K. Kei, A. Kouchner, S. Marka, Z. Marka, J. Rollins, C. D. Ott, T. Pradier, and A. Searle. Joint Searches Between Gravitational-Wave Interferometers and High-Energy Neutrino Telescopes:. Science Reach and Analysis Strategies. *Int. J. Mod. Phys. D*, 18:1655–1659, 2009.

[280] B. Baret, I. Bartos, B. Bouhou, A. Corsi, I. D. Palma, C. Donzaud, V. V. Elewyck, C. Finley, G. Jones, A. Kouchner, S. Márka, Z. Márka, L. Moscoso, E. Chassande-Mottin, M. A. Papa, T. Pradier, P. Raffai, J. Rollins, and P. Sutton. Bounding the time delay between high-energy neutrinos and gravitational-wave transients from gamma-ray bursts. *Astropart. Phys.*, 35(1):1 – 7, 2011.

[281] Imre Bartos, Chad Finley, Alessandra Corsi, and Szabolcs Márka. Observational constraints on multimessenger sources of gravitational waves and high-energy neutrinos. *Phys. Rev. Lett.*, 107:251101, Dec 2011.

[282] Shin'ichiro Ando et al. Multimessenger astronomy with gravitational waves and high-energy neutrinos. *Rev. Mod. Phys.*, 85(4):1401–1420, 2013.

[283] E. Waxman and J. Bahcall. High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physical Review Letters*, 78:2292–2295, March 1997.

[284] E. Waxman and J. N. Bahcall. Neutrino Afterglow from Gamma-Ray Bursts: 10^{18} eV. *Astrophys. Journal*, 541:707, October 2000.

[285] P. Mészáros and E. Waxman. TeV Neutrinos from Successful and Choked Gamma-Ray Bursts. *Phys. Rev. Lett.*, 87(17):171102–+, October 2001.

[286] P. Mészáros, S. Kobayashi, S. Razzaque, and B. Zhang. High Energy Photons, Neutrinos and Gravitational Waves from Gamma-Ray Bursts. In R. Ouyed, editor, *Beaming and Jets in Gamma Ray Bursts*, page 30, 2002.

[287] S. Razzaque, P. Mészáros, and E. Waxman. High Energy Neutrinos from Gamma-Ray Bursts with Precursor Supernovae. *Phys. Rev. Lett.*, 90(24):241103–+, June 2003.

[288] S. Razzaque, P. Mészáros, and E. Waxman. Neutrino tomography of gamma ray bursts and massive stellar collapses. *Phys. Rev. D*, 68(8):083001, October 2003.

[289] S. Ando and J. F. Beacom. Revealing the Supernova Gamma-Ray Burst Connection with TeV Neutrinos. *Phys. Rev. Lett.*, 95(6):061103, August 2005.

[290] S. Razzaque, P. Mészáros, and E. Waxman. High Energy Neutrinos from a Slow Jet Model of Core Collapse Supernovae. *Mod. Phys. Lett. A*, 20:2351–2367, 2005.

[291] M. G. Aartsen, R. Abbasi, Y. Abdou, M. Ackermann, J. Adams, J. A. Aguilar, M. Ahlers, D. Altmann, J. Auffenberg, X. Bai, and et al. First Observation of PeV-Energy Neutrinos with IceCube. *Physical Review Letters*, 111(2):021103, July 2013.

[292] S. Adrián-Martínez, I. A. Samarai, A. Albert, M. André, M. Anghinolfi, G. Anton, S. Anvar, M. Ardid, T. Astraatmadja, J.-J. Aubert, and et al. A first search for coincident gravitational waves and high energy neutrinos using LIGO, Virgo and ANTARES data from 2007. *JCAP*, 6:8, June 2013.

[293] B Baret et al. Multimessenger sources of gravitational waves and high-energy neutrinos: Science reach and analysis method. *Journal of Physics: Conference Series*, 363(1):012022, 2012.

[294] H. B. J. Koers and R. A. M. J. Wijers. Enhanced high-energy neutrino emission from choked gamma-ray bursts due to meson and muon acceleration. *ArXiv e-prints*, 711, November 2007. arXiv:0711.4791.

[295] S. Horiuchi and S. Ando. High-energy neutrinos from reverse shocks in choked and successful relativistic jets. *Physical Review D*, 77(6):063007, March 2008.

[296] J Ahrens et al. Sensitivity of the IceCube detector to astrophysical sources of high energy muon neutrinos. *Astroparticle Physics*, 20(5):507 – 532, 2004.

[297] I. Bartos, B. Dasgupta, and S. Márka. Probing the structure of jet-driven core-collapse supernova and long gamma-ray burst progenitors with high-energy neutrinos. *Physical Review D*, 86(8):083007, October 2012.

[298] B. Baret, I. Bartos, B. Bouhou, E. Chassande-Mottin, A. Corsi, I. Di Palma, C. Donzaud, M. Drago, C. Finley, G. Jones, S. Klimenko, A. Kouchner, S. Márka, Z. Márka, L. Moscoso, M. A. Papa, T. Pradier, G. Prodi, P. Raffai, V. Re, J. Rollins, F. Salemi, P. Sutton, M. Tse, V. Van Elewyck, and G. Vedovato. Multimessenger science reach and analysis method for common sources of gravitational waves and high-energy neutrinos. *Physical Review D*, 85(10):103004, May 2012.

[299] T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J.Phys.A*, A9:1387–1398, 1976.

[300] A. Vilenkin and E. Shellard. *Cosmic strings and other Topological Defects*. Cambridge University Press, 1994.

[301] I. B. Zeldovich. Cosmological fluctuations produced near a singularity. *Mon. Not. R. Astron. Soc.*, 192:663–667, sep 1980.

[302] Mark Wyman, Levon Pogosian, and Ira Wasserman. Bounds on cosmic strings from WMAP and SDSS. *Phys.Rev.*, D72:023513, 2005.

[303] Neil Bevis, Mark Hindmarsh, Martin Kunz, and Jon Urrestilla. Fitting CMB data with cosmic strings and inflation. *Phys.Rev.Lett.*, 100:021301, 2008.

[304] Richard Battye, Bjorn Garbrecht, and Adam Moss. Tight constraints on F- and D-term hybrid inflation scenarios. *Phys.Rev.*, D81:123512, 2010.

[305] Richard Battye and Adam Moss. Updated constraints on the cosmic string tension. *Phys.Rev.*, D82:023521, 2010.

[306] P.A.R. Ade et al. Planck 2013 results. XXV. Searches for cosmic strings and other topological defects, 2013.

[307] Andrei D. Linde. Hybrid inflation. *Phys.Rev.*, D49:748–754, 1994.

[308] Edmund J. Copeland, Andrew R. Liddle, David H. Lyth, Ewan D. Stewart, and David Wands. False vacuum inflation with Einstein gravity. *Phys.Rev.*, D49:6410–6433, 1994.

[309] G.R. Dvali, Q. Shafi, and Robert K. Schaefer. Large scale structure and supersymmetric inflation without fine tuning. *Phys.Rev.Lett.*, 73:1886–1889, 1994.

[310] Saswat Sarangi and S.H. Henry Tye. Cosmic string production towards the end of brane inflation. *Phys.Lett.*, B536:185–192, 2002.

[311] Rachel Jeannerot, Jonathan Rocher, and Mairi Sakellariadou. How generic is cosmic string formation in SUSY GUTs. *Phys.Rev.*, D68:103514, 2003.

[312] Edward Witten. Cosmic superstrings. *Physics Letters B*, 153(4???5):243 – 246, 1985.

[313] Edmund J. Copeland, Levon Pogosian, and Tanmay Vachaspati. Seeking String Theory in the Cosmos. *Class.Quant.Grav.*, 28:204009, 2011.

[314] N. Kaiser and A. Stebbins. Microwave anisotropy due to cosmic strings. *Nature*, 310:391–393, 1984.

[315] Amy S. Lo and Edward L. Wright. Signatures of cosmic strings in the cosmic microwave background, 2005.

[316] E. Jeong and G. F. Smoot. The Validity of the Cosmic String Pattern Search with the Cosmic Microwave Background. *The Astrophysical Journal Letters*, 661(1):L1, 2007.

[317] Jon Urrestilla, Neil Bevis, Mark Hindmarsh, and Martin Kunz. Cosmic string parameter constraints and model analysis using small scale Cosmic Microwave Background data. *JCAP*, 1112:021, 2011.

[318] J. Dunkley, R. Hlozek, J. Sievers, V. Acquaviva, P.A.R. Ade, et al. The Atacama Cosmology Telescope: Cosmological Parameters from the 2008 Power Spectra. *Astrophys.J.*, 739:52, 2011.

