LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	**LIGO–T1400054** **VIR-0176A-14**

The LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics (2014-2015 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 2014/11/21

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 into four LSC-Virgo working groups: compact binary coalescences (CBC), generic transients (Burst), continuous waves (CW) and stochastic gravitational-wave background (SGWB). Each of these groups pursues distinct astrophysical sources of gravitational waves with different methods. Joint teams formed from the members of two or more working groups exist, where the science suggests overlap between sources or methods. In addition to these astrophysics groups, the Detector Characterization group, which also collaborates with the detector commissioning teams, works to improve search sensitivity by identifying and mitigating instrumental noise sources that limit the sensitivity to astrophysical signals.

The *LSC-Virgo White Paper on Gravitational Wave Searches and Astrophysics* describes the goals, status, and plans of the data analysis teams in the LSC and Virgo, as well as plans and status of Detector Characterization, calibration and development and maintenance of software tools and the management of software and computing resources. This document, which is revised and updated every year, is intended to facilitate the understanding of the science that we are pursuing, the prioritization of our objectives, the development plan and resource needs. The 2014-2015 version offers a significant re-structure from previous year, resulting from discussions in the LSC-Virgo Data Analysis Council, with four components:

1. **Chapter 1** is the executive summary for this white paper, outlining for each search group a mission statement and a ranked list of priorities both in terms of scientific targets and tasks in preparation for Advanced LIGO and Advanced Virgo.
2. **Chapter 2** lists papers and milestones from the period May 2013-May 2014.
3. **Chapter 3** is a collection of search plans for Advanced LIGO and Advanced Virgo, with details on developments and resources. The plans can also be read as standalone documents and will be updated on a 6-month timescale.
4. **Chapter 4** includes ongoing analysis on initial detector data or exploratory work that do not yet fit the search plan template.

We refer to the Advanced Detector Era (ADE) as the epoch of Advanced LIGO and Advanced Virgo science data acquisition, currently scheduled to start in the second half of 2015. The schedule of science run is formulated by the LSC-Virgo Joint Running Plan Committee, which includes representatives from the laboratories, the commissioning teams and search groups, and the plans are formulated according to the schedule outlined in the document "Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories" (arXiv:1304.0670).

Epoch	Estimated Run Duration	$E_{\rm GW} = 10^{-2} M_\odot c^2$ Burst Range (Mpc)		BNS Range (Mpc)		Number of BNS Detections	% BNS Localized within	
		LIGO	Virgo	LIGO	Virgo		5 deg^2	20 deg^2
2015	3 months	40 – 60	–	40 – 80	–	0.0004 – 3	–	–
2016–17	6 months	60 – 75	20 – 40	80 – 120	20 – 60	0.006 – 20	2	5 – 12
2017–18	9 months	75 – 90	40 – 50	120 – 170	60 – 85	0.04 – 100	1 – 2	10 – 12
2019+	(per year)	105	40 – 80	200	65 – 130	0.2 – 200	3 – 8	8 – 28
2022+ (India)	(per year)	105	80	200	130	0.4 – 400	17	48

Table 1: Summary of a plausible observing schedule, expected sensitivities, and source localization with the advanced LIGO and Virgo detectors, which will be strongly dependent on the detectors' commissioning progress. Extracted from arXiv:1304.0670

1.1 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 spacetimes of compact object mergers generate gravitational waves that will reveal the physics of strong-field gravity. The gravitational waves from mergers involving black holes will allow us to explore General Relativity and the nature of gravity. With densities of matter inaccessible to terrestrial experiments, mergers involving neutron stars hold the key to understanding the equation of state of nuclear matter. This is a crucial piece of fundamental physics missing from our understanding of the universe. 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 a maximum distance of ~ 400 Mpc, neutron star–black hole (NSBH) binaries to ~ 1 Gpc, and stellar-mass binary black holes (BBHs) at distances over 2 Gpc. LIGO and Virgo conduct their searches jointly giving us a three-detector network that can be used to localize sources on the sky. 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 is designed to identify gravitational wave signals from compact binary sources in the detector data, measure the waveform parameters, and use detected signals study the nature of gravity and the astrophysics of nature's most compact objects. Detection and parameter measurement of CBC signals is carried out by the joint LSC-Virgo CBC working group. The CBC group's science program 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.

The LSC charter states that the mission of the LSC is *to detect gravitational waves, use them to explore the fundamental physics of gravity, and develop gravitational wave observations as a tool of astronomical discovery.* We prioritize CBC tasks both in view of this mission statement and in terms of their chronological importance in the advanced detector era. Accordingly, *highest priority* is given to the ability to detect the most probable CBC sources with the LIGO and Virgo detectors, followed by science with the first detections (*high priority*), and finally the detection of sources which in the current astrophysical paradigm may have a low rate within the detectors' volume reach, but whose discovery would nevertheless be of great astrophysical importance (*priority*).

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 binary neutron stars, stellar mass binary black holes, neutron star–stellar mass black hole binaries, and the detection of gravitational waves from compact binaries in coincidence with an externally triggered short-hard gamma ray burst. The CBC group must be ready to measure the masses and spins of sources on first detection, to provide rate estimates for detected sources, and to provide rapid significance measurements of externally triggered gamma-ray burst events.

   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 priorities**

   Once CBC sources have been detected, a significant amount of astrophysics can be extracted from the observed gravitational waves. Preparing to extract this information accurately is a high priority for the CBC group. This includes: precise measurement of masses and spins to understand the properties of compact objects and their formation; determination of the neutron star equation of state; tests of the genuinely strong-field dynamics of spacetime, a regime which can only be probed with direct gravitational wave detection; cosmological studies without the need for a cosmic distance ladder; and accurate measurement of coalescence rates for CBC sources.

   Although a coincident electromagnetic (EM) counterpart is not required to detect gravitational waves from compact binary coalescence, the coincident detection of an electromagnetic counterpart with a CBC event would add significant astrophysical information to our discoveries. Preparing for detections of EM counterparts to gravitational wave events is a high priority for the CBC group. This includes: development of low latency analysis pipelines, low-latency data quality assessment, low-latency significance estimation, sky localization for CBC sources, and preparations for joint gravitational wave/EM observations in the advanced detector era.

3. **Priorities**

   Priorities include expanding the CBC search to 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.

CBC group science plans are organized by the physics of the sources, although there is overlap in the way that searches are performed. Each science plan provides an abstract, scientific justification, publication plan, technical requirements and development plan and the resources required. The organization of the science plans is as follows:

- search for binary neutron star coalescences;
- search for stellar mass binary black hole coalescences;
- search for neutron star–black hole coalescences;
- search for intermediate mass black hole binary coalescences (joint with the Burst working group);
- search for gravitational wave counterparts to gamma-ray bursts (joint with the Burst working group).

1.2 Searches for Generic Transients, or Bursts

The Burst group's mission is to detect gravitational wave transients, or *bursts*, and gain new information on transient signal populations and emission mechanisms of astrophysical objects, as well as to test theories of gravity. The group aims to extract a broad range of observational results from early data from the Advanced gravitational wave detector network, building on the online and offline analysis experience and infrastructure developed for first generation interferometric detectors.

Some gravitational wave progenitors, such as supernovae, involve complex physics and dynamics for which no robust signal model currently exists. Other sources, such as the merger of intermediate-mass black holes, produce gravitational wave transients which appear as short bursts in the data. Therefore, the Burst Group implements a variety of methods to find transients in the data that are inconsistent with the baseline noise. These methods identify coincident excitations between multiple detectors to discriminate between gravitational waves and noise fluctuations. In a few special cases when an accurate signal model is available, such as for cosmic string cusps or neutron star or black hole ringdowns, a search can be done using matched filtering with a bank of templates. Otherwise, gravitational wave bursts can be identified in the strain data as excess-power localized events in the time-frequency domain.

Although burst search algorithms are designed to detect a wide range of signals, their tuning and interpretation benefits 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 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 increases 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.

In preparation for the Advanced Detector Era, the Burst Group have prioritized their activities into 3 categories. Burst activities that are listed as *highest priorities* are activities that must be performed leading up to the Advanced Detector Era and must be on time, targeting the widest possible survey of the gravitational wave transient sky or serendipitous astrophysical phenomena. Activities listed as *high priorities* are activities which will boost the scientific return in the early years of the Advanced Detector Era, targeting deeper searches for selected source classes or the most promising extensions of the searched transient parameter space with respect to the highest priority activities. Finally, additional *priorities* are activities that will achieve a high science potential according to the current astrophysical and cosmological framework once the Advanced detectors reach design sensitivity.

1. **Highest priority**

   The Burst group is focused on an eyes wide open approach to detecting gravitational wave bursts in the Advanced Detector Era. The goal is to identify the presence of gravitational wave transients in data acquired by the Advanced detectors over as broad a parameter space as possible, under the assumption that the gravitational wave signal searched for is transient in nature. The Burst group must also be ready to extract astrophysical information from any detected signal. Therefore, the highest priority activities will contribute to:

   - a statement on the transient gravitational wave sky; performing population studies if we have

several detections, reporting on a rare-event detection significance if we have one candidate or producing an upper limit on the rate of gravitational wave bursts if there is no detection;
- astrophysical interpretation of any detected signals, including the development of methods for signal characterization and parameter estimation;
- prompt burst analysis, trigger production and sky localization, for 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- or high-energy neutrinos.

2. **High priorities**

   In addition to the eyes wide open approach, the Burst group will perform searches guided by information from various astrophysical sources. The additional astrophysical information reduces the parameter space over which the search must be performed, leading to a reduction in the false alarm rate and, consequently, an improvement in the search sensitivity. These activities will boost the scientific return in the early years of the Advanced Detector Era and should, therefore, be carried out with a high priority. They are:

   - searches for gravitational wave bursts from intermediate mass binary black holes, binary black holes with eccentric orbits and intermediate mass ratio inspirals;
   - 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;
   - searches for gravitational wave bursts originating from cosmic strings.

3. **Priorities**

   Additional priorities are activities that will have the most potential for scientific impact once the Advanced detectors have reached design sensitivity. Such activities 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.

The Burst working group's wide parameter search for gravitational wave bursts across the entire sky is detailed in the All-Sky Burst Search plan. This search plan lays out the search's scientific justification, publication plan, technical requirements, development plan and resources required. Further to the All-Sky Burst Search plan, the Burst group has also formulated search plans for intermediate mass black hole binaries and gravitational wave counterparts to gamma ray bursts, jointly with the CBC working group. Additionally, members of the Burst, Continuous Wave and Stochastic groups have been working together on searches for long-lived gravitational wave transients.

1.3 Searches for Continuous Wave Signals

The LSC/Virgo Continuous Waves (CW) Group aims to detect and measure GW 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 GW 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 GW 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 actual continuous GW emission and because electromagnetic astronomers have detected less 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 GW 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 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. 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 known binary sources with unknown source frequency must make similar sensitivity tradeoffs, and all-sky searches for unknown binary sources 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 priorities**
    - 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 and PSR J1751–305. 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 third object for a sharp X-ray periodicity 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 serious prospect of discovery.

2. **High priorities**

   - Targeted searches for known pulsars for which the spindown limit is unlikely to be beaten.
   - 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.
   - Spotlight searches of interesting sky patches, *e.g.,* the galactic center or other star-forming regions.
   - 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.

3. **Priorities**

   - Supernova post-birth search in our galaxy or a nearby galaxy.
   - Other long-lived periodic signals (*e.g.,* sidereal effects, non-sinusoidal periodicity)

There is some overlap in 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 Gravitational Wave Backgrounds

The prime objective of the Stochastic Gravitational Wave Background (SGWB) working 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 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; recent results from the BICEP2 experiment provide evidence for the existence of primordial gravitational waves in the form B-modes. 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 $\Omega_0 < 6.1 \times 10^{-6}$ 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. 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.

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 successfully used to search for gravitational waves from Sco 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, Sco X-1). The radiometer limits on Sco X-1 from initial LIGO remain the most constraining to date, and the Group continues to develop the search, in collaboration with the Continuous Wave Group.

The SGWB group has developed a pipeline to search for long-lived $\sim 10\,\mathrm{s}$–$1000\,\mathrm{s}$ gravitational-wave transients using cross-correlation. Long-lived gravitational-wave transients may be produced following the collapse of a massive star through rotational instabilities in the protoneutron star remnant or from instabilities in (or fragmentation of) the resultant accretion disk. A single long-lived gravitational-wave transient detection could provide an unparalleled glimpse into the moments following stellar collapse and the birth of a neutron star or black hole. Other scenarios have been proposed as well, including protoneutron star convection, rotational instabilities in merger remnants, and eccentric binary systems. In collaboration with the Bursts Group, the stochastic group has searched for long-lived gravitational-wave transients coincident with long gamma-ray bursts; the first long-lived gravitational-wave transient all-sky search is ongoing.

The SGWB group maintains an active role in detector characterization efforts. Current/recent work includes: noise line hunting, environmental monitors and subsystem coherence studies, correlated noise (Schumann resonance) studies, realtime data-quality monitor development, and participation in the hardware injection subgroup.

1. **Highest priority**

   The highest priority of the SGWB group are its two flagship searches: isotropic and directional. 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 is an import as a tool for distinguishing between different sources of the stochastic background and a mature, long-established analysis, providing targeted radiometer results and of the gravitational-wave sky maps.

2. **High priority**

   We designate as high priority activities that are carried out in order to extend the science possible with stochastic observations. Current/recent work in this vein includes mock data challenges, modeling of different sources of the stochastic background, and parameter estimation. The group continues to support the development of a pipeline to look for non-standard polarization modes in the stochastic background. We also designate as high priority searches for long-lived transients, which will be increasingly coordinated with other search groups. In parallel, we are making preparations for long transient searches with advanced detectors. Finally, we designate as high-priority the detector characterization efforts that pertain directly to the stochastic background searches, including a fast time-domain cross correlation engines (with possible applications to transient searches), a real-time data-quality monitor, an environmental monitoring system based on the long transient search software (which has recently made its mark as helpful commissioning tool), and investigations into correlated noise between different detector sites.

3. **Priorities**

   We designate as priorities activities, which are in a earlier stage of development than those listed above. The SGWB group is interested in the development of techniques for measuring the non-Gaussianity of the stochastic background, exploring the deployment of a robust non-Gaussian analysis pipeline. The group supports the continued development of folded data products in which a full science run's worth of data is collapsed into one sidereal day. Finally, the SGWB group supports an ongoing effort to further upgrade the stochastic pipeline, to improve in-code documentation, and to improve the user experience while leaving the bulk of the code intact. This project is being carefully coordinated so that, at no point, we will find ourselves dependent on a untested, unreviewed pipeline.

There is overlap in the stochastic group's search for 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 Stochastic 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

LIGO: LIGO's sensitivity to gravitational-wave signals is limited by noise from the instruments and their environment. In the early runs of Advanced LIGO, anticipated signals will be near the search detection thresholds, and so the success of detection, the vetting of candidate signals, and the accuracy of parameter estimation will rely heavily on the quality of the data searched and the collaboration's knowledge of the instrument and environment. The LIGO Detector Characterization group is therefore focused on working together with the astrophysical search groups and the detector groups to (i) deliver the data quality information necessary to clean the data sets, veto false positives, and allow candidate follow up for gravitational-wave searches and (ii) characterize the early Advanced LIGO detectors to help to identify data quality issues early enough that they can be mitigated in the instruments to improve future instrument and search performance.

Search Data Quality: LIGO data contain non-Gaussian components such as noise transients and quasi-periodic lines that have a negative impact on astrophysical searches. Compact Binary Coalescence sources are the most likely to be detected in early Advanced LIGO and transient noise in the detector data can mimic or mask their transient signals, limiting their detection probability and the accuracy of the source parameters recovered. The detection probability for burst searches is affected even more by the rate of transient noise in the detectors. To minimize these negative effects, LIGO data must be cleaned of transient data quality issues. The primary forms of data quality information that must be delivered to the astrophysical search groups are: *state segments* that indicate which data should be analyzed, based on the state of the instrument and its calibration; *veto segments* that indicate periods of poor quality data; and *data quality triggers* that identify short durations where the data are likely to contain a non-astrophysical disturbance. Searches will use veto segments and data quality triggers to either ignore problematic data or to reduce confidence in any search triggers associated with these times. For continuous-wave and stochastic backgrounds searches, frequency bins that are contaminated by non-astrophysical disturbances must be identified and removed, and low-level, broadband contamination from correlated magnetic noise must be mitigated.

Early aLIGO Characterization: During initial LIGO, detector characterization was primarily done by looking at artifacts in the detector outputs or search results and trying to reconstruct their cause from the thousands of interferometer channels. Throughout aLIGO installation, the Detector Characterization group has worked with the detector groups to identify and resolve issues in the early aLIGO subsystems related to glitch and noise contamination, channel signal fidelity and robustness, etc. This work has led to some early data quality improvements and helped to train a wider pool of scientists who are familiar with the instruments. Continued work aims to expedite aLIGO detections by ensuring that the detectors are well understood and that data quality issues are aggressively pursued. In 2014-2015, this work will be very important as the LIGO detectors transition from installation to running and intense commissioning begins to identify limits to sensitivity and robustness in the instruments.

1. **Highest priority.** The highest priority of the LIGO Detector Characterization group is to provide data quality information to the LSC-Virgo search groups that designate what data should be analyzed, remove egregious data quality issues, and identify periods/frequencies of poor data quality.

2. **High priorities.** Characterize early Advanced LIGO to improve future searches and inform the search data quality production. Complete a well-understood and documented physical environmental monitoring system. Develop and maintain the software infrastructure required to provide useful data quality information to online searches.

3. **Priorities.** Develop improved methods to uncover the causes of the noise transients which most impact the searches, with the goal of mitigating them or producing vetoes.

To accomplish these priorities, the LIGO Detector Characterization is organized two sub-groups: one devoted to search data quality and one devoted to instrument characterization.

Virgo

The search for gravitational signals requires a careful monitoring of the detector's response and the mitigation of noise sources. Noise mitigation, spectral lines identification, glitch reduction and vetoes production are, in coordination with commissioning and data analysis teams, the main tasks of the Virgo detector characterization. We also aim at a general understanding of noise features: its stationarity and Gaussianity as well as characterization of spectral lines or noise couplings. The detector characterization work is an interconnection between the commissioning team tracking down any limitation of the detector's sensitivity, the calibration team maintaining the calibration and timing accuracy to an acceptable level for GW searches, the Virgo data quality (VDQ) and Noise Studies teams providing noise information and vetoes to the data analysis groups and commissioning team. Along the previous scientific runs and commissioning periods, Virgo detector characterization already provided several tools, studies and vetoes which impacted positively both commissioning activity and astrophysical searches.

Search Data Quality: We aim to base future glitch studies and online vetoes on a standard trigger generator and on an efficient hierarchical veto algorithm. We develop, in common with LIGO, a database of data quality segments (DQSEGDB) which will be dedicated to transient signals for Burst and CBC search groups. We plan to use it also to identify families of glitches during commissioning. Other data quality needs are more specific to CW or Stochastic search groups and focus on identifying noise source contributions to spectral lines or non stationary and non linear features. We plan to use automatic spectral lines identification tools already well tested and the maintenance of a lines database which can be used by CW search group to veto the data and by the commissioning to find hints of noise sources.

Early AdvVirgo Characterization: We plan to start noise and glitch studies on each commissioned subsystem as soon as it starts to provide usable signals. This will be done in close collaboration with subsystem coordinators. To provide early information about the detector status, the noise level and the glitches impact, we have developed specific tools able to monitor in-time or online the various signals acquired, a standard trigger generator and an identification of possible correlation between the gravitational signal and the auxiliary signals. To help noise characterization, we developed also noise monitoring tools plugged in the framework NMAPI (Noise Monitor Application Programming Interface), most of them already tested on past Virgo Scientific Runs. These tools give information on spectral lines, non stationary or non linear noise and glitches rate. We developed a new pipeline (SiLeNTe) which finds the rank for non linear noise coupling between gravitational and auxiliary signals.

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 priorities.** 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 spectral lines database catalogue, identification of non stationary lines and a software infrastructure to provide useful online data quality information.

3. **Priorities.** Develop improved methods to uncover the paths and the sources of the noise transients which most impact the searches. Setup automated classification tools for noise sources.

To accomplish these priorities, Virgo Detector Characterization is organized in two groups: one devoted to search data quality and glitch studies (VDQ) and one devoted to lines identification and noise hunting (Noise Studies).

1.6 Data Calibration

LIGO

Calibration of the LIGO interferometers is critical to the success of the searches and to the confidence associated with their results. It 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. Among the critical calibration activities to be done, one can emphasize:

- 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's strain channel 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 to the interpretation of results from data analysis pipelines. 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.

The Burst and CBC groups work with the subgroup to provide transient waveforms and to determine suitable injection rates. The CW group selects the parameters for neutron star signals, which persist throughout the science run. The SGWB group typically carries out one or two $\approx$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. The results of these studies are reported back to the hardware injection team so that adjustments can be made.

While most injections are made known to the LSC, there are also blind injections. Blind injections are carried out by a separate team. However, the hardware injection group is in charge of maintaining the blind injection infrastructure, nearly identical to the regular injection infrastructure, and provide training. We aim to strengthen the collaboration between the hardware injection group and the blind injection team, e.g., by sharing one team member who can facilitate the flow of information between the two teams.

1.8 Computing and Software

LIGO

In the LIGO Computing model, data is acquired at the LIGO Observatories where it is processed locally to characterize the detector and allow timely identification and correction of problems with the instruments. Calibrated data is transferred over a dedicated low-latency network to the Tier-1 center at Caltech where it is processed to provide alerts to other astrophysicists within minutes of data acquisition. All the data is transferred over the LIGO Data Replicator (LDR) network to the archive at Caltech; calibrated data is transferred to the Tier-2 data centers where deep searches are run in batch mode.

During initial LIGO (iLIGO), the LIGO Scientific Collaboration (LSC) successfully established a collaboration controlled, distributed computing environment–the LIGO Data Grid (LDG)–to carry out prototyping, testing, and production data analysis activities. The current plan builds on that and leverages new opportunities to integrate shared resources, such as XSEDE, with the fabric of the LDG to accommodate our computational requirements as advanced LIGO (aLIGO) commissioning gives way to continuous observing runs in 2018. In this context, we refer to the computing facilities at the LIGO Observatories as Tier-1-Observatories; the computing facilities at Caltech as Tier-1-Caltech; shared resources (yet to be identified, but possibly including XSEDE) as Tier-2-Shared; and other dedicated computing facilities as Tier-2-Other.

Data analysis is assumed to proceed at the pace of data acquisition. This leads to the need for captive computing resources when the instruments are acquiring observational data. The location of hardware resources follows from the science drivers and the acceptable latencies between data acquisition and analysis for each of them. In particular, the Tier-1 resources (at Caltech and the Observatories) are used for low-latency and day-latency computing that will run essentially continually from 2015-2018. It is hoped that Tier-2-Shared facilities will carry out the deep, production searches and provide resources for rapid-turnaround data exploration and pipeline development. Full details of

The **highest priorities** for software and computing directly support O1 readiness. The following list of task flow directly from the science requirements stated elsewhere: a) provide an accurate determination of the computing requirements and optimizations for O1; b) identify and/or deploy hardware resources to support O1 science; c) deploy authorization infrastructure and tools for collaboration with electromagnetic partners; d) deploy and support a build and test environment that allows the scientific review teams' assessment of the correct behavior of software to be used in future science statements; e) define and deploy improved computing environment for detector monitoring; f) support and improve the low-latency data transfer network software and services; g) complete, deploy and support the new data-quality database; h) deploy and support segment generation, aggregation, publication and verification services; i) support data publication, replication and verification on the ligo-data-replication network; j) support application layer I/O and framework libraries including framecpp, gds, lalsuite, glue, pyCBC, gstlal, etc; k) support software development related to the Online Detector Characterization (ODC) infrastructure; l) support and enhance monitoring tools to provide insight into the status of computing and analyses; m) support and enhance the network data server infrastructure (NDS2); n) support and enhance the Gravitational-wave Candidate Event Database (GraCEDb) and associated tools; o) provide support to package and distribute software via the LSCSoft repositories; p) support the LIGODataFind services including data discovery; q) support and enhance the Channel Information System; r) support and enhance the data-quality summary service; t) support and enhance LIGO Data Viewer and LDVweb.

A detailed description of LIGO commputing and software can be found in the LSC Computing Plan (LIGO Document T050053).

Virgo

The Virgo Computing Model and the Implementation Plan describe in detail the Advanced Virgo computing infrasturcture. The purpose of Advanced Virgo Computing Model is to design, implement and maintain the computing infrastructure necessary to achieve Virgo's scientific goals. In this spirit, we make use of earlier Virgo Computing developments, but at the same time it is necessary to enlarge the range of computing activities and improve the various solutions.

Virgo computing and storage resources are part of the European Grid Infrastructure, as such are shared resources and are scattered around in various Computer Centers of the collaborating and / or supporting institutions. Virgo data analysis activities has also made use of LDG resources in the past.

Motivated by the above facts we define the following highest priority ones for the forthcoming period:

Highest priority tasks **a.)** clarify the and / or make more precise the exact computing needs of the various LIGO - Virgo scientific data analysis pipelines relevant for Virgo computing resources **b.)** ensure that the analysis pipelines that are ported and/or optimized to new, shared resources (such as for example XSEDE) remains executable on Virgo computing resources **c.)** ensure that all or at least big majority of the scientific data analysis pipelines can be executed on Virgo resources **d.)** make sure that pipeline developers are aware and in fact using new technologies which could possibly make their code more efficient and fast (such as for example GPUs, more efficient compilers, vectorisation, etc) **e.)** deploy and support a build and test environment that allows the scientific review teams' assessment of the correct behavior of software to be used in future science statements; **f.)** enable wider usage of Pegasus-based workflows inside the collaboration **g.)** develop closer integration of data transfer solution and file catalogs used by the offline pipelines **h.)** Create a new Virgo web page which can be used more efficiently by collaboration members, as well as by external visitors **i.)** install and maintain a nightly build or a continous integration server which ensures that software packages are always in a usable, clean state in the repository **j.)** upgrade to more advanced revision control system, such as for example git **k.)** replace our current (CMT based) software configuration and installation to a solution which is compatible with Linux standars **l.)** set up file catalogs in Computer Centers which are compatible with ligo_data_find queries so that pipelines using these tools can run **m.)** simplify the various authentication methods used, set up and integrate a new authentication system which is compatible with the LIGO solution **n.)** make it possible for LIGO colleagues to submit jobs to Virgo resources more easily **o.)** examine the various cloud and virtualisation solutions and perform evaulation studies how those can be useful for data analysis activities

2 Previous Accomplishments (2013-2014)

2.1 Burst Working Group

The following papers have appeared in journals in the past year:

S5 Long-duration GWB GRB search [1]: We searched for long-duration (10-1000 s) GW transients associated with long GRBs detected by the *Swift* satellite during the LIGO S5 run. No candidate events were found, and distance exclusion limits were set using signal models for accretion disk instabilities. The search was developed within the Stochastic group. This paper has been published in PRD.

S6/VSR2,3 optical image analysis [2]: We present the results of a search for optical transients associated with eight candidate GW events identified by LIGO and Virgo during the S6/VSR2,3 runs. No optical transient was found with a convincing association with any of these candidates, and none of the GW triggers showed strong evidence of astrophysical nature. This paper has been published in ApJS.

S5,6/VSR1,2,3 cosmic strings [3]: We present a search for GWs from cosmic string cusps in data collected by LIGO and Virgo between 2005 and 2010, with over 625 days of live time. We find no evidence of GW signals from cosmic strings. From this, we derive new constraints on cosmic string parameters, which complement and improve existing limits from previous searches for a stochastic background of GWs from cosmic microwave background measurements and pulsar timing data. This paper has been published in PRL.

S6/VSR2,3 IMBHB [4]: We present a search for gravitational waves from merging intermediate mass black hole binaries (IMBHB) performed on data from the S6/VSR2,3 runs. No significant candidate was found. We placed upper limits on the coalescence-rate density of non-spinning IMBHBs with total masses between 100 $M_\odot$ and 450 $M_\odot$ and mass ratios between 0.25 and 1. This was a *joint Burst-CBC effort*. This paper has been published in PRD.

Astrowatch/GEO GRBs [5]: We present a search for gravitational waves associated with associated with 129 GRBs observed satellite-based gamma-ray experiments between 2006 and 2011. This search used data from GEO600 and one of the LIGO or Virgo detectors to analyse GRBs not included in previous searches. No significant candidate was found. This paper has been published in PRD.

2.2 Compact Binary Coalescences Working Group

The following paper has appeared in a journal in the past year:

Parameter estimation paper [6]: Full authorlist methods paper on comparison and validation of CBC parameter estimation with Nested Sampling, MCMC, and MultiNest using S6/VSR2-3 data. This paper has been published in PRD.

NINJA-2 paper [7]: Results from a blind injection challenge whereby hybridized numerical relativity waveforms were injected into data that was recolored to early advanced detector era LIGO and Virgo noise curves. This was a *joint effort involving Numerical Relativity scientists*, external to LV Collaborations, and *CBC and Burst groups*. This paper has been published in CQG.

S5-S6/VSR2-3 ringdown search [8]: Intermediate mass black holes were searched for by looking for possible ringdown signals. No significant events were found, and mass dependent upper limits were placed. This paper has been published in PRD.

S5/S6 IPN GRBs [9]: Triggered search for GW associated with 223 GRBs found by the InterPlanetary Network. No GW signals were found, and where possible lower bounds on the distances were set. This was a *joint CBC-Burst effort*. This paper has been published in PRL.

2.3 Continuous Waves Group

The following papers have been published or submitted for publication in journals since May 2013:

S6/VSR2,4 targeted [10]: Targeted search in S6 and VSR2-4 data for nearly 200 known pulsars "Gravitational waves from known pulsars: results from the initial detector era". This paper has been published in ApJ.

S5 directed [11]: Directed search in S5 data for unknown isolated neutron stars near the galactic center "A directed search for continuous Gravitational Waves from the Galactic Center" This paper has been published in PRD.

S5 Hough all-sky [12]: All-sky search in S5 data for unknown isolated neutrons stars "Refinement and application of a Hough all-sky search for continuous gravitational waves to the full LIGO S5 data". This paper has been published in CQG.

VSR1 F-statistic all-sky [13]: All-sky search in VSR1 data for unknown isolated neutrons stars "Implementation of an F-statistic all-sky search for continuous gravitational waves in Virgo VSR1 data". This paper was submitted to CQG.

S6/VSR2-3 TwoSpect all-sky binary [14]: All-sky search in S6 / VSR2-3 data for unknown binary neutrons stars "First all-sky search for continuous gravitational waves from unknown sources in binary systems". This paper was submitted to PRD.

In addition, the following searches were completed and are under collaboration review:

1. All-sky search in S6 data for unknown neutron stars in binary systems.
2. Narrow-band search in VSR4 data for the Crab and Vela pulsars.
3. Directed search for unknown isolated neutron stars in 10 supernova remnants.

The following searches completed production running and now are in post-processing:

1. All-sky broadband search in S6 data for unknown isolated neutron stars (Powerflux pipeline).
2. All-sky mid-frequency search in S6 data for unknown isolated neutron stars (Einstein@Home pipeline).
3. All-sky low-frequency search in VSR2/4 data for unknown isolated neutron stars (Frequency Hough pipeline).
4. Spotlight search in S6 data for unknown isolated neutron stars in a spur of the Orion spiral arm.
5. Directed search in S6 data for isolated neutron stars in globular cluster NGC 6544.