[319] Thibault Damour and Alexander Vilenkin. Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows. *Phys.Rev.*, D71:063510, 2005.

[320] S. Olmez, V. Mandic, and X. Siemens. Gravitational-Wave Stochastic Background from Kinks and Cusps on Cosmic Strings. *Phys.Rev.*, D81:104028, 2010.

[321] Mark G. Jackson, Nicholas T. Jones, and Joseph Polchinski. Collisions of cosmic F and D-strings. *JHEP*, 0510:013, 2005.

[322] T. Kibble. *The formation and evolution of cosmic strings*, chapter Cosmic strings – An overview. Cambridge Univ. Press, gibbons, g. and hawking, s edition, 1990.

[323] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys.Rev.Lett.*, 85:3761–3764, 2000.

[324] David P. Bennett and Francois R. Bouchet. Evidence for a Scaling Solution in Cosmic String Evolution. *Phys.Rev.Lett.*, 60:257, 1988.

[325] Christophe Ringeval, Mairi Sakellariadou, and Francois Bouchet. Cosmological evolution of cosmic string loops. *JCAP*, 0702:023, 2007.

[326] Jose J. Blanco-Pillado, Ken D. Olum, and Benjamin Shlaer. The number of cosmic string loops. *Phys.Rev.*, D89:023512, 2014.

[327] V. Berezinsky, B. Hnatyk, and A. Vilenkin. Superconducting cosmic strings as gamma-ray burst engines, 2000.

[328] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cusps and kinks on cosmic strings. *Phys.Rev.*, D64:064008, 2001.

[329] J. Aasi et al. The characterization of Virgo data and its impact on gravitational-wave searches. *Class.Quant.Grav.*, 29:155002, 2012.

[330] T. Accadia et al. Virgo Document VIR-0128A-12, 2012.

[331] P. Binetruy, A. Bohe, T. Hertog, and D. A. Steer. Gravitational wave signatures from kink proliferation on cosmic (super-) strings. *Phys.Rev.*, D82:126007, 2010.

[332] Xavier Siemens, Jolien Creighton, Irit Maor, Saikat Ray Majumder, Kipp Cannon, et al. Gravitational wave bursts from cosmic (super)strings: Quantitative analysis and constraints. *Phys.Rev.*, D73:105001, 2006.

[333] Kipp C. Cannon. A Bayesian coincidence test for noise rejection in a gravitational-wave burst search. *Class.Quant.Grav.*, 25:105024, 2008.

[334] Tristan L. Smith, Elena Pierpaoli, and Marc Kamionkowski. A new cosmic microwave background constraint to primordial gravitational waves. *Phys.Rev.Lett.*, 97:021301, 2006.

[335] B.P. Abbott et al. An Upper Limit on the Stochastic Gravitational-Wave Background of Cosmological Origin. *Nature*, 460:990, 2009.

[336] Frederick A. Jenet, G.B. Hobbs, W. van Straten, R.N. Manchester, M. Bailes, et al. Upper bounds on the low-frequency stochastic gravitational wave background from pulsar timing observations: Current limits and future prospects. *Astrophys.J.*, 653:1571–1576, 2006.

[337] E Thrane, S Kandhasamy, C D Ott, et al. Long gravitational-wave transients and associated detection strategies for a network of terrestrial interferometers. *Phys. Rev. D*, 83:083004, 2011.

[338] A. Corsi and P. Mészáros. Gamma-ray burst afterglow plateaus and gravitational waves: multi-messenger signature of a millisecond magnetar? *Astrophys. J.*, 702:1171, 2009.

[339] A. L. Piro and C. D. Ott. Supernova fallback onto magnetars and propeller-powered supernovae. *Astrophys. J.*, 736:108, 2011.

[340] A. L. Piro and E. Thrane. Gravitational Waves from Fallback Accretion onto Neutron Stars. *Astrophys. J.*, 761:63, 2012.

[341] E Thrane and M Coughlin. Searching for gravitational-wave transients with a qualitative signal model: seedless clustering strategies. *Phys. Rev. D*, 88:083010, 2013.

[342] E. Thrane and M. Coughlin. Seedless clustering in all-sky searches for gravitational-wave transients. *Phys. Rev. D*, 89:063012, 2014.

[343] K Kiuchi, M Shibata, P J Montero, and J A Font. Gravitational waves from the Papaloizou-Pringle instability in black hole-torus systems. *Phys. Rev. Lett.*, 106:251102, 2011.

[344] M. H. P. M. van Putten. Proposed Source of Gravitational Radiation from a Torus around a Black Hole. *Phys. Rev. Lett.*, 87:091101, 2001.

[345] M. H. P. M. van Putten. Gravitational Waveforms of Kerr Black Holes Interacting with High-Density Matter. *Astrophys. J. Lett.*, 684:91, 2008.

[346] A. L. Piro and E. Pfahl. Fragmentation of Collapsar Disks and the Production of Gravitational Waves. *Astrophysical Journal*, 658:1173, April 2007.

[347] J Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Phys. Rev. D*, 88:122004, 2013.

[348] C. D. Ott. The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26:063001, 2009.

[349] R. M. O'Leary, B. Kocsis, and A. Loeb. Gravitational waves from scattering of stellar-mass black holes in galactic nuclei. *Mon. Not. R. Ast. Soc.*, 395:2127, 2009.

[350] S. Sutton, I. S. Heng, G. Prodi, and E. Chassande-Motin. Search Plan: All-Sky Burst Search. *LIGO-T1400192*, 2014.

[351] Eric Thrane, Vuk Mandic, and Nelson Christensen. Detecting very long-lived gravitational-wave transients lasting hours to weeks. 2015.

[352] T. Prestegard, E. Thrane, et al. Identification of noise artifacts in searches for long-duration gravitational-wave transients. *Class. Quantum Grav.*, 29:095018, 2012.

[353] `http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/stochastic/trunk/stamp2/`.

[354] T. Prestegard for the STAMP all-sky working group. The STAMP all-sky search for long-transient gravitational waves, 2014. `https://dcc.ligo.org/LIGO-G1400176`.

[355] The LSC and Virgo. All-sky search for long-duration gravitational-wave transients with ligo, 2014. `https://dcc.ligo.org/LIGO-P1400138`.

[356] B. P. Abbott, R. Abbott, F. Acernese, et al. Search for gravitational-wave bursts associated with gamma-ray bursts using data from ligo science run 5 and virgo science run 1. *The Astrophysical Journal*, 715(2):1438, 2010.

[357] M. S. Briggs, V. Connaughton, K. C. Hurley, P. A. Jenke, A. von Kienlin, A. Rau, X.-L. Zhang, The LIGO Scientific Collaboration, Virgo Collaboration: J. Abadie, B. P. Abbott, and et al. Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3. *Astrophysical Journal*, 760:12, November 2012.

[358] S. Adrián-Martínez et al. A First Search for coincident Gravitational Waves and High Energy Neutrinos using LIGO, Virgo and ANTARES data from 2007. *Journal of Cosmology and Astroparticle Physics*, 2013(06):008, 2013.

[359] J. Abadie, B. P. Abbott, T. D. Abbott, R. Abbott, M. Abernathy, C. Adams, R. Adhikari, C. Affeldt, P. Ajith, B. Allen, and et al. Implications for the Origin of GRB 051103 from LIGO Observations. *Astrophysical Journal*, 755:2, August 2012.

[360] J. Aasi et al. Search for gravitational waves associated with gamma-ray bursts detected by the Inter-Planetary Network. *Physical Review Letters*, 113:011102, 2014.

[361] J. Aasi et al. Methods and results of a search for gravitational waves associated with gamma-ray bursts using the GEO600, LIGO, and Virgo detectors. *Physical Review D*, 89:122004, 2014.

[362] Kipp C. Cannon. Efficient algorithm for computing the time-resolved full-sky cross power in an interferometer with omnidirectional elements. *Phys.Rev.*, D75:123003, 2007.

[363] Mark Edwards. On the search for intermediate duration gravitational waves using the spherical harmonic basis. 2013. https://dcc.ligo.org/LIGO-P1300227.

[364] P. Sutton et al. X-Pipeline All-Sky Burst Search: Optimisation. 2015. `https://dcc.ligo.org/T1500033`.

[365] N. Cornish and T. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Submitted to Class. Quantum Grav., LIGO-P1400185*, 2014.

[366] LSC/Virgo Burst Analysis Group. LALInference_Burst. *https://wiki.ligo.org/Bursts/LIB/WebHome*, 2014.

[367] Michael Coughlin, Patrick Meyers, Shivaraj Kandhasamy, Eric Thrane, and Nelson Christensen. Chasing 5-sigma: Prospects for searches for long-duration gravitational-waves without time slides. 2015.