Finally, the following major CW group efforts are also under way:

- An Einstein@Home search in S6 data for unknown isolated neutron stars in supernova remnants.
- A 4-stage mock data challenge for detection of isolated neutron stars, where targeted, directed and all-sky pipelines all participated. Stage 3 is nearing completion.
- A mock data challenge for detection of the low-mass X-ray binary Scorpius X-1. Six pipelines are participating. Results are expected in late spring 2014.

2.4 Stochastic Group

The following stochastic searches has been approved for journal submission (following final internal circulation):

S6-VSR23 isotropic paper: "Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009-2010 LIGO and Virgo Data." (LIGO-P1300154)

S5 H1H2 paper: "Searching for stochastic gravitational waves using data from the two co-located LIGO Hanford interferometers." (LIGO-P1000112)

In addition, following stochastic searches are nearing completion:

- STAMP all-sky search (joint with burst).
- S6-VSR23 directional search.

2.5 Detector Characterization

LIGO:

Publications: The following projects were completed and are now under collaboration review.

1. Characterization of the LIGO detectors during their sixth science run.
2. Environmental Influences on the LIGO detectors during the 6th science run.

Intangibles: Our most important accomplishments in the past year have been intangible: forging better ties with instrumentalists and building knowledge, creating sub-groups on data quality and instrumental characterization, and building documentation and infrastructure related to the aLIGO subsystems and data quality.

Summary pages: We have set up a summary page service at LHO and LLO that presents key figures of merit and state information from the aLIGO detectors and subsystems and figures of merit for the astrophysical searches. These automated pages will serve as the primary resource for data quality investigations.

Online Detector Characterization (ODC): A front-end real-time status monitoring system has been implemented and tested for the pre-stabilized laser, input mode-cleaner, and seismic isolation subsystems.

Transient monitoring: We set up improved automated transient monitoring tools (Omicron and Excess Power), integrated these with correlation-finding tools (hveto and UPV), and used these to identify sources of transient noise in the LIGO Engineering Runs.

Physical Environmental Monitoring System installation: A large fraction of the Advanced LIGO Physical Environmental Monitoring System was installed and tested at the LIGO Hanford Observatory.

LIGO Channel Activity Monitor (LigoCAM): The Advanced LIGO detectors record hundreds of thousands of channels, and many of these are used in servo control loops that are important to the performance of the instruments. A new monitor was developed to automatically detect large changes in the behavior of recorded channels, due, for example to sensor failure or cables being unplugged, reducing the time and effort required to find "dead" channels.

Source of combs in seismic isolation system identified: Together with on-site instrument scientists, we helped identify that not-synced oscillators used for the seismic isolation systems in aLIGO were responsible for strong combs of periodic lines in the platform motion. Because this problem was identified early, it will be completely fixed before the first aLIGO science data is taken.

Virgo:

Omicron online: A faster and improved version of the Omega transient signal search algorithm has been set up and tested with LIGO and Virgo data. It provides triggers output in ROOT files and is based on a Virgo library called GWOLLUM, dealing with data quality segments and triggers. Specific web pages have been developed to allow in-time monitoring of the triggers. Omicron has been adapted to run online within the Virgo online software environment. This online version of Omicron has been tested successfully during the year 2013.

Improved UPV algorithm: The Use Percentage Veto (UPV) algorithm has been improved to take into account the frequency dependence of the use percentage. Run over the triggers provided by Omicron, it will be the main online vetoes provider. At the beginning of 2014, tests have shown that a hierarchical version of the UPV algorithm allows good veto efficiency while using only a few auxiliary channels. Monitoring tools are provided along with UPV, showing either the efficiency and deadtime of vetoes of a given auxiliary channel or a matrix of UPV results made over several auxiliary channels. This version of UPV provides also useful information about glitch frequency up-conversion.

Data Quality Segments Database (DQSEGDB): During year 2013, we have defined, documented and set up a common LIGO-Virgo database which will contain all the data quality flags produced online and offline in Advanced Detectors. The structure of this mySQL database has been defined, a server and a client have been set up and are currently under test. A web interface is in preparation.

Monitoring tools: A large upgrade of the dataDisplay, an offline and online monitoring tool, has started at the end of 2013. A Data Quality flags performance estimator is available since April 2013 and several monitoring web pages based on UPV and Omicron results are available since the end of 2013.

3 Search Plans for the Advanced Detector Era

We collect here the plans formulated by the astrophysics search groups for Advanced LIGO and Advanced Virgo. These plans connect the science case, the criteria for publication, software development plans and resource needs for each search. The goal of these plans are transparent information sharing and the ability to look at the broad picture of the science we are pursuing as a Collaboration. Search plans descriptions and progress reports will allow expert, but not-too-involved, eyes to review the plans and flag possible effort duplications across groups and/or weak science cases.

Sections §3.11–§3.22 were appended in November 2014, in order of submissions.

Contents

In addition, the search groups are involved in exploratory work which is meant for search plan but is not as mature, and in some cases concluding searches of 2005-2010 data; those activities will be described in §4.

3.1 All-Sky 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.

We present the plans for a search for gravitational-wave bursts using the advanced LIGO - advanced Virgo network (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.

A key part of this search is to estimate 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 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 burst searches have focused on signal durations of up to approximately one second, the advanced detector search will extend to durations of order 100 seconds or more. Potential sources of GW bursts include long 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 be as effective as 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]. For this reason, 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.

---

[1] We note that the rate of binary neutron star detections, which is the best-modelled GW source for the advanced detectors, is uncertain to three orders of magnitude; see §3.2.

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. [15]

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 search will be the reference used to benchmark other burst searches. 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) [16, 17, 18], X-PIPELINE [19], and STAMP [20]. Of these, only CWB has been used to date for all-sky 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) [21], which processes detector data streams separately ("incoherently") and looks for coincident transients; LAL Inference Burst (LIB), which also searches incoherently [22] to trigger a coherent MCMC followup; and BAYESWAVE [23], which performs Bayesian MCMC followups. In light of these activities, we currently envision the following:

1. The main search pipeline will be the updated version of coherent WaveBurst, CWB2G (see §3.1.5).
2. We intend to adopt a second pipeline with sensitivity overlapping that of CWB2G. The second pipeline will increase the robustness of the search by providing an independent analysis of the data. In particular, by being subject to different systematics, a second analysis is 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. Finally, we note that historically the direct comparison of two pipelines has spurred improvements in each. The proposed criteria for the adoption of a second pipeline are outlined in §3.1.5.
3. We will explore the feasibility of searching for longer-duration bursts ($\sim$100 s), which is motivated by possible sources such as instabilities in accretion disks or proto-neutron stars. This extension of the search parameter space may be done in CWB2G or by the second pipeline.

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 §3.1.5.

All science data with two or more detectors operating at reasonable sensitivity will be searched, to maximise chances of a detection, since a minimum of two detectors is required to reject background noise fluctuations. The search will identify and rank candidate events, using GW correlation tests and data quality checks to reduce the noise background. The time-lag method will be used to assess the background and thus the significance of candidate events. As in S6/VSR2,3, an online minute-latency analysis will be used to provide alerts of significant events to electromagnetic observatories. The definitive analysis will be performed offline using the best calibration, data quality, and other relevant information available.

Preliminary offline analyses will also be performed on a regular basis using online data quality and calibration information, to refine the low-latency investigations on event reconstruction, confidence, and detection efficiency. Intermediate offline reruns will likely be needed when significant upgrades of the calibration and data quality are released.

3.1.4 Publication plan

We expect to publish one all-sky observational results paper following each advanced detector science run. Based on the current run schedule [24], 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 §3.1.5.

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 §3.1.5. 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.

3.1.5 Technical requirements and development plan

Main pipeline:
The CWB2G pipeline is under development, starting from the offline version and porting with minimal differences to the online version. 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).

CWG2G development is now in the late stages. An intermediate online version of CWB2G was tested in Engineering Run 5. The final version will be tested in ER6. The offline CWB2G was finalized in late 2013. The pipeline review is underway and should be completed by 2014. Ongoing tests include the re-analysis of all initial LIGO-Virgo data by summer 2014. 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 density statistic [3].

Second pipeline:
As discussed in §3.17.3, several additional pipelines are under active development: STAMP, X-PIPELINE, EXCESSPOWER, LIB, and BAYESWAVE. STAMP has been used to look for long-duration GWs from gamma-ray bursts. A feasibility study of all-sky analysis of initial LIGO data is currently underway, and advanced time-frequency clustering techniques are being explored [20]. X-PIPELINE has been used for a number of externally triggered burst searches, and is currently being extended to an all-sky pipeline using a fast correlator engine based on spherical radiometry [25, 26]. 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 [27]. The 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 CWB as well as simulated signals, both in ER5 and S6/VSR2,3 data.

The qualification of a second pipeline will rely on a performance comparison against CWB2G. This will involve the analysis of all S6/VSR2,3 data with 2 or more detectors operating, over the full 60 Hz – 4000 Hz band, using the same simulations, data quality, and calibrations as in the published analysis [15][4]. Emphasis will also be placed on how the second pipeline extends the science reach of the all-sky analysis, such as through superior sensitivity in some areas of parameter space (e.g., for long-duration signals) or better sky localisation capability.

Mock Data Challenge (Simulation) infrastructure:
GRAVEN/BURSTMDC, the traditional burst infrastructure to produce software injections, will be maintained both to produce simulation frames for pipelines lacking internal simulation engines, and as a reference with which to validate simulation engines which are internal to other pipelines. Its present capabilities are satisfactory, with no significant further development planned and almost all of its infrastructure reviewed.

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[5]. 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.

---

[3] See http://www.virgo.lnl.infn.it/Wiki/index.php/Main_Page.

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

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

3.1.6 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.

Review:
The CWB2G review has started, and is expected to take approximately one year. Given the relevance of the pipeline the review team is made up of 5 reviewers. After addressing the pipeline review, the plan is to review the re-analysis of past S5-6/VSR1-3 data, and of the new methods for the future all-sky search. After the successful qualification of the second pipeline, a similar review effort will be needed.

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 [28]. 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 to be in the range $0.7\,\mathrm{M}_\odot$ to $2.7\,\mathrm{M}_\odot$ with a mean mass of $\sim 1.4\,\mathrm{M}_\odot$ [29], though the lower value of $0.7\,\mathrm{M}_\odot$ comes from an imprecise measurement of a single system that is also consistent with a higher mass. NSs in BNS systems have a narrower observed mass distribution of $(1.35 \pm 0.13)\,\mathrm{M}_\odot$ [29]. 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 [30, 31]. Lower mass Fe cores are predicted to lead to NSs with masses almost as low as $1\mathrm{M}_\odot$ [32].

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 [33], however the fastest known NS spin is 0.4 [34]. 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 [35], whereas a search for more rapidly spinning systems would require the use of waveforms that capture the effect of spin [36, 37, 35]. BNS waveforms for systems with small spin are well modelled by post-Newtonian theory [38] at seventh order beyond the leading-order orbital phase [39].

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 [40]. 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 [41, 42]. Low-latency pipelines [43, 44, 45, 46, 47, 48] 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 [49], 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 [50]. 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 [51]. With a few tens of detections, it may be possible to significantly constrain the EOS by combining information from many coalescence events [52], although this requires the incorporation of waveform developments to better model effects beyond the level of current post-Newtonian templates to avoid systematic errors [53, 54]. In cases where an EM counterpart can be identified, further information can be used to understand the physics of the merger [55, 56].

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 [57]. Even the binary pulsar J0737-3039 [58] 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 [59, 60]. Such studies will provide stronger tests of general relativity than any that have been performed to date [61].

3.2.3 Search Description

During the initial-detector era we developed algorithms to search for BNS signals coincident in a detector network [62, 63, 64, 65, 66, 67]. The BNS search is conducted with matched filtering [63, 67, 45], 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 [49]. Search results were published in a series of papers for the six LIGO and three Virgo "science runs": S1 [68], S2 [69], S3 and S4 [70], S5/VSR1 [71, 72, 73] and S6/VSR2,3 [74]; 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 [69, 70, 71, 72, 73, 74]. 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. For this reason, we will conduct a BNS-only search in the future. 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 entire plausible NS mass range of $0.9\,\mathrm{M}_{\odot}$ 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 [34, 33]. Although such low spins are insignificant from a search point of view [35], 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 [45, 46, 47, 48] and MULTI-BAND TEMPLATE ANALYSIS (MBTA) will run on distinct computing facilities to provide redundancy.
2. A traditional offline search pipeline, AHOPE, 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 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 [75], 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 [76] 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 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.

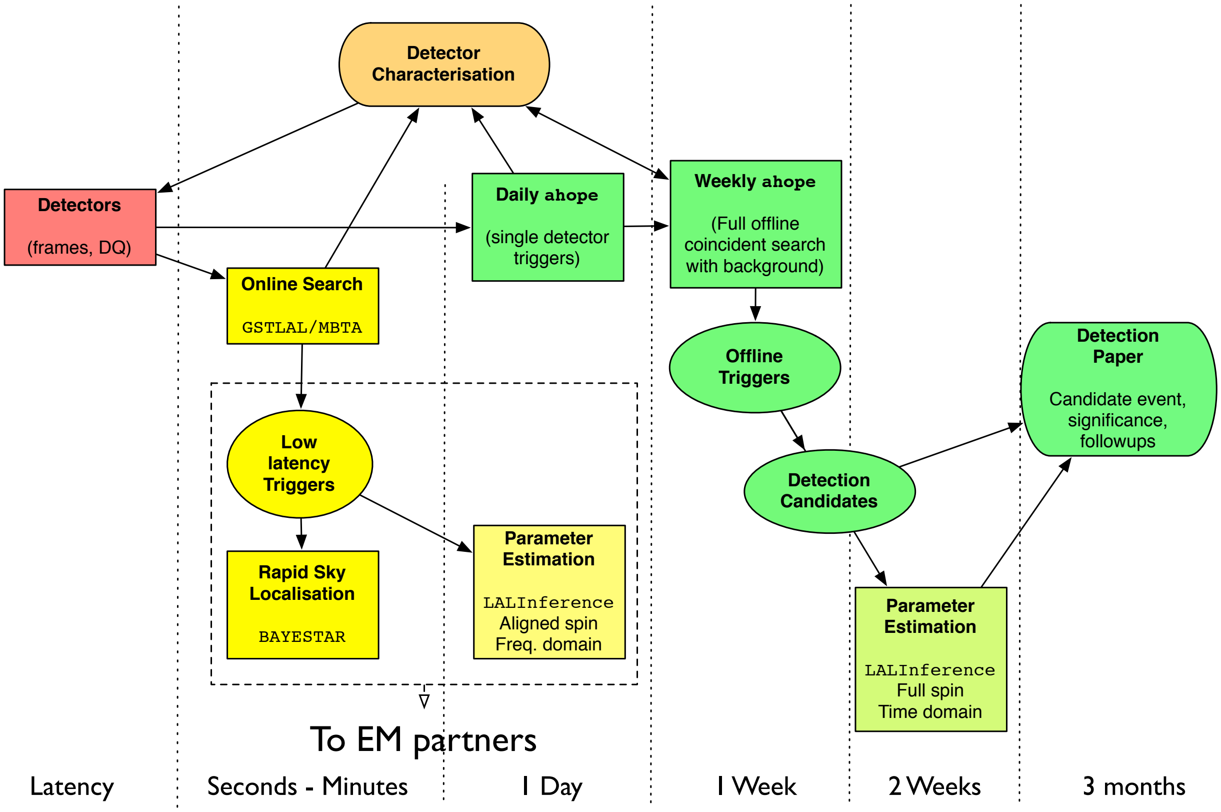

Figure 1: Overview of BNS search, identifying the core components for online and offline analyses, their latencies, and the interdependency between them. The low-latency components (in yellow) will provide rapid triggers for EM followup, whereas the higher latency part (in green) will provide final results 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. `ahope`, an advanced LIGO offline compact binary analysis pipeline. `LALInference`, a parameter estimation pipeline written with in the LAL library.

3.2.4 Publication plan

In the past, we have published the results of BNS searches together with the results from NSBH and stellar-mass BBH [69, 70, 71, 72, 73, 74]. 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 [75] 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 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. 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 detections, 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.

3.2.5 Technical requirements and development plan

The LSC and Virgo have proposal for early observing run scenarios is described in Ref. [78]. 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/adVirgo, 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

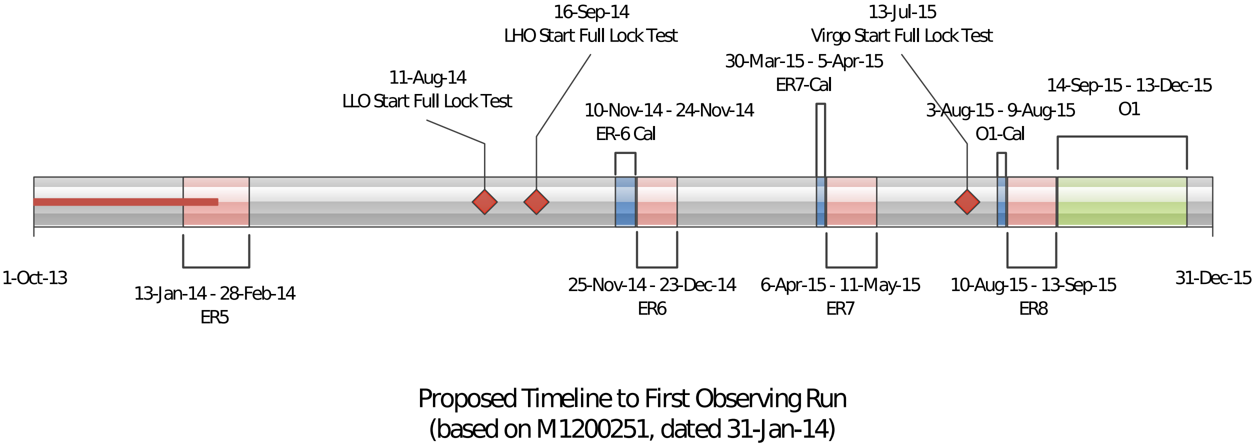

Figure 2: Proposed time line to first observing run (available as LIGO-G1301309 [77]).

close coupling between the search group and the detector group is critical for a timely completion of the analysis [74]. 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 $\chi \leq 0.05$ [35]. 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 $\chi \leq 0.4$ [35]. Full parameter estimation requires the inclusion of spin effects in the waveforms and template banks using models such as SpinTaylorT4. Required actions include:

- an MDC to test the efficacy of aligned-spin BNS searches in real data as a followup to [35]
- 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.

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 [67], 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, more modular offline pipeline called AHOPE, as well as an accompanying python analysis library, PYCBC, 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 AHOPE/PYCBC.

MBTA and GSTLAL-CBC are specifically designed for advanced detector BNS searches. GSTLAL-CBC review has begun with a review committee convened. Preparations for the AHOPE/PYCBC review have begun within the CBC group and we will request a review committee be convened in Spring 2014.

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. Required actions include:

- a common MDC to assess the sensitivity of the pipelines and ensure they meet performance targets.
- a method to combine results from multiple pipelines and/or searches on overlapping parameter spaces.

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. We need to decide:

- if the veto category system used in S5/S6 will go forward.
- how to best encode real-time DQ information into the online h(t) files.
- how to veto different types of glitches: do we veto triggers or excise the glitch from h(t)?

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:

- Understand false alarm probability uncertainties, and their implications for detection confidence.
- Understand the impact of removing zero-lag coincidences from the background vs. leaving them in.
- Understand systematic and statistical uncertainties in rate estimations and their behavior in the presence of one or more signals.
- We have code to compute null-result rate upper limits, and recently new techniques have been implemented to estimate rates both in the null case and in the detection case. The latter need to be validated and reviewed.

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. We need to:

- formalize the use of GraceDB metadata ("group", "type", etc.).
- stabilize and document the data product formats for releasing information on candidate events.

Parameter Estimation (PE) and Testing GR.
To accurately and rapidly extract the parameters of a signal, a Bayesian inference package called LALINFERENCE has been developed. This package can recover the physical parameters of the signal, and it was reviewed and successfully used in previous science runs [49]. Further developments are required to fully exploit the advanced detector data. Required actions include:

- Increase parameter estimation performance below 40 Hz.
- Ensure robustness against calibration errors and noise artifacts.
- GR tests will follow using a population or loud signals to test different models. Testing GR pipelines must be tuned to compute background in realistic data.

3.2.6 Resources

Calibration. We will require calibrated data over the sensitive band of the detectors (from the low-frequency wall to a few kHz). An initial estimate is that calibration accuracy to 5% in amplitude and $5\,\mathrm{deg}$ in phase will be adequate for early science runs where any signals are likely to be low in amplitude. We would like to have calibration to the required accuracy prior to the start of any science run.

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 ER7. 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

Compact binaries of black holes (BH) resulting from the collapse of massive stars are one of the most promising sources for the first detection of gravitational waves; we may see tens of signals per year in the Advanced detector network. Black holes formed by stellar collapse could be up to tens of times more massive than the Sun. We are searching for gravitational waves from systems where two such black holes orbit round each other and finally merge into one.

The first detection of a stellar-mass binary black hole (BBH) would prove the existence of a so far unobserved class of binary, constituting a major discovery. The information in the mass parameters of the first few detections could distinguish between astrophysical models of compact object populations. Knowledge of merger rates from current astrophysical observations of black holes and theoretical modelling is extremely uncertain, so improved upper limits can contain new information about how massive stars and black holes evolve.

Some of the questions we aim to answer by searching Advanced detector data 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, and how does the collapse of one 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 simulations of the formation and evolution of massive binary stars (population synthesis) and through dynamical modeling of dense stellar clusters. BBH coalescence rates are very uncertain: the rate estimates for a $(10+10)\mathrm{M}_\odot$ BH binary range from 10^{-4} to $0.3\,\mathrm{Mpc}^{-3}\mathrm{Myr}^{-1}$ [28]. Current observational rate upper limits from LIGO and Virgo are only a factor $\sim$4 above the upper end of this range [74].

The mass distribution of Galactic stellar mass BH is estimated in [79, 80, 81], and X-ray observations yield BH masses $5 \leq M_\bullet/\mathrm{M}_\odot \leq 20$, confirmed with dynamical mass measurements for 16 BHs. An apparent lack of BH ("mass gap") in the range 3–5 $\mathrm{M}_\odot$ [79, 80, 82] has been ascribed to the supernova explosion mechanism [83, 84]. Objects initially forming as neutron stars may also accrete material during a common envelope phase and collapse to low-mass, near-maximally spinning black holes [85]. The highest well-determined neutron star masses are close to $2\,\mathrm{M}_\odot$ [86, 87], thus binaries may contain BH upwards of this mass.

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 [88, 89]. 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 $\sim 80\,\mathrm{M}_\odot$ [84, 90]. Common envelope binary evolution [91] may reduce the maximum expected component mass and total mass to $\lesssim 100\,\mathrm{M}_\odot$ [92], however stellar BH with mass above $100\,\mathrm{M}_\odot$ are conceivable [93], overlapping the range associated with IMBH formed by repeated mergers. There is no direct observational limit on BBH mass ratios; current models of field binary evolution favor $q \lesssim 5$ [92] (field binaries are those where progenitor stars evolve into black holes in isolation, without interaction with other bodies).

X-ray observations of accreting black holes indicate a fairly uniform distribution of spins over the entire range allowed by general relativity, $0 \leq S/m^2 \leq 1$ [94, 95, 96, 97, 98, 99, 100]; both low ($\sim$0.1) [101] and high (> 0.85) values [102] are represented. The microquasar XTE J1550-564 [103] and population synthesis models [104] indicate small spin-orbit misalignment in field binaries. For massive field binary progenitors, the common envelope phase and mass transfer [105] are expected to cause strong correlations between spins and masses of the two BHs in field binaries [106]. 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 [107]. A detection with moderately high SNR may yield strong evidence for at least one spinning component, independently establishing the existence of black hole spin in GR.

Several detections will begin to establish an observational distribution over chirp mass, and, possibly, spin or mass ratio, further constraining models of BBH populations [108]. 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 [109, 110] 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 PN expansion, numerical relativity and BH perturbation theory. Establishing a deviation from the predictions of GR, for example the Kerr description of BH, would be a major discovery.

3.3.3 Search Description

Pipeline In initial LIGO/Virgo the "ihope" pipeline [67] 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 $M \leq 25\,\mathrm{M}_\odot$ [71, 72, 73, 74], and a "high-mass" pipeline using inspiral-merger-ringdown (IMR) templates to search for signals with $25 \leq M \leq 100\,\mathrm{M}_\odot$ [111, 112]. These searches yielded upper limits that were a factor $\sim$10 higher than optimistic rate predictions.

In [7] it was shown that the existing "ihope" pipeline with no additional changes would still be able to detect BBH signals in early advanced LIGO data. However, ihope cannot accommodate technical developments envisaged for Advanced searches: different methods of template generation including the use of spinning templates; handling large numbers of templates and high volumes of data output; better treatment of boundaries in parameter space; and more accurate modelling of the signal and noise event distributions (see *Technical requirements and development plan* below). Thus we plan to use "pycbc", a more modular and adaptable implementation of the coincident matched filter search, as code base for the analysis. A pycbc-based pipeline is also expected to run offline BNS and NSBH searches in Advanced data. Running an online low-latency pipeline such as gstlal-CBC is possible but not a priority as stellar BBH mergers are not expected to be accompanied by EM emission.

Mass and spin ranges BBH may have a very wide range of component masses. We consider 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 BH formed via NS accretion and collapse. The upper mass limit of the search will be primarily determined by the efficacy of the matched filters and of signal-based vetoes, as fewer cycles are seen in the detector. There is no clear division between the maximum possible stellar BBH mass and the lower end of the IMBH range. Based on previous experience, a search can be performed up to a total mass of at least $50\,\mathrm{M}_\odot$ without further development, such that the loss of sensitivity due to non-Gaussian noise is small; beyond this point, depending on Advanced data quality, more sophisticated vetoes and/or ranking statistics may be required to optimize the analysis. Mock Data Challenges (MDCs) comparing pipeline efficiency with IMBH searches using large sets of simulated signals will guide the eventual choice of upper mass boundary. BBH search output may be also used within a method of combining search pipelines to be developed for IMBH sources.

We may also restrict the maximum mass ratio if asymmetric templates with significant waveform uncertainties (beyond ∼10:1) contribute strongly to the noise background.

Possible black-hole spins cover the whole range $0 \leq \left|S/m^2\right| \leq 1$; pipelines using non-spinning templates may miss a significant number of BBH signals with large component spins. We are currently developing an analysis pipeline using aligned-spin templates to improve sensitivity to such signals. 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.

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. The same analysis pipeline may search for two or more different sources (e.g., NSBH and BBH). 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 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 [75, 113], we will consider events with false alarm probabilities $\gtrsim 4\sigma$ ($\approx 6 \times 10^{-5}$) 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.

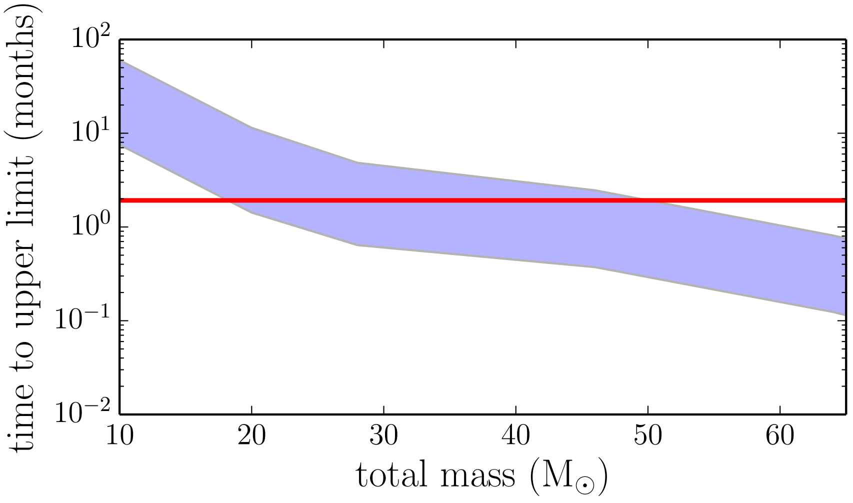

Figure 3: Amount of search time needed in early advanced LIGO to place upper limits that would constrain optimistic population synthesis models (“Submodel A” in [91]) of the rate of BBH coalescence, as a function of total mass. The red-line region indicates the estimated search time in the first observing run (O1) assuming it is 3 months long with a coincident duty cycle between Hanford and Livingston of 64% [114]. The blue-shaded region indicates the time to upper-limit assuming the optimistic (lower line) and pessimistic (upper line) “Early" PSDs in [114]. 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 $\gtrsim 20\,\mathrm{M}_\odot$.

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 should also be thoroughly tested and 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 false alarm probabilities $\lesssim 4\sigma$ but $\gtrsim 3\sigma$ (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 CBC 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” described in [91] would be ruled out for certain equal-mass, non-spinning systems with just 3 months of run time predicted for the first aLIGO and AdV 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 CBC 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.

3.3.5 Technical requirements and development plan

Waveforms

Accurate and computationally efficient waveform models for BBH coalescences are crucial to the scientific output of the search. They are used as matched filters for detection, and as simulated signals for tuning the pipeline parameters and event ranking statistic. When estimating the source parameters of a detected coalescence signal, waveforms are subtracted from the data to assess the relative likelihoods of different parameter values. Finally, in estimating astrophysical coalescence rates waveform models are used to quantify the sensitivity to the search to a realistic population of signals. 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. Requiring development and/or review in the short term are SEOBNRv2 [115], an aligned-spin waveform model usable over a broad range of component masses and spins; IMRPhenomP, a fast approximate representation of precessing binaries [116]; and "surrogates" using reduced order modelling, an approximation technique offering speedup up to a factor of $O(10^3)$ [117], to assist applications where computational cost is dominated by waveform generation. In the longer term SEOBNRv3, a model for generic precessing BBH, is to be implemented.

Analysis pipeline

The data analysis pipeline begins with the layout of matched filter templates over the parameter space to be searched. For non-spinning signals, existing template placement codes [62, 118] are able to cover the space limiting the largest possible loss in SNR due to bank discreteness to a predetermined value (usually 3%) [119, 120]. However we are pursuing an aligned-spin bank for Advanced data, which is expected to improve the search sensitivity to spinning systems; using aligned-spin template waveforms will yield up to 45% improvement in detection efficiency for systems with aligned spin, compared to a search that does not include spin effects [121]. In [35] a placement algorithm using a PN metric for aligned-spin BNS systems was described. Preliminary results for the BBH region show that this metric can achieve the required minimal match for the expected sensitivity curve in the first observing run of aLIGO (O1). The sensitivity of the detectors is expected to improve through commissioning with each successive observing run in aLIGO. Assuming that in O1 we will not be sensitive to signals below 30 Hz, we will need $\sim 100\,000$ templates to carry out an aligned-spin search. We estimate that we will need $\sim 250\,000$ and $\sim 630\,000$ templates if the lower frequency cutoff improves to 20 and 10 Hz in successive runs, respectively.

It has been shown that a search using templates including aligned-spin waveforms will yield up to 45% improvement in detection efficiency *for systems with aligned spin*, compared to a search that does not include spin effects [121]. Studies on BNS and NSBH systems indicate that a search using aligned-spin waveforms will also perform better at detecting generic binary black hole systems than a non-spinning search [35, 122]. Quantifying these effects is a current priority in the BBH group.

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 [116], and SEOBNRv3 [123] when available — expected end 2014). 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 [120, 119]. The effect of higher harmonics in spinning signals currently cannot be quantified due to the lack of suitable IMR waveform models.

Pipeline

Current searches regenerate template banks approximately every 40 minutes. This is done to capture changes in the sensitivity of the instrument over this timescale. Due to the slow generation time of SEOBNRv2 templates, this will not be feasible for Advanced detectors: generating all of the waveforms in a template bank on a single core would take between 1 and 10 days. We will investigate whether it is possible to use a fixed bank for much longer periods of time, perhaps even for an entire analysis interval ($\sim$ 1 week). If so, then we can parallelize waveform generation.

As with past searches, we demand that candidate events be coincident across multiple detectors; currently there are two methods available, "ethinca" [66] which allows template of different masses to be coincident, accounting for parameter correlations, and "exact" match which requires the same template in all detectors. MDC studies indicate exact match has a slight sensitivity advantage over ethinca in some cases, in particular high masses where the PN metric becomes increasingly inaccurate. Technical issues include that ethinca does not currently account for component spins; conversely, exact match requires unclustered triggers to be stored, significantly increasing the cost of performing time slide background estimation.

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 Hz, the shortest BBH templates ring down at $\sim$ 100 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 indicators of 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 CBC search is identifying a detection candidate with very low false alarm rate. Large numbers of time shifts may be performed to produce background samples comparable to zerolag, however the conditions for the resulting estimate to be reliable/useful for detection purposes need to be understood, and more computationally efficient methods of performing slides may be required. Developments should also be applicable to BNS and NSBH searches.

"Slideless" methods have been proposed which are more computationally efficient and may reach smaller false alarm probabilities over a given data set. These work by using the combined probability density functions of single detector triggers to estimate the probability of getting a multiple-detector trigger with a given significance [124, 125]. However it is not clear if their results are comparable to slides, particularly at high masses where the noise distribution and coincidence test are harder to model.

The efficiency of a search in non-Gaussian noise depends on the choice of ranking statistic, which is a function of the recovered parameters of candidate events. Previous searches have used simple, hand-tuned functions, however these will be far from optimal at high masses, where signal-based chisq tests [64] become ineffective. Variation over the parameter space has so far been dealt with by dividing into a small number of bins, however efficiency and background estimation should be improved if the variation can be modelled by a smooth function, which is likely for lower-mass systems.

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 [126] 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 [127]. 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 [128, 129, 130, 131, 132, 133, 117]. To be usable within the same computational constraints, surrogate SEOBNR models should be comparable to SpinTaylorT4 in the cost of generation.

Testing GR

Bayesian methods for testing whether the waveforms of any detected signals are consistent with the expectations from GR have been developed [134] and are demonstrated for BNS systems [61]. 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 GR with BBH systems is demonstrated. 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.

3.3.6 Resources

Calibration requirements The calibration requirements for BBH searches are similar, or less stringent, than for BNS since the range of frequencies spanned by BBH signals is the same or less. An initial estimate is that calibration accuracy to 5% in amplitude and $5\,\mathrm{deg}$ in phase from the low-frequency wall to $\sim 2\,\mathrm{kHz}$ will be adequate for early observing runs, where any signals are likely be low in amplitude. In order to meet the target schedule for releasing candidate detections, calibration should be in place 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 1 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 within DetChar and some from the CBC 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 CBC 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. Some codes have already begun the review process and others will start review in Summer/Fall of 2014. 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, DetChar 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 [82], and whether black holes preferentially have large spins [102]. For certain configurations the neutron star will tidally disrupt during merger and this can give rise to an observable electromagnetic counterpart. 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.

Acheiving 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 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 [135]. 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 $< 10\%$ [136, 137].

With ~ 10 NSBH observations it will be possible to more closely probe the underlying physics [138]. 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 [139, 140, 141, 142, 143] and could help to distinguish whether NSBH systems are predominantly formed by dynamical capture or stellar evolution of field binaries [144, 145, 146, 147]. Additionally, the magnitude of the spin of the BH would provide information on the amount of accretion experienced by the BH [101, 85, 148].

Fully general-relativistic numerical simulations of NSBH systems have been performed (for e.g. [149, 150, 151, 152, 153, 154]) 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 [155]. 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 [156, 155, 157]. 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 [158, 159]. Additionally, galaxy host identification of EM counterparts can allow us to better constrain their formation process, as with (the possibly identical) short GRB events [160]. 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 [142, 141].

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 [161, 162, 52, 54]. Investigation is needed to understand if this kind of study is better facilitated with NSBH signals than with BNS signals.

In the absence of a detection of NSBH GW signals, astrophysically relevant constraints can be placedon 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 [144, 145, 146, 147]. Though no NSBH systems are known to exist, one likely progenitor has been observed [163]. 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 Mpc^{-3} Myr^{-1} [164]. A “pessimistic” estimate is given as 6×10^{-4} Mpc^{-3} Myr^{-1} and an “optimistic” estimate as 1 Mpc^{-3} Myr^{-1} [164]. 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 $\sim 1.4\,\mathrm{M}_\odot$ [29, 165], 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 $(1.35{\pm}0.14)\,\mathrm{M}_\odot$ [29, 165]. 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 [166, 167]. Lower mass Fe cores are predicted to lead to neutron stars with masses almost as low as $1\,\mathrm{M}_\odot$ [32].

The mass distribution of Galactic stellar mass black holes is estimated in [168, 79, 169], 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$\,\mathrm{M}_\odot$ [168, 82] has been ascribed to the supernova explosion mechanism [83, 84]. 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 [88, 170]. 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 $\sim 80\,\mathrm{M}_\odot$ [84, 171]. Common envelope binary evolution may reduce the maximum expected component mass and total mass to $\lesssim 100\,\mathrm{M}_\odot$ [91], however stellar BH with mass above $100\,\mathrm{M}_\odot$ are conceivable [93], 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 [94, 95, 96, 97, 98, 99, 100]. Some black holes observed in X-ray binaries have very large dimensionless spins (> 0.7), while others could have much lower spins (~ 0.1) [101]. 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 [102]. Black hole spins are only constrained by the relativistic bound $\chi = S/m^2 \leq 1$. The microquasar XTE J1550-564 [103] and population synthesis models [104] 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 [138], 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 [172] It is possible for neutron stars to be spun up through accretion to much higher spins that will persist until merger [34], 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 [173, 63, 174, 175, 176, 127]. 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 [177, 178, 179]. Search results were published in a series of papers: S3 and S4 [180], S5/VSR1 [181, 72, 182] and S6/VSR2,3 [74]. 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 still needed however, including the inclusion of spin effects in search templates and extending search and parameter estimation pipelines to incorporate the longer waveform templates required.

To date, the NSBH search has been performed in conjunction with CBC searches for BNS and stellar-mass BBH [180, 181, 72, 182, 74]. 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 envisage 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 [173]. 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 [38]. 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 $0.2 - 300\mathrm{yr}^{-1}$ [183]. This assumes an isotropic in direction and uniform in magnitude spin distribution and assumes Gaussian noise. 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. 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 [184]. 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 [136, 185]. 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 [186, 183]. We will use these developments for placing template banks for NSBH searches with Advanced LIGO (aLIGO) and Advanced Virgo (AdV). 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. 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 [183]. However, previously tested algorithms incorporating precession increased false alarm rates and thus did not show an overall increase in search sensitivity [179, 187]. Therefore, our initial search proposal is for a search using only aligned-spin template waveforms (ie. not including precessional effects), while we develop and assess potential new methods for efficiently incorporating precession.

The search will target NSBH inspirals and mergers comprised of component masses spanning the entire plausible NS mass range of $0.9\mathrm{M}_{\odot}$ to $3\mathrm{M}_{\odot}$. Note that while the error bars on observations of NS masses extend as low as $0.7\mathrm{M}_{\odot}$ [29], 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 $3\,\mathrm{M}_{\odot}$. The reliability of waveform models, along with studies and tests on simulated data, will guide the choice of the upper bound on the BH mass. Current NSBH waveform models with high-mass black holes ($\gtrsim 15\,\mathrm{M}_{\odot}$) are not yet reliable enough to make meaningful tests of search performance [188]. However, ongoing work from the numerical relativity and analytical relativity community, focused on this region of parameter space, should yield improved waveforms, which will allow us to assess the efficacy of searches with black hole mass $> 15M\,\mathrm{M}_{\odot}$. 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 [183].

The NSBH search will be conducted using three classes of pipeline. One will be a traditional offline search pipeline, using XSEDE computer facilities. 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. 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` [43] and LSC-based GSTLAL-CBC [189, 46, 47, 48], 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 [74], however they can be computationally inefficient. Therefore new “slideless” techniques are being developed and will be assessed for use in NSBH searches [190, 125].

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 [127]. To check for consistency, at least two methods will be used to sample the high dimensional parameter space: Markov Chain Monte Carlo [191, 192] and nested sampling [193]. Additionally, several waveform models will be used, including the precessing post-Newtonian models used in [183, 188], the best-available spinning effective one-body (EOB) [194, 195, 123], and so-called “phenomenological” spinning inspiral-merger-ringdown (IMR) waveforms [196, 116]. This redundancy will check for consistency among waveform models and sampling methods, as well as quantify any potential systemic biases due to waveform model uncertainty. We note that 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 [197]. 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 Publication plan

In the past, we have published non-detection papers of NSBH searches together with the results from BNS and BBH searches [180, 181, 72, 182, 74]. 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 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 [114]. 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 occured near the end of the LIGO’s S6 and Virgo’s VSR3 [74, 75]. 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 [75] 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 If 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 [113, 198]. 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 NSBH 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 (ie., 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.

3.4.6 Technical requirements and development plan

A proposed roadmap for the development of the aLIGO and AdV detectors towards their design sensitivity is given in [114]. 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 have begun to use real detector data and we expect this to increase as aLIGO and AdV are commissioned. 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 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 BNS 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 [114].

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 will be 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 ($\|\vec{S}_{\rm BH}\|/m_{\rm BH}^2 \gtrsim 0.5$) [188]. From we imply that a similar discrepancy exists between the models and waveforms produced by Nature [188]. 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” [123] and “IMRPhenomP” [116] 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 and IMRPhenomP precessing IMR waveform models.
- 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: `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 [177, 199, 179, 187], though such methods have yet to demonstrate an improvement in sensitivity on initial LIGO and Virgo data [179, 187]. These methods should be applied and tested on the ERs and MDCs, but aligned-spin search development is the priority.

Early observation runs are only expected to have good sensitivity above ~ 30 Hz, so initial testing should focus on this region. As the advanced detectors are commissioned during the initial years of operation the pipelines should be tuned for running with low-frequency sensitivity. This will be a priority issue for the BNS group, and the NSBH group will utilize developments in that context.

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 `ahope`, `MBTA` or GSTLAL-CBC 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 NSBH pipelines will report high significance triggers to the same alert system for EM observers used by the BNS search. 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 ariving in the search pipeline. Parameter estimation pipelines have already demonstrated the ability to analyze candidate events in initial LIGO and Virgo data [127]. 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 (for all except EOB models, which could take months). Computational cost is not an immediate barrier for NSBH parameter estimation with any of the currently available waveform families except for the EOB models. 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). Because the aligned-spin IMR SEOBNRv1 [194] model is prohibitively expensive for use with parameter estimation, we must consider not just accuracy but also computational efficiency when implementing improved waveform models. Efforts are underway to use reduced order modeling to create computationally-efficient surrogate models [128, 129, 130, 131, 132, 133, 117] for expensive waveform families such as SEOBNRv1, 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 [200] and are demonstrated for BNS systems [61]. 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.

3.4.7 Resources

Calibration requirements NSBH searches will require calibrated data over the sensitive band of the detectors (from the low-frequency wall to a few kHz). An initial estimate is that calibration accuracy to 5% in amplitude and $5\,\mathrm{deg}$ in phase will be adequate for early observation runs where any signal we are lucky enough to see will very likely be low amplitude. Identifying calibration requirements is one of the priorities for future MDCs in the BNS group, and will apply directly to NSBH searches. This will require work with the calibration team prior to the run to ensure that realistic models are being used to test sensitivity to calibration errors. Following this, detailed feedback will be provided to the calibration team. Calibration of the data should be at the required accuracy prior to the start of any observation run.

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 arise, 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 In order to perform a search for NSBH signals in Advanced LIGO data in the zero-detuned high-power optical configuration, and a lower matched-filter frequency of 10Hz, requires roughly 2100000 templates to cover the space with a minimal match of 97%. This assumes the parameter space described previously; namely black hole mass $\in (3.0, 15.0)\,\mathrm{M}_{\odot}$, neutron star mass $\in (0.9, 3.0)\,\mathrm{M}_{\odot}$, black hole dimensionless spin $\in (0, 1)$ and restricted to be aligned, or anti-aligned with the orbit and neutron star spin $\in (0, 0.05)$. We do not expect Advanced LIGO to have this low frequency sensitivity in early observation runs [114]. If we instead use a lower matched-filter frequency of 20 and 30 Hz we find that 560000 and 210000 templates respectively are required. 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 $\in (1.0, 3.0)\,\mathrm{M}_{\odot}$, 670000 templates would be required for a 10Hz lower matched-filter frequency.

Computational Infrastructure and Services The following computational infrastructure and services must be available, tested and working to successfully perform searches for BBH 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 search for precessing binaries [177]. 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. [177] 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 aligned 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 neutron star bank. Significantly compounding this is the fact that the filters described in Ref. [177] 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. [177] 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 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 processing 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 ~ 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. Therefore, to perform exhaustive parameter estimation on a single NSBH event will require 1,600 CPU cores for 16 days. In the unlikely event of multiple NSBH events within this period it would be preferable to have 1,600 cores for *each* event.

Review The review of several search components will overlap between source groups: i.e., search algorithms (3 pipelines), waveforms, rate 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. A collection of joint short GRB 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 search. 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 also are expected to be GW sources: compact binary coalescences for short GRBs, with gamma-ray duration < 2 s and harder spectra, and collapsars for long GRBs, which typically last > 2 s and have softer spectra. The cosmological distance of GRBs remains the principal challenge for a GW detection associated to GRBs with advanced detectors.

Cosmological short GRB 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 in GW with advanced detectors to ~ 400 Mpc for NS-NS or ~ 1 Gpc for NS-BH at a rate of $\sim 1\,\mathrm{yr}^{-1}$ for each [201]. 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).

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 a full understanding. 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 $\sim 1\,\mathrm{yr}^{-1}$ [202, 203]. The hypothesized population of low-luminosity long GRBs [204] would also have a rate of $\sim 1\,\mathrm{yr}^{-1}$.

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 [205] 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-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 [206].

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 [207]. 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 can make few % measurement of the Hubble constant H_0 [208]. 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 [209] 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 $\sim 40\%$ [209].

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. 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 incorporate features in the search that would not 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 large, the NS is swallowed whole and no GRB can be formed. Taking this into account may allow us to reduce the number of templates searched.

Both searches — for short and long GRBs — will be promptly initiated, within about five to ten 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 event, most likely within minutes, 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 has the advantage of being very low latency, and 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, long or short, and the COHPTF binary merger search on any GRBs that are identified as short, or possibly short. These searches will be initiated as soon as possible after the GRB alert, and we will endeavor 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. The final results will be available within a few days.
- We will run an archive 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 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 [210], 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) We plan to publish with a goal of ∼2 weeks after the GRB. This would apply to a (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.

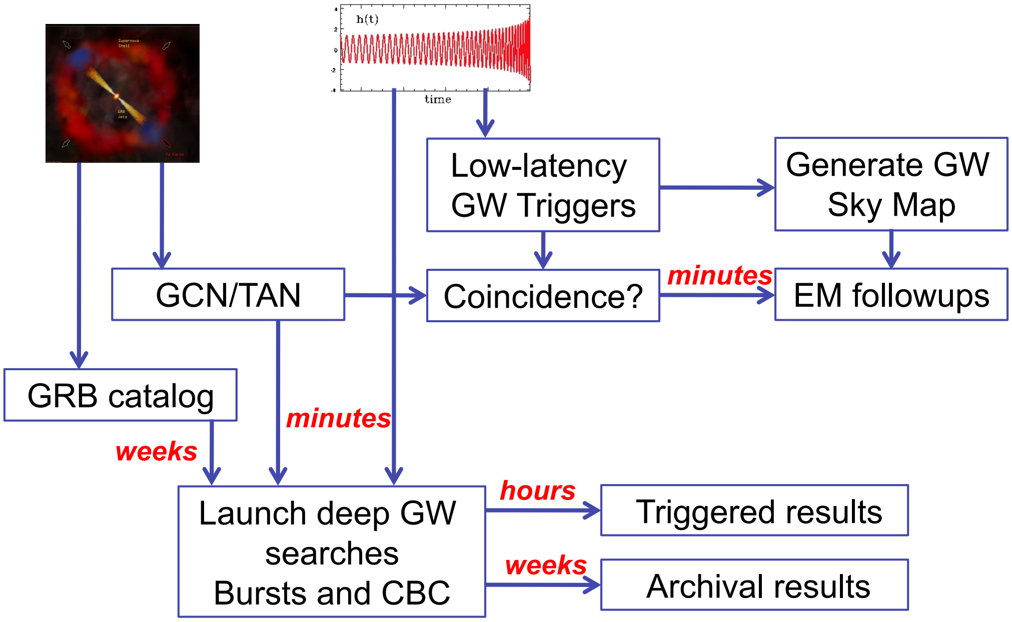

Figure 4: Flow of gamma-ray and GW data for the GRB searches. There are three main time scales, with each having an associated search: Low-latency (minutes), prompt triggered (hours), archival (weeks).

3.5.5 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 (∼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. (∼hour,or *prompt*, and ∼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 yet been reviewed. Its performance on the next engineering run will be used to decide which version will be used in the first science run and whether a review of the modifications is needed.

- **CBC.** is a fully-coherent search employing matched filtering to target NS-NS and NS-BH waveforms in association with GRBs. 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 and the low-frequency cutoff has been reduced to 15 Hz. This latter feature works, although template bank sizes with 2018 sensitivity curves will be much larger than the current ones. Running with larger template banks should be tested. Finally, work has started to allow the code to run on GRBs where only short stretches of data are available, and to use longer stretches of data when that is available.

- **Promp triggered search infrastructure.** The coherent Bursts (X-PIPELINE) and CBC (COHPTF) 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 [211]. 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 [212]. Stochtrack is not at present reviewed, but it is a relatively small subroutine called by otherwise reviewed STAMP infrastructure.

- **GRB plateaus.** An analysis technique is in development to target secular bar-mode signals of duration ∼1000 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.

3.5.6 Resources

Number of events

The computing needs of prompt and archival triggered searches are driven by the numbers of GRBs in the samples. We will assume that the GRB rates for the early ADE running will be the same as that for 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. The IPN detections are not conveyed promptly via the GCN, so are only considered for the archival searches. Table 2 summarizes these rates.

Table 2: 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 analysis with medium latency (∼15 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 within ∼ 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, cohPTF 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 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: 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, ligolw_segment_query, 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 pipeline the expected review will depend on the configuration selected as the multivariate configuration, which has not yet been reviewed. Otherwise, most of the modifications done recently have been reviewed for different analysis projects.
- For cohPTF, review will be needed on the different upgrades foreseen, see *Technical Requirements and Development Plan*.
- The core of the STAMP pipeline have been already reviewed for previous searches. Some parts will also be reviewed for on-going projects on burst and stochastic side. We anticipate short reviews 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 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 [213]. These IMBHs may form binaries, either when two or more IMBHs are formed in the same cluster [214], or as a result of a merger of two clusters each of which contains an IMBH in the suitable mass range [215]. 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 [216]. IMBH binaries could also form as a result of evolution of isolated binaries with very high initial stellar masses [217].

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 [194], 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 [218].

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 [219, 218] and the CBC ringdown search [220]. 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 §3.6.5.

The coherent WaveBurst (cWB) burst search is well established. Two IMBHB burst searches have been performed with data from initial LIGO and Virgo [219, 218]. 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 [220], while binaries with a total mass up to $100\,\mathrm{M}_\odot$ were searches for with full inspiral-merger-ringdown searches [111, 221]. 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. 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 [121]. However, further work will be necessary (see §3.6.5) 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 [222], which has been already used in the burst IMBHB search [218]. 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 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 §3.6.5 and §3.6.6. The S6 blind injection manuscript, LIGO-P1000146-v16 [75], 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.

3.6.5 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 EOBNRv2, EOBNRv2HM, SEOBNR, and IMRPhenomP.

Mock Data Challenges. To prepare for aLIGO/AdV runs, the Burst and CBC group plan to use 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 will be used to test pipeline performance and data analysis infrastructure. The pipeline performance will be evaluated based on the False Alarm rate Density (FAD) figure of merit [222]. The burst group also plans to re-run the upgraded CWB2G pipeline on the S5/S6/VSR1-4 data for comparison with the previous burst analysis. The MDC studies will 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.

3.6.6 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 3.6.5, 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. We have identified calibration requirements as one of the key outcomes of MDC 2, scheduled for late 2014. Following this, we will provide more detailed feedback to the calibration team on the level of calibration accuracy required for the IMBHB 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: 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). It is envisioned that the review of all of these components could begin in the second half of 2014, so that the majority of the analysis is reviewed prior to the beginning of the run. 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 All-sky Searches for Isolated Spinning Neutron Stars

3.7.1 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.

3.7.2 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 [223, 224, 225], magnetic deformations [226, 227], unstable r-mode oscillations [223, 228, 229], and free precession [230]. A review of these emission mechanisms can be found in [231]. 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 [232], although it is likely that only a tiny fraction would both be rotating fast enough to be accessible to LIGO [233] and remain bound to the galaxy over the age of the galaxy [234]. Only $\sim$ 2000 radio or x-ray galactic pulsars have been discovered so far [235]. 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 [231]. Note, however, that a later simulation analysis found significantly lower expectations that depend on the assumed source frequency and ellipticity [236]. 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, once a corresponding electromagnetic counterpart were identified, ideally, in multiple bands.

3.7.3 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. 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 $0 < -\dot{f}_{\rm GW} < 10^{-8}$ Hz s^{-1}. 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.

PowerFlux [237, 238] applies a stack-slide approach [239] 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 [240, 241].

The **Einstein@Home** [242] distributed computing resources enable the use of much longer coherence times ($\sim$day in recent searches) in an all-sky search than can be supported by the LIGO and Virgo computing clusters. 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 Hough thresholding algorithm [242] 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.

The Sky Hough method [243] 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 [244].

The hierarchical Frequency Hough method [245, 246, 247, 248] 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 [249] 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.

All pipelines, with the exception of Einstein@Home, share comparable sensitivity and computational cost; the mock data challenge described in §3.14.5 will yield a quantitative comparison amongst them. We do not foresee significant improvements in the performance of these pipelines before the Advanced Detector Era.

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 [240], which build on top of the PowerFlux infrastructure; and 2) a cross-correlation method [250] 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 an ongoing 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 §3.14.5, 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 [240, 241] used in the full-S5 PowerFlux paper [251] and the

"Nomad" approach [252] used in the full-S5 Einstein@Home paper [253]. 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[6]. 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 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

---

[6] 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.

past as targets for narrow-banding, it is not at all clear that narrow-banding is truly sensible.)

3.7.4 Publication Plan

Following the observing scenarios document [254], 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. 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 within a few months of the end of data collection, before the next run begins. We plan to analyze only $\sim$3 months of data for these quick-look searches, even when more data is available, given the diminishing returns on sensitivity and rapidly increasing computational costs for longer observation times.

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 pubic. 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 §3.14.3.

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	*arxiv:1402.4974*	VSR1 search review completed

Table 3: Summary of CW all-sky pipelines, including published observational results from previous data runs and the current search and review status.

3.7.5 Technical requirements and development plan

Pipelines status and review

PowerFlux has been used in published searches on LIGO S4 data [237] and S5 data [255, 251], 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 [256], based on a "universal statistic." Review of the S6 results is well under way, with publication expected in summer 2014. The tagged pipeline from this review serves as a reviewed baseline for the ADE.

Einstein@Home has been used in published searches on LIGO S4 data [242] and S5 data [257, 253], 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 has not yet been carried out. Review of some recent code improvements is under way as part of the review of a directed search for the galactic center. Review of the most recent full E@H all-sky pipeline is planned for completion in 2014.

The **Sky Hough** pipeline has been used in published or soon-to-be-published searches in LIGO S2, S4 and S5 data [258, 237, 12]. 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 summer 2014, 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 recently submitted for publication. It is planned to run the pipeline on VSR2 and VSR4 data in 2014. The pipeline has been tagged for default ADE use. The future addition of global correlations component would require additional review.

Mock Data Challenge The performances of the various pipelines used in this search will be 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 (Fall 2012) – 20 injections (upper limits and yes/no detection success)
- Stage 2 (Jan-Sept 2013) – 200 injections (same as stage 1 plus efficiency curves)
- Stage 3 (Oct 2013 - July 2014) – 2400 injections (same as stage 2, with better statistics)
- Stage 4 (Aug-Oct 2014) – Remaining 2400 injections up to 2000 Hz (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 2014. 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.

At present, the mock data challenge is based on S6 data sets, but will be extended to include VSR2/VSR4 data sets. In addition, engineering run data sets with full or nearly full interferometers will also be used once those become available with substantial livetime, in 2014.

3.8 Targeted Searches for Gravitational Waves from Known Pulsars

3.8.1 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.

3.8.2 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 [223, 224, 225], magnetic deformations [226, 227], unstable r-mode oscillations [228, 223, 229], and free precession [230]. A review of these emission mechanisms can be found in [231].

Here, we focus on the targeted search for gravitational waves from known neutron stars that are also radio or X-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. 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 [231]

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

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) [259]. 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) [260].

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 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 [REF]. 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.

3.8.3 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 pipeline use complementary techniques to reject out-of-band signals and for statistical interpretation.

Time domain Bayesian Pipeline. The time-domain Bayesian method has been applied successfully to data from the first six LSC science runs [261, 262, 263, 264, 265, 259] and to the Virgo VSR2/VSR4 runs [266, 259]. The method is described in detail in [267]; the inclusion of binary system parameters is described in [268]. 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 downsampling of the raw data. Currently this takes about 25 min per pulsar per detector per day of data. We are investigating 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 should give a very significant speed-up when processing multiple targets.

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 [249], 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 [269]. 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 [266] 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 [259].

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 [270, 266]. 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 [266] 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 [259].

3.8.4 Publication Plan

Following the observing scenarios document [254], 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 a 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.

3.8.5 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* in press (S6/VSR2/VSR4)	Review complete	
Time Domain Matched Filter	*ApJ* 737 (2011) 93 (VSR2) *ApJ* in press (S6/VSR2/VSR4)	Review complete	
Signal Fourier 5-compoent	*ApJ* 737 (2011) 93 (VSR2) *ApJ* in press (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 [REFs]. Additions and refinements to the pipeline are reviewed as they come into use. 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 the end of 2014.

The new parameter estimation code has already been validated against the current code using both fake data and simulated signals [271, 272]. 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 review too is expected to finish by the end of 2014 and the code refinement 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 [273]. Moreover, a new method for computing upper limits in the frequentist framework has been developed, which overcomes some problems of the standard frequentist methods [274], 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.

Mock Data Challenge
The performances of the various 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. 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.

3.9 Search for an Isotropic Stochastic Gravitational Wave Background

3.9.1 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 $\Omega_0 < 6.1 \times 10^{-6}$ at 95% confidence [275]. 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.

3.9.2 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 [276], phase transitions in the early universe [277, 278], cosmic strings [279, 280, 281, 282], and pre-Big Bang models [283, 284]. In astrophysical scenarios, the gravitational-wave signals from more ordinary objects, such as neutron star binaries, combine to produce a stochastic signal (see [285] for a review). Astrophysical backgrounds may arise from core collapses to neutron stars or black holes [286, 287, 288, 289], rotating neutron stars [290, 291] including magnetars [292, 293, 294, 295], phase transition [296, 297] or initial instabilities in young neutron stars [298, 299, 300, 299], compact binary mergers [301, 302, 303, 304, 305, 306] and compact objects around supermassive black holes [307, 308]. 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 [303]. 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. The recent measurement of the B-mode polarization in the cosmic microwave background by the BICEP 2 experiment [309] claims detection of the gravitational waves produced during inflation. 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 $\sim 10^9$–$10^{10}\,\mathrm{GeV}$ in the early universe [310]. 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 [311]; the statistical nature of the signal will also provide information [303]. 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 $\Omega_{\mathrm{gw}}(f) \approx 10^{-15}$, 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 [312, 313]. Similarly, there are regions of parameter space where we may be able to detect cosmic strings [314, 3]. 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 [315]. 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 non-standard polarization [316]. While general relativity allows only for two kind of tensorial polarizations, a generic prediction of extended gravitational theories, such as scalar-tensor ones [317, 318], $f(R)$ gravity [319, 320], bimetric [321] and massive [322] 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 [323, 324, 314, 325]. 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 remaining initial detector isotropic searches [325, 326] are undergoing review and are expected to be submitted for Executive Committee approval soon. A search for extended polarizations is also ongoing, and is expected to be ready for a first review at the end of this year. Searches with advanced detectors are expected to benefit from dramatically improved sensitivity.

3.9.3 Search Description

The primary goal of the isotropic search is to estimate the energy density of the stochastic background:

$$\Omega_{\mathrm{GW}}(f) \equiv \frac{1}{\rho_c} \frac{d\rho_{\mathrm{GW}}}{d\ln f}\,, \tag{3}$$

where ρ_{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 [327, 328], which has served as the basis for all previous LIGO/Virgo stochastic searches, e.g., [323, 324, 314, 325, 326]. The expected signal-to-noise ratio for a stochastic search, using two detectors I and J, is given by

$$\mathrm{SNR} = \frac{3H_0^2}{10\pi^2}\left(2T\int_0^\infty df\,\gamma_{IJ}^2(f)\frac{\Omega_{GW}^2}{f^6 P_I(f)P_J(f)}\right)^{1/2} \tag{4}$$

where H_0 is the present value of the Hubble expansion rate, $\gamma_{IJ}(f)$ is the overlap reduction function (see [328]), 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[7]. 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 [329].

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 $\Omega_{\mathrm{gw}}(f)$ given some assumed power law $\Omega_{\mathrm{gw}}(f) \propto f^\alpha$. 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 [314, 311]. Also, we will use the results to place constraints on specific models such as cosmic strings [314] and astrophysical backgrounds from binary coalescence [311]. 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 [311].

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.

---

[7] http://tinyurl.com/kgbffx7

3.9.4 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 [330], 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.

3.9.5 Technical Requirements and Development Plan

The stochastic pipeline is Matlab-based code that lives in the matapps repository[8]. The main function is stochastic.m.

---

[8] http://tinyurl.com/kpmyruv

The stochastic pipeline can be traced back to S1 [329]. 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., [323, 324, 314, 325, 326]). 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.

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[9]. 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 to be tested on ER5 data and then run real time during ER6. 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 is currently in progress in order to validate the different aspects of the search; no cause for concern has been displayed by this exercise. A stochastic detection checklist exists, but it is being updated for the advanced detector era.

3.9.6 Resources

Code infrastructure. The stochastic isotropic search utilizes the Matlab-based stochastic pipeline, which lives in matapps[10]. 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

---

[9] https://svn.ligo.caltech.edu/svn/sgwb/trunk/stochmon

[10] http://tinyurl.com/kpmyruv

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 [331]. 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 [332]. The primary sources of these correlated fields are Schumann resonances [332]. 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 [314] 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.

3.10 Directional Search for Persistent Gravitational Waves

3.10.1 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 $2 - 20 \times 10^{-50}$ strain^2 Hz^{-1} and $5 - 35 \times 10^{-49} \mathrm{strain}^2$ Hz^{-1} 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.

3.10.2 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.

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 [276], mechanisms that terminate inflation and may give contributions at high frequencies [312, 313], phase transitions in the early universe [277, 278], cosmic strings [279, 280, 281, 282], and pre-Big Bang models [283, 284]. While in most models the SGWB is predicted to be isotropic, there are mechanisms that could introduce anisotropy [282, 333]. Astrophysical backgrounds may arise from binary mergers [315, 334, 301], core-collapse supernovae [335, 336], neutron-star excitations [337, 296], persistent emission from neutron stars [338, 339], and compact objects around supermassive black holes [307, 308]. 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).