[368] E. Thrane for the The Stochastic Group. Calibration requirements for stochastic searches in the advanced detector era, 2014.
`https://dcc.ligo.org/LIGO-T1200562`.

[369] E. P. Mazets, S. V. Golentskii, V. N. Ilinskii, R. L. Aptekar, and I. A. Guryan. Observations of a flaring X-ray pulsar in Dorado. *Nature*, 282:587–589, December 1979.

[370] Y. T. Tanaka, T. Terasawa, N. Kawai, A. Yoshida, I. Yoshikawa, Y. Saito, T. Takashima, and T. Mukai. Comparative Study of the Initial Spikes of Soft Gamma-Ray Repeater Giant Flares in 1998 and 2004 Observed with Geotail: Do Magnetospheric Instabilities Trigger Large-Scale Fracturing of a Magnetar's Crust? *Astrophysical Journal Letters*, 665:L55–L58, August 2007.

[371] T. Terasawa, Y. T. Tanaka, Y. Takei, N. Kawai, A. Yoshida, K. Nomoto, I. Yoshikawa, Y. Saito, Y. Kasaba, T. Takashima, T. Mukai, H. Noda, T. Murakami, K. Watanabe, Y. Muraki, T. Yokoyama, and M. Hoshino. Repeated injections of energy in the first 600ms of the giant flare of SGR1806 - 20. *Nature*, 434:1110–1111, April 2005.

[372] O. Benhar, V. Ferrari, and L. Gualtieri. Gravitational wave asteroseismology reexamined. *Physical Review D*, 70(12):124015–+, December 2004.

[373] G. L. Israel, T. Belloni, L. Stella, Y. Rephaeli, D. E. Gruber, P. Casella, S. Dall'Osso, N. Rea, M. Persic, and R. E. Rothschild. The discovery of rapid x-ray oscillations in the tail of the sgr 1806-20 hyperflare. *The Astrophysical Journal Letters*, 628(1):L53, 2005.

[374] A. L. Watts and T. E. Strohmayer. Detection with RHESSI of High-Frequency X-Ray Oscillations in the Tail of the 2004 Hyperflare from SGR 1806–20. *Astrophysical Journal Letters*, 637:L117–L120, February 2006.

[375] N. Andersson, J. Baker, K. Belczynski, S. Bernuzzi, E. Berti, L. Cadonati, P. Cerdá-Durán, J. Clark, M. Favata, L. S. Finn, C. Fryer, B. Giacomazzo, J. A. González, M. Hendry, I. S. Heng, S. Hild, N. Johnson-McDaniel, P. Kalmus, S. Klimenko, S. Kobayashi, K. Kokkotas, P. Laguna, L. Lehner, J. Levin, S. Liebling, A. MacFadyen, I. Mandel, S. Marka, Z. Marka, D. Neilsen, P. O'Brien, R. Perna, J. Read, C. Reisswig, C. Rodriguez, M. Ruffert, E. Schnetter, A. Searle, P. Shawhan, D. Shoemaker, A. Soderberg, U. Sperhake, P. Sutton, N. Tanvir, M. Was, and S. Whitcomb. The transient gravitational-wave sky. *Classical and Quantum Gravity*, 30(19):193002, October 2013.

[376] B. Abbott et al. Search for gravitational-wave bursts from soft gamma repeaters. *Physical Review Letters*, 101(21):211102, 2008.

[377] B. P. Abbott et al. Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm. *Astrophysical Journal Letters*, 701:L68–L74, August 2009.

[378] J. Abadie et al. Search for Gravitational Wave Bursts from Six Magnetars. *The Astrophysical Journal Letters*, 734(2):L35, 2011.

[379] J. Abadie et al. Search for gravitational waves associated with the august 2006 timing glitch of the vela pulsar. *Physical Review D*, 83(4):042001, Feb 2011.

[380] B. Abbott and for the LIGO Scientific Collaboration. Implications for the Origin of GRB 070201 from LIGO Observations. *Astrophysical Journal*, 681:1419–1428, 2008.

[381] B. Abbott et al. Search for gravitational wave radiation associated with the pulsating tail of the SGR 1806–20 hyperflare of 27 December 2004 using LIGO. *Physical Review D*, 76:062003, 2007.

[382] M. Shibata. Constraining Nuclear Equations of State Using Gravitational Waves from Hypermassive Neutron Stars. *Physical Review Letters*, 94(20):201101, May 2005.

[383] L. Baiotti, B. Giacomazzo, and L. Rezzolla. Accurate evolutions of inspiralling neutron-star binaries: Prompt and delayed collapse to a black hole. *Physical Review D*, 78:084033, Oct 2008.

[384] K. Hotokezaka, K. Kyutoku, H. Okawa, M. Shibata, and K. Kiuchi. Binary neutron star mergers: Dependence on the nuclear equation of state. *Physical Review D*, 83(12):124008, June 2011.

[385] A. Bauswein, H.-T. Janka, K. Hebeler, and A. Schwenk. Equation-of-state dependence of the gravitational-wave signal from the ring-down phase of neutron-star mergers. *Physical Review D*, 86(6):063001, September 2012.

[386] N. Stergioulas, A. Bauswein, K. Zagkouris, and H.-T. Janka. Gravitational waves and non-axisymmetric oscillation modes in mergers of compact object binaries. *MNRAS*, 418:427–436, November 2011.

[387] K. Hotokezaka, K. Kiuchi, K. Kyutoku, T. Muranushi, Y. Sekiguchi, M. Shibata, and K. Taniguchi. Remnant massive neutron stars of binary neutron star mergers: Evolution process and gravitational waveform. *Physical Review D*, 88(4):044026, August 2013.

[388] A. Bauswein, N. Stergioulas, and H.-T. Janka. Revealing the high-density equation of state through binary neutron star mergers. *Physical Review D*, 90(2):023002, July 2014.

[389] J. Clark, A. Bauswein, L. Cadonati, H.-T. Janka, C. Pankow, and N. Stergioulas. Prospects for high frequency burst searches following binary neutron star coalescence with advanced gravitational wave detectors. *Physical Review D*, 90(6):062004, September 2014.

[390] S. Dall'Osso, B. Giacomazzo, R. Perna, and L. Stella. Gravitational waves from massive magnetars formed in binary neutron star mergers. *ArXiv e-prints*, July 2014.

[391] C. Ott. Analytic CCSN Waveforms: 1. Barmode. *LIGO-T1000553*, 2013.

[392] K. Kotake. Multiple physical elements to determine the gravitational-wave signatures of core-collapse supernovae. *Comptes Rendus Physique*, 14:318, 2013.

[393] C. D. Ott. TOPICAL REVIEW: The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26:063001, 2009.

[394] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Class. Quant. Grav.*, 26(20):204015, October 2009.

[395] J. Logue, C. D. Ott, I. S. Heng, P. Kalmus, and J. Scargill. Inferring Core-Collapse Supernova Physics with Gravitational Waves. *Physical Review D*, 86(4):044023, 2012.

[396] E. Abdikamalov, S. Gossan, A. M. DeMaio, and C. D. Ott. Measuring the angular momentum distribution in core-collapse supernova progenitors with gravitational waves. *Physical Review D*, 90(4):044001, 2014.

[397] W. J. Engels, R. Frey, and C. D. Ott. Multivariate regression analysis of gravitational waves from rotating core collapse. *Physical Review D*, 90(12):124026, December 2014.

[398] J. W. Murphy, C. D. Ott, and A. Burrows. A Model for Gravitational Wave Emission from Neutrino-Driven Core-Collapse Supernovae. *Astrophysical Journal*, 707:1173, 2009.

[399] B. Müller, H.-T. Janka, and A. Marek. A New Multi-dimensional General Relativistic Neutrino Hydrodynamics Code of Core-collapse Supernovae. III. Gravitational Wave Signals from Supernova Explosion Models. *Astrophysical Journal*, 766:43, 2013.

[400] K. N. Yakunin, P. Marronetti, A. Mezzacappa, S. W. Bruenn, C.-T. Lee, M. A. Chertkow, W. R. Hix, J. M. Blondin, E. J. Lentz, O. E. Bronson Messer, and S. Yoshida. Gravitational waves from core collapse supernovae. *Class. Quantum Grav.*, 27:194005, October 2010.

[401] T. Fischer, S. C. Whitehouse, A. Mezzacappa, F.-K. Thielemann, and M. Liebendörfer. The neutrino signal from protoneutron star accretion and black hole formation. *Astronomy & Astrophysics*, 499:1–15, 2009.

[402] C. D. Ott, C. Reisswig, E. Schnetter, E. O'Connor, U. Sperhake, F. Löffler, P. Diener, E. Abdikamalov, I. Hawke, and A. Burrows. Dynamics and Gravitational Wave Signature of Collapsar Formation. *Phys. Rev. Lett.*, 106:161103, 2011.

[403] S. E. Gossan et al. Observing Gravitational Waves from Core-Collapse Supernovae. *In preparation, LIGO-P1400233*, 2015.

[404] J. Aasi et al. Targeted search for Gravitational-Waves from Core-Collapse Supernovae in initial LIGO and Virgo data. *To be submitted to Phys. Rev. D., LIGO-P1400208*, 2014.

[405] G. Pagliaroli, F. Vissani, E. Coccia, and W. Fulgione. Neutrinos from Supernovae as a Trigger for Gravitational Wave Search. *Phys. Rev. Lett.*, 103(3):031102, July 2009.

[406] J. F. Beacom and P. Vogel. Can a supernova be located by its neutrinos? *Physical Review D*, 60(3):033007, 1999.

[407] Pietro Antonioli, Richard Tresch Fienberg, Fabrice Fleurot, Yoshiyuki Fukuda, Walter Fulgione, et al. SNEWS: The Supernova Early Warning System. *New J.Phys.*, 6:114, 2004.

[408] S. Klimenko. Coherent WaveBurst 2G Update. *LIGO-G1400993*, 2014.

[409] C. D. Ott, E. Abdikamalov, P. Mösta, R. Haas, S. Drasco, E. P. O'Connor, C. Reisswig, C. A. Meakin, and E. Schnetter. General-relativistic Simulations of Three-dimensional Core-collapse Supernovae. *Astrophysical Journal*, 768:115, May 2013.

[410] C.L. Fryer, D.E. Holz, and S.A. Hughes. Gravitational Wave Emission from Core Collapse of Massive Stars. *Astrophysical Journal*, 565:430, 2002.

[411] IceCube. http://icecube.wisc.edu.

[412] T. Gaisser and F. Halzen. Icecube. *Ann. Rev. Nuc. Part. Sc.*, 64(1):101, 2014.

[413] Large Volume Detector, `http://www.bo.infn.it/lvd`.

[414] C. Vigorito [LVD Collaboration]. Galactic supernovae monitoring at LVD. *Nuc. Phys. B Proc. Suppl.*, 221:410, December 2011.

[415] Borexino, `http://borex.lngs.infn.it`.

[416] G. Alimonti et al. [BOREXINO Collaboration]. The Borexino detector at the Laboratori Nazionali del Gran Sasso. *Nuc. Inst. Meth. Phys. Res. A*, 600:568, March 2009.

[417] F Acernese et al. Search for gravitational waves associated with grb 050915a using the virgo detector. *Classical & Quantum Gravity*, 25(22):225001, 2008.

[418] F. Acernese et al. Gravitational wave burst search in the Virgo C7 data. *Classical & Quantum Gravity*, 26:085009, 2008.

[419] Thomas S. Adams. Detector Characterisation and Searches for Gravitationa l Waves Using GEO600. 2014. https://dcc.ligo.org/LIGO-P1400088.

[420] M. Ikeda et al. Search for Supernova Neutrino Bursts at Super-Kamiokande. *Astrophysical Journal*, 669:519, November 2007.

[421] Super-Kamiokande, `http://neutrino.phys.washington.edu/superk`.

[422] S. Mukherjee and L. Salazar. New Method for Enhanced Efficiency in Detection of Gravitational Waves from Supernovae Using Coherent Network of Detectors. *LIGO-P1400116*, 2014.

[423] C. Plapous, C. Marro, and P. Scalart. Improved signal-to-noise ratio estimation for speech enhancement. *Trans. Audio, Speech and Lang. Proc.*, 14(6):2098–2108, November 2006.

[424] Thomas S. Adams, Duncan Meacher, James Clark, Patrick J. Sutton, Gareth Jones, et al. Gravitational-Wave Detection using Multivariate Analysis. *Phys.Rev.*, D88:062006, 2013.

[425] E. O'Connor and C. D. Ott. Black Hole Formation in Failing Core-Collapse Supernovae. *Astrophysical Journal*, 730:70, April 2011.

[426] D. Thornton, B. Stappers, M. Bailes, B. Barsdell, S. Bates, N. D. R. Bhat, M. Burgay, S. Burke-Spolaor, D. J. Champion, P. Coster, N. D'Amico, A. Jameson, S. Johnston, M. Keith, M. Kramer, L. Levin, S. Milia, C. Ng, A. Possenti, and W. van Straten. A Population of Fast Radio Bursts at Cosmological Distances. *Science*, 341:53–56, July 2013.

[427] E.F. Keane. Fast Radio Bursts - I: Initial Cogitation, 2014.

[428] V. Berezinsky, B. Hnatyk, and A. Vilenkin. Superconducting cosmic strings as gamma-ray burst engines. *Baltic Astronomy*, 13:289, 2004.

[429] Yun-Wei Yu, Kwong-Sang Cheng, Gary Shiu, and Henry Tye. Implications of fast radio bursts for superconducting cosmic strings, 2014.

[430] Tanmay Vachaspati. Cosmic sparks from superconducting strings. *Physical Review Letters*, 101(14):141301, Sep 2008.

[431] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys. Rev. Lett.*, 85(18):3761–3764, Oct 2000.

[432] V. M. Lipunov and I. E. Panchenko. Pulsars revived by gravitational waves. *Astronomy & Astrophysics*, 312:937, August 1996.

[433] M.S. Pshirkov and K.A. Postnov. Radio precursors to neutron star binary mergings. *Astrophysics and Space Science*, 330(1):13–18, 2010.

[434] B. M. S. Hansen and M. Lyutikov. Radio and X-ray signatures of merging neutron stars. *MNRAS*, 322:695, April 2001.

[435] J. Moortgat and J. Kuijpers. Gravitational waves in magnetized relativistic plasmas. *Phys. Rev. D*, 70:023001, Jul 2004.

[436] J. Moortgat and J. Kuijpers. Indirect visibility of gravitational waves in magnetohydrodynamic plasmas. *Proceedings of the XXII Texas Symposium on Relativistic Astrophysics, Stanford University, Palo Alto, USA, December 13-17 2004*, pages 1225–1230, 2005.

[437] Lucia M. Franco, Bennett Link, and Richard I. Epstein. Quaking neutron stars. *The Astrophysical Journal*, 543(2):987, 2000.

[438] K.D. Kokkotas and B.G. Schmidt. Quasi-normal modes of stars and black holes. *Living Rev. Rel.*, 2:2, 1999.

[439] J. Middleditch, F. E. Marshall, Q. D. Wang, E. V. Gotthelf, and W. Zhang. Predicting the starquakes in psr j0537-ĂŞ6910. *The Astrophysical Journal*, 652(2):1531, 2006.

[440] Abbott, B. and others. Search for Gravitational Wave Bursts from Soft Gamma Repeaters. *Physical Review Letters*, 101:211102, 2008.

[441] E.F. Keane, M. Kramer, A.G. Lyne, B.W. Stappers, and M.A. McLaughlin. Rotating Radio Transients: new discoveries, timing solutions and musings. *Mon. Not. R. Astron. Soc.*, 415:3065, 2011.

[442] A.G. Lyne et al. Unusual glitch activity in the RRAT J1819?1458: an exhausted magnetar? *Mon. Not. R. Astron. Soc.*, 400:1439, 2009.

[443] P. Weltevrede, S. Johnston, and C.M. Espinoza. The glitch-induced identity changes of PSR J1119-6127. *AIP Conference Proceedings*, 1357:109, 2011.

[444] J. Boyles, R. S. Lynch, S. M. Ransom, I. H. Stairs, D. R. Lorimer, M. A. McLaughlin, J. W. T. Hessels, V. M. Kaspi, V. I. Kondratiev, A. Archibald, A. Berndsen, R. F. Cardoso, A. Cherry, C. R. Epstein, C. Karako-Argaman, C. A. McPhee, T. Pennucci, M. S. E. Roberts, K. Stovall, and J. van Leeuwen. The green bank telescope 350 mhz drift-scan survey. i. survey observations and the discovery of 13 pulsars. *The Astrophysical Journal*, 763(2):80, 2013.

[445] V. Predoi, J. Clark, T. Creighton, E. Daw, S. Fairhurst, I. S. Heng, J. Kanner, T. Regimbau, P. Shawhan, X. Siemens, P. Sutton, A. Vecchio, D. White, and G. Woan. Prospects for joint radio telescope and gravitational-wave searches for astrophysical transients. *Classical and Quantum Gravity*, 27(8):084018, April 2010.

[446] J. Aasi et al. Search for gravitational waves associated with gamma-ray bursts detected by the Inter-Planetary Network. 2014.

[447] J. Aasi et al. Methods and results of a search for gravitational waves associated with gamma-ray bursts using the GEO600, LIGO, and Virgo detectors. *Phys.Rev.*, D89(12):122004, 2014.