When run in narrowband mode, the search focuses on a 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.) More generally, 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 binary neutron stars 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 [312, 313]. 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 [314, 340]. 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 [315], although parts of the parameter space in magnetar- or pulsar-based models will also be within reach [341].

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 [342, 343]. 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 ≈2× better strain sensitivity to Sco X-1 than previous radiometer measurements [343], 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).

3.10.3 Search Description

Similarly to the isotropic search, the anisotropic SGWB search estimates the energy density of the stochastic background, but keeps the directional information [343]:

$$\Omega_{\mathrm{GW}}(f) \equiv \frac{1}{\rho_c}\frac{d\rho_{\mathrm{GW}}}{d\ln f} = \frac{2\pi^2 f^3}{3H_0^2}\int d\hat{\Omega}\, H(f)\, P(\hat{\Omega}) \tag{5}$$

where ρ_{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

$$P(\hat{\Omega}) = \eta(\hat{\Omega}_0)\delta^2(\hat{\Omega}, \hat{\Omega}_0), \tag{6}$$

and in the spherical harmonic decomposition (SHD) algorithm we assume that the signal can be written as a superposition of spherical harmonics

$$P(\hat{\Omega}) = \sum_{lm} P_{lm} Y_{lm}(\hat{\Omega}). \tag{7}$$

In either basis, we can compute the "dirty map" X_ν and the corresponding Fisher matrix $\Gamma_{\mu\nu}$:

$$X_\nu = \sum_{ft} \gamma_\nu^*(f,t)\frac{H(f)}{P_1(f,t)P_2(f,t)}C(f,t) \tag{8}$$

$$\Gamma_{\mu\nu} = \sum_{ft} \gamma_\mu^*(f,t)\frac{H^2(f)}{P_1(f,t)P_2(f,t)}\gamma_\nu(f,t) \tag{9}$$

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 [343]. Likelihood maximization then leads to the estimators of the angular content in the two bases:

$$\hat{\eta}_{\hat{\Omega}} = \left(\Gamma_{\hat{\Omega}\hat{\Omega}}\right)^{-1} X_{\hat{\Omega}} \tag{10}$$

$$\hat{P}_{lm} = \sum_{l'm'} (\Gamma^{-1})_{lm,l'm'} X_{l'm'} \tag{11}$$

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 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[11]. 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.

3.10.4 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 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 [342, 343].

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 .

---

[11] http://tinyurl.com/kgbffx7

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 is currently participating in a CW-led mock data challenege to study the sensitivity of different pipelines to gravitational waves from Sco X-1. We will continue to coordinate this search with related CW searches.

3.10.5 Technical Requirements and Development Plan

The stochastic pipeline is Matlab-based code that lives in the matapps repository[12]. The main function is stochastic.m. The stochastic radiometer pipeline can be traced back to S4 [342], with the SHD algorithm added in the S5 analysis [343]. 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.

3.10.6 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. Minor 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 [332]. The primary sources of these correlated fields are Schumann resonances [332]. Having documented this potential source of correlated noise, we will be careful to determine if correlated

---

[12] http://tinyurl.com/kpmyruv

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 [344].) 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 [314] 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.

3.11 Directed Searches for Gravitational Waves from Isolated Neutron Stars

3.11.1 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.

3.11.2 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 [223, 224, 225], magnetic deformations [226, 227], unstable r-mode oscillations [228, 223, 229], and free precession [230]. A review of these emission mechanisms can be found in [231].

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 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 4. 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 centers of globular clusters. 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 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.

Table 4: 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)
1.9+0.3		174846.9−271016	8.5	0.1
18.9−1.1		182913.1−125113	2	4.4
93.3+6.9	DA 530	205214.0+551722	1.7	5
111.7−2.1	Cas A	232327.9+584842	3.4	0.3
189.1+3.0	IC 443	061705.3+222127	1.5	3
266.2−1.2	Vela Jr.	085201.4−461753	0.2	0.69
266.2−1.2	Vela Jr.	085201.4−461753	0.75	4.3
291.0−0.1	MSH 11−62	111148.6−603926	3.5	1.2
347.3−0.5		171328.3−394953	0.9	1.6
350.1−0.3		172054.5−372652	4.5	0.6

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.

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)[345]

$$h_0^{\rm age} = \frac{1}{D}\sqrt{\frac{5GI_{zz}}{8c^3\tau}} \tag{12}$$

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:

$$h_0^{age} = 1.2\times10^{-24}\left(\frac{3.4kpc}{D}\right)\sqrt{\left(\frac{I_{zz}}{10^{45}{\rm gcm}^2}\right)\left(\frac{300{\rm years}}{\tau}\right)} \tag{13}$$

In designing a search, it is customary [345] 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.

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.)

3.11.3 Search Description

At present two mature approaches are available for directed searches. The first, used for the S5 Cas A search [346] and for the S6 SNR search, is based on computing the $\mathcal{F}$-Statistic[249] 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 [347] 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 [11] 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 5 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 [252] used in the full-S5 Einstein@Home paper [253]. 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 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

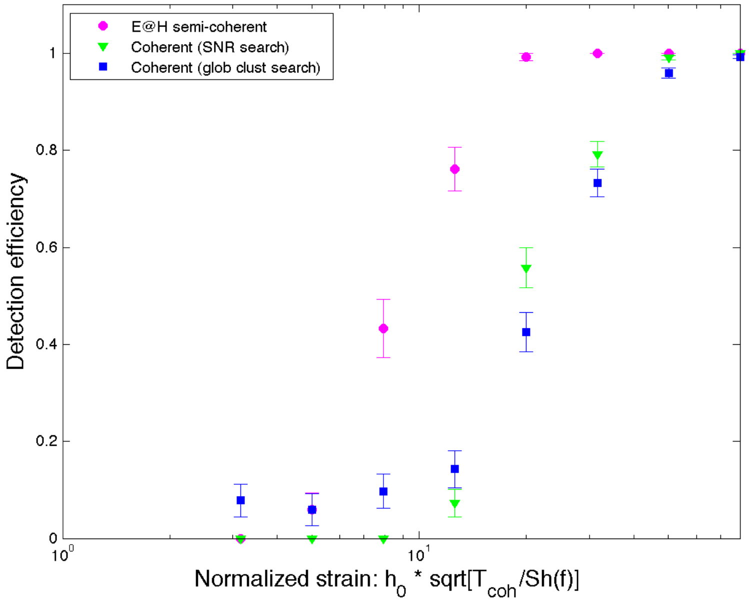

Figure 5: Approximate detection efficiencies of coherent pipelines with (glob cluster) and without (supernova remnant) resampling; and of a semicoherent directed search pipeline, based on mock data challenge studies. The semi-coherent results are taken from a cluster emulation of what can be achieved with Einstein@Home, using a 20 mHz search band and ignoring $\ddot{f}$.

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 [348] 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.

3.11.4 Publication Plan

Following the observing scenarios document [254], 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, 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[13]. 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.

3.11.5 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 Q1 2015	

Coherent Pipeline
The first iteration of this pipeline was originally reviewed for the S5 Cas A search [346]. 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 [11]. An updated

---

[13] 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.

version of the pipeline, including a search over non-zero $\ddot{f}$, will begin review in Q4 2014, with completion expected in Q1 2015.

Mock Data Challenge
The performances of the various 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.

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

3.12.1 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.

3.12.2 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 [349, 350, 351, 352, 229, 223, 228, 353]. 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 [349, 350, 223]:

$$h_{\rm torque} \quad \sim \quad (5\times10^{-27})\sqrt{\left(\frac{600\ {\rm Hz}}{f_{\rm GW}}\right)\left(\frac{\mathcal{F}_{\rm x}}{10^{-8}\ {\rm erg/cm^2/s}}\right)} \tag{14}$$

where $\mathcal{F}_{\rm x}$ 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:

$$h_{\rm torque}^{\rm ScoX1} \quad \sim \quad (3\times10^{-26})\sqrt{\left(\frac{600\ {\rm Hz}}{f_{\rm GW}}\right)} \tag{15}$$

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) [354]. It has been suggested [355] 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 [223]. 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 [356] indicate the slowing is not dramatic. Another proposed mechanism is excitation of r-modes in the fluid interior of the star [351, 223, 352, 228], with both steady-state emission and cyclic spinup-spindown possible [353, 357].

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 [358]. We envision a program in which the search for Sco X-1 is given highest priority, 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.

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 [359] or that the r-modes are restricted to the neutron star crust and hence gravitationally much weaker than core r-modes [360]. 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.

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 [361, 362, 363]. 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 [363] 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 [10]. 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 $f_{\rm GW} = 2f_{\rm Rot}$, 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, 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.

3.12.3 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 [354].

Seven existing or potential pipelines in all are under consideration for these searches (one of which is a Stochastic Search Group pipeline). Four of these pipelines are mature enough to have participated fully in the first Sco X-1 mock data challenge [364]: Sideband, TwoSpect, Polynomial and Radiometer (Stochastic Search Group pipeline). Based on these initial MDC results, Polynomial appears to offer no significant advantage for a directed binary search (but is being pursued for all-sky binary searches). TwoSpect and Radiometer have similar sensitivities when analyzing a year of data, while Sideband has somewhat 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 [365]. The semicoherent pipelines described here are expected to be robust against plausible fluctuations. Sideband, TwoSpect and Radiometer all have modest computational cost and seem appropriate to include in the suite described below.

Three 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: CrossCorr (similar to the Radiometer pipeline, but with frequency demodulation and longer correlation coherence times); PowerFlux (developed for isolated-star searches, but adaptable to binary searches); and an Einstein@Home-based F-Statistic search. As of October 2014, the CrossCorr pipeline is rapidly approaching maturity and shows promise to be more sensitive than all of the first four pipelines, albeit at higher computational costs. That additional cost has not yet been empirically quantified at high signal frequencies, but should grow faster than f^2; if CrossCorr is run to target frequencies, it will likely need to offset the computing cost increase by operating with a shorter coherence time at those frequencies, mitigating somewhat its sensitivity advantage. It is difficult to make concrete plans using any of the three developing pipelines at this time, but by summer 2015, CrossCorr is likely to be ready for use in O1. 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 three search pipelines that are currently mature and have comparable sensitivities: Sideband, Radiometer and TwoSpect. We also describe the CrossCorr pipeline, which is less mature, but whose preliminary results indicate that it could be the most sensitive method available for use in O1.

The Sideband pipeline [366, 365] is based on the venerable $\mathcal{F}$-Statistic [249] 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 $M \sim 4\pi f_{\rm gw} a \sin i/c$ frequency resolvable sidebands where $f_{\rm gw}$ 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 [367, 368]. 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 [366]. Search results have been obtained from a ten-day stretch of S5 data and reviewed, with submission for publication expected in November 2014. 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[342] 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 [364] across frequency bins.

The present search program is quite mature, having been used previously for S4 and S5 searches [342, 343], 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[369, 14] 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 [369] 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. 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 polarizaed source.

An all-sky binary search (in S6 and VSR2-3 data) over a ~500-Hz band was recently published [14], which included a low-frequency search for Sco X-1. A dedicated S6 search for Sco X-1 up to ~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.

The Cross-Correlation (CrossCorr) pipeline[250] 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 separated in time by less than some maximum lag time T_{lag}, the pipeline can be tuned to trade off sensitivity versus computing cost.

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.

A methods paper is in progress[370] describing applications of the search method to LMXBs in particular, including an approximate parameter space metric, reduction of spectral leakage effects, and the impact

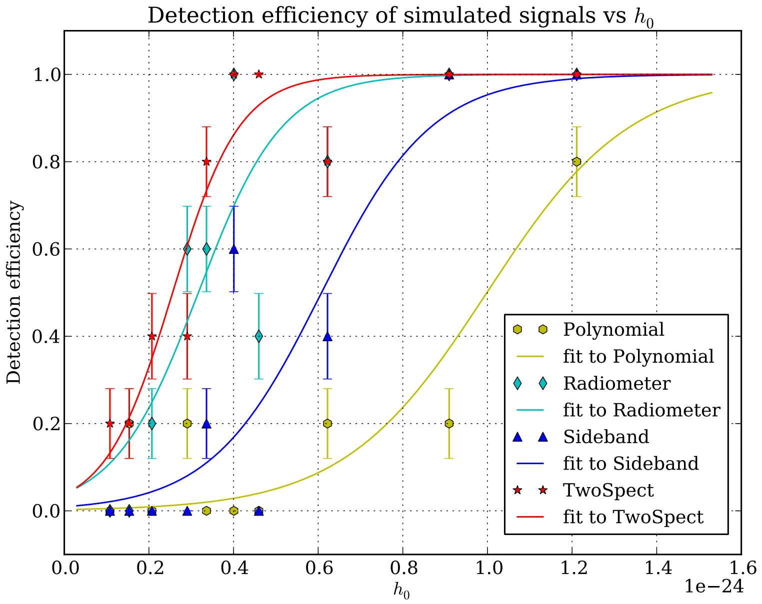

Figure 6: Approximate detection efficiencies of four participating directed binary search pipelines in the first stage of the Sco X-1 mock data challenge.

of unknown inclination and polarization angles for the neutron star.[14] The CrossCorr search has been run on simulated data, including the data of the Sco X-1 mock data challenge (see below), with encouraging indications about its sensitivity. Typical lag times used were $T_{\text{lag}} \leq 2400\,\text{s}$.

Preliminary comparison of pipelines

As described below (section 3.12.5), an ongoing **mock data challenge** is evaluating the pipeline performances (detection and upper limits strain sensitivity, parameter estimation), with a method paper to be completed by December 2014 [364]. Figure 6 shows the detection efficiency of the four 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, the sensitivities of the top three pipelines are comparable enough and different enough in their tolerances for astrophysical uncertainties that it is sensible to support all three pipelines in O1 running. (As noted earlier, Polynomial has been optimized for all-sky binary searches.) 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 CrossCorr pipeline is currently being tested on the "open" signals of the Sco X-1 MDC, with a full run on the "blind" signals expected to begin by the end of October 2014. (The MDC included 50 "open" signals whose parameters were published initially, and 50 "blind" signals whose parameters were initially kept secret. Those parameters have been released, and comparisons used in Figure 1, but the CrossCorr team is endeavoring to remain "self-blinded" as much as practically possible, by confining initial tests to the open

---

[14] See for example https://dcc.ligo.org/LIGO-G1300509/public

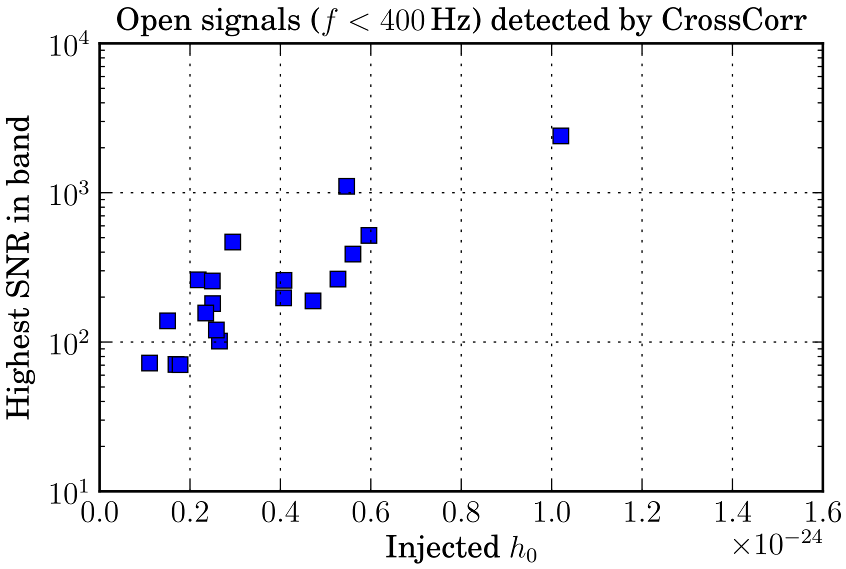

Figure 7: Strength of detection of MDC signals by CrossCorr pipeline. While these all represent strong detections, some caveats apply when comparing these preliminary results to the efficiencies in figure 6. The CrossCorr results were obtained for a different set of injected signals (open rather than blind), and thus particulars such as inclination angle and orbital parameters which may affect the detectability will be different. The CrossCorr results were only for f below 400 Hertz, rather than over the whole frequency band. Computing cost considerations (shorter SFTs as well as a higher density of templates in parameter space) will necessitate running CrossCorr with a shorter coherence time at higher frequencies. Finally, these large single-template SNRs do not take into consideration multiple-template trials factors, errors in power spectrum estimation in the presence of strong signals, or non-Gaussianities on the tail of the sampling distribution, although none of these effects would be large enough to reduce the lowest effective SNR below about 60.

signals. The 19 open signals in the band $50\,\text{Hz} < f < 400\,\text{Hz}$ have been analyzed so far, using 600 second SFTs and lag times of 2400 and 1200 seconds for different regions of the orbital parameter space. All 19 have been strongly detected, with peak SNR values of over 70. (The expected SNR of the cross-correlation statistic scales like h_0^2, and the optimally combined statistic is normalized to have a theoretical variance of unity and assumed to obey a normal distribution thanks to the Central Limit Theorem.) Figure 7 illustrates the recovered SNR values for those signals. Note that other signals in the MDC data have lower h_0 values. We anticipate that the CrossCorr method will transition from detection to non-detection when h_0 is about $1/3$ of the quietest value in this set of 19 signals, or at about $h_0 \sim 5 \times 10^{-26}$ for this simulated year-long 3-detector run at design sensitivity, with coherence times of order 2000 seconds.

A search for the open signals in the frequency range $400\,\text{Hz} < f < 800\,\text{Hz}$ is underway. In order to capture the orbital Doppler modulation at these higher frequencies, shorter (300 s) SFTs are needed, which requires a reduction in T_{lag} to keep the computing cost fixed. The initial run is being done with lag times of 1140 and 840 seconds.

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 14, 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 also go in the other direction too, with accretion disk models that maintain torque balance without significant GW radiation[371]. 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

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. 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 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[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.

---

[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 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)	MDC analysis underway; Review not yet begun
Sideband	(Expect S5 paper submission - Nov 2014)	S5 pilot search review completed
Radiometer	*PRD* 76 (2007) 082003 (S4) *PRL* 107 (2011) 0271102 (S5)	S6 / VSR2-4 search under way
TwoSpect	*PRD* 90 (2014) 062010 (S6/VSR2) (all-sky binary pipeline)	S6 Sco X-1 search under way

Table 5: Summary of directed binary pipelines, including published observational results from previous data runs, and the current search and review status.

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.

3.12.4 Publication Plan

Following the observing scenarios document [254], 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 §3.12.3.

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.

3.12.5 Technical requirements and development plan

Pipelines status and review

Sideband has been used in a search of LIGO S5 data, the review of which was completed recently, with a mature paper draft recently circulated to the collaboration and with submission for publication expected in November 2014.

Radiometer has been used in published searches of LIGO S4 data [342] and S5 data [343], with an S6 / VSR2-4 search nearing completion in fall 2014. Modest enhancements are planned to correct for signal leakage across bins at high frequencies, which will require additional pipeline review.

TwoSpect has been used in its original, all-sky implementation in a published search of LIGO S6 and Virgo VSR2 data [14]. Modest review will be needed for recent minor modifications for use in a directed search, tested in the Sco X-1 MDC [364] and being used in a LIGO S6 search. More substantial review will be needed for ongoing enhancements to support coherent detector combination, explicit polarization searching and orbital phase exploitation that improve sensitivity to known binary systems.

CrossCorr has a functional lalapps search code which is being used for MDC analyses, but a bit of technical cleanup remains to be done. Post-processing and follow-up procedures are undergoing rapid development in the context of the MDC. Pilot analyses of real data from S5/S6 and/or ER7 are planned for December 2014/January 2015, with review readiness targeted for late January 2014.

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 is 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} 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 6. A methodological paper describing the results of the first stage of the MDC is in preparation [364].

The next stage of the MDC, planned for early 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.

3.13 Search for transients from Cosmic Strings

3.13.1 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.

3.13.2 Scientific Justification

A cosmic network of strings may form as a result of phase transitions in the early Universe [372]. When a U(1) symmetry is broken in multiple causally disconnected spacetime regions, one-dimensional topological defects, i.e., strings, are expected to form [373]. For a long time, cosmic strings were considered candidate sources for structure formation in the early Universe [374]. Cosmic microwave background (CMB) experiments, however, have shown that cosmic strings can only contribute up to a few percent of the overall anisotropies observed [375, 376, 377, 378, 379]. 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 [380, 381, 382, 383, 384]. Cosmic strings produced in string theory motivated models (dubbed "cosmic superstrings") have received much attention since they could provide observational signatures of string theory [385, 386].

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 [387, 388, 389] and from simulations of string-sourced CMB anisotropies [375, 376, 377, 378, 390, 391]. 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 [379] 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 [392, 393]. When two string segments meet, they exchange partners or intercommute with a probability p. For superstrings, the reconnection probability can be less than unity ($10^{-4} < p < 1$ [394]) 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 [395] 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 [396].

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. [397] suggested that the size of loops is set by gravitational backreaction and $\alpha \leq \Gamma G\mu$, where $\Gamma \sim 50$ [373]. 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 $\alpha \lesssim 1$ [398, 399]. We parametrize $\alpha = \varepsilon\Gamma G\mu$ with $\varepsilon < 1$ following the convention of Ref. [392].

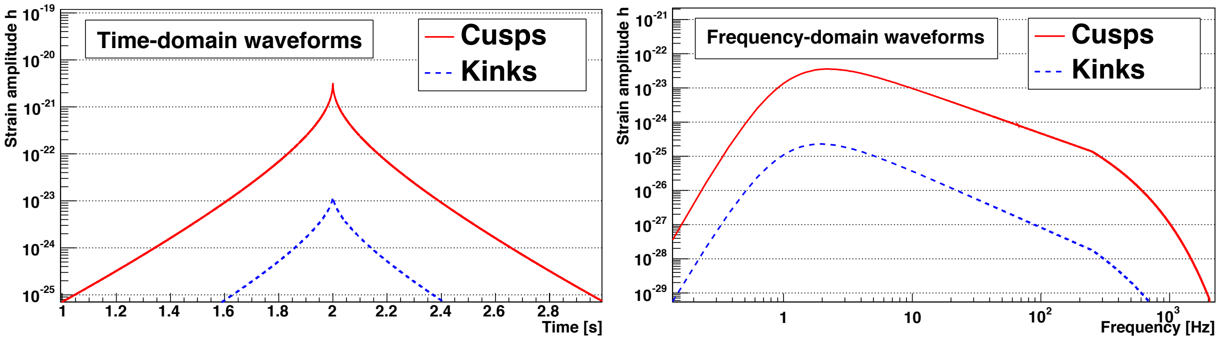

Figure 8: Time-domain (left) and frequency-domain (right) waveforms for gravitational waves produced by cusps (red plain curve) and kinks (blue dashed curve). These waveforms have been computed with the same set of cosmic string parameters: $G\mu = 10^{-8}$, $\varepsilon = 10^{-5}$, $z = 1$ and $f_h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

The possibility of direct detection of GW bursts from cosmic string cusps was first suggested in 2000 by Berezinsky *et al.* [400]. Shortly after, Damour and Vilenkin showed that the stochastic GW background generated by oscillating loops is strongly non-Gaussian [396]. Occasional sharp bursts of GWs produced by cusps are expected to stand out above the stochastic background [396, 401, 392]. 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 8 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 [27]. The next generation of ground-based GW detectors will probe the cosmic string parameter space further. The improved sensitivity of Advanced LIGO [402] and Advanced Virgo [403] 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. 8) but a proliferation mechanism could offer a production rate several order of magnitude larger [404] see Fig. 9). In the future, it is planned to search for GWs associated to both cusps and kinks and to publish combined results.

3.13.3 Search Description

The search for GW bursts from cosmic strings begins with a matched-filter analysis of strain data from each detector separately [405]. It consists of projecting the whitened data onto an overpopulated [16] 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 > 3.6 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 [174].

To discriminate true signals from background events, we apply the multivariate technique described in Ref. [406], which uses a set of simulated GW events and typical noise events to statistically infer the

---

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

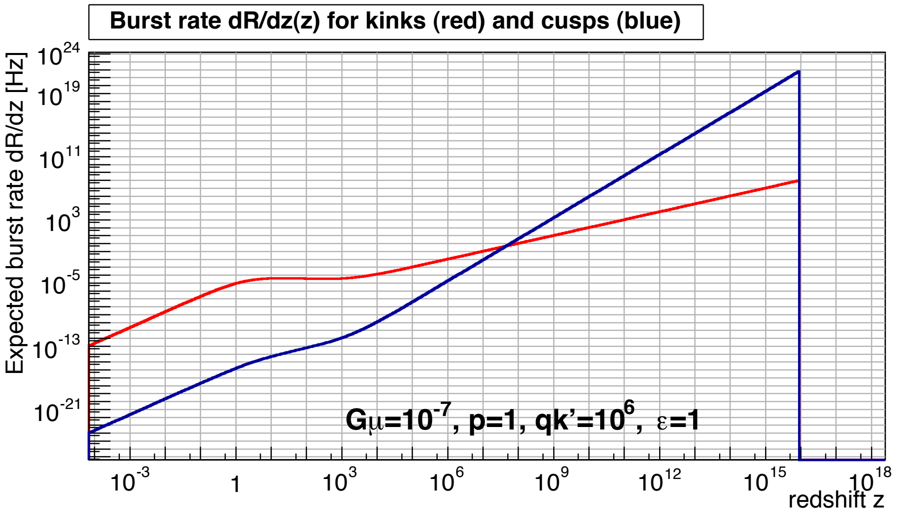

Figure 9: Predicted rate of GW bursts produced by cosmic string cusps and kinks. The proliferation mechanism, parameterized by $qk' \gg 1$, favors the kink event rate over a wide range of redshifts.

probability for a candidate to be signal or noise.

Figure 10 displays the region of the cosmic string parameter space that is excluded by the analysis of initial detectors' data [27]. 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 [407, 408] by a factor 3 for intermediate ε values. It nicely complements existing limits provided by pulsar timing experiments for large ε [409, 378] and by the LIGO stochastic search in the very small loop regime [408].

3.13.4 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.

3.13.5 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 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

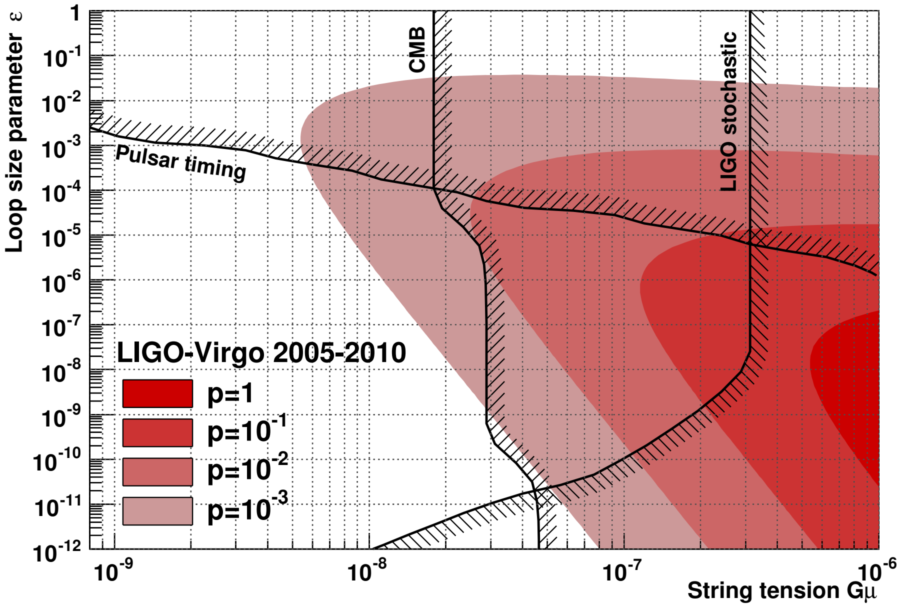

Figure 10: This plot presents existing constraints on cosmic string parameters: the string tension $G\mu$, the loop size parameter ε and the probability p that two string segments interact when they meet. The S5-6/VSR1-2-3 analysis is able to reject the regions filled in red. For comparison, other constraints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

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

3.14.1 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.

3.14.2 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 [223, 224, 225], magnetic deformations [226, 227], unstable r-mode oscillations [223, 228, 229], and free precession [230]. A review of these emission mechanisms can be found in [231]. 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 [232], although it is likely that only a tiny fraction would both be rotating fast enough to be accessible to LIGO [233] and remain bound to the galaxy over the age of the galaxy [234]. Only $\sim$ 2000 radio or x-ray galactic pulsars have been discovered so far [235].

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 [233]. 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% [235].

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 [14], 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.

3.14.3 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[14, 369] 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 [237]. 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 [14], 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** [410] 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

$$f_{\rm ssb} = f_{\rm gw}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right) \tag{16}$$

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

$$\frac{df_{\rm ssb}}{dt} = f_{\rm gw}\gamma\left(1 - \frac{d\vec{v}}{dt}\cdot\frac{\vec{n}}{c}\right) + \frac{df_{\rm gw}}{dt}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right). \tag{17}$$

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. [239] 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} $\mathrm{Hz/s^2}$.

The Polynomial pipeline has been implemented and tested on simulated data with white noise (see S. van der Putten's Ph.D. thesis [410]) and has been further tested in the Scorpius X-1 mock data challenge. 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 [411], 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 [250] 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 [240, 241, 252].

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 3.14.2. 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 6: Summary of existing CW all-sky binary pipelines, including published observational results from previous data runs and the current search and review status.

3.14.4 Publication Plan

Following the observing scenarios document [254], 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 §3.14.3. 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[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. 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.

3.14.5 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 [14]. 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

---

[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 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 [364].

Polynomial too has participated in the Sco X-1 mock data challenge and is preparing for a review to begin in 2015.

3.15 Search for Eccentric Binary Black Holes

3.15.1 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 $e >\sim 0.05$ [412], 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.

3.15.2 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 [413] suggests that our galactic nucleus should have ~ 2000 BHs and ~ 400 NSs in the central 0.1 pc. In [414, 415], 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 [414].

Dynamical capture binaries may also form in globular clusters (GCs) that undergo core collapse [416, 417]. In [418], 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 [419, 420, 421] and CO mergers around supermassive BHs in galactic nuclei [422], as well as in coevolved or dynamically formed BH-NS or NS-NS binaries [423]. Efforts to understand this mechanism in the general-relativistic regime are ongoing (see e.g. [424]), and the event rates of these systems are not well known (though see [425]).

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 [426], $r_{p,i} \approx 0.57 r_{p,f}(1+e_f)e_f^{-12/19}[1+O(e_f^2)]$. 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.

3.15.3 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. However, the cWB2G pipeline will not provide estimates for the source parameters. In order to actually measure the parameters of a signal candidate, the BayesWave pipeline will be used, which is well suited for the analysis of eBBH signals.

BayesWave coherently models transient events in LIGO data as a linear combination of wavelet basis functions [427]. 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 [428]. 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 $f(t)$ [426]. 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 [429] and is naturally extensible to eBBH searches.

3.15.4 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 §3.15.5 and §3.15.7. 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.

3.15.5 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 [430] 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 [428]. 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 [431]. 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.

3.15.6 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.

3.15.7 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 both CWB2G and BayesWave is nearly complete. The remaining development items (multi-chirp reconstruction and improvements in multiresolution analysis for CWB2G, and reconstruction of eBBH waveforms in 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. All 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 will start by 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.