[448] J. Abadie, B. P. Abbott, R. Abbott, T. Accadia, F. Acernese, R. Adhikari, P. Ajith, B. Allen, G. Allen, E. Amador Ceron, and et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys. Rev. D*, 81(10):102001, May 2010.

[449] E. Petroff, E. F. Keane, E. D. Barr, J. E. Reynolds, J. Sarkissian, P. G. Edwards, J. Stevens, C. Brem, A. Jameson, S. Burke-Spolaor, S. Johnston, N. D. R. Bhat, P. Chandra, S. Kudale, and S. Bhandari. Identifying the source of perytons at the Parkes radio telescope. *ArXiv e-prints*, April 2015.

[450] M. S. Briggs et al. Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3. *Astrophysical Journal*, 760:12, November 2012.

[451] L. Bildsten. Gravitational radiation and rotation of accreting neutron stars. *Astrophys.J.*, 501:L89, 1998.

[452] G. Ushomirsky, C. Cutler, and L. Bildsten. Deformations of accreting neutron star crusts and gravitational wave emission. *Mon.Not.Roy.Astron.Soc.*, 319:902, 2000.

[453] B. J. Owen. Maximum elastic deformations of compact stars with exotic equations of state. *Phys.Rev.Lett.*, 95:211101, 2005.

[454] C. Cutler. Gravitational waves from neutron stars with large toroidal B fields. *Phys.Rev.*, D66:084025, 2002.

[455] A. Melatos and D.J.B. Payne. Gravitational radiation from an accreting millisecond pulsar with a magnetically confined mountain. *Astrophys.J.*, 623:1044–1050, 2005.

[456] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio, et al. Gravitational waves from hot young rapidly rotating neutron stars. *Phys.Rev.*, D58:084020, 1998.

[457] N. Andersson, K. D. Kokkotas, and N. Stergioulas. On the relevance of the r mode instability for accreting neutron stars and white dwarfs. *Astrophys.J.*, 516:307, 1999.

[458] D.I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon.Not.Roy.Astron.Soc.*, 331:203, 2002.

[459] B. Abbott et al. Coherent searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Phys.Rev.*, D76:082001, 2007.

[460] R. Narayan. The birthrate and initial spin period of single radio pulsars. *Astrophysical Journal*, 319:162–179, August 1987.

[461] D. R. Lorimer. Binary and millisecond pulsars. *Living Reviews in Relativity*, 8(7), 2005.

[462] J. M. Cordes and D. F. Chernoff. Neutron Star Population Dynamics. II. Three-dimensional Space Velocities of Young Pulsars. *Astrophysical Journal*, 505:315–338, September 1998.

[463] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility Pulsar Catalogue. *Astronomical Journal*, 129:1993–2006, April 2005. See also `http://www.atnf.csiro.au/research/pulsar/psrcat/`.

[464] B. Knispel and B. Allen. Blandford's Argument: The Strongest Continuous Gravitational Wave Signal. *Phys.Rev.*, D78:044031, 2008.

[465] B. Abbott et al. All-sky search for periodic gravitational waves in LIGO S4 data. *Phys.Rev.*, D77:022001, 2008.

[466] V. Dergachev. Description of PowerFlux Algorithms and Implementation, 2005. LIGO-T050186, available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[467] P. R. Brady and T. Creighton. Searching for periodic sources with LIGO. 2. Hierarchical searches. *Phys.Rev.*, D61:082001, 2000.

[468] V. Dergachev. On blind searches for noise dominated signals: a loosely coherent approach. *Class.Quant.Grav.*, 27:205017, 2010.

[469] V. Dergachev. Loosely coherent searches for sets of well-modeled signals. *Phys.Rev.*, D85:062003, 2012.

[470] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca, et al. The Hough transform search for continuous gravitational waves. *Phys.Rev.*, D70:082001, 2004.

[471] B. Abbott et al. The Einstein@Home search for periodic gravitational waves in LIGO S4 data. *Phys.Rev.*, D79:022001, 2009.

[472] J. Aasi et al. Application of a Hough search for continuous gravitational waves on data from the fifth LIGO science run. *Class.Quant.Grav.*, 31:085014, 2014.

[473] Llucia Sancho de la Jordana and Alicia M. Sintes. A chi**2 veto for continuous wave searches. *Class.Quant.Grav.*, 25:184014, 2008.

[474] S. Frasca, P. Astone, and C. Palomba. Evaluation of sensitivity and computing power for the Virgo hierarchical search for periodic sources. *Class.Quant.Grav.*, 22:S1013–S1019, 2005.

[475] P. Astone, S. Frasca, and C. Palomba. The short FFT database and the peak map for the hierarchical search of periodic sources. *Class.Quant.Grav.*, 22:S1197–S1210, 2005.

[476] C. Palomba, P. Astone, and S. Frasca. Adaptive Hough transform for the search of periodic sources. *Class.Quant.Grav.*, 22:S1255–S1264, 2005.

[477] F. Antonucci, P. Astone, S. D'Antonio, S. Frasca, and C. Palomba. Detection of periodic gravitational wave sources by Hough transform in the f versus f(.) plane. *Class.Quant.Grav.*, 25:184015, 2008.

[478] P. Jaranowski, A. Krolak, and B. F. Schutz. Data analysis of gravitational - wave signals from spinning neutron stars. 1. The Signal and its detection. *Phys.Rev.*, D58:063001, 1998.

[479] J. Aasi et al. Implementation of an F-statistic all-sky search for continuous gravitational waves in Virgo VSR1 data. *Class.Quant.Grav.*, 31:165014, 2014.

[480] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan. Cross-correlation search for periodic gravitational waves. *Phys.Rev.*, D77:082001, 2008.

[481] J. Abadie et al. All-sky Search for Periodic Gravitational Waves in the Full S5 LIGO Data. *Phys.Rev.*, D85:022001, 2012.

[482] M. Shaltev and R. Prix. Fully coherent follow-up of continuous gravitational-wave candidates. *Phys.Rev.*, D87(8):084057, 2013.

[483] J. Aasi et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys.Rev.*, D87(4):042001, 2013.

[484] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. Preprint arXiv:1304.0670, 2013.

[485] B. Abbott et al. All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data. *Phys.Rev.Lett.*, 102:111102, 2009.

[486] V. Dergachev. Novel universal statistic for computing upper limits in an ill-behaved background. *Phys.Rev.*, D87(6):062001, 2013.

[487] B. Abbott et al. First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform. *Phys.Rev.*, D72:102004, 2005.

[488] B.P. Abbott et al. Einstein and Home search for periodic gravitational waves in early S5 LIGO data. *Phys.Rev.*, D80:042003, 2009.

[489] J. Aasi et al. Gravitational waves from known pulsars: results from the initial detector era. *Astrophys.J.*, 785:119, 2014.

[490] B.J. Owen. Detectability of periodic gravitational waves by initial interferometers. *Class.Quant.Grav.*, 23:S1–S8, 2006.

[491] J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, T. Adams, and et al. Narrow-band search of continuous gravitational-wave signals from Crab and Vela pulsars in Virgo VSR4 data. *Physical Review D*, 91(2):022004, January 2015.

[492] B. Abbott et al. Setting upper limits on the strength of periodic gravitational waves using the first science data from the GEO 600 and LIGO detectors. *Phys.Rev.*, D69:082004, 2004.

[493] B. Abbott et al. Limits on gravitational wave emission from selected pulsars using LIGO data. *Phys.Rev.Lett.*, 94:181103, 2005.

[494] B. Abbott et al. Upper limits on gravitational wave emission from 78 radio pulsars. *Phys.Rev.*, D76:042001, 2007.

[495] B. Abbott et al. Beating the spin-down limit on gravitational wave emission from the Crab pulsar. *Astrophys.J.*, 683:L45–L50, 2008.

[496] B.P. Abbott et al. Searches for gravitational waves from known pulsars with S5 LIGO data. *Astrophys.J.*, 713:671–685, 2010.

[497] J. Abadie et al. Beating the spin-down limit on gravitational wave emission from the Vela pulsar. *Astrophys. J.*, 737:93, 2011.

[498] Dupuis, R.J. *Bayesian Searches for Gravitational Waves from Pulsars*. PhD thesis, University of Glasgow, 2004. LIGO-P050014-00-R.

[499] M. Pitkin and G. Woan. Binary system delays and timing noise in searches for gravitational waves from known pulsars. *Phys.Rev.*, D76:042006, 2007.

[500] P. Jaranowski and A. Krolak. Search for gravitational waves from known pulsars using $\mathcal{F}$ and $\mathcal{G}$ statistics. *Class.Quant.Grav.*, 27:194015, 2010.

[501] P. Astone, S. D'Antonio, S. Frasca, and C. Palomba. A method for detection of known sources of continuous gravitational wave signals in non-stationary data. *Class.Quant.Grav.*, 27:194016, 2010.

[502] M. Pitkin, C. Gill, J. Veitch, E. Macdonald, and G. Woan. A new code for parameter estimation in searches for gravitational waves from known pulsars. *J.Phys.Conf.Ser.*, 363:012041, 2012.

[503] Gill, C. D. *Searching for gravitational waves from pulsars*. PhD thesis, University of Glasgow, 2012.

[504] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba. Coherent search of continuous gravitational wave signals: extension of the 5-vectors method to a network of detectors. *J.Phys.Conf.Ser.*, 363:012038, 2012.

[505] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba, 2012. poster at GWPAW12.

[506] K. Wette et al. Searching for gravitational waves from Cassiopeia A with LIGO. *Class.Quant.Grav.*, 25:235011, 2008.

[507] J. Abadie et al. First search for gravitational waves from the youngest known neutron star. *Astrophys.J.*, 722:1504–1513, 2010.

[508] Pinkesh Patel, Xavier Siemens, Rejean Dupuis, and Joseph Betzwieser. Implementation of barycentric resampling for continuous wave searches in gravitational wave data. *Phys.Rev.*, D81:084032, 2010.

[509] J. Aasi et al. A directed search for continuous Gravitational Waves from the Galactic Center. *Phys.Rev.*, D88:102002, 2013.

[510] C. Chung, A. Melatos, B. Krishnan, and J. T. Whelan. Designing a cross-correlation search for continuous-wave gravitational radiation from a neutron star in the supernova remnant SNR 1987A. *Mon. Not. R. Astron. Soc.*, 414:2650, 2011.

[511] R.V. Wagoner. Gravitational radiation from accreting neutron stars. *Astrophys. J.*, 278:345, 1984.

[512] J. Papaloizou and J.E. Pringle. Gravitational radiation and the stability of rotating stars. *Mon. Not. R. Ast. Soc.*, 184:501, 1978.

[513] Nils Andersson. A New class of unstable modes of rotating relativistic stars. *Astrophys.J.*, 502:708–713, 1998.

[514] J.L. Friedman and S.M. Morsink. Axial instability of rotating relativistic stars. *Astrophys. J.*, 502:714, 1998.

[515] Y. Levin. Runaway Heating by r-Modes of Neutron Stars in Low-Mass X-Ray Binaries. *Astrophys. J.*, 517:328, 1999.

[516] G.B. Cook, S.L. Shapiro, and S.A. Teukolsky. Rapidly rotating neutron stars in general relativity: Realistic equations of state. *Astrophys. J.*, 424:823, 1994.

[517] D. Chakrabarty et al. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature*, 424:42–44, 2003.

[518] M. Vigelius and A. Melatos. Gravitational-wave spin-down and stalling lower limits on the electrical resistivity of the accreted mountain in a millisecond pulsar. *Astrophys. J.*, 717:404, 2010.

[519] P. Arras and et al. Saturation of the r-Mode Instability. *Astrophys. J.*, 591:1129, 2003.

[520] A. L. Watts, B. Krishnan, L. Bildsten, and B. F. Schutz. Detecting gravitational wave emission from the known accreting neutron stars. *MNRAS*, 389:839–868, September 2008.

[521] N. Andersson, D.I. Jones, and W.C.G. Ho. Implications of an r mode in xte j1751?305: mass, radius and spin evolution. *Mon. Not. R. Ast. Soc.*, 442:1786, 2014.

[522] U. Lee. Excitation of a non-radial mode in a millisecond X-ray pulsar XTE J1751-305. *Mon. Not. R. Ast. Soc.*, 442:3037, 2014.

[523] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell. String axiverse. *Physical Review D*, 81(12):123530, June 2010.

[524] A. Arvanitaki and S. Dubovsky. Exploring the string axiverse with precision black hole physics. *Physical Review D*, 83(4):044026, February 2011.

[525] Hirotaka Yoshino and Hideo Kodama. Probing the string axiverse by gravitational waves from Cygnus X-1, 2014.

[526] J. Aasi et al. Gravitational waves from known pulsars: results from the initial detector era. *Astrophys.J.*, 785:119, 2014.

[527] H.J. Bulten et al. Gravitational waves from Sco X-1 : A comparison of search methods. in preparation.

[528] J. Aasi et al. Directed search for gravitational waves from Scorpius X-1 with initial LIGO data. *Physical Review D*, 91(6):062008, March 2015.

[529] J.T. Whelan et al. A Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1. in preparation, LIGO Publication LIGO-P00142 (2015).

[530] L. Sammut, C. Messenger, A. Melatos, and B. J. Owen. Implementation of the frequency-modulated sideband search method for gravitational waves from low mass x-ray binaries. *Physical Review D*, 89(4):043001, February 2014.

[531] S. M. Ransom, J. M. Cordes, and S. S. Eikenberry. A New search technique for short orbital period binary pulsars. *Astrophys.J.*, 589:911–920, 2003.

[532] C. Messenger and G. Woan. A Fast search strategy for gravitational waves from low-mass X-ray binaries. *Class.Quant.Grav.*, 24:S469–S480, 2007.

[533] B Abbott et al. Upper limit map of a background of gravitational waves. *Phys. Rev. D*, 76:082003, 2007.

[534] B Abbott et al. Directional limits on persistent gravitational waves using LIGO S5 science data. *Phys. Rev. Lett.*, 107:271102, 2011.

[535] E. Goetz and K. Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Class.Quant.Grav.*, 28:215006, 2011.

[536] The LIGO Scientific Collaboration, the Virgo Collaboration, J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, T. Accadia, F. Acernese, K. Ackley, et al. First all-sky search for continuous gravitational waves from unknown sources in binary systems. *Phys. Rev.*, D90:062010, May 2014.

[537] B. Haskell, N. Andersson, C. D'Angelo, N. Degenaar, K. Glampedakis, W. C. G. Ho, P. D. Lasky, A. Melatos, M. Oppenoorth, A. Patruno, and M. Priymak. Gravitational waves from rapidly rotating neutron stars. *ArXiv e-prints*, July 2014.

[538] S. van der Putten. *Thermal lensing in Virgo and Polynomial search: an all-sky search for gravitational waves from spinning neutron stars in binary systems*. PhD thesis, Free University Amsterdam, 2011. `http://www.nikhef.nl/pub/services/newbiblio/theses.php`.

[539] A. Vicere and M. Yvert. First tests of a moment autocorrelation method to detect gw from pulsars in binary systems – virgo technical note vir-0292c-13, October 2013. Virgo technical note VIR-0292C-13 (October 2013).

[540] J. Aasi et al. Improved Upper Limits on the Stochastic Gravitational–Wave Background from 2009–2010 LIGO and Virgo Data. *prl*, 113:231101, 2014.

[541] E. W. Kolb & M. S. Turner. *The Early Universe*. Westview Press, 1994.

[542] A. A. Starobinskii. Spectrum of relic gravitational radiation and the early state of the universe. *JETP Lett.*, 30, 1979.

[543] R. Bar-Kana. Limits on Direct Detection of Gravitational Waves. *Phys. Rev. D*, 50, 1994.

[544] T. W. B. Kibble. Topology of cosmic domains and strings. *J. Phys.*, A9, 1976.

[545] T. Damour & A. Vilenkin. Gravitational radiation from cosmic (super)strings: bursts, stochastic background, and observational windows. *Phys. Rev. D*, 71, 2005.

[546] S. Olmez, V. Mandic, and X. Siemens. Gravitational-Wave Stochastic Background from Kinks and Cusps on Cosmic Strings. *Phys. Rev. D*, 81:104028, 2010.

[547] S. Olmez, V. Mandic, and X. Siemens. Anisotropies in the Gravitational-Wave Stochastic Background. *J. Cosmol. Astropart. Phys.*, 2012:009, 2011.

[548] A. Buonanno. Spectrum of relic gravitational waves in string cosmology. *Phys. Rev. D*, 55, 1997.

[549] V. Mandic & A. Buonanno. Accessibility of the Pre-Big-Bang Models to LIGO. *Phys. Rev. D*, 73, 2006.

[550] T. Regimbau. The astrophysical gravitational wave stochastic background. *Research in Astronomy and Astrophysics*, 11:369–390, April 2011.

[551] A. Buonanno, G. Sigl, G. G. Raffelt, H.-T. Janka, and E. Müller. Stochastic gravitational-wave background from cosmological supernovae. *Physical Review D*, 72(8):084001, October 2005.

[552] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni. Gravitational waves from the first stars. *Physical Review D*, 73(10):104024, May 2006.