3.16 Search for Intermediate-mass-ratio Coalescences

3.16.1 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 are more burst-like than lower mass systems, sweeping through a relatively narrow range of frequencies, and might therefore be vetoed by current techniques and these systems lie in a regime where little modelling of the gravitational waveforms has been done. 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 with tuning for the appropriate mass range and a template based search with gstlal. The performance of both techniques will be assessed using the best available models in the IMRI mass range.

3.16.2 Scientific Justification

There is some observational evidence that IMBHs of a few hundred solar masses may exist generically in globular clusters [213]. 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 radiation emission takes over and drives the binary to coalesce [432]. This gradual hardening mechanism means that the binary is likely to have sufficient time to circularize. In [432], an upper limit on the Advanced LIGO IMRI rate of $\sim 30 \mathrm{yr}^{-1}$ 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 [217], 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. A $1.4 + 100M_{\odot}$ 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 [433]. 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. Once the low-frequency sensitivity reaches $\sim$ 10Hz, IMRIs will have strong potential for fundamental physics [434]. 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 difference in 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 [114]. 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.

3.16.3 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 IMBH 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 large difference in mass ratio, the effect of precession accumulates over the large number of observed waveform cycles and therefore precessing templates may be needed [435, 436]. The template bank will therefore ultimately need to employ non-precessing spinning waveform models to increase coverage and it may 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 etc. 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 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 [222], which has been already used in the burst IMBHB search in initial-detector data [218], 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 [437]. Successful parameter estimation of IMRI-like systems, including mass ratio measurements, has been demonstrated using the same (EOBNR) templates for injection and recovery. 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.

3.16.4 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., [111]). 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.

3.16.5 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 [438]. Several waveform models are now available which make predictions in the IMRI range — EOBNR(v2) [439, 440], SEOBNR(v2) [115], Huerta & Gair [437], Callister & Gair [441], 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 [442] 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 and should be finished by end of Q4 2014.

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 vetoes 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.

3.16.6 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 3.6.5, 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.

3.17 High-energy Neutrino Multimessenger Analysis

3.17.1 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.

3.17.2 Scientific Justification

Many of the most violent and interesting phenomena producing GW transients are also thought to be sources of high-energy neutrinos [443, 444, 445, 446, 447, 448]. These non-thermal, GeV-PeV neutrinos neutrinos are thought be produced in relativistic outflows driven by central engines also responsible for GW emission [449, 450, 451, 452, 453, 454, 455, 456]. 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 [448]. A particularly interesting development is the recent detection of astrophysical PeV-energy neutrinos [457], 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 [447]. For compact binary mergers, which are the likely progenitors of short GRBs, we expect GW emission detectable out to $\sim 450\,\mathrm{Mpc}$ 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 $\sim 200\,\mathrm{Mpc}$ 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 $\sim 10^{51}\,\mathrm{erg}$ for some sub-photospheric and collisionless shock models.

Search results can be used to constrain joint emission models [447]. 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 [443, 444, 445, 446, 447, 448, 458, 459]. (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 [451, 460, 461] 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 [446, 447, 459].

The IceCube neutrino detector [462], 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 [447, 458], while the projected constraints [447] and expectations (e.g., [463]) suggest that multimessenger GW+high-energy neutrino searches will be a fruitful direction of research during ADE.

3.17.3 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 [464, 446, 459]. 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 [464].

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 [464].

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 [464] 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.

3.17.4 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.

3.17.5 Technical requirements and development plan

The ADE search will use the multimessenger search pipeline developed for joint searches with IceCube [464]. 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.

3.17.6 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.

3.18 Searches For Transient Gravitational Waves From Isolated Neutron Stars

3.18.1 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 of creating an entirely new field of 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.

3.18.2 Scientific Justification

Galactic Magnetar Hyperflares
Isolated NSs exhibit a variety of violently energetic and unpredictable phenomena such as X-ray and gamma-ray flaring in magnetars and abrupt changes in rotation rate seen in pulsar glitches. While these phenomena are not well understood, it is reasonable to expect that any sudden and localized energy release, such as that observed in a giant flare in a soft-gamma repeater (SGR), 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 [465]. 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 $\sim$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 [466, 467]. While torsional modes themselves are not likely to be particularly efficient GW emitters, if the 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., [468]) 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 [469, 470, 471], pulsar glitches [472] and potentially nearby ($\sim$ 1 Mpc) short γ-ray bursts where the burst may be an extragalactic magnetar hyperflare [473, 474]. 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 [475]. 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. We therefore plan only to pursue GW analyses for events believed to be giant flares (aka ‘hyperflares’) from unambiguously Galactic magnetars, or similar events, such as the 2004 flare in SGR 1806-20.

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.* 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

Another candidate event for GW emission from NS oscillations is the end-state of binary neutron star coalescence. 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., [476, 477, 478, 479]. 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 eventially 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 [480]. A number of studies [479, 481, 482], 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, however, this source is only detectable to $\mathcal{O}(10)$ Mpc (see [483] for a comprehensive detectability study of the short burst following the merger), requiring a serendipitously nearby coalescence. 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 [484]. 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.

3.18.3 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 [19, 485], 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 [486, 487, 474, 488, 9, 5]. 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., $\sim 1-4\,$Kz) and long-duration ($\sim 100-1000\,$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 [489] which has previously been used to search for long-lived GW signals associated with GRBs [490]. 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 LALINFERENCE 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*[18].

---

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

X-PIPELINE will be used to target short-duration bursts from 10's of Hz up to $\sim$ 4 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 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 ($\sim$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 ($\sim$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.

Detector Characterization, Calibration Requirements & Other Common Search Aspects

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

These analyses will be rerun offline as needed after any significant changes to GW data calibration and data quality. Since it is proposed that the analysis is only run after the event is confirmed as being in galactic in origin or following a BNS GW inspiral detection, it is unlikely that any reruns will be required as a result of further electromagnetic observations, as the sky location and timing information will already be known.

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 hyper-flares and GW detections from binary neutron star coalescence are rare and will be of sufficient intrinsic 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.

3.18.4 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.

Confident detection. Confident detection would be comprised of a detection candidate with significance at $> 3\sigma$, accompanied by a robust waveform reconstruction (i.e., narrow posterior measurements, free of artefacts) and spectral analysis.

Marginal detections 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.

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[19]. 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 identcal 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.

3.18.5 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 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[20] 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 [483] 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 [484] and [491], 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.

---

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

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

3.18.6 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.

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.

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.

3.18.7 Person Power

- 0.2 FTE requested for review of simulation infrastructure and sanity checks for run configurations.
- 0.2 FTE is needed for monitoring and evaluating extraordinary astrophysical events and online trigger studies/checks. This estimate is common to all NS activities.
- 0.8 FTE is needed for development and testing of short burst followups of BNS inspirals, including simulation infrastructure development.
- 0.5 FTE is needed for tuning and testing of magnetar-targeted X-pipeline analyses, including simulation infrastructure development.
- 1.0 FTE is needed for for development and testing for STAMP-based analyses of magnetar signals, including simulation infrastructure development.

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

3.19.1 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 be able to pick 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 × 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 bouyant 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 an excellent chance of detecting the GWs from a core-collapse supernova in the Milky Way or in the Magellanic Clouds. The chance of detection 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.

3.19.2 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 galactic (or SMC/LMC) 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 (as summarized, e.g., in [492, 493]) 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-driven mechanism or the strong signal from rotating core collapse can be linked to magnetorotational explosions (e.g., [494, 495]). GWs from rotating core collapse can be used to measure the angular momentum of the collapsing core [496, 497]. 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., [498, 499, 500]). An abrupt end of GW and neutrino emission would unambiguously herald the formation of a black hole (e.g., [501, 502]).

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., [493]), sensititivity estimates based on real re-colored noise with projected O1 sensitivity and a two-detector network [503], and preliminar detection upper limits from the ongoing search for GWs from distant CCSNe in S5/A5/S6 data [504]. All these studies show that a galactic CCSN at a fiducial distance of $10\,\mathrm{kpc}$ should be detectable by an L1–H1 O1 detector configuration even if pessimistic emission scenarios are considered. A non-detection would put highly significant constraints on the degree of asymmetry of the supernova engine. Detection of a CCSN at SMC/LMC distance will require 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., [505]). 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 [506]. Information about these neutrino detections will be disseminated with low latency by the SuperNova Early Warning System (SNEWS, `http://snews.bnl.gov/`, [507]) and can be used to set a tight ($\sim 4\,\mathrm{min}$, [503]) 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 [19]. 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, [508]). 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 [504] and by the sensitivity study of [503]. CCSNe are prime galactic science for Advanced LIGO and Advanced Virgo. CCSNe occuring at distances beyond the Magellanic Clouds 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., [499, 509, 493]). However, extreme emission scenarios associated with longer-lasting ($\sim$ one to two seconds) barmode instabilities (e.g., [510]) or fragmenting accretion disks [511] could be constrained for more distant CCSNe [503].

3.19.3 Search Description

SNEWS-triggered search: The goal of the SNEWS-triggered analysis is to provide an 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 [19] and will complement the all-sky online analysis with Coherent WaveBurst 2G (described in the all-sky burst search proposal, [512]). 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 [504]. 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. In addition, 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 [513].

A nearby CCSN will produce a prominent signal in the global array of neutrino detectors such as Super-Kamiokande [514, 515], Borexino [516, 517], and LVD [518, 519]. In preparation for such an event, the neutrino community has an established alert system known as SNEWS [507]. SNEWS will provide an automated email alert of "GOLD" events to registered users with an estimated latency of five minutes or less.[21] The best pointing accuracy will be approximately $5^\circ - 20^\circ$ from Super-Kamiokande [506], but this information may not be immediately available at the time of the alert. We therefore envision performing 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 [19, 485] and offline with X-PIPELINE and CWB2G. Both pipelines can operate in a single-detector mode in the case only one Advanced LIGO detector is operational at the time of the SNEWS trigger. X-PIPELINE [19, 485] is a pipeline designed for GW follow-up of astrophysical triggers. X-PIPELINE has been used in a number of LSC-Virgo observational results papers in searches triggered by gamma-ray bursts and high-energy neutrinos [486, 487, 474, 488, 9, 5]. 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 [492, 493]. 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 line as needed after any significant changes to calibrations, data quality, or as further information on the supernova is received. In particular, we forsee re-running the analysis as improved sky position information becomes available.

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, [508]), but preparation and checking will require human input and intervention. 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 2014 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 determined using the CWB2G-reconstructed detector responses. For this, code freeze will happen by the March 2015 LVC meeting. (2) A CWB2G plug-in that implements the noise reduction procedure [520] 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., [521]). If mature, this procedure will be used and reviewed by the March 2015 LVC meeting.

Parameter Estimation Follow-Up: Upon an identification of a detection candidate by X-PIPELINE (or

---

[21] The false alarm rate for SNEWS alerts is estimated to be less than one per century [507].

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 LALInference_Burst (LIB, [522]), BAYESWAVE [523], MAXENT [524], and CWB2G [508] algorithms. It is expected that LIB, BAYESWAVE, and CWB2G will be reviewed by the start of O1.
- We will employ the Supernova Model Evidence Extractor (SMEE, [495]) 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 will use the method of Abdikamalov *et al.* [496] to infer the total angular momentum of the progenitor star's inner core. 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 (lead: Caltech graduate student Sarah Gossan).

3.19.4 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. 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.

3.19.5 Technical Requirements and Development Plan

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 are:

- Adapt the online triggering system to receive and parse SNEWS alerts.
- Test the X-PIPELINE job monitoring by online system and update as needed.
- Implement and test sending LV alerts of detected GW events, with emphasis on the estimated sky position.
- Explore what choices of X-PIPELINE analysis parameters will maximize the speed and sensitivity of the search. Investigate and characterize X-PIPELINE's performance on single-detector data.

- 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 [525].)

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 [504]. 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 implementation and testing of a ring-shaped sky mask;
- investigation and characterization of CWB2G for the case of single-detector data;
- the implementation and testing of a CWB2G module of the noise reduction procedure developed by the University of Texas at Brownsville group [520]. 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.

Parameter Estimation Follow-Up: At this point, the LIB, BAYESWAVE, and SMEE parameter estimation / model selection pipelines have reached relative maturity. 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 angular momentum estimation [496], 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 [526]. The major milestones for parameter estimation are:

- 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 study (paper draft by the end of December 2014).
- Ready a matched-filtering based angular momentum estimation pipeline [496] that can be run on X-PIPELINE and CWB2G detection candidates.
- Further test, bring to maturity, and review basic signal characteristic extraction and signal reconstruction with the optimal set of the LIB, BAYESWAVE, MAXENT, and CWB2G codes. Note that MAXENT is unlikely to be reviewed by O1 and development of it has currently stalled.
- Develop and test an advanced parameter estimation pipeline based on GW spectrograms to infer progenitor physics such as postbounce accretion rate and compactness (Gossan PhD project).

3.19.6 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*.

Person-power.

- 1 FTE is needed for code development, testing, and integration of the X-PIPELINE SNEWS-triggered search.
- 2 FTE is needed for the CWB2G development and tuning (shared FTEs with the optical search).
- Approximately 3 FTE are needed in total for readying the parameter estimation framework for GWs from the next nearby CCSN.
- 0.5 FTE will be needed for review of the X-PIPELINE SNEWS-triggered search.
- 1 FTE will be needed for review of the parameter estimation codes.

3.20 Search for transients in coincidence with Fast Radio Bursts

3.20.1 Abstract

Since the publication in summer 2013 of four Fast Radio Bursts (FRBs) identified in Parkes Telescope data [527] 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 already signed with HTRU.

3.20.2 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 [528] 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 [529] and have been suggested as a possible emission mechanism for FRBs [530, 531]. A cosmic string cusp may emit gravitational waves with a $f^{-4/3}$ frequency dependence up to a cutoff frequency [532], potentially at frequencies and amplitudes detectable by ground-based interferometers [405, 27].

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 [533] 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 [534]. A third possible mechanism is the radiation at radio frequencies as a result of magnetospheric interactions [535].

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 [536]. 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 [537].

Single neutron stars - Most models of gravitational 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 [528], 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 [538], or from other asteroseismic phenomena resulting in shifting of the neutron star's crust [539]. 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 [465]. 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. [540] 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 [472]. Models for neutron star asteroseismic phenomena similar to those under discussion have also motivated previous gravitational wave searches in coincidence with SGR flares [541].

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 [542, 543, 544]. We therefore consider single neutron stars as a possible source of coincident GW and radio transient events.

3.20.3 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 and NS-BH 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 statistical significance of at least 0.01 would trigger follow-up procedures as laid out on the burst “detection checklist”. The candidate may be rejected before the entire list is executed.

3.20.4 Publication Plan

No publications are planned regarding specific FRBs or sets of FRBs in O1 regarding upper limits set on FRB triggers. “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 specific analysis.

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 [545] and appear to occur at galactic distance scales based on observed dispersion measures. 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.

3.20.5 Resources

Computing needs - Requirements on a per-trigger basis are similar to the burst analysis of a Swift GRB. We estimate that 300 CPU-days 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) 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 [546]. 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.

3.21 Contributions of long transient gravitational waves to the stochastic background search

3.21.1 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.

3.21.2 Scientific Justification

The scientific rationale for this search is based on the following observation: one or more long-lived transient signals 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 transient gravitational-wave signal, lasting $\approx 10\,\mathrm{hr}$, corresponding to 1% of the total coincident data during O1, assuming a (previously typical) duty cycle of 50%. A $10\,\mathrm{hr}$-long signal, with a stochastic signal-to-noise ratio [327] of SNR $= 50$, would produce an SNR $= 5$ signal through dilution in the full O1 stochastic search. However, the same signal would produce only a SNR $= 0.26$ 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. Clearly, individual transient signals longer than $10\,\mathrm{hr}$ could influence the stochastic measurement with even smaller stochastic SNR. Similarly, gravitational wave signals much shorter than $10\,\mathrm{hr}$ are unlikely to influence the stochastic measurement, unless they are very frequent during O1 [315, 547].

Most previous work on long-duration transients has focused on signals in the regime of $\approx$10–3000 s. However, longer signals are possible as well, including very long-lived signals (potentially lasting weeks) coincident with gamma-ray bursts [548] and individual neutron stars (including signals from type-I bursts and glitching/accreting millisecond pulsars) [549]. The detectability of such very long signals is highly uncertain. 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 [332, 344].

The scientific rationale for a search for very-long transients, spanning hours to weeks, was first explored in [549]. The authors of [549] 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 [550, 551, 552], Alfvén oscillations from giant magnetic flares (also lasting days to months) [553, 554], emission from free precession (with a damping time possibly lasting from weeks to years) [555, 556, 557], magnetic instabilities in newborn neutron stars (lasting days) [558], and gravitational-waves from r-modes [559, 560]. In addition, generic rotational instabilities in newborn neutron stars, potentially powered by fallback accretion [561, 562], may persist on a timescale of hourss [548]. Finally, it is worthwhile to be prepared for a surprise: a very long lived transient signal not predicted by modelers [549].

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 a deep understanding of 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). It is currently not clear how STAMP (which we propose using for this search plan) compares to CW pipeline(s). In any case, we will work with the CW group to understand potential complementarity and adjust the search accordingly.

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 $\Omega_{\rm gw}$ [327] 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 [489, 563, 490, 564, 565]. STAMP is a Matlab-based code package that resides in the Matapps svn repository [566]. STAMP works by cross-correlating data from two detectors to produce cross-power spectrograms [489]. Gravitational-wave signals appears as tracks of brighter-than-usual spectrogram pixels. STAMP employs a user-specified clustering algorithm (there are a few options [489, 567, 564, 565, 568]) in order to identify statistically significant clusters of pixels. One of these clustering algorithms, which we shall refer to below, is called Stochtrack [564]. 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 [490].

A project to carry out an all-sky search with initial LIGO data is complete [569]; 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 [569]. The S5 dataset has been analyzed and no candidate events were observed. A paper draft describing the search is available on the DCC [570].

There has also been significant progress recently improving the sensitivity of STAMP for advanced detectors. In particular, recent work shows that a technique called “seedless clustering” can dramatically improve the sensitivity of searches for long-lived signals (at increased computational cost) by a factor of up to two in distance (7.4 in volume) compared to seed-based methods and depending on the waveform [564, 565]. The new seedless clustering algorithm code 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 [565] 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 [571]. 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 [565].

We will also study the distribution of intermediate-duration triggers lasting $\approx 10\,\mathrm{s}$–$10\,\mathrm{hr}$ in order to constrain their potential contribution to the energy density. We will attempt to use trigger distributions generated by all transient pipelines (e.g., from CBC searches for binary neutron stars) applied to the same data. While transient phenomena on this time scale are of potential interest to the Burst and CBC Groups, rather than focusing on the loudest candidate(s) we will focus on understanding the distribution of the zero-lag triggers more generally. This is because, as noted above, signals shorter than $\approx 10\,\mathrm{hr}$ could contribute to the stochastic measurements only if they are sufficiently frequent during O1. We will work within the Burst and CBC Groups to coordinate any overlapping efforts in the $10\,\mathrm{s}$–$10\,\mathrm{hr}$ regime. We note that the CBC and Burst groups actively study the distribution of their triggers, especially pertaining to how it can provide information on the noise background. We will also coordinate with Burst and CBC Groups to ensure that zero-lag transient distributions used here are not unblinded in a way that may bias ongoing Burst and CBC searches.

3.21.3 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).

3.21.4 Technical requirements and development plan

Pipeline. The baseline plan is to carry out the analysis with STAMP [489, 563, 490]. STAMP (and Stochtrack) are Matlab-based code. They are part of Matapps [566]. Most of the STAMP code was reviewed for the published initial LIGO targeted analysis [490], 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 [564, 565] 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 [564, 565]. 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 [563, 490]. This is attributable to the long time scales probed by the search as well as the fact that STAMP uses cross-correlation [563]. Subsequent investigations have applied seedless clustering techniques to recolored initial LIGO noise leading to comparable results [564, 565]. There are no special requirements for calibration beyond what has already been requested by the Stochastic Group [331].

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 [564, 565]. 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 [572, 573] and follow-up/visualization of CBC triggers [568]. We expect continued development and maintenance of STAMP will be broadly useful for Stochastic Group activities and the wider LSC/Virgo community.

4 Ongoing searches on Initial Detector Data and activities not ready for a Search Plan

Here we summarize activities on initial detector data and work that is not ready for a mature search plan.

4.1 Burst Working Group

The Burst Group activity is distributed across eight science teams:

- **All Sky – All Time:** the most general search for GW transients, which is also the starting point for certain targeted source analyses;
- **Gamma-Ray Bursts (GRB):** GWs associated with gamma-ray burst triggers;
- **High Energy Neutrinos (HEN):** GWs coincident with high energy neutrino events;
- **Supernova (SN):** GW searches triggered by optical or low-energy (MeV) neutrino observations;
- **Neutron Star Physics (NS):** GWs from isolated neutron stars, including searches triggered by magnetar flares, pulsar glitches and follow-ups of binary neutron star mergers for post-merger GW signals from surviving hypermassive neutron stars;
- **Other Electromagnetic Counterparts (EM):** GWs associated with X-ray, UV, optical, radio counterparts, using searches triggered by EM events other than the specific sources mentioned above, or else triggered by GW events and followed up with EM observations;
- **Binary Black Holes (BBH):** searches for GWs from intermediate mass inspirals and eccentric BBH systems, with close ties with the CBC group;
- **Exotica:** alternative theories of gravity, cosmological defects (cosmic strings), bursts with memory.

This organization is centered on science targets and sources, rather than analysis techniques. Although the ultimate decision on the science goals belongs to the full Burst Group, each team is charged with formulating the science case for its analysis, proposing the methods and defining a coherent publication strategy. Each team is responsible to lay out a timeline with readiness milestones, using the engineering runs schedule to plan on testing, review accomplishments and report to the whole Burst Group. Moreover each team has to prepare a Search Plan for each planned search.

New searches or teams will be embraced by the group only once they have proven their potential with astrophysical motivation and a viable implementation plan. The team should consider the likely signals; what astrophysical interpretation(s) should be pursued for a non-detection or for a detection; the suitability of existing analysis tools and techniques; what specific improvements need to be made; and what new tools or techniques (if any) need to be developed. Simulations and astrophysical input will help decide whether a specific source requires a targeted analysis or a different tuning and interpretation of the standard all-sky and externally triggered searches.
In the following we summarize the ongoing activities which are not yet described in mature Search Plans.

4.1.1 [All-Sky] All-sky searches for long-duration transients

An all-sky (untriggered) search for long-lived gravitational waves. In collaboration with the stochastic group, we are pursuing an all-sky search for long-duration transients; see Section 4.3. STAMP (Stochastic Transient Analysis Multi-detector Pipeline) is an extension of the stochastic pipeline. STAMP analyzes the same cross-correlation data products used in a stochastic search in order to identify long-lived transient signals lasting $\gtrsim 10\,\mathrm{s}$. The STAMP All-Sky (STAMP-AS) analysis is an untriggered search looking for long duration $\sim 10-600\,\mathrm{s}$ GW transients [569]. Possible astrophysical sources for such signals are accretion disk instabilities, instabilities in nascent neutron stars, eccentric black hole binary systems, and neutron stars quasi-periodic oscillations. To carry out this search, several enhancements have been made with respect to the original STAMP code such as a faster clustering algorithm, allowing the code to look for a signal everywhere in the sky even without investigating every sky direction, and a more efficient mechanism to cross-correlate data. A dedicated STAMP-AS pipeline has been run on $\approx 1\,\mathrm{yr}$ of coincident Hanford-Livingston data from S5, including 200 time slides for background estimation and a thorough sensitivity study for a dozen test waveforms. (The STAMP-AS team has also performed extensive data-quality investigations.) At present, the analysis team is ready to open the box and to begin a review.

Advanced detector era. The STAMP code suite has reached a mature stage. The first analysis targeting GRBs has been published [490]. The more computationally challenging STAMP all-sky analysis is nearly complete; see above. The improved sensitivity of aLIGO, coupled with recent improvements to STAMP clustering algorithms, should allow us to probe long-duration transient models out to astrophysically interesting distances: $\approx 600\,\mathrm{Mpc}$ for extreme models emitting $E_{\mathrm{GW}} = 0.1 M_{\odot}$ of energy in GWs [564]. Recent work has demonstrated that seedless clustering can improve the sensitivity of both targeted and all-sky searches [564, 565] by a factor of up to two in distance. Other work has investigated new techniques for parameter estimation [574] and specialized searches targeting compact binary systems [568] (envisioned as a follow-up to matched filter triggers). The ADE triggered STAMP search will be applied to interesting electromagnetic triggers including gamma-ray bursts and supernovae. The ADE all-sky STAMP search will be run on all high-sensitivity data.

4.1.2 [HEN] High-Energy Neutrino Multimessenger Analyses

Many of the most violent and interesting phenomena producing GW transients are also thought to be sources of high-energy neutrinos (HENs) [575, 444, 445, 446, 447, 576, 488]. These non-thermal, GeV-PeV neutrinos neutrinos are thought be produced in relativistic outflows driven by central engines also responsible for GW emission [577, 578, 579, 452, 453, 454, 580, 456]. Both long and short gamma ray bursts (GRBs), core-collapse supernovae with fast rotating cores, and highly magnetized neutron stars (magnetars) are thought to produce GWs and HENs that may be detectable out to relevant distances [576].

There are multiple scientific benefits of simultaneously observing GWs and high energy neutrinos from a common source:

- **High-confidence detection** – Both GWs and neutrinos are weakly interacting messengers, and so far there has been no confirmed detection of either sources of cosmic origin. The combined information from GW and HEN observatories can greatly enhance our confidence in a joint detection [575, 444, 445, 446, 447, 576, 488, 459]. In particular, a comparison [447] of joint GW+HEN searches using advanced GW detectors and the completed km^3 IceCube detector to the reach of independent searches using the same detectors concluded that, while the main advantage of joint searches is increased sensitivity for the actual detection of sources, joint searches will provide better constraints than independent observations if, upon non-detection they result in an increased exclusion distance by at least a factor $\sim f_b^{1/3}$ compared to independent searches, where f_b is the neutrino beaming angle. This study derived the first observational constraints on common GW-HEN sources using initial LIGO, Virgo, and IceCube data, and also projected population constraints from joint searches with advanced GW and HEN detectors.
- **New probe of the depths of violent astrophysical phenomena** – GWs and HENs 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 HEN flux is reflective of the presence of hadrons in the relativistic outflow generated and driven by the central engine. 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, their surrounding, and the nature of the outflows. For example, it

has recently been demonstrated [463] that the energy-dependence of the onset time of neutrino emission in advancing relativistic jets can be used to extract information about the supernova/gamma-ray burst progenitor structure. Together with observed GWs, this would provide information on the inner density of the progenitor beneath the shock region ($\sim 10^{10}$ cm for mildly relativistic jets). In favorable conditions, very few neutrinos, in coincidence with a GW signal, would be sufficient to provide important information, and/or to differentiate between progenitor types.

- **Prospect of common sources dark in gamma rays** – The emission of HENs is tightly connected to the presence of high energy photons (gamma rays) in the outflow. There are specific cases where the source optical thickness is large and prevents the gamma-rays to escape from the source. One of the most interesting prospects of joint GW - high energy neutrino searches are common sources that are dark in gamma rays. One of the prominent such sources are choked GRBs [579, 460, 461] or low-luminosity GRBs [581, 582, 583, 584, 585, 586, 587, 588]. These sources are difficult to detect with electromagnetic observatories, and hence provide an exciting opportunity to joint GW - HEN searches that can discover them and/or constrain their population [446, 447, 459]. Further, it is plausible that previously unanticipated sources or mechanisms can be discovered and studied with joint searches.

Currently operating HEN observatories include IceCube [462], a cubic-kilometer detector at the South Pole, and ANTARES [589] in the Mediterranean sea. ANTARES is proposed to be upgraded to cubic-kilometer detector (called KM3NeT) in the coming years [590]. A third HEN detector is operating in lake Baikal and has been proposed to be upgraded [591].

There have been coincident data taking periods between initial GW and HEN detectors in the last few years, providing datasets that have already been used to derive population constraints on joint sources [447]. This includes the first coincident search for GWs and HENs, using the S5/VSR1 data and the partial ANTARES detector in its 5-string configuration. The analysis uses the directional distribution and the time of arrival of HENs to trigger a GW follow-up analysis, similar to the analysis used for GW follow-up searches of GRBs. This analysis has been completed and published [488].

The joint analysis of the LIGO S5/S6, Virgo VSR1/VSR2/VSR3 and IceCube 22/59/79 has been finished. The analysis was based on the GW-HEN baseline search method [464, 459]. It takes into account the significance of the GW and HEN signals, calculated using the excess energy in the GW datastream (see e.g. [592]), and the reconstructed neutrino energy and neutrino flux (i.e. number of coincident neutrinos). The analysis also took into account the directional probability distributions of the GW and HEN signals, as well as the *a priori* source distribution using the observed distribution of blue luminosity in the universe[22]. The joint search used the coincidence time window derived in [446]. The search found no significant coincident signals. It used the results to obtain upper limit estimates of multimessenger GW+HEN sources, as well as to derive the projected sensitivity increase of searches in the advanced GW detector era. The analysis was designed and carried out to be applicable to searches in the advanced detector era, and will be used for this purpose once GW data becomes available.

The joint analysis of ANTARES 12-line and LIGO-Virgo S6-VSR2/3 data is under review. A layer of improvement to the baseline analysis has been developed and applied to this data set [220]. It consists of a modification of the GW event generation algorithm which allows the full use of all GW data (currently restricted periods when three GW detectors are operating).

Searches in the advanced GW detector era will be carried out using the operating HEN detectors (IceCube and possibly Antares/Km3NET). They will use the searches and search parameters developed for earlier searches. For example, the maximum time difference between the arrivals of the observed GW trigger and HEN events [446], one of the key parameters of the joint GW+HEN search algorithm, will be usable for advanced searches as well. Here, a too small time window might exclude some potential sources, while a too large time window would unnecessarily increase the false alarm rate and the computational cost.

---

[22] I.e. the analysis assumes that the source distribution follows the blue-luminosity distribution of galaxies.

Low-latency multimessenger GW+HEN searches will constitute an interesting new direction for the advanced detector era. Both GW and HEN 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 HENs can arrive prior to the onset of electromagnetic emission from sources such as GRBs, joint GW+HEN events may be primary targets for electromagnetic follow-up searches.

In short, GW+HEN observational results have already proved to produce exciting scientific results [447, 488], while the projected constraints [447] and expectations (e.g., [463]) suggest that multimessenger GW+HEN searches will be a fruitful direction of research during the advanced detector era.

4.1.3 [SN] Supernovae as Astrophysical Triggers

Core-collapse supernovae are interesting candidates as gravitational-wave sources and can be studied in conjunction with both neutrino and optical messengers. The theoretically expected GW signal is most likely very weak (e.g., [493, 593, **?**]), but a galactic core-collapse supernova would be detectable even in pessimistic models. Extreme emission scenarios can be constrained with advanced LIGO out to a few few Mpc.

Most optical triggers carry the burden of a large uncertainty on the derived event time (order of several hours or more), making the GW data analysis task challenging due to large backgrounds and varying detector duty cycles. Well-known sky locations are a significant aid to the analysis. **A near-term goal is completion of the optical supernova search in data taken by the initial detectors**, which is underway and currently at the review stage. It targets four S5/A5/S6 core-collapse supernovae with will constrained explosion times: SN 2007gr, SN 2008ax, SN 2008bk, SN 2011dh. These core-collapse supernovae exploded in host galaxies $\lesssim 12\,\mathrm{Mpc}$ from Earth and the GW search may be able to constrain the most extreme GW emission models. The enhanced reach of advanced detectors will be able to constrain a more significant swath of the model space. We will propose an advanced-detector era supernova search upon completion of the current search.