[553] S. Marassi, R. Schneider, and V. Ferrari. Gravitational wave backgrounds and the cosmic transition from Population III to Population II stars. *Monthly Notices of the Royal astronomical Society*, 398:293–302, September 2009.

[554] X.-J. Zhu, E. Howell, and D. Blair. Observational upper limits on the gravitational wave production of core collapse supernovae. *Monthly Notices of the Royal astronomical Society*, 409:L132–L136, November 2010.

[555] T. Regimbau and J. A. de Freitas Pacheco. Cosmic background of gravitational waves from rotating neutron stars. *Astronomy and Astrophysics*, 376:381–385, September 2001.

[556] P. A. Rosado. Gravitational wave background from rotating neutron stars. *Physical Review D*, 86(10):104007, November 2012.

[557] T. Regimbau and J. A. de Freitas Pacheco. Gravitational wave background from magnetars. *Astronomy and Astrophysics*, 447:1–7, February 2006.

[558] E. Howell, T. Regimbau, A. Corsi, D. Coward, and R. Burman. Gravitational wave background from sub-luminous GRBs: prospects for second- and third-generation detectors. *Monthly Notices of the Royal astronomical Society*, 410:2123–2136, February 2011.

[559] S. Marassi, R. Ciolfi, R. Schneider, L. Stella, and V. Ferrari. Stochastic background of gravitational waves emitted by magnetars. *Monthly Notices of the Royal astronomical Society*, 411:2549–2557, March 2011.

[560] C.-J. Wu, V. Mandic, and T. Regimbau. Accessibility of the stochastic gravitational wave background from magnetars to the interferometric gravitational wave detectors. *Physical Review D*, 87(4):042002, February 2013.

[561] G Sigl. Cosmological gravitational wave background from phase transitions in neutron stars. *J. Cosmol. Astropart. Phys.*, JCAP04:002, 2006.

[562] J. C. N. de Araujo and G. F. Marranghello. Gravitational wave background from neutron star phase transition. *General Relativity and Gravitation*, 41:1389–1406, June 2009.

[563] V. Ferrari, S. Matarrese, and R. Schneider. Stochastic background of gravitational waves generated by a cosmological population of young, rapidly rotating neutron stars. *Monthly Notices of the Royal astronomical Society*, 303:258–264, February 1999.

[564] X.-J. Zhu, X.-L. Fan, and Z.-H. Zhu. Stochastic Gravitational Wave Background from Neutron Star r-mode Instability Revisited. *Astrophysical Journal*, 729:59, March 2011.

[565] E. Howell, D. Coward, R. Burman, D. Blair, and J. Gilmore. The gravitational wave background from neutron star birth throughout the cosmos. *Monthly Notices of the Royal astronomical Society*, 351:1237–1246, July 2004.

[566] A. J. Farmer & E. S. Phinney. The gravitational wave background from cosmological compact binaries. *Mon. Not. R. Ast. Soc.*, 346:1197, 2003.

[567] X.-J. Zhu, E. Howell, T. Regimbau, D. Blair, and Z.-H. Zhu. Stochastic Gravitational Wave Background from Coalescing Binary Black Holes. *Astrophysical Journal*, 739:86, October 2011.

[568] P. A. Rosado. Gravitational wave background from binary systems. *Physical Review D*, 84(8):084004, October 2011.

[569] S. Marassi, R. Schneider, G. Corvino, V. Ferrari, and S. Portegies Zwart. Imprint of the merger and ring-down on the gravitational wave background from black hole binaries coalescence. *Physical Review D*, 84(12):124037, December 2011.

[570] C. Wu, V. Mandic, and T. Regimbau. Accessibility of the gravitational-wave background due to binary coalescences to second and third generation gravitational-wave detectors. *Physical Review D*, 85(10):104024, May 2012.

[571] X.-J. Zhu, E. J. Howell, D. G. Blair, and Z.-H. Zhu. On the gravitational wave background from compact binary coalescences in the band of ground-based interferometers. *Monthly Notices of the Royal astronomical Society*, 431:882–899, May 2013.

[572] L Barack & C Cutler. Confusion noise from LISA capture sources. *Phys. Rev. D*, 70:122002, 2004.

[573] G Sigl & J Schnittman & A Buonanno. Gravitational-wave background from compact objects embedded in AGN accretion disks. *Phys. Rev. D*, 75:024034, 2007.

[574] M Maggiore. Gravitational Wave Experiments and Early Universe Cosmology. *Phys. Rept.*, 331:283, 2000.

[575] V Mandic, E Thrane, S Giampanis, and T Regimbau. Parameter Estimation in Searches for the Stochastic Gravitational-Wave Background. *Phys. Rev. Lett.*, 109:171102, 2012.

[576] N. Barnaby, E. Pajer, and M. Peloso. Gauge Field Production in Axion Inflation: Consequences for Monodromy, non-Gaussianity in the CMB, and Gravitational Waves at Interferometers. *Phys. Rev. D*, 85:023525, 2012.

[577] J L Cook and L Sorbo. Particle production during inflation and gravitational waves detectable by ground-based interferometers. *Phys. Rev. D*, 85:023534, 2012.

[578] B. Abbott et al. Searching for Stochastic Gravitational Waves with LIGO. *Nature*, 460:990, 2009.

[579] J. Aasi et al. Constraints on cosmic strings from the LIGO-Virgo gravitational-wave detectors. *Physical Review Letters*, 112:131101, 2014.

[580] C. Wu, V. Mandic, and T. Regimbau. Accessibility of the gravitational-wave background due to binary coalescences to second and third generation gravitational-wave detectors. *Phys. Rev. D*, 85:104024, 2012.

[581] Kazuhiro Hayama and Atsushi Nishizawa. Model-independent test of gravity with a network of ground-based gravitational-wave detectors. *Phys. Rev. D*, 87:062003, Mar 2013.

[582] C. Brans and R.H. Dicke. Mach's principle and a relativistic theory of gravitation. *Phys.Rev.*, 124:925–935, 1961.

[583] Y. Fujii and K. Maeda. *The Scalar-Tensor Theory of Gravitation*. Cambridge Monograph on Mathematical Physics. Cambridge University Press, Cambridge, 2002.

[584] Thomas P. Sotiriou and Valerio Faraoni. f(R) Theories Of Gravity. *Rev.Mod.Phys.*, 82:451–497, 2010.

[585] Antonio De Felice and Shinji Tsujikawa. f(R) Theories. *Living Rev. Relativity*, 13(3), 2010.

[586] Matt Visser. Mass for the graviton. *General Relativity and Gravitation*, 30(12):1717–1728, 1998.

[587] Claudia de Rham, Gregory Gabadadze, and Andrew J. Tolley. Resummation of massive gravity. *Phys. Rev. Lett.*, 106:231101, Jun 2011.

[588] B Abbott et al. Searching for a Stochastic Background of Gravitational Waves with the Laser Interferometer Gravitational-Wave Observatory. *Astrophys. J.*, 659:918, 2007.

[589] J. Abadie et al. Upper limits on a stochastic gravitational-wave background using LIGO and Virgo interferometers at 600-1000 Hz. *Phys. Rev. D*, 85:122001, 2012.

[590] J. Aasi et al. Searching for stochastic gravitational waves using data from the two colocated LIGO Hanford detectors. *prd*, 91:022003, 2015.

[591] Bruce Allen and Joseph D. Romano. Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities. *Phys.Rev.*, D59:102001, 1999.

[592] N Christensen. Measuring the Stochastic Gravitational Radiation Background with Laser Interferometric Antennas. *Phys. Rev. D*, 46:5250, 1992.

[593] B Abbott et al. Analysis of first LIGO science data for stochastic gravitational waves. *Phys. Rev. D*, 69:122004, 2004.

[594] D Meacher, E Thrane, and T Regimbau. Statistical properties of astrophysical gravitational-wave background. *Phys. Rev. D*, 89:084063, 2014.

[595] D. Meacher et al. A Mock Data and Science Challenge for Detecting an Astrophysical Stochastic Gravitational Wave Background with Advanced LIGO and Advanced Virgo. *In Preparation*, 2015. `https://dcc.ligo.org/LIGO-P1500078`.

[596] J Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories, 2013. http://arxiv.org/abs/1304.0670.

[597] E Thrane, N Christensen, and R Schofield. Correlated magnetic noise in global networks of gravitational-wave interferometers: observations and implications. *Phys. Rev. D*, 87:123009, 2013.

[598] A Ain, P Dalvi, and S Mitra. Fast gravitational wave radiometry using data folding. 2015. `http://xxx.tau.ac.il/abs/1504.01714`.

[599] E Thrane, S Mitra, N Christensen, V Mandic, and A Ain. All-sky, narrowband, gravitational-wave radiometry with folded data. *Accepted in Phys. Rev. D*, 2015.