Supernova triggers from detectors sensitive to low energy (up to tens of MeV) neutrinos can be used in GW searches as well. For example, a core-collapse supernova near the galactic center is expected to produce a large flux of ~ 8000 detected neutrinos in the Super-Kamiokande detector [594] with a pointing accuracy of 4°. Unlike photons, which can take up to a $\sim$day to break out, neutrino bursts and gravitational waves mark the moment of core collapse, and are expected to be coincident in time to $\lesssim 1$ s (most likely, GW and neutrino emission will set in within milliseconds of each other [493, 595]). The expected strong neutrino signal for a galactic core-collapse supernova would provide excellent timing and good pointing, thus allowing an improved sensitivity gravitational-wave burst search, similar to that employed for GRBs. For extragalactic supernovae, the neutrino signature would in general provide timing but not pointing. At the distance of Andromeda, the expected flux of detected neutrinos in Super-Kamiokande would fall off to $\mathcal{O}(1)$. In this case, joint neutrino-GW time-coincident searches would substantially increase detection probability and decrease the false-alarm rate.

A proposed joint search for GWs and low-energy neutrinos from core-collapse supernovae on archival (and future) data has been approved by the LSC and Virgo in early 2014. This is accompanied by an MoU commonly agreed with the neutrino experimental groups of Borexino, LVD and IceCube. The Super-Kamiokande collaboration is not intending to be part of this effort at this time. An exchange of sub-threshold events from the neutrino and gravitational-wave detectors from recent data taking is expected to take place in order to develop joint search strategies, including assessing the improvement in sensitivity, detection confidence and coincidence livetime coverage the joint neutrino-gravitational-wave detector network will provide. We expect such work to pave the way for the real-time participation of gravitational-wave detector to the supernovae early warning system (SNEWS) once LIGO and Virgo return to science running in 2015 and beyond. **Our near term goal is to establish the joint, low-threshold search methods for core-collapse supernovae using archival data from LIGO, Virgo and the neutrino partners in this effort**

and to characterize such search in terms of improvement in sensitivity.

In the case of a supernova GW detection, likely from a Milky-Way or Large/Small-Magellanic-Cloud core-collapse supernova accompanied by a neutrino signal, the two urgent question will be (1) what can the signal teach us about the physics underlying the core-collapse supernova explosion mechanism, and (2) what can the signal tell us about the about the characteristics (e.g., core rotation rate, core structure) of the star that is collapsing and exploding? Both questions are best approached with model selection and parameter estimation algorithms. The supernova group is in the process of developing multiple approaches to extract physics from a detected supernova GW signal.

The Supernova Model Evidence Extractor (SMEE), a Bayesian pipeline, is capable of deciding between neutrino-driven, magnetorotational, and acoustic-mechanism core-collapse supernova explosions [495]. **Our near term goal is to prepare and characterize SMEE for the realistic multiple-detector aLIGO detection scenario of gravitational-waves from the next nearby supernova.**

In [496], it was shown that for the case of rotating core collapse, a matched-filtering approach can be used to *measure* the angular momentum of the inner part of the collapsing core to $\sim$30% accuracy and to constrain the degree of differential rotation. This work was followed up by the Bayesian principle-component regression (PCR) parameter estimation approach of [596] and by the frequentist multivariate approach of [497]. The multivariate approach is a particularly powerful addition to our set of parameter estimation tools, because it for the first time allows the prediction of rotating core collapse waveforms on the basis of astrophysical parameters (e.g., degree if differential rotation, angular momentum in the core). **Our near term goal is to incorporate the new multivariate model of [497] in a Bayesian parameter estimation framework that combines the strengths of SMEE and the work of [596]. Moreover, we plan to approach the parameter estimation problem for GWs from neutrino-driven convection, which introduces stochastic components to the GW time series.**.

In the intermediate term, we intend to incorporate a basic ability to identify supernova-like signals and estimate parameters / select models into a broader low-latency burst parameter estimation pipeline. **A short term goal is to study the ability of the existing pipelines MaxEnt [524], coherent WaveBurst 2G [17], and BayesWave [523] to reconstruct supernova waveforms from noisy data.** Subsequently, we will apply a model selection technique similar to SMEE and compare its performance on reconstructred waveforms to boosted decision tree algorithms.

4.1.4 [NS] GWs as probes of Neutron Star Physics

Isolated neutron stars in our galaxy may be sources of detectable GW burst via a number of processes. Searches for GWs from these systems are carried out in coordination with the Stochastic and Continuous Waves Groups as appropriate.

Burst Followups of Binary Neutron Star Inspirals

The likely outcome of binary neutron star coalescence is the formation of a metastable and possibly hyper-massive neutron star [476, 597, 598, 599, 600, 601, 477, 602, 480, 603, 478, 604, 605, 479, 606, 481, 607, 608]. Transient non-axisymmetric deformations in the post-merger remnant lead to a short duration ($\sim$ 10-100 ms) burst of high-frequency ($\sim$ 2-4 kHz) GWs. Although this merger/post-merger signal may only be detectable to $\sim$ 10 Mpc, yielding rather low detection rates ($\sim$ 1 per century at aLIGO design sensitivity), the opportunity for increased detection confidence by association with the inspiral signal and the potential for constraining the EOS of matter at supra-nuclear density [479, 606, 608] via an independent and complementary channel to related inspiral measurements make the post-merger signal from BNS systems a potential source of high-profile burst science.

Investigations are underway to develop a search strategy using traditional burst detection pipelines and newer Bayesian parameter estimation methods.

Triggers from Magnetars

Past externally triggered searches have looked for GWs associated with bursts from magnetar candidates – soft gamma repeaters (SGRs) and anomalous X-ray pulsars (AXPs) – following two distinct search strategies. The first strategy looks for GW emission associated with the prompt gamma-ray burst from individual magnetar bursts and giant flares. An emphasis is placed on GW emission from the damping of f-modes in the magnetar which may be excited in the burst, but the entire detector band is searched up to 3 kHz. This strategy has been used to analyze over 1400 electromagnetic SGR and AXP triggers, including the 2004 giant flare [469, 471]. The second strategy looks for the collective signature of weak GW bursts from repeated flares from a single SGR source, stacking the GW data to dig deeper into the noise at the expense of additional but plausible model dependence. This strategy was used to analyze the 2006 SGR 1900+14 storm [470]. New strategies, such as that used to search for long-duration bursts from GRBs [211], are being explored to increase sensitivity to potentially interesting lower-frequency modes which might be rung up in association with flares. These can be tested with initial LIGO-Virgo data, possibly leading to an electromagnetically triggered archival search for long duration GW bursts associated with magnetar flares.

Triggered r-mode Searches

Investigations into excitations of pulsars (post-glitch) and accreting millisecond pulsars (post-flare) have shown that gravitational waves emitted at r-mode frequencies are a potential source for second and third generation gravitational wave detectors.

A search strategy has been formulated with members of the Continuous Waves working group to search for these long-lasting transient signals. Trial runs on simulated data have shown the search strategy capable of identifying the presence of long transient signals at strengths estimated by previous feasibility studies. The next step is to perform a comprehensive characterisation of the search using both simulated data and data acquired by interferometers. Efforts will also be made to translate the outputs from search codes into astrophysically interesting constraints on the source parameters.

4.1.5 [EM] Other Electromagnetic Counterparts and Follow-up Observations

The previous subsections described several scenarios in which astrophysical systems are expected to emit electromagnetic (EM) radiation along with GW bursts, and the substantial benefits of joint observations. Here, we consider other EM counterparts where the nature of the source is unclear, and/or where the EM signature is present but was not initially detected by surveys.

Radio burst triggered GW searches

Bursts of radio waves ("Fast radio bursts", or FRBs) with large dispersion measures suggesting an extragalactic origin have been detected by a few radio telescopes [609, 610, 527, 611], but their origin remains mysterious. Some are strongly suspected to have a terrestrial origin [612, 613], but the question of how such highly dispersed signals are created is currently a subject of intense debate. It is plausible that at least *some* FRBs are from genuine astrophysical events. Since energetic events which produce detectable GWs are likely to feed some of the released energy into EM emission, radio transient counterparts are quite possible. The radio transient either could be a short, prompt pulse or else a longer-duration afterglow (rising and peaking days, weeks, or months later, depending on the radio frequency band in which it is observed.) A short transient would result from coherent emission, e.g. from some plasma excitation, while a radio afterglow could be synchrotron emission from an expanding jet or fireball.

Prompt radio bursts are natural candidates for joint analysis with GW data, especially because there are theoretical models in which various mechanisms give rise to a prompt pulse of radio emission from some putative gravitational wave sources, e.g. compact binary coalescence scenarios or cosmic string cusps [531]. Therefore, the Burst Group science program includes GW burst searches targeting the times and sky locations of reported radio bursts recorded by radio telescopes such as Green Bank [614], Arecibo, LWA, and potentially others.

An interesting aspect of follow-up of radio pulse transients is that multi-channel radio receivers and digital signal processing provide a "dispersion measure" which indicates the integrated electron density through which the radio signal traveled. Although the dispersion delay can be many seconds or minutes, de-dispersion yields the time that the radio pulse would have arrived if it had not been dispersed, allowing us to better predict when the gravitational wave should have arrived at our detectors. Tasks for this area include:

- Identify interesting radio transients, taking into account whether they appear to be truly astrophysical
- Consider the detectability of possible GW sources that could have produced the radio transients
- Develop search techniques with appropriate coherent/coincident analysis conditions
- Complete searches
- Consider what conclusions can be drawn from positive and negative search results
- Formulate a good plan for joint radio-GW searches in the advanced detector era

EM follow-ups of GW event candidates

Telescopes cannot cover the whole sky at all times with good sensitivity. Current EM survey projects typically have lengthy cadence times, or are considerably less sensitive than their more directed counterparts [615]; therefore it is quite possible that the EM signature from an interesting event will be missed because no telescope of an appropriate type is looking in the right direction at the right time. The GW detector network, on the other hand, is effectively an all-sky monitor for highly energetic astrophysical events. Thus there is a strong motivation to point optical, X-ray and/or radio telescopes in response to potential GW triggers and thereby attempt to catch an associated transient EM signal which would otherwise be missed.

For example, because the gamma rays from GRBs are believed to be beamed (and because even the Fermi Gamma-ray Burst Monitor covers only half the sky), there may be short-lived "orphan afterglows" from nearby GRB progenitors which are going undetected [616]. Other possible sources of joint EM/GW emission include decaying neutron star matter ejected during merger [617, 618, 619, 620] and supernovas [**?**]. Some more details on these sources and their expected observational signatures can be found in [621, 615]. If such an event is detected by the network of GW interferometers, then the reconstructed sky position can in principle be used to point one or more telescopes and catch the afterglow before it fades away.

Like externally triggered GW searches, GW-triggered EM follow-up observations can be beneficial in two ways. First, they would help establish the event as astrophysical in nature, and may increase the reach of the interferometer network in the case of a GW event that is not quite strong enough to stand out from the background. Additionally, having an associated EM signal increases the astrophysical information that can be mined from a GW event [622, 623, 624]. Note that both prompt and delayed follow-up observations, as well as simply checking against lists of transients identified by other surveys, are appropriate to catch different possible light curves.

Scientific prospects for EM follow-up observations with specific observing strategies have been studied over the past several years [621, 625, 201, 626], and a first complete EM follow-up projected was implemented and carried out during the S6/VSR2-3 science run [627]. For the Advanced Detector era, we will again search for burst signals with low latency as the data is collected (complementing the low-latency CBC search effort which is part of the binary neutron star search plan outlined in section 3.2). Burst candidates identified in the data will be inserted into the GraceDB database, with reconstructed sky maps. Signifi-

cant events will be selected and packaged for alerts to participating observing partners, enlisted through the LSC-Virgo partnership protocols that are currently being worked out. Tasks for this area include:
1) Generating GW alerts

- Low-latency analysis — included in the all-sky burst search plan (section 3.1)
- Event selection
- Position reconstruction (and studies)
- Fast parameter estimation (and studies)
- Study effect of calibration errors
- Information storage and monitoring
- VOEvent packaging
- Human validation procedures (as needed)
- Transport to observing partners

2) Supporting the usage of GW alerts

- Software for unpacking and interpreting VOEvents
- Provide test events for partners to practice with
- General observation planning tools ("observing calculator")
- Galaxy catalogs: general usage options, formatting, integration with GW sky maps, usefulness studies

3) Using GW alerts to make and analyze EM observations

This is the "external" component of the EM follow-up campaign, to be done by observing partners and by some LSC-Virgo members using their other research time.

4) Working with joint GW-EM observing results and analysis

- Bulletin Board system: receive/collect EM observations
- Bulletin Board system: display EM observations
- Event coincidence conditions
- Joint confidence calculation
- Joint parameter estimation
- Interpretation of GW data in the context of EM counterpart info
- Publication readiness

4.1.6 [BBH] Explore Binary Black Holes in a Wide Parameter Space

Burst searches are usually un-modeled or weakly modeled to be open to the broad spectrum of sources. In particular, they can be sensitive to a wider class of CBC sources than dedicated inspiral template searches. The burst approach is robust against possible discrepancies between theory and nature, albeit at the cost of reduced search sensitivity and increased false alarm rates.

The goal of the burst BBH searches is to explore as wide as possible the parameter space of BBH sources which may not be accessible by the CBC matched-filtering searches due to the lack of complete or accurate template banks. In principle, the existing all-sky pipelines can be used for such searches. However, for specific source signatures a better reach can be achieved by introduce a weak-model constraint that allows to better suppress the false alarm rate, but still preserves the robustness of un-modeled burst search. For instance, such dedicated searches for intermediate mass BBH (100–400$M_\odot$) have been conducted on the initial LIGO and Virgo data, by enforcing an elliptical polarization constraint [628, 629]. In preparation for the first detection of BBH sources in a wide range of BBH parameters with arbitrary spin configurations and high eccentricities, the burst BBH working group identifies the following science targets, pursued in collaboration with the other working groups.

Conduct dedicated BBH searches. The Burst group is working on several dedicated BBH searches: a) Intermediate Mass Black Hole Binaries (IMBHB) (a search for these systems is part of the aLIGO search plans [see Section 3.6]), b) Intermediate Mass Ratio Inspirals (IMRI) and c) Eccentric Binary Black Holes (EBBH), which is discussed below.

The scientific motivation of the IMRI search is discussed in details in Section 4.2.1. The IMRI sources is challenging to detect both for the CBC and the burst algorithms. In the first case, it is very hard to obtain accurate IMRI waveforms required for construction of the CBC template banks. In the second case, it is hard to efficiently detect low chirp mass ($M \lesssim 50M_{\odot}$) waveforms, which span some part of the IMRI parameter space. In the burst group it has been demonstrated that the upgraded Coherent WaveBurst algorithm (cWB2G) is capable to achieve the performance close to a matched filter for IMRI sources with the chirp mass down to $10M_{\odot}$, which should cover the range of chirp masses expected for IMRI sources. The IMRI search plan for advanced detectors is in preparation.

The EBBH sources may be formed by dynamical interactions in galactic nuclei containing a supermassive black hole (SMBH), or in globular clusters [630]. If stellar-mass or intermediate-mass black holes form a dense population around SMBHs, the probability of close encounters of two black holes could be high [631]. Such encounters may lead to the formation of binary black hole systems. The initial eccentricity of such binary system is likely to be close to unity and remain large all the way to the merger. The merger may happen within hours, and such short lived systems are expected to have a unique GW signature: a series of short bursts. There are no accurate eccentric binary black hole waveforms with the eccentricity greater than 0.05 available at this time and the burst searches may be the only way to detect and study binary sources with high eccentricity. The search for eccentric binary sources is being performed on the initial LIGO data with the cWB2G algorithm. The EBBH search plan for advanced detectors is in preparation.

Detection with confidence. Most of anticipated BBH signals are expected to have relatively short duration (few seconds or less) and can be easily confused with the instrumental/environmental artifacts (glitches). Identification and rejection of such false events is a serious challenge in the burst analysis. Therefore, advances in the detector characterization (§5.1,**??**) and data regression are critical for the burst BBH searches. To improve the BBH detection efficiency and reduce false alarm rate, the BBH searches employ model constraints. The group will improve existing constraints on the GW polarization states and develop new robust BBH constraints based on the time-frequency pattern recognition. A special concern is a production of large background samples (by at lest a factor of 10 compared to the S5/S6 burst analysis) and better estimation of the false alarm rates, Background production is very CPU intensive and require the development of efficient analysis algorithms.

Development and validation of astrophysical waveforms. An ongoing collaborative effort with the CBC group is studying the use of the full coalescence BBH waveforms, including the inspiral, merger, and ringdown phases. The same set of phenomenological and EOBNR waveforms [632, 439, 633, 634] is being analyzed by burst search methods, inspiral and ring-down matched filtering to compare their detactability across the BBH parameter space. We will also pursue active development of the faithful BBH waveforms with spin and eccentricity [630, 631, 635, 636]) covering the BBH parameter space as wide as possible. Studies of such waveforms will help to devise new detection techniques and provide important guidance for the interpretation of the search results. The astrophysical BBH waveforms will be used to quantify the reach of the BBH searches and in the parameter estimation studies.

Coordinate reconstruction. Sky localization is a challenging problem for the BBH analysis, particularly for sources with high masses. Such BBH sources are expected to produce a signal at low frequency where the triangulation capabilities of detector networks are affected by the diffraction limit: ($\lambda/d \sim 1$, where λ is the GW wavelength and d is the network base length. For these reasons it is important to develop coherent network reconstruction algorithms which use advantage of the antenna polarization coverage. The burst BBH working group concentrates on the un-modeled (robust) sky localization in close coordination with the CBC group and the astrophysical follow up effort.

Waveform reconstruction and source parameter estimation. In preparation for a detection, the working groups need to improve parameter estimation techniques One of the main priorities of the BBH working group is a weakly-modeled reconstruction of the BBH waveforms, Such analysis can identify the polarization state of the BBH system and determine such source parameters as the binary inclination angle and eccentricity. The reconstructed waveforms can be compared to known models for extraction of other source parameters (component masses, etc.). Progress towards waveform reconstruction has already been made via coherent techniques, but more progress is needed to compare a candidate to waveform parameters. Bayesian [637] and MCMC [191] techniques are currently being explored by the CBC Group as well as members of the Burst Group. We are open to exploration of new techniques which may prove useful for reconstruction of BBH waveforms.

4.1.7 [EXOTICA] Beyond Standard Theories and Sources

A prototype analysis looking for cosmic (super)string signatures has been performed on S4 data [638]; this search was extended to the remaining of LIGO and Virgo data collected through 2010 and over 625 days of livetime. The paper publication with the result of the search with the initial detectors has appeared in the literature in 2014 [27]. No event was found consistent with a GW signal from cosmic strings. From this, new constraints on cosmic string parameters were derived. These complement and improve existing limits from previous searches for a stochastic background of GWs from cosmic microwave background measurements and pulsar timing data. In particular, if the size of loops is given by the gravitational backreaction scale, we place upper limits on the string tension $G\mu$ below 10^{-8} in some regions of the cosmic string parameter space. We will resume searches for cosmic strings with the start of the advanced detectors' data-taking.

Search strategies for alternative theories of gravity have already started being investigated within the Burst Group. For instance, a preliminary study has shown that the standard all-sky burst search pipeline has good sensitivity for GW bursts with a scalar polarization component (rather than the standard tensor polarizations), while a modified pipeline can achieve slightly better sensitivity for such signals [639]. A supernovae-triggered search for scalar GW burst signals, using a modified version of the pipeline normally used in the GRB triggered searches, was also explored. Future work in this area may lead to a dedicated search or searches specifically targeting alternative GW signals, although a more complete advanced detector network will likely be needed to reliably distingiush the scalar and tensor modes. Tests of general relativity based on the relative timing of GW and EM signals, or on detailed source dynamics revealed in GW waveforms, are also possible.

A search for bursts with memory may complement searches for generic bursts or from binary systems. A formal proposal for a development project/search for such signatures was presented to the bursts working group in September 2012[23]. We have identified the potential sources, including predicted energy and strain scales associated with them. We have no major progress to report following the formulation of that search proposal and until now. What is foreseen in the 1-2 year horizon are proof-of-concept/prototypical searches: we plan to adopt existing general burst-search methods as well templated search methods like the one used in the search for cosmic (super)strings in order to look for memory in archival data first. This includes adding functionality within LAL for generation of such signal morphologies. The emphasis is mainly for the advanced detector era, but any significant progress until then in terms of methods and proof-of-concept work using data already in hand may lead to publication. Building phenomenological waveforms and developing injection machinery in order to benchmark existing burst methods are of the top priority and near term goals.

---

[23] `https://dcc.ligo.org/LIGO-G1200974`

4.2 Compact Binary Coalescences Working Group

4.2.1 Intermediate-mass-ratio inspirals (IMRIs)

Gravitational-wave sources:
Intermediate-mass black holes (IMBHs) with a mass between ~ 100 and $\sim 10^4$ solar masses occupy the mass range between stellar-mass and massive black holes (see [213] for a detailed review). There is growing but still ambiguous evidence for IMBH existence, including observations of ultra-luminous X-ray binaries (e.g., [640]). A number of formation mechanisms have been proposed, which may lead to the existence of IMBHs of a few hundred solar masses in globular clusters. Such IMBHs are likely to form binaries which can be hardened through dynamical interactions to the point of merging through gravitational-wave radiation reaction. Alternatively, compact binaries involving IMBHs may form through the evolution of isolated binaries composed of very massive stars [217]. If IMBHs in this mass range are generic, then Advanced LIGO and Virgo can detect gravitational waves from their coalescence [28]. Compact binary mergers involving IMBHs can be divided into two categories: (i) *mergers of IMBH binaries*, and (ii) *intermediate-mass-ratio inspirals.* Searches for merging IMBH binaries are included in the advanced detector search plans [see 3.6]; here, we focus on intermediate-mass-ratio inspirals (IMRIs).

If an IMBH exists in an environment where it is surrounded by neutron stars and stellar-mass black holes, such as a globular cluster, IMRIs of compact objects into IMBHs become likely sources of gravitational waves [434]. Advanced LIGO and Virgo could detect tens of such events per year [432]. The dominant capture mechanism is likely to involve gradual binary hardening via three-body interactions, meaning that the binary eccentricity will be very low in the GW detector band [432]. Spins of IMBHs that grow primarily through such minor mergers should not exceed ~ 0.3 [641]. For systems with very asymmetric mass ratios, the power emitted during the merger and ringdown portions of the waveform will be suppressed, so the late inspiral may dominate the signal-to-noise ratio.

Waveforms:
Accurate waveforms for IMRIs are particularly difficult, because the PN approximation is expected to fail for binaries that spend so many cycles close to the ISCO, while the IMRI mass ratio is not extreme enough for the mass-ratio expansion to be accurate [438]. EOBNR waveforms may be accurate in this regime; other approaches include hybrid waveform families [437]. Waveform accuracy studies are made difficult by the absence of accurate numerical-relativity waveforms at IMRI mass ratios, although there have been promising recent advances in numerical simulations of IMRIs [642]. Nonetheless, good overlaps between the hybrid waveforms and inspiral-only EOBNR waveforms indicate that existing waveforms may allow an IMRI search without extreme loss of sensitivity [643].

In the future, waveforms including precessing spins, inspiral-merger-ringdown phases, higher harmonics, and possibly eccentricity will be highly desirable for the development of searches and interpretation of search results.

Detection and parameter estimation:
The following are some of the special challenges presented by the detection and parameter estimation of binaries involving IMBHs, and the available avenues of addressing them:

- The low frequency of sources involving IMBHs places a greater emphasis on low-frequency detector performance. (*DQ, Searches*)
- The possible semblance of IMRIs to instrumental artifacts, lead to stringent data quality requirements (see Section **??**). (*DQ, Searches*)
- How can a dedicated IMRI search incorporate waveform uncertainty? Further studies of detection efficiency using injections from one waveform family and templates from another are necessary, including considerations of signal-based (e.g., chi-squared) vetoes. (*Searches, Waveforms*)

- The parameter space spanned by IMRI sources can be also explored with (coherent) burst searches (see Section 4.1.6), which do not rely on accurate GW waveforms.

- The accuracy of parameter inference, including mass and spin measurements, needs to be explored in the IMRI parameter space (see Section **??**). This includes challenges with parameter estimation when using templates with potentially significant systematic uncertainties. (*PE, Waveforms*)

IMBH Science:
A single detection of a $100+\ M_{\odot}$ system could provide the first unambiguous confirmation of the existence of IMBHs. This alone would be a major discovery.

Further detections could allow us to investigate the prevalence of IMBHs in globular clusters and cluster dynamics.

IMBHs could provide particularly exciting ways of testing general relativity (see Section **??**). For example, independent measurements of the IMBH mass quadrupole moment from IMRI gravitational waves would probe the IMBH spacetime structure [434, 644]. Ringdown studies of IMBHs could similarly test whether IMBHs are really Kerr black holes [645].

4.2.2 Determining the significance of CBC events

The rates and significance subgroup has chosen to use a series of mock-data challenges (MDCs) to evaluate the performance of the currently available methods for both event significance and astrophysical event rate estimation. Unlike other MDCs within the CBC group, these datasets simulate âĂIJtriggersâĂİ, the outputs of search pipelines rather than the direct output of the interferometers. This allow us far more control over, and understanding of, the statistical distribution of these triggers which can be composed of arbitrary populations of background (non-astrophysical) and foreground (astrophysical) events. These MDCs are being designed to probe and resolve the following priority issues regarding event significance and event rate estimation relevant to all CBC searches.

- The effect of removal or non-removal of zero-lag triggers when performing background estimation. This is the focus of the first MDC (known as the Hamlet test).

- The accuracy of event significance estimates, an understanding of their associated uncertainties and how this impacts our aim of confident detection.

- The ability of our algorithms to correctly report on the significance of multiple detectable astrophysical events.

- Our ability to perform the most accurate astrophysical event rate estimation in both detection and null-detection scenarios.

The first MDC is complete (as of early June 2014) and results are being compiled into a publication to be circulated and presented at the Autumn 2014 LVC meeting.

4.3 Continuous Waves Group

Rapidly rotating neutron stars are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band.[24] These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [223, 224, 225], magnetic deformations [226, 227], unstable r-mode oscillations [228, 223, 229], and free precession [230], all of which operate differently in accreting and non-accreting stars. We present a review of these emission mechanisms in [231]. Indirect upper limits on gravitational wave emission inferred from photon astronomy are more optimistic for non-accreting stars, but according to present theories accreting neutron stars are more likely to be emitting at or near the indirect limits.

The sources for which we search fall into four broad categories: non-accreting known pulsars for which timing data is available, non-accreting known stars without timing data, unknown isolated stars, and accreting stars in known binary or stars in unknown binary systems. For each type of source, we know or can infer properties of the source population; and for particular stars, there are indirect upper limits on gravitational wave emission which LIGO or Virgo must beat in order to claim a novel result. From our point of view, each type of object presents a distinct data analysis challenge which is directly constrained by more conventional astronomical observations. In particular, as most of our searches are computationally limited, their sensitivities are directly dependent on the constraints that come from conventional astronomy. As a result of our computational limitations we support a variety of search codes, each optimised for a different portion of parameter space. Where possible these code share common libraries and are cross-checked on fake signals injected into the detectors.

The breadth of investigation is fundamental to our search method. Given the large uncertainties in neutron star demographics (only $\sim$2000 of 10^8-10^9 neutron stars in the galaxy have been detected), evolution, and structure, we cannot confidently predict which type of source will provide our first continuous-wave discovery. Prudence demands an eyes-wide-open approach and enough flexibility to exploit unexpected waveform models. We do adhere, however, to certain priorities in allocating resources (scientists and computers) to different searches. Specifically, we place the highest priority on targeted searches for known pulsars (especially those for which the spin-down limit is achievable – see below) and on all-sky searches for unknown isolated pulsars. Plans for these two types of searches are described in section 3. Below we describe some remaining ongoing (spring 2014) all-sky searches in initial detector data, along with exploratory work to refine these searches.

Plans for three other types of searches in advanced detector data are in preparation: 1) directed searches for isolated stars of unknown spin frequency in a known location; 2) directed searches for binary stars of unknown spin frequency in a known location; and 3) all-sky searches for unknwon binary stars. Much exploratory work and pipeline development for these searches are underway, as described below, and results of two ongoing mock data challenges (isolated and binary stars) will be used in formulating concrete search plans later this year for the advanced detector era.

For each specific search effort described below, the section header states whether the work concerns an ongoing search in initial-detector data or exploratory pipeline development (or both).

4.3.1 Non-accreting known pulsars

We include in this source type all objects for which pulsations are observed in radio, X-ray, or other electromagnetic radiation, with the exception of accreting millisecond pulsars. Photon astronomy can tell us precisely the sky positions, frequencies, and frequency changes of these objects, meaning that our analyses need search only a small parameter space and are not computationally limited. Photon astronomy also sets an upper limit on the gravitational wave strain we could see from a known pulsar, assuming that all of the

---

[24] We use the term "neutron star" broadly, keeping in mind that some such stars may contain quark matter or other exotica.

observed spin-down is due to gravitational waves. In terms of the distance D, gravitational wave frequency f_{gw} and its time derivative $\dot{f}_{\mathrm{gw}}$, this indirect limit is [231]

$$h_{\mathrm{IL}} = 8.1 \times 10^{-25} \left(\frac{1\,\mathrm{kpc}}{D}\right) \left(\frac{-\dot{f}_{\mathrm{gw}}}{10^{-10}\,\mathrm{Hz/s}} \frac{100\,\mathrm{Hz}}{f_{\mathrm{gw}}}\right)^{1/2} \left(\frac{I}{10^{38}\mathrm{kgm}^2}\right)^{1/2}. \tag{18}$$

Here I is the star's moment of inertia, as estimated by theory but not directly observed, and could be higher than the fiducial value by a factor of up to 3. 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 S5-S6 data and the Virgo VSR2-4 data has beaten this indirect "spin-down limit" by a factor of 10 for the Crab pulsar (59.45 Hz) and by a factor of 3 for the Vela pulsar (22.38 Hz). Other pulsars for which the spin-down 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) [260], which should all be spin-down-accessible with advanced detectors.

The discussion above assumes gravitational wave emission from a triaxial neutron star, with the electromagnetic and gravitational wave components rotating as one unit. The astrophysical return from detecting such emission would be the first ever measurement of the difference between the two (equatorial) components of the inertia tensor. This in turn would give 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. Other emission mechanisms include free precession, excited modes of oscillation of the fluid, and the spin-down of a multi-component star. The astrophysical returns from detection of such wave generation could be considerable, potentially giving 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. This means that searches for such waves require careful thought in order to pick out a range of parameter space (i.e., the wave frequency and its time derivatives) that is both astrophysically reasonable and computationally feasible to search over. As described below (4.3.1), 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. Clearly, 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 are those for gravitational wave emission from pulsars of known position, rotation frequency, spin-down rate, and binary orbital parameters where necessary. This additional information greatly reduces the size of the parameter space over which we must search, and allows us to perform a fully coherent search for signals over all the available data. 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 lowest signal sensitivities achievable by LIGO and Virgo.

Three different pipelines are in current use for targeted searches: 1) a time-domain Bayesian method used in previous LIGO searches; 2) a new Fourier-domain method with Wiener filtering and deconvolution of amplitude modulation; and 3) a new matched filter method based on the $\mathcal{F}$-statistic and (new) G-statistic. These three methods are described in section 3. Below we discuss some ongoing enhancements of two of these pipelines.

4.3.1.1 Narrowband Searches for Known Pulsars – Exploratory

We know of several physical mechanisms that could cause the frequency of a neutron star's emitted gravitational waves to differ slightly from the typically assumed $2f_{\rm rot}$, with $f_{\rm rot}$ being the rotation frequency. Using estimates of the maximum frequency difference the $\mathcal{F}$-statistic search method has been used to search a small frequency and frequency derivative parameter space for the Crab pulsar with nine months of data from the S5 run [264]. This search has also been performed on 28 days of S6 and VSR2 data.

The nested sampling parameter estimation code developed for the targeted known pulsar search can also be extended to search narrow frequency bands. We will study its performance for a range of bandwidths (and ranges of frequency derivatives) using the multi-source mock dataset that we have developed. As well as allowing the search to span different emission scenarios, this will also allow previously poorly timed pulsars to be included in searches.

In the longer term, these narrow-band search methods will be explored as a follow-up step to semi-coherent all-sky searches.

4.3.1.2 Narrow-band search with the signal Fourier 5 components method – Ongoing search

The targeted search method based on signal Fourier five components, described in the search plans, has been extended to narrow-band searches, where a possible small mismatch between two times the EM frequency and the GW frequency is taken into account by exploring a range of frequencies and spin-down values around those derived from EM observations. A methodological paper has been written (PRD 89, 062008, 2014). The method has recenty been applied to a narrow-band for Crab and Vela pulsars in the data of Virgo VSR4 run. An observational paper, containing the results of this analysis, is nearly completed as well as the method review. In the future the method will be applied to the data of advanced detectors for potentially interesting sources.

4.3.1.3 Compound signal searches for pulsars – Exploratory

We will investigate new methods to detect the collective signal from multiple known pulsars. One such method is Siegel's test for compound periodicities in a time series [646]. Adapted to gravitational-wave data, this might hold the promise that such a compound signal from multiple sources can have a higher probability of detection than for the strongest single-source signal alone. We shall also consider the collective detectability of certain known populations of pulsars. Finally we shall investigate the detectability of compound signals from unknown populations among candidates obtained in all-sky searches.

4.3.2 Non-pulsing non-accreting neutron stars and favorable directions

This type of search includes point sources, such as central compact objects in supernova remnants, as well as highly localized regions, such as the innermost parsec of the galactic center or the core of a globular cluster. Photon astronomy can provide sky positions for these objects, but since no pulses are observed, the external measurements cannot provide us with frequencies or spin-down parameters. Since we must search over many frequencies and spin-down parameters, sky-survey positional errors (such as from ROSAT) are too large: we require arcminute accuracy or better to keep down the computational cost of a long integration time and thus a deep search. Although no f and $\dot{f}$ are observed for these objects, we can still define an indirect limit we need to beat. If we assume an object has spun down significantly from its original frequency and that this spin-down has been dominated by gravitational wave emission, we can rewrite Eq. (18) as

$$h_{\rm IL} = 2.3 \times 10^{-24} \left(\frac{1\,{\rm kpc}}{D}\right) \left(\frac{10^3\,{\rm yr}}{\tau_{\rm sd}}\right)^{1/2} \left(\frac{I}{10^{38}\,{\rm kg\,m^2}}\right)^{1/2} \tag{19}$$

in terms of the spin-down age $\tau_{\rm sd}$, which can be inferred in various ways.

Initial LIGO can beat this upper limit for several objects of this type, including the youngest – the object in supernova remnant Cas A ($\tau_{\rm sd} = 326\,$yr, $h_{\rm IL} = 1.2 \times 10^{-24}$) – and the closest, Vela Junior ($D > 200\,$pc, though the precise value is uncertain). Several more objects have indirect limits attainable with advanced detectors, including the remnant of Supernova 1987A ($h_{\rm IL} = 3.2 \times 10^{-25}$). However this putative neutron star is only 25 years old and would require a search over a large parameter space including six or seven frequency derivatives. Searches over small sky areas, such as the Galactic center region (single "pixels") are computationally the same as searches for known point sources. In addition to having published upper limits for Cassiopeia A [346] and the Galactic center [**?**] using LIGO S5 data, we are nearing completion of searches in more sensitive S6 data for Cas A and O(10) other supernova remnants, and a search in the direction of globular cluster NGC 6544. We are collaborating with several photon astronomers on constructing more complete target lists of point sources and small areas, both for initial and advanced LIGO (see section 4.3.2).

The first search for a source with no timing (Cas A) used the $\mathcal{F}$-statistic code with a single integration of $\mathcal{O}(10)$ d. Our estimate of computational cost and sensitivity [345] showed that this is enough to start beating indirect upper limits on some sources. For young sources even such a short integration time requires up to the second frequency derivative; thus metric-based methods and code for tiling parameter space in multiple dimensions are important to reduce the computational cost. The more recent S6 searches for Cas A, for other SNRs and for NGC 6544 have also used a single coherence time.

In the future, however, we expect to rely more on semi-coherent, hierarchical searches, using many different observation spans of shorter coherence time, an approach used in the recently published Galactic center search [**?**] and now being used in an Einstein@Home search for Cas A and other SNRs. Another pipeline enhancement uses barycentric resampling of the data during preprocessing to improve computational efficiency. This resampling method was first exercised in a LIGO S5 search for the Calvera central compact object [347] and has been more recently used in the search for NGC 6544. This improvement will be incorporated in the coming year into the Einstein@Home directed searches. These pipeline improvements will allow us to search a significant fraction of the full ADE data sets, rather than data spans of only $\mathcal{O}(10)$ d.

4.3.2.1 Coherent directed searches – Ongoing search

The coherent S5 search for Cas A [346] is being updated to S6 data and extended to supernova remnants G1.9+0.3, G18.9–1.1, G93.3+6.9, G189.1+3.0 (IC 443), G266.2–1.2 (Vela Junior, two searches), G291.0–0.1, G347.3–0.5, and G350.1-0.3, for a total of ten searches and nine objects. These remnants, like Cas A, contain non-pulsing candidate neutron stars, except for G1.9+0.3 which has now taken Cas A's place as the youngest known supernova remnant in the galaxy and is small enough to be searched with one sky position. The update includes the SSE2 enhancements to floating-point instructions on Intel processors, resulting in nearly a factor 3 speed-up over the standard $\mathcal{F}$-statistic code. All of the searches are done, no plausible signals have been found, preliminary upper limits have been set, and the internal review is underway. The sensitivities of these coherent searches, integrating over time spans from 5 days to several weeks of S6 data, beat the indirect limits on gravitational-wave emission over frequency bands of at least 100 Hz and in one case nearly 2 kHz. Publication is planned in fall 2014.

In addition, a search based on time domain resampling [347], following the template placement method of the Cas A search [346], has been carried out for isolated stars in globular cluster NGC 6544. The search pipeline is based on an Einstein@Home implementation of the resampling algorithm. Code development was coordinated with the E@H team, to permit a common code base for a variety of future searches. Production running is complete with broadband upper limits obtained and outlier follow-up nearly complete. Review will begin in spring 2014, with a publication planned in late summer 2014.

A search for continuous gravitational waves from unidentified Fermi-satellite gamma-ray sources at high galactic latitudes is under exploration, where the relatively coarse resolution from Fermi requires gridding a small patch of the sky. This opportunistic search for gamma-ray sources with the spectral characteristics of known gamma-ray pulsars and with small time variability, would use relatively short coherence times – $\mathcal{O}(5\text{ d})$ and barycentric resampling.

4.3.2.2 Supernova 1987A using the cross-correlation technique – Exploratory

As described elsewhere, the semi-coherent excess power methods are more robust than the fully coherent searches. This is because they demand phase coherence of the signal only over the coherent integration time, which is much shorter than the total observation duration. This reduction in the minimum coherence time has the added advantage of significantly reducing the computational cost. It is possible to reduce this coherence time even further by using cross-correlations between data from multiple detectors. In the case when we correlate coincident data from two detectors, the minimum coherence time is just the light travel time between the two detectors. In the general case, we can correlate data streams collected at arbitrary times from two distinct detectors, and also from the same detector at distinct times. The details of this method, which is a generalization of methods used previously in the stochastic "radiometer" search [647, 648], can be found in [250]. The main feature of this generalization is the presence of a free parameter, the minimum coherence time required of the signal, which can be tuned depending on the desired sensitivity, robustness and computational cost.

The starting point for this search is a set of SFTs of duration T_{sft} covering a total observation time T_{obs} followed by: i) a choice of the minimum coherence time $T_{\text{coh}-\text{min}}$ which is used to create pairs of SFTs, ii) a computation of the cross-correlation statistic for each pair for a given set of pulsar parameters, and iii) calculating a weighted linear combination of the various cross-correlations, with the weights chosen to maximize the sensitivity exactly as in the PowerFlux or the weighted Hough searches. Many of the existing standard CW searches can be viewed as special cases of this scheme. The standard PowerFlux search corresponds to considering only self correlations of the SFTs, a full coherent search corresponds to considering all possible SFT pairs, and the hierarchical search is an intermediate case with $T_{\text{obs}} \gg T_{\text{coh}-\text{min}} \gg T_{\text{sft}}$. (This is however a computationally inefficient way of calculating the coherent statistic, for which it is better to use the existing $\mathcal{F}$-statistic.) The choice of $T_{\text{coh}-\text{min}}$ will be set by a tradeoff of computational time with sensitivity, but we expect to have $T_{\text{coh}-\text{min}}$ be at least a few times T_{sft} in order to take advantage of the flexibility of the method.

One object to which this semi-coherent method can be applied in a directed search is a neutron star in the supernova remnant SN1987A [348]. In searching for such a young object, searching over frequency derivatives can be prohibitive because of the need to search over higher derivatives. It turns out that the search space can be narrowed by using a physical model for the frequency evolution: $\dot{\nu} = Q_1\nu^5 + Q_2\nu^n$. The first term is the usual term due to gravitational wave emission while the second term represents all other effects (ideally, for electromagnetic braking, one would expect a braking index of $n = 3$, but in practice one observes a range $1 < n < 3$ for the ~ 10 objects whose n can be measured by absolute pulse numbering). With this model, and using $T_{\text{coh}-\text{min}} \approx 1$ hr, it turns out that the computational cost becomes manageable for astrophysically interesting parameter ranges, namely $B \lesssim 10^{11}$ G and $\epsilon \gtrsim 10^{-4}$, where B is the magnetic field strength and ϵ in the ellipticity [348].

A pipeline to perform a cross-correlation search for isolated neutron stars has been written and tested and undergone significant review. The pipeline has been calibrated via Monte-Carlo simulations for pure noise and injections, with and without averaging over the inclination and polarization angles, and with and without spin down. A partial search for SN 1987A over $\sim 30\%$ of the parameter space has been done with S5 data. In the coming year, the search will be completed on S5, S6 and/or VSR2-3 data, with an eye towards producing a publication before Advanced LIGO data begins to flow. As described in Sec 4.3.4, a re-organized version of the pipeline is being developed in parallel, to address some of the shortcomings of the original pipeline (e.g., the need to keep track of the detection statistic for every point in parameter space rather than using a "toplist" of the most significant points). That new pipeline is currently constructed specifically for the directed search for a neutron stars in a known binary system, but will eventually be generalized to include the (f_0, Q_1, Q_2, n) parameter space described here for young isolated neutron stars.

4.3.2.3 Semi-targeted search using stroboscopic resampling – Exploratory

In general, the correction of the Doppler effect due to Earth motion depends on the source sky direction and frequency. Since the parameters are often unknown, a large computational effort is required to correct for any possible direction and emission frequency. A correction technique independent of the frequency is used in a pipeline based on stroboscopic resampling. The antenna proper time is accelerated or slowed down by deleting or duplicating in a timely manner single samples of the digitized signal in order to keep the reference clock synchronized with the source clock, within an accuracy given by the inverse of the sampling frequency f_s (several kilohertz) [649]. The removal (or the duplication) of the samples takes place typically each few seconds. The list of samples to be removed or duplicated (named *mask*) is thus not huge and can be easily computed by simple geometrical consideration. As detailed in [649] the mask corresponding to a given direction is provided by the times when the antenna crosses one of the equi-phase parallel planes fixed in the space, perpendicular to the wave vector and each at a distance c/f_s from the next one. Each "crossing time" is computed by the scalar product of the antenna velocity and the wave direction (in practice by a few operations each second of data).

The maximum phase error due to the non-perfect synchronization is given by $2\pi f_0/f_s$ where f_0 is the signal expected frequency and f_s is the sampling one. As a reminder, a phase error around a few tenths of rad is small enough to guarantee that almost all the signal energy is recovered around the main frequency. It is thus important to resample the data working at the Virgo data acquisition frequency (20 kHz) in order to use the method effectively up to several hundred Hz. This frequency independence makes the method very appealing for sources where the direction is well fixed, but the emission frequency is uncertain (semi-targeted search). The pulsar spin-down is taken into account by properly shifting the equi-phase target plane during the acquisition time. As a consequence, a single mask requires specifying both the direction and the spin-down value of the source. The Einstein delay and the Shapiro effect can be also easily computed without any significant additional computational cost.

We have developed an analysis pipeline and are applying it to VSR2 data. The Earth ephemeris is computed by using the Roma 1 group PSS routine. In just a few minutes the ephemeris and Einstein delay data are computed and stored for the entire VSR2 period with a sampling time of a few seconds (enough to approximate Earth motion with enough accuracy).

Starting from the ephemeris, another routine computes the masks for a set of directions and spin-down values. The computation time was tested not to exceed a few 10^{-8} of the integration time, per each mask (i.e., per each direction and spin-down).

In parallel the antenna data, already cleaned from non-stationary events by usual PSS techniques, is pass-band filtered around the signal expected frequency. The bandwidth must be large enough to contain all the sidebands produced by Doppler and spin-down. Several tens of operations per sample are necessary in the data filtering. The final cost will be evaluated after implementation, but we expect to work with a computing time around 10^{-4} – 10^{-3} of the integration time.

During the signal decimation, different masks can be applied in parallel to the filter output (at signal sampling frequency). Very light buffers are produced at the downsampling frequency (inverse of the filter band) for FFT spectral analysis. Usual statistical analysis for peak identification will be adopted in the final step of the pipeline.

Since the Doppler correction (computation of masks and their parallel application in decimation of the filtered data) is negligible, the optimization strategy for the semi-targeted search is straightforward. We need only choose the width of the pass-band filter ("slice"). Indeed this choice determines the downsampling factor, thus the length of the buffers governing the FFT computation time. Finally we must multiply the time required for the previous operation (filtering and parallel FFTs) times the number of slices required to cover all of the interesting detection band. The optimization of the pass-band filter width, obtained minimizing the total computation time, depends on the analysis to be performed. Many tests have been performed on simulated data assuming different antenna orbits, spin-down values, sampling frequencies and source frequencies. In all cases, the expected phase-locking and peak reconstruction accuracy has been found. Similar tests have been performed injecting signals in the VSR1 data. All the results are described in Torre's graduation thesis [650], or (more in summary) in [651]. The resampling of the data requires less than 10^{-5} of the investigated time (for a single direction and spin-down on a few Hz band), that is negligible with respect to the time required to read the HF stream (of the order of 10^{-4}). A method to read the data and apply the resampling technique directly to the down-sampled data (making negligible the computing time for reading data) is in progress.

A methods paper was published in Phys. Rev. D [652].

The amplitude modulation will be taken into account using a matched filtering in the frequency domain, in a way similar to the one developed by the Rome group. We are currently implementing a full pipeline for this purpose, and testing it on hardware injections in VSR2 data.

After the validation of the pipeline we plan to apply it to the observation of the supernova remnant RX J0852.0-4622. Currently the former Ph. D. student dedicated to this (Oriella Torre) had to suspend her involvement, so we had a reduction in manpower. Our plan is to complete the analysis anyway. Our expectation is to have it ready for the review at the end of the current year (2014).

4.3.2.4 Semi-targeted searches for "transient CW signals" – Exploratory

This section concerns searches for gravitational wave signals longer than those traditionally considered by the Burst group, but shorter than those traditionally considered by the CW group. A code capable of performing such searches was developed by Giampanis & Prix [653]. As an astrophysical application of this, the possibility of the excitation of r-modes in rotating neutron stars has been considered. Two different kind of sources have been identified:

- **Glitching pulsars** which are mainly young pulsars (although one millisecond pulsar has been observed to glitch).
- **Type I bursts** from accreting millisecond pulsars.

Some thought has already been given to the astrophysics of carrying out such searches, in terms of wave frequencies, durations and detectability [654]. Given that the r-mode's frequency is known with a precision of $\sim 20\%$, searches will need to be carried out in a frequency band defined by this uncertainty. This means that for young pulsars like the Vela Pulsar ($f_{\rm mode} \sim$ 15 Hz) a $f_{\rm band} \sim$ 3 Hz is required, while for faster spinning stars like 4U 1608−522 ($f_{\rm mode} \sim$ 825 Hz) a $f_{\rm band} \sim$ 165 Hz is required.

In the case of pulsar glitches, an assumption that $E_{\rm mode} = E_{\rm glitch}$ was made, and in a similar way, for the Type 1 Bursts $E_{\rm mode} = E_{\rm burst}$ was assumed, where $E_{\rm mode}$ is the energy deposited in the r-mode, $E_{\rm glitch}$ an estimate of the energy involved in a glitch, and $E_{\rm burst}$ an estimate of the electromagnetic luminosity of an X-ray burst.

The detectability of the signal also depends upon the signal decay timescale, which in turn depends upon the poorly constrained dissipation mechanisms (probably related to shear viscosity) that act to suppress the r-mode. One such mechanism is the related to the existence of a viscous boundary layer at the star's crust-core interface, as studied by Levin & Ushomirsky [655]. Regardless of the details of the mechanism, we can assume that the dissipation is strong enough to make all observed millisecond pulsars and low-mass X-ray binaries stable.

Using simple estimates, it was found that, if the Levin and Ushomirsky dissipation model [655] is considered, the detection of gravitational waves associated with glitches in young pulsars is not plausible with 3rd generation gravitational waves detectors [654]. On the other hand, glitching millisecond pulsars may be significantly easier to detect, although such glitches may be extremely rare. Attention is now being focused on type I X-ray bursts in accreting millisecond pulsars. Some of the fastest such pulsars may in fact reside close to the threshold for instability, resulting in relatively long-lived signals and potentially interesting levels of gravitational wave emission. The recent observation of an X-ray oscillation during an outburst of an low-mass X-ray binary may even be an example of such a phenomenon [656].

To gain a more accurate picture of the feasibility of such searches, the transient gravitational waves search code [653] is now being tested to recover signals with timescales assumed to lie in some narrow range of frequency and damping time. Efforts are now focused on (i) Completing the extraction of injected transient signals from realistic data, searching over a wider range of frequencies and damping times. (ii) The assembly of a methods paper, combining the astrophysical motivation with details of the search method. (iii) Assessing what can be learnt regarding the source parameters in the event of a successful detection.

4.3.2.5 Directed search using $\mathcal{F}$-statistic – Exploratory

In order to allow for fast, dedicated searches for GWs from astrophysically motivated regions of the sky (e.g., supernova remnants and star-forming regions, globular clusters) over large frequency and frequency derivatives range, we have implemented an already existing and tested all-sky version of the $\mathcal{F}$-statistic method on the graphical processor unit (GPU) hardware accelerators. Since the massively parallel GPU architecture allows for substantial speed-up (preliminary tests show speed-ups larger than 30 times, as compared to the previous CPU version) and provides sufficient operating memory, it - in principle - alleviates some previous computational constraints, like the length of two sideral days for the time-series data interval used in the FFT computation. Presently, a version for the Nvidia GPUs and Nvidia/CUDA programming language is being tested - we anticipate the optimized production code will be fully functional in the beginning of the second half of 2014.

Meanwhile, we are researching the possibilities of using the GPUs together with the CPU version of the code, allowing for massive parallelization in the cluster environment, parallel calculations with several GPUs at once, and planning for the expansion of the $\mathcal{F}$-statistic toolset with the ability to analyze the periodic GW signal at both once and twice the spin frequency of the star [657], as well as for the analysis of the data from a network of detectors.

4.3.2.6 Einstein@Home semi-coherent directed searches – Ongoing search

In mid 2013 we launched a new class of searches on the Einstein@Home platform, targeting known or suspected compact objects that are well-localized on the sky, but have unknown rotation frequency and spindown. These "directed" semi-coherent searches on LIGO S6 data will focus on the list of astrophysical targets that was initially collected for fully-coherent directed searches (see Sec. 4.3.2 and Sec. 4.3.2). We aim for a near-optimal distribution of computing power, both over the different astrophysical targets, as well as in the choice of semi-coherent pipeline parameters (length and number of segment, template-bank mismatches). Semi-coherent directed searches are expected to be the most sensitive searches of this kind, and will beat the indirect spindown-limit for most of these targets over a wide frequency range of up to $\sim$1kHz. In practice we have started with a single directed search for Cas A, lasting for about a year, which is expected to end in August 2014. This will be followed by several searches over further targets, either individually or "bundled" into groups of targets. In total we expect to spend about 2 years of Einstein@Home computing time on these directed searches. Detailed publication plans are yet to be decided.

4.3.2.7 Other targets – Exploratory

We are collaborating with several astronomers on constructing lists of interesting targets for further directed searches, i.e., targets where LIGO and Virgo can beat the indirect limits on gravitational-wave emission. Apart from Cas A there are of order ten candidate non-pulsing neutron stars with indirect limits on h_0 beatable with S5 or S6/VSR2 coherent and semi-coherent searches. There are also several small, young supernova remnants (such as SN 1987A) and pulsar wind nebulae where the neutron star is not seen. Other small sky regions (further discussed below) also can be targets of this type of search. Examples include regions of massive star formation such as the galactic center and massive young clusters containing magnetars such as Westerlund 1. Globular clusters are not likely to contain young neutron stars, but some old neutron stars are known to possess planets and debris disks. Frequent perturbations in the dense environment of a cluster core could trigger bombardment episodes, and a star with an impact-related deformation counts as rejuvenated for purposes of a gravitational-wave search. Considering interaction timescales, most of the best targets are nearby, dense clusters such as NGC 6544. However 47 Tuc's interaction timescale is short enough to make it an attractive target even though it is further away; and furthermore the first GLAST/Fermi results show considerable high-energy diffuse emission which is likely related to neutron star activity in the relatively recent past.

It is useful to maintain ties because X-ray and gamma-ray astronomers are finding many point source neutron star candidates in recent years, and thus it is likely that the interesting target list for LIGO will expand substantially even before advanced LIGO. Examples include Fermi point sources and HESS TeV gamma-ray sources that are followed up with Chandra and XMM-Newton X-ray observations, yielding pulsar wind nebulae and sometimes the neutron stars themselves.

A paper describing an interesting list of targets in in preparation and will be submitted for publication this year. As the advanced detector era approaches, these studies will be extended.

4.3.3 All-sky searches for isolated neutron stars

These are objects which have not been previously identified at all, and thus we must search over various possible sky positions, frequencies, and frequency derivatives. They are believed to constitute the overwhelming majority of neutron stars in the Galaxy, but most of them are not believed to be good sources for LIGO or Virgo. 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

$$h_{\mathrm{IL}} = 4 \times 10^{-24} \left(\frac{30\,\mathrm{yr}}{\tau} \right)^{1/2}, \tag{20}$$

where τ is the mean time between supernovae in the Galaxy. The latest and most detailed derivation of this upper limit is given in [231]. Note, however, that a more recent simulation analysis finds significantly lower expectations that depend on the assumed source frequency and ellipticity [236].

It is useful here to briefly explain the computational challenge that must be overcome for these searches. The parameter space for blind searches for weak signals from unknown isolated neutron stars is very large. The number of templates N_p, required to cover the entire sky, a large frequency band, and a range of spin-down parameters and using data which spans a duration T, is roughly proportional to T^5. The computational cost therefore scales as $\sim T^6$. In fact, for any reasonable volume of parameter space, N_p becomes so large that using our existing coherent integration code and using the full computational power of our largest computational platform `Einstein@Home` running for a few months, it is not possible to consider values of T larger than a few days. Even if we were able to speed up our coherent demodulation algorithm by, say, a factor of 100, T would increase only by a factor of $100^{1/6} \approx 2.2$. On the other hand, we require T to be a few months to have a realistic chance of detection. The situation is, naturally, even more demanding for neutron stars in binary systems.

For this reason, different methods using a combination of coherent and semi-coherent techniques have been designed. The basic idea is to break up T into smaller segments which are analysed coherently, and to stitch together these segments using a semi-coherent technique. Outlier candidates are then followed up. The sophistication and automation of the follow-ups have improved in recent analyses and offer the promise of lowering detection thresholds significantly in some search pipelines.

Five all-sky pipelines have been used in recent years for carrying out all-sky searches: 1) a quick-look semi-coherent method known as PowerFlux using incoherent sums of strain spectral power from many 30-minute "Short Fourier Transforms" (SFTs), 2) a multi-interferometer Hough transform method starting from 30-minute SFTs, 3) a hierarchical algorithm using `Einstein@Home`, based on phase-preserving demodulation over many ∼day long intervals, followed by a semi-coherent step (see below), 4) a hierarchical method, developed in Virgo, based on the alternation of coherent and incoherent steps; and 5) an $\mathcal{F}$-statistic-based search also developed on Virgo data. These five methods are described in section 3. Below we discuss some ongoing searches with several of the pipelines, along with development work to improve sensitivity or computational efficiency.

In addition, two new methods are under development that offer greater robustness against uncertainty in the source model: 1) a "loosely coherent" method [240] using the PowerFlux infrastructure; and 2) a cross-correlation method [250] which provides a smooth bridge between semi-coherent and coherent methods, with the possibility of parameter tuning to improve sensitivity over semi-coherent methods while maintaining robustness.

4.3.3.1 PowerFlux method – Ongoing search

For a monochromatic, constant-amplitude sinusoidal wave in Gaussian noise, summing the strain power from M short Fourier transforms improves the sensitivity (strain value for a fixed signal-to-noise ratio) by a factor $M^{1/4}$. In contrast, a coherent search based on a single Fourier transform over the entire M intervals gives a sensitivity that improves like $M^{1/2}$. One strong advantage of the semi-coherent methods is their robustness against unknown source phase disturbances, such as from frequency glitches due to starquakes.

The searches we must perform are more complicated than simple power sums. Frequency and amplitude modulations that depend on source direction are relatively large. The frequency modulations arise from the motion of the detectors with respect to the source, with components due to the Earth's rotation ($v/c \sim 10^{-6}$) and to its orbital motion ($v/c \sim 10^{-4}$). The amplitude modulation arises from the changing orientation of the interferometer arms with respect to the source direction as the Earth rotates. As a result, an all-sky search requires a set of finely spaced templates on the sky with varying corrections for these modulations. In general, the number of templates required for a given coverage efficiency scales like the square of the source frequency.

Over the last several years, we have explored three related methods for incoherent strain power summing: StackSlide [239], Hough [243, 258], and PowerFlux [238]. These methods take different approaches in summing strain power and in their statistical methods for setting limits, but their performances are quite similar. Because PowerFlux has been found to yield somewhat better efficiency than the other two methods for most frequency bands, it has been chosen as the quick-look semi-coherent algorithm used on data from the S5 and S6 science runs. An article based on applying all three methods to the S4 data was published in early 2008 [237] in Physical Review D.

In short, PowerFlux computes from many thousands of 30-minute 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, like Hough or StackSlide, 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 ~0.56 mHz spacing and limits presented separately for 0.25 Hz bands.

A short publication based on an improved PowerFlux search over the first 8 months of S5 data was published in Physical Review Letters in early 2009 [255]. These results cover the frequency range 50-1000 Hz and negative spin-down as large as 5×10^{-9} Hz/s. The present PowerFlux program permits deeper searches for coincident candidates among multiple interferometers than in S4 and applies tighter coincidence requirements between candidates in the H1 and L1 interferometers, which allows setting lower SNR thresholds for followup of candidates.

A series of major improvements to computational efficiency were made to facilitate a PowerFlux run over the full S5 data while keeping memory requirements within the bounds of LIGO processors and keeping total computational time for the first-pass search within a half-year. A two-interferometer power sum was used, together with coincidence between H1 and L1, to push deeper into the noise than before.

In parallel, a "loosely coherent" follow-up was added directly to PowerFlux, one that allows for slow drifts or modulations in phase from one SFT epoch to the next [240]. It offers the possibility of a "controlled zoom" of interesting candidates and reduces the chances of missing a true signal because it doesn't quite fit the template for a long coherent search. Upper limits for the 50-800 Hz band based on a search of the full 2-year S5 data were produced, and outliers followed up with the new loose coherence step. These results were published in 2012 [251].

Preliminary S6 results have been obtained, with follow-up of outliers in progress. A large number of spectral artifacts seen in S6 data (associated with the output mode cleaner) make this follow-up more laborious than expected. A novel "universal" statistic [256] has been developed to cope with the large number of non-Gaussian S6 bands in both H1 and L1 data. Other recent improvements included a coherent IFO-sum option for each SFT to gain further SNR [658]. Publication of final S6 results is expected in fall 2014.

4.3.3.2 `Einstein@Home` All-Sky Searches

– Ongoing search

Overview: `Einstein@Home` is a public distributed-computing project in which users contribute the spare CPU cycles on their computers for gravitational wave searches. Thus far, it has been used in the all-sky wide parameter space searches for CW sources. It was launched in February 2005, and since then it has built up a user-base of over 200 000 active users; it currently delivers more than ∼ 1 Pflops of continuous computing power. This is by far the largest computational platform available to the LSC and it is also one of the largest public distributed projects of its kind in the world. The project is targeted towards making a detection. So far it has analysed LIGO data from S3, S4, S5 and S6.

S5GC1 search: A significant improvement with respect to the Hough transform hierarchical search was claimed in [659] by exploiting global correlations in source parameter space. An engineering run (S5GCE) to test this improvement was completed in spring 2010, and a new production run (S5GC1) over the full S5 data has recently completed. A preliminary study of the sensitivity of this run shows that it is in fact less sensitive than the S5 run whose results were published in February 2013 (H.B. Eggenstein, P. Leaci and M.A. Papa). However this run covers a range of frequencies above 1190 Hz which were not searched before. The plan is to analyse the results from this band and publish a short report only for the frequency band not covered by the S5R5 search.

S6 Bucket search: A production run over the most sensitive S6 data (in the "bucket") begun in 2011, where the relatively low frequencies (50-450 Hz) permit searching with a longer coherence time (90 segments of 60 hours each) and up to spin-downs corresponding to an age of 600 years. The run completed, and a similar run was launched in 2012 which returns candidates ordered according to a different statistic from used in the first bucket search. In particular, the results are ranked according to a "line veto" statistic that, broadly speaking, decreases the significance of candidates with signal-inconsistent average multi versus single-IFO F-stat values. This run also completed, at the end of 2012. The systematic post processing of these runs has begun in 2013 and it is ongoing. Apart from an improved sensitivity and hence a publication of the results of this search, a comparison between the candidates of the different runs will allow us to characterise the performance of the different detection statistics that have been developed.

Data selection and run set-up procedures: The best exploitation of the limited computing resources for a given parameter space is an optimization problem that depends on the noise level of the data and on their gaps relative to all the usable detectors. Up to now each `Einstein@Home` run has required an ad hoc iterative study of the set-up parameters. The team is trying to quantify through figures of merit the relevant choice parameters and to set up automatic procedures to define run parameters. An additional factor that has entered the run-set up decisions up to now in different, and not always consistent ways, is the prior chance to make a detection as a function of parameter space. We would like to develop schemes that at least highlight the priors underlying the run set-up choices.

Speed-up of the F statistic search: Incorporate the resampling-based code developed by Chris Messenger in the `Einstein@Home` hierarchical pipeline. Progress has been made in 2011-2012 on this. We foresee integrating the resampling technique in our mainstream searches but we are not sure that we will be able to achieve this during this year.

Signal-based veto procedures: A “Line Veto statistic” has been developed in order to discard spurious disturbances related to instrumental “line” artifacts from the top-list of returned candidates, which will improve the sensitivity of searches. The S6 Bucket runs have used this for the first time and, as described above, the post processing of the results will allow to characterise the performance of this statistic and determine a solid tuning method for its free parameters.

Clustering of correlated candidates, detection tests over portions of the parameter space Based on the final results of ongoing post processing efforts, identification of correlated candidates will be studied and possibly incorporated in the hierarchical `Einstein@Home` pipeline allowing to save follow-up cycles. For the same purpose, i.e. identification of interesting regions in parameter space, we will also investigate the application of standard hypothesis testing methods, for example the Bayes ratio. We plan to address this in the context of the S6 bucket post processing. For directed searches a clustering scheme was designed and successfully used in the Galactic centre search.

re-write of GCT code In 2011-2012 the GCT search code routines were modified to allow the use of 2nd order spin-down to expand the population of sources that we can search for. Recent investigations in the context of setting up the S6 Einstein@Home directed searches have uncovered performance issues with the most current implementation that have been addressed with ad-hoc fixes for the S6 directed-search runs. For these we are confident that the used GCT scheme still achieves good detection efficiencies.

New all-sky hierarchical code The GCT scheme in its current implementation is not suited for all-sky searches over periods of a year or more. This is a major drawback and makes it compelling for us to develop an alternative stack-slide (semicoherent sum) code. Whereas there are drawbacks to the Hough code, at the time of writing it is still competitive with the GCT implementation for year-long searches up to first order spindown.

Sensitivity estimates and comparisons We maintain and develop an analytical framework to estimate the sensitivity of the complex search procedures that the `Einstein@Home` performs as well as other searches, most notably the Powerflux search. Such framework successfully predicted the S5R5/R6 $h_0^{90\%}$ ULs.

Performance benchmarking We will use the reference MDC set of the CW group to characterise the performance of the Einstein@Home searches through the standard figures of merit used by the other CW group teams.

Automatic organisation and look-up of results : We have set up a dedicated machine to assimilate the E@H results as they come in, organise them in their ultimate directory and file structure and produce statistical summary plots that will help us diagnose issues early on in the run.

Identification of instrumental artefacts: In the S6 Bucket results we identify frequency bands that host enhanced values of the detection statistic and try and correlate with disturbances in the input data.

Automatic follow-up procedures: Based on the final results of ongoing optimization studies for a hierarchical search (see below), automatic follow-up procedures will be implemented in the `Einstein@Home` core code in the long term, ultimately permitting improved sensitivity (cf Sec. 4.3.3).

Support and maintenance: The `Einstein@Home` servers and project require continual maintenance in the form of software updates, message board interaction with users, publicity, maintenance of server hardware, maintenance, repair and extension of the BOINC libraries, bug tracking and elimination, etc.

Automatization of work-unit generator for different searches: Currently much work and specialized expertise is required in order to set up a new BOINC project, or even to prepare and launch a new run in an existing project such as `Einstein@Home`. Some of the key steps required are a "workunit generator" that needs to be implemented (coded in C++ against the BOINC library), together with a validator and an assimilator. The science application needs to be installed on the server, together with various setup steps required on the server in order to prepare the scheduler and the database. Work has now begun on a project to make this increasingly easier and more "user-friendly", allowing users to set up new runs or even whole projects "on the fly".

Relation to the "Grid":

BOINC is a general computing platform that is able to leverage huge computing power from a pool of heterogeneous computing resources in a fault-tolerant and robust way. In this it achieves an important goal that is also part of various "grid" initiatives. If one can create a flexible and simple interface, similar to that of condor, say, to this powerful infrastructure, one could leverage the massive pool of LSC computing clusters or other "grid" resources in a more transparent and flexible way than is currently possible.

4.3.3.3 Followup-searches to confirm or veto CW signal candidates – Exploratory

Better theoretical understanding and the development of software tools is required in order to deal efficiently with the follow-up of interesting candidates from incoherent search pipelines. The previously developed fully coherent follow-up method based on numerical optimization using a Mesh Adaptive Direct Search algorithm (NOMAD) has been extended to include higher-order spindown and to follow-up candidates from searches for unknown binary systems. However further work is required to optimize and streamline the search software. A methods paper about the application of this grid-less follow-up procedure to S5R5 Einstein@Home search outliers has been written and is in the process of final approval and preparation for publication. An alternative coherent and completely deterministic procedure was used to follow up low significance candidates of the targeted Galactic center search. Investigations along both lines will continue. As the final goal is to move the coherent follow-up step onto the hosts participating in the Einstein@Home search, given a recent speed-up of the grid-based search approach utilizing Graphics Processing Units we will investigate a possible hybrid (grid / grid-less) follow-up procedure. In order to minimize the follow-up search volume we will answer a question of central interest, namely how applicable is the Fisher matrix formalism to candidates with moderate signal-to-noise ratio at low confidence (the lower the confidence, the smaller the follow-up parameter space).

4.3.3.4 Instrumental line-robust statistics for wide-parameter searches – Exploratory

Recent Einstein@Home runs are using not only the standard $\mathcal{F}$-statistic long established for CW searches, but also new line-robust statistics (LR-statistics). This can increase sensitivity in data bands affected by near-monochromatic disturbances in a single detector, the *line* artifacts. The LR-statistic, based on an extended noise model, was published as a methods paper.

A simple version has been used in the S6LV1 run, and more advanced versions, matching the published methods paper [660], is currently in use for the joint S6bucket/S6LV1 post-processing and the S6Directed runs.

Current work focusses on improving and generalizing the line-robust statistics. For example, a change in model priors can improve the performance for data from detectors with very different sensitivities, resulting in "weighted" LR-statistics. Modified line models, based on unmodulated sinusoids or on typical structures in parameter-space which have been identified during the joint S6bucket/S6LV1 post-processing, will also be investigated, as well as application of similar methods to time-domain disturbances.

4.3.3.5 Hierarchical Hough search – Ongoing search

The All-Sky search, as already described in past WPs [245, 246, 247, 248] for unknown neutron stars, is carried out in the Rome group with the "Frequency Hough Transform method", a transformation from the time/observed frequency plane to the source frequency/spin-down plane. A detailed method paper (following a previous one in which the basis of the method were outlined) is going to be submitted (May 2014). We have so far focused on the analysis of Virgo VSR2-VSR4 low frequency region, from 10 to 128 Hz, where the good Virgo sensitivity gives the opportunity to exploit a band never analyzed before for All-Sky searches, in particularly below 50-60 Hz. The presence of several disturbances has conducted to the development and implementation of various cleaning steps. The review of the method and software is nearly finished. The production of candidates (coarse and refined) for VSR2 and VSR4 data has been completed, search sensitivity and upper limits computed in the band 10-128 Hz for each run separately. Candidate clustering, to identify and consolidate groups of candidates that are reasonably due to a single "cause" (normally a disturbance), has been done and we are now working on the coincidences among candidates found in the two runs. Surviving candidates will be subject to various verification steps, to discard them or to furtherly increase confidence in detection. These verification consists in the application of various criteria not directly to the coincidences but, rather, to the candidates that originated them or even to the peaks in the peakmap. The possible instrumental origin of candidates will be studied also with the aid of detector experts. We are starting to write a observational paper containing the results of VSR2 and VSR4 low frequency analysis, which should be ready in summer 2014. Later, we plan to run some semi-directed search, e.g. toward the Galactic center, with small spin-down age (e.g. 100 years).

4.3.3.6 Follow-up of hierarchical Hough candidates – Exploratory

To allow exploring deeper in the detector noise, a follow-up procedure will be applied to lower-threshold candidates that survive after the coincidences and verification step and that cannot be ascribed to an instrumental origin. It will consist in the analysis of a small portion of the parameter space around each candidate with a longer coherence time. This implies the construction of a new set of longer FFTs and, possibly, a further application of the Hough transform procedure. We still have to design the details of the follow-up procedure. We aim at completing this work by spring 2015.

4.3.3.7 Loosely coherent search – Ongoing search

A new method called “loose coherence” is based on the notion of allowing slow drifts or modulations in phase from one SFT to the next, as described above for following up PowerFlux outliers. But, at least for small sky regions, the method can be applied from scratch in a dedicated pipeline to identify those outliers. A stand-alone program has been developed to permit “spotlight” searches for narrow regions in frequency, frequency derivative and sky location [241]. Production running began in fall 2012 for two spotlight directions in the galactic plane with overdensities of stars, associated with spiral arm spurs. Publication of the results is expected in late 2014.

4.3.3.8 Cross-correlation search – Exploratory

The cross-correlation method has been described previously. Its present uses are as a directed search for known neutron stars, either isolated (Sec 4.3.2) or in binary systems (Sec 4.3.4). It could in principle also be used for an all-sky search complementing the existing semi-coherent searches, but such plans are in the future.

4.3.4 Accreting and unknown binary neutron stars

For this class of source the gravitational radiation is thought to be powered by ongoing or previous accretion onto the neutron star. In this scenario, first proposed by [661], the neutron star achieves an equilibrium state whereby the angular momentum fed to it through accretion is balanced with the angular momentum radiated away through gravitational waves. This argument and its history is summarized in [231]. The resulting indirect limit can be put in terms of X-ray flux F_x and spin frequency $f_{\rm rot}$ as

$$h_{\rm IL} = 5 \times 10^{-27} \left(\frac{300\,{\rm Hz}}{f_{\rm rot}}\right)^{1/2} \left(\frac{F_x}{10^{-8}\,{\rm erg\,cm^{-2}\,s^{-1}}}\right)^{1/2}. \tag{21}$$

At present we divide the known accreting neutron stars into two groups: the low-mass X-ray binaries (LMXBs) and the accreting millisecond X-ray pulsars (AMXPs). Sources from both groups consist of a neutron star in orbit around a lower mass companion object from which it is accreting matter. From a data analysis perspective the key difference is that for the majority of the ∼85 known LMXBs the spin frequency of the neutron star is unknown but thought to lie in the range ∼200 Hz – 1 kHz and for the 7 known AMXPs the spin frequency (equal to the pulsar frequency) is known to high accuracy. This difference makes searches for the LMXBs a far more challenging task that of the AMXPs. Note that there are 10 LMXBs for which type-I thermonuclear bursts are seen from which the spin frequency can be constrained to within ∼1 Hz.

Another important difference comes from the indirectly measured time-averaged accretion rates which are typically at, or within a factor of a few of, the Eddington limit for the LMXBs. The AMXPs exhibit accretion rates lower by a factor of 10 – 100 in comparison. This difference, according to Wagoner’s arguments, makes the LMXBs likely to be stronger gravitational wave emitters than the AMXPs.

To date we have published a single coherent analysis on the accreting neutron star in the LMXB Sco X-1 using S2 data [231]. This was an exercise in wide multi-dimensional parameter space matched filtering and due to the rapid increase of search templates with observation time, the search was computationally limited to an observation time of only 6 h. Sco X-1, although the brightest X-ray source in the sky and consequently, also likely to also be the brightest continuous gravitational wave source, is typical of the LMXBs. As such it is clear that incoherent hierarchical search strategies need to be developed in order to maximise the search sensitivity, given the volume of parameter space we need to search and the computational resources. In this spirit, an incoherent search approach, based on the "radiometer" cross-correlation technique developed within the stochastic background group was applied to S4 data to set an upper-limit on radiation from Sco X-1 [647]. This result was recently updated with an S5 search, which also included a search in the directions of the galactic center and SN1987A [648]. A modification of the cross-correlation method to include information from the CW signal model, which allows cross-correlation of data from different times as well as different detectors, was proposed in [250]. This allows for a "tunable" semicoherent method where sensitivity can be increased by allowing the correlation of data further separated in time, at the expense of a greater needed resolution in parameter space. At present there is an ongoing study into the relative merits of each of the possible search strategies available within the LVC for targeting Sco X-1. A paper comparing and contrasting the methods, including the sideband, cross-correlation, twospect and radiometer searches is under way, with publication expected in summer 2014.

Finally, we are exploring new methods to carry out an all-sky search for unknown neutron stars in binary systems. Because the unknown orbital parameters increase the parameter space enormously, it is expected that only relatively insensitive methods using short coherence times will be feasible.

4.3.4.1 Sideband search for known binary systems – Ongoing search

The GWs from a continuously emitting source in a binary system will be received at a ground–based detector as a frequency and amplitude modulated signal. For known binary sources such as the low–mass X–ray binaries (LMXBs) we can remove the effects of the detector motion and maximize over the unknown amplitude modulation parameters through barycentric corrections and the use of the $\mathcal{F}$-statistic. The remaining time-dependent frequency modulation, due to the binary Doppler motion of the source, allows us to decompose the signal into the 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 $M \sim 4\pi f_{\rm gw} a \sin i/c$ frequency resolvable sidebands where $f_{\rm gw}$ 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 could be extracted by incoherently summing together the $\mathcal{F}$-statistic at each sideband frequency [367, 368]. 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 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), however, the sensitivity 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 to which this search is most suited. This includes the Z and atoll sources (rather than the accreting millisecond X-ray pulsars) which have known sky position, and for some, a reasonably well known orbital period. The remaining orbital parameters, semi-major axis, time of passage through the ascending node, 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 preliminary results have been obtained from a ten day stretch of S5 data. This data length has been chosen so that a monochromatic signal can be assumed for the duration observation time span. The expected sensitivity of this search will become astrophysically interesting (i.e., will start challenging accretion balance upper-limits) for advanced LIGO and specifically for the LMXB source Sco X-1.

A methods paper detailing the search was published in early 2014 [**?**], and the associated results paper is currently being drafted, with the long-suspended review planned to resume in mid 2014.

4.3.4.2 Cross-correlation searches for known binary systems – Exploratory

The cross-correlation search described in section 4.3.2 can also be applied to a search for binary systems at known sky locations, such as Sco X-1. The parameter space is three-dimensional, consisting of the gravitational wave frequency and the two unknown binary orbital parameters (e.g., projected semimajor axis and binary orbital phase), so a semi-coherent cross-correlation search with a short coherence time should allow a search using a manageable number of templates. This search should allow the use of more data than in the fully-coherent short-time search done in [231], and a more sensitive search than the incoherent cross-correlation search done in [648]. A methods paper [662] is in preparation which describes the expected sensitivity and computational cost of this search, as well as some further enhancements to the cross-correlation method (e.g., a parameter space metric, and the effects of windowing and leakage [663]).

A reorganized version of the LAL/LALApps cross-correlation code has been written, which addresses some of the shortcomings of the original pipeline (e.g., the need to keep track of the detection statistic for every point in parameter space rather than using a "toplist" of the most significant points), and which includes in its parameter space the orbital parameters for a neutron star in a binary system. It also incorporates the phase metric for the cross-correlation search when choosing its parameter space grid, and allows for combining signal from multiple frequency bins to avoid the effects of leakage [663]. A basic version of this pipeline is currently in place and being tested on simulated data. In the next year it will be refined and generalized to include a wider class of parameter space.

We are applying the pipeline in a search for radiation from Sco X-1. This will initially be done as part of the Sco X-1 Mock Data Challenge, described in Sec 4.3.4, and subsequently performed on S6/VSR2 data, with an eye towards having the pipeline fully developed by the time advanced detectors come on line.

4.3.4.3 TwoSpect searches for known binary systems – Exploratory & Ongoing search

The TwoSpect program, developed originally for all-sky binary searches and described below, has been recently updated to permit more optimized and sensitive searches for known binary systems, such as Sco X-1. The primary updates were eliminating the hierarchical first stage of the search, necessary in the all-sky search but harmful to sensitivity and exploitation of orbital parameters when known. This revised version of TwoSpect will be used in an opportunistic search for Sco X-1 and XTE J1751-305, while further pipeline improvements are explored, including coherent combination of different interferometer data streams and exploitation of orbital phase information.

4.3.4.4 Einstein@Home searches for known binary systems – Exploratory

The aim of this project consists of mainly analytically estimating the sensitivity of a directed search for continuous wave sources in binary systems by using Einstein@Home. For this purpose it is relevant to perform accurate template mismatch investigations to study the parameter-space metric for such kind of sources, which has been derived in [664].

By using the term "directed" we refer to searches where only the sky position of the source is supposed to be known with enough accuracy, while its frequency, frequency derivatives and orbital parameters are unknown. The parameter space we want to sieve is defined by the ranges of uncertainty on the frequency and signal orbital parameters. When we investigate such parameter space to search for the true signal parameters, we resort to the concept of the parameter-space metric, firstly described in [665, 666, 667], which allows us to define a measure of distance, thus establishing how to sample within the space. Hence, by following this geometrical approach, we can place templates within the parameter space in such a way that they will neither be too sparse to fail to spot a real signal nor too fine to waste computational power. In this way, any real signals will be recovered with some maximum tolerable loss in signal-to-noise ratio, which is directly related to the notion of mismatch [667, 668].

We recomputed the coherent and semichoherent parameter-space metric components given in [664] in the case of both a coherent segment length much larger and shorter than the source orbital period. In order to test the validity extent of such parameter-space metrics for circular, low- and high-eccentric orbits, we developed a very fast and efficient machinery, able to perform multiple software injections for known continuous wave sources in binary systems all at once, and to search for these signals and the corresponding mismatched parameter-space-points by computing the $\mathcal{F}$-statistic [669] values. In this process we generalized one of the routines originally written in octave by Karl Wette, using the SWIG software development tool to perform injections of isolated continuous wave sources, to the case of continuous wave sources located in binary systems.

All the numerical simulations we plan to perform will make use of the following LALSuite software package applications [670]: `Makefakedata_v4`, `ComputeFStatistic_v2` and `PredictFStat` to estimate the expected $\mathcal{F}$-statistic values for a given simulated continuous wave signal.

The details of the developed method and algorithm, as well as the achieved results, will be presented in a methodological paper, which we anticipate to be published in early 2015.

This accomplishment will pave the way to carry out a real Einstein@Home search for continuous wave sources in known binary systems in the upcoming advanced detector era.

4.3.4.5 Polynomial search for unknown binary systems – Exploratory

As discussed above, searches for unknown binaries present formidable computing challenges. 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_{ssb} detected in the solar system barycenter may be modeled as

$$f_{\text{ssb}} = f_{\text{gw}}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right) \tag{22}$$

with f_{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

$$\frac{df_{\text{ssb}}}{dt} = f_{\text{gw}}\gamma\left(1 - \frac{d\vec{v}}{dt}\cdot\frac{\vec{n}}{c}\right) + \frac{df_{\text{gw}}}{dt}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right). \tag{23}$$

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 orbital, 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 $0.002 \times f_{\rm gw}/s$.

In order to accommodate such large frequency shifts, a new search algorithm is being developed. An extension of the coherent search method with extra parameters to describe the orbital motion of the neutron star is not computationally feasible (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). A hierarchical search method like the stack-slide or Hough transform methods as discussed in Ref. [239] is also not promising, since the short FFT database must have a time length below about 25 s in order to keep the strength of the gravitational wave in 1 bin. As an alternative, we propose to apply 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). 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. For binary systems with orbital periods of the order of 4000 s, the coherence time is limited to about 500 s for this reason. However, for such waves the frequency could spread over hundreds of frequency bins in a 500 s Fourier transform, hence the proposed set of filters should give a sizeable improvement over stack-slide or Hough-transform techniques that start from a short FFT base. Searches for binary systems with larger orbital periods may be applied with a larger coherence time.

If a correlation between a filter and the data exceed 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 easily. We are currently developing this analysis strategy and the algorithms. Analysis of the Virgo and LIGO data with this set of filters could set an upper limit on the existence of gravitational waves within 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.

For this search, the code has been implemented and tested on simulated data with white noise. The results have been published in S. van der Putten's thesis [410]. Polynomial Search is also part of the Sco X-1 pipeline comparison project discussed in section 4.3.4.

Currently, we are working on a method paper and we plan to have the code reviewed soon, in order to analyse existing Virgo and LIGO data.

4.3.4.6 TwoSpect search for unknown binary systems – Ongoing search

The TwoSpect search is a hierarchical method for detecting unknown continuous wave sources from binary systems. The goal of the TwoSpect search is to probe regions of the large parameter space of pulsars in binary systems without exhausting the existing computational resources available. It seems unlikely that the search will have the sensitivity to make a detection in S5 or S6/VSR2-4 data, but since accreting neutron stars in binary systems are the best candidates to have large ellipticities, carrying out a search is prudent.

The TwoSpect method [369] relies on computing two successive power spectra of the calibrated strain data channel, hence the name TwoSpect. First, we take a power spectrum of the time series data, where the coherence time for the first power spectrum depends on the region of parameter space we wish to cover. For shorter-period binary systems, we use a shorter coherence time for each SFT. We make these choices to ensure the signal remains in one bin during most of each SFT interval. We then demodulate the SFTs 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 which are either rough approximations of a CW signal from a binary system (less computations required) or a more detailed approximation (more computations required). This two-stage pipeline acts as a filter to find the best candidates for a deeper search. We also use a spectrum folding algorithm known as Incoherent Harmonic Summing (IHS) developed by the radio pulsar community. This algorithm can provide a threshold filter for deciding whether or not to carry out a template calculation for a putative set of source parameters.

The first search (in S6 and VSR2-3 data) over a ∼500-Hz band was recently completed and reviewed. A paper presenting the first-ever all-sky upper limits on binary sources is expected to be submitted for publication in May 2014.

4.3.4.7 Sco X-1 pipeline comparison – Exploratory

The low-mass X-ray binary Scorpius X-1 is potentially our most luminous source of continuous gravitational wave radiation. Within the continuous wave and stochastic groups there exist a number of search pipelines that are sensitive to this source. We have begun the process of comparing the characteristics and sensitivities of these pipelines with regards to Sco X-1 through the use of a mock-data-challenge. Those participating in the comparison project have been meeting regularly since the end of 2012 to define and carry out the challenge. First results were obtained in May 2014, with a short-authorlist article describing the study planned for summer 2014.

4.3.5 Support infrastructure

There is critical infrastructure needed to support the various CW searches. New software injections will be an important element in evaluating the performance of search pipelines. Systematic detection and identification of instrumental or environmental spectral lines, particularly wandering lines, has been and will continue to be necessary to eliminate outliers in many searches. In this section we describe ongoing and planned work to address these support needs.

4.3.5.1 Mock data challenge and software injections

An important new element in evaluating search pipeline performance is the creation of simulated data with several thousand software injections for pulsars at randomly chosen sky locations and with frequencies and frequency derivatives spanning our search ranges. These standardized data sets should enable us to compare with more statistical precision and more systematically the sensitivity and robustness of the various pipelines in detecting pulsars in blind searches or in reconstructing source parameters in targeted searches. The outcomes of these comparisons might well lead to the abandonment (or at least de-prioritization) of some search pipelines.

We created 6 000 parameter files for artificial pulsars and generated signals from these over the period of S6. The signals have been added to a replica of the S6 frames and distributed to the LSC clusters. These reference software injections are being used to compare and contrast the performance of search pipelines in an ongoing, four-stage mock data challenge. Results from stage 1 of the MDC (∼20 injections) were reported by a subset of targeted, directed and all-sky isolated-star pipelines in December 2012 and January 2013. A larger subset of pipelines participated in stage 2 (∼200 injections) in summer 2013. Stage 3 will conclude in spring 2014 (∼2200 injections), and stage 4 will conclude in fall 2014 (∼2200 *blind* injections).

In preparation for the advanced detector era engineering runs are being performed to simulate detector data. In the first three engineering runs, ER1-3, two simulated pulsars were created and injected into the dataset. These have been successfully extracted using the known pulsar search method.

4.3.5.2 CW detector characterization

The CW Group has a strong interest in monitoring (and mitigating) instrumental spectral lines (see the detector characterization chapter) with low latency. In preparation for S6 the nascent "F-Scan" infrastructure developed during S5 for generating high-resolution spectrograms to detect wandering lines visually was automated and expanded to auxiliary channels and to Virgo channels. These spectrograms and data files proved invaluable in quickly spotting and tracking down instrumental lines in S6 data. The F-Scan output has also been mined for spectral coincidences between interferometers and between strain and auxiliary channels. Further improvements included dedicated F-scans for bands relevant to special pulsars, such as the Crab and Vela. An F-Scan version with further enhancements has been run on H1 and L1 PSL data, on H2 One Arm Test data in summer 2012, H1 Half-Interferometer (HIFO) Y arm data in summer 2013 and is being run in 2014 on H1 HIFO X and L1 HIFO X and Y. Full-interferometer studies will be carried out when the data becomes available.

In addition, the offline auxiliary-channel coherence studies used in following up on S4 and S5 pulsar candidate outliers have continued for S6/VSR2 analyses. Special attention has been given to narrow spectral bands around a half dozen known pulsars for which the spin-down limit may be accessible in S6/VSR2/VSR4 and to bands containing outliers coincident between H1 and L1. A new technique using spectral correlations is also under development.

On the Virgo side, a new infrastructure called NoEMi has been developed for identifying/mining stationary or slowly wandering spectral lines. It was run offline on VSR2 data, and in real-time on VSR3 and VSR4 data. During science runs NoEMi publishes daily summary pages, with plots and lists of the found noise lines, on the Virgo monitoring web pages. In particular, for the objectives of the CW targeted search, noise lines found at the same frequencies of the known pulsars are highlighted in the summary pages. More details on NoEMi can be found in [671].

Lines found by NoEMi are stored in a database which is made accessible via a web interface. The database allows to add user-defined information (metadata) to each identified line.

NoEMi will be used as the main noise line monitoring tool for the Advanced Virgo, LIGO and GEO detectors. NoEMi will run on the "local" computing centers (thus distributing the computing load and storage requirements) and send the formatted data to a centralised (cluster of) databases. As a preliminary test-bench of the upgraded framework, NoEMi was installed and run at Caltech on the full S6 LIGO data and then commissioned at Hanford to run daily on the H2 One Arm Test data in summer 2012. This infrastructure too has been run on the HIFO data at Hanford and Livingston as arms have been commissioned.

4.3.6 Scientific goals and plans for the advanced-detector era

As discussed above, most of the searches carried out by the CW Search Group are computationally limited. We simply cannot search all of the parameter space we wish to search. We must compromise and make strategic choices, mindful that we have no "guaranteed" sources, even in the advanced detector era. Keeping our eyes open to a variety of potential CW sources is prudent in current searches and will remain so. Consequently, we will maintain (and likely expand) a broad suite of search approaches.

That said, we have in mind a number of high-priority searches we want to ensure are carried out promptly in the advanced-detector era, as sufficient data becomes available. For some searches we want more than one independent search pipeline in place, both to reduce the chance of error (software bug or user error) and to allow better assessment of systematic uncertainties, especially in the event of a discovery.

We commit as a group to producing journal-ready papers within the first two years of the advanced detector era on the following "flagship" subjects, with highest priority given to the first two:

1. Targeted search for $>$100 known pulsars (radio, X-Ray, γ-ray).
2. All-sky search for unknown isolated neutron stars ($\sim$10-2000 Hz) and spindown range $0 < \dot{f}_{\rm GW} < 10^{-8}$ Hz s^{-1}.
3. Narrowband search for the targeted pulsars above, in which a slight frequency mismatch between electromagnetic and gravitational wave mismatch is permitted.
4. Directed search for interesting objects & sky locations assumed to be isolated, such as Cassiopeia A (if pulsations not yet detected), SN1987A, and isolated stars near the galactic center and in nearby globular clusters.
5. Directed search for known binary sources without pulsations, such as Sco X-1, and assumed binary sources near the galactic center and in nearby globular clusters.

In addition, we will strive to be in a position to produce journal-ready papers within the first two years on the following subjects:

1. All-sky search for unknown stars in binary systems
2. Search for transients from known stars exhibiting electromagnetic glitches

As will become clear below, we already have a multitude of search pipelines for some of the searches we intend to carry out, at a variety of developmental stages. Hence in preparing to meet our science goals for the advanced-detector era, the CW group will undergo a period of consolidation and stabilization of the flagship search pipelines in the various high-priority search categories discussed above.

The performance of the various pipelines will be evaluated via a standardized set of software signal injections into a parallel stream of real data from the LIGO and Virgo interferometers. There will be several thousand injections of isolated, binary, and glitching pulsars distributed uniformly in frequency across the detection band and randomly in sky location, orientation, spindown, and other parameters that apply to binary or glitching sources. Pipelines will be evaluated according to their success in recovering these signals and in the precision of their source parameter estimations.

We believe that evaluating standard figures of merit for sensitivity, parameter space coverage, robustness and speed on this simulation data set will reveal which pipelines can be safely abandoned and which should receive the efforts needed for stabilization in the advanced-detector era.

Mature plans for targeted and all-sky isolated-star searches can be found in section 3. In the following, we address for the other flagship searches and for each "desirable" search the following issues:

- the analysis tools (pipelines) needed and the readiness of those tools for use, including review status;
- the interpretation tools needed in the event of a discovery, although none of these searches is guaranteed a discovery; and
- expected milestone dates for being ready (where milestones can be specified with confidence).

Flagship Searches:

1. *Targeted search for $>$100 known pulsars.* (see section 3)
2. *All-sky search for unknown isolated neutron stars.* (see section 3)

3. *Narrowband search for the targeted pulsars, allowing EM/GW mismatch.*

    - We plan to have at least one stable and proven narrowband pipeline to search for signals only slightly mismatched in frequency from that implied by electromagnetic pulsations. One such pipeline was used in a published search near the Crab frequency in early S5 data. That pipeline has been adapted to use Doppler resampling code, to gain in computational speed, code which was largely reviewed in the context of the recent Calvera search (aborted when out-of-band pulsations observed by x-ray astronomers).
    - In the event of a discovery, it should be straightforward to "zoom in" to determine source parameters well enough for detection with targeted search pipelines, in which case SNR would likely increase, and parameter estimation prove to be powerful.
    - We expect to produce a journal-ready publication using the current (or slightly improved) pipeline on the time scale of summer 2014. That pipeline should be ready for the advanced detector era with a snapshot taken. Software injections will be used to evaluate its performance.

4. *Directed search for interesting isolated objects & sky locations.*

    - We plan to have at least one stable and proven directed-search pipeline for isolated, well modeled neutron stars (known direction, but unknown frequency) based on long coherence times. Recent or ongoing searches for Cas A, Calvera, the galactic center and nearby globular clusters, all based on the $\mathcal{F}$-statistic, provide prototype approaches. For stars satisfying the nominal source model, Doppler resampling will be valuable in reducing the computational cost of these searches. While more than one resampling method has been used within the group, none has been completely reviewed yet. In addition, we plan to have a pipeline robust against source modeling errors, one based on cross correlation methods, now under development for a SN 1987A search.
    - In the event of a discovery, zooming in is, again, likely to prove effective and permit powerful parameter estimation.
    - We expect at least one Fstat-based directed, isolated-star searches to be fully reviewed by summer 2014 with published results. Review of the cross-correlation pipeline will likely complete in 2015. Again, software injections will be used to evaluate the performance of all directed-search pipelines.

5. *Directed search for binary sources in known directions.*

    - We plan to have at least three independent, stable and proven pipelines able to search for Sco X-1 and other LMXBs. The current Sideband and TwoSpect pipelines, along with the Radiometer pipeline (used in stochastic GW searches), provide defaults, but we expect other pipelines (Cross-correlation, Polynomial, PowerFlux and an Einstein@Home Fstat-based search) to become available for Sco X-1 application, too. Given Sco X-1's relatively high likelihood of producing detectable radiation and its likely phase wandering during accretion, it is prudent to attack it with a variety of algorithms.
    - In the event of a discovery, zooming in could prove challenging for accreting binaries, but at the same time, if a source is strong enough to be seen over short coherence times, then it may be possible to track time-varying accretion effects directly, which could yield