[600] D Talukder, E Thrane, S Bose, and T Regimbau. Measuring neutron-star ellipticity with measurements of the stochastic gravitational-wave background. *Accepted for publication in Phys. Rev. D*, 2014. `http://arxiv.org/abs/1404.4025`.

[601] T Regimbau & B Chauvineaux. A stochastic background from extra-galactic double neutron stars. *Class. Quantum Grav.*, 24:627, 2007.

[602] E Howell et al. The gravitational wave background from neutron star birth throughout the cosmos. *Mon. Not. R. Ast. Soc.*, 351:1237, 2004.

[603] V Ferrari & S Matarrese & R Schneider. Gravitational wave background from a cosmological population of core-collapse supernovae. *Mon. Not. R. Ast. Soc.*, 303:258, 1999.

[604] V Ferrari & S Matarrese & R Schneider. Stochastic background of gravitational waves generated by a cosmological population of young, rapidly rotating neutron stars. *Mon. Not. R. Ast. Soc.*, 303:258, 1999.

[605] T Regimbau & J A de Freitas Pacheco. Cosmic background of gravitational waves from rotating neutron stars. *Astron. Astrophys.*, 376:381, 2001.

[606] T Regimbau & J A de Freitas Pacheco. Gravitational wave background from magnetars. *Astron. Astrophys.*, 447:1, 2006.

[607] J. Aasi et al. Constraints on cosmic strings from the LIGO-Virgo gravitational-wave detectors. *Phys.Rev.Lett.*, 112:131101, 2014.

[608] C. Wu, V. Mandic, and T. Regimbau. Accessibility of the stochastic gravitational wave background from magnetars to the interferometric gravitational wave detectors. *Phys. Rev. D*, 87:042002, 2013.

[609] E Thrane, N Christensen, R M S Schofield, and A Effler. Correlated noise in networks of gravitational-wave detectors: subtraction and mitigation. *Phys. Rev. D*, 90:023013, 2014.

[610] E Thrane, V Mandic, and N Christensen. Detecting very long-lived gravitational-wave transients lasting hours to weeks. *Phys. Rev. D*, 91:104021, 2015.

[611] R Prix, S Giampanis, and C Messenger. Search method for long-duration gravitational-wave transients from neutron stars. *Phys. Rev. D*, 84:023007, 2011.

[612] C A van Eysden and A Melatos. Gravitational radiation from pulsar glitches. *Class. Quantum Grav.*, 25:225020, 2008.

[613] M F Bennett, C A van Eysden, and A Melatos. Continuous-wave gravitational radiation from pulsar glitch recovery. *Mon. Not. R. Ast. Soc.*, 409:1705, 2010.

[614] A G Lyne, S L Shemar, and F G Smith. Statistical studies of pulsar glitches. *Mon. Not. R. Ast. Soc.*, 315:534, 2000.

[615] K Kashiyama and K Ioka. Magnetar asteroseismology with long-term gravitational waves. *Phys. Rev. D*, 83:081302, 2011.

[616] Y Levin and M van Hoven. On the excitation of f modes and torsional modes by magnetar giant flares. *Mon. Not. R. Ast. Soc.*, 418:659, 2011.

[617] P Jaranowski, A Kr'olak, and B F Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58:063001, 1998.

[618] D I Jones and N Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Ast. Soc.*, 331:203, 2002.

[619] L Gualtieri, R Ciolfi, and V Ferrari. Structure, deformations and gravitational wave emission of magnetars, 2010. Prepared for 19th International Conference on General Relativity and Gravitation (GR19), Mexico City, Mexico, July 5-9, 2010, arxiv/1011.2778.

[620] S Dall'Osso, S N Shore, and L Stella. Early evolution of newly born magnetars with a strong toroidal field. *Mon. Not. R. Ast. Soc.*, 398:1869, 2009.

[621] N Andersson and K D Kokkotas. The r-mode instability in rotating neutron stars. *Int. J. Mod. Phys. D*, 10:381, 2001.

[622] R Bondarescu, S A Teukolsky, and I Wasserman. Spinning down newborn neutron stars: Nonlinear development of the r-mode instability. *Phys. Rev. D*, 79:104003, 2009.

[623] A Melatos and M Priymak. Gravitational radiation from magnetically funnelled supernova fallback onto a magnetar. *Astrophys. J.*, 794:170, 2014.

[624] F Camilo, S M Ransom, S Chatterjee, S Johnston, and P Demorest. PSR J1841-0500: A Radio Pulsar That Mostly Is Not There. *Astrophys. J.*, 746:63, 2012.

[625] L Bildsten. Gravitational radiation and rotation of accreting neutron stars. *Astrophys. J. Lett.*, 501:L89, 1998.

[626] A Arvanitaki, M Baryakhtar, and X Huang. Discovering the qcd axion with black holes and gravitational waves. *Phys. Rev. D*, 91:084011, 2015.

[627] T. Prestegard and E. Thrane. Burstegard: a hierarchical clustering algorithm. *LIGO DCC*, page L1200204, 2012. `https://dcc.ligo.org/cgi-bin/DocDB/ShowDocument?docid=93146`.

[628] M. Coughlin, E. Thrane, and N. Christensen. Detecting compact binary coalescences with seedless clustering. *Phys.Rev.*, D90(8):083005, 2014.

[629] P. Meyers, M. W. Coughlin, and J. Luo. Investigating Environmental Noise Using the Stochastic Transient Analysis Multi-Detector Pipeline (STAMP-PEM), 2014. `https://dcc.ligo.org/LIGO-G1400354`.

[630] M. Coughlin for the LIGO Scientific Collaboration and the Virgo Collaboration. Identification of long-duration noise transients in LIGO and Virgo. *Class. Quantum Grav.*, 28:235008, 2011.

[631] LSC Detector Characterization Group home page. `https://wiki.ligo.org/foswiki/bin/view/DetChar/WebHome`.

[632] J. Zweizig. Data Monitor Tool Project .
`http://www.ligo.caltech.edu/~jzweizig/dmt/DMTProject/`.

[633] EPICS home page. `http://www.aps.anl.gov/epics/index.php`.

[634] Home page for the omega pipeline:.
`https://trac.ligo.caltech.edu/omega/`.

[635] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21(20):S1809, 2004.

[636] D. M. Macleod, S. Fairhurst, B. Hughey, A. P Lundgren, L. Pekowsky, J. Rollins, and J. R. Smith. Reducing the effect of seismic noise in LIGO searches by targeted veto generation. *Classical and Quantum Gravity*, 29(5):055006, 2012.

[637] T. Isogai, LIGO Scientific Collaboration, Virgo Collaboration, et al. Used percentage veto for LIGO and virgo binary inspiral searches. *Journal of Physics: Conference Series*, 243(1):012005, 2010.

[638] J. R Smith, T. Abbott, E. Hirose, N. Leroy, D. MacLeod, J. McIver, P. Saulson, and P. Shawhan. A hierarchical method for vetoing noise transients in gravitational-wave detectors. *Classical and Quantum Gravity*, 28(23):235005, 2011.

[639] Advanced LIGO Subsystem Characterization Matrix.
`https://wiki.ligo.org/DetChar/AligoSubsystemMatrix`.

[640] D. Sigg, R. Bork, and J. Zweizig. Detector Characterization and Global Diagnostics System of the Laser Interferometer Graviational-wave Observatory (LIGO). In *Proceedings of the Ninth Marcel Grossman Conference*, 2000.

[641] ligoDV home page.
`https://www.lsc-group.phys.uwm.edu/daswg/projects/ligodv.html`.

[642] mDV is available from the 40m svn.

[643] J. Smith, A. Lundgren, N. Christenson, F. Robinet, C. Pankow, D. MacLeod, J. McIver, and J. Zweizig. DetChar Trigger Handling Conventions, LIGO Document T1300468-v4. Technical report, LIGO, 2013. https://dcc.ligo.org/LIGO-T1300468.

[644] LIGO calibration web page.
`http://blue.ligo-wa.caltech.edu/engrun/Calib_Home/`.

[645] R. Adhikari *et al.*, “Calibration of the LIGO Detectors for the First LIGO Scientific Run”, LIGO-T030097 (2003);
G. Gonzalez *et al.*, “Calibration of the LIGO Detectors for S2”, LIGO-T040060 (2004);
G. Gonzalez *et al.*, “Calibration of the LIGO Detectors for S3”, LIGO-T050059 (2005);
A. Dietz *et al.*, “Calibration of the LIGO detectors for S4”, LIGO-T050262-01-D (2006).

[646] Y. Aso et al. Accurate measurement of the time delay in the response of the ligo gravitational wave detectors. *Classical and Quantum Gravity*, 26(5):055010, 2009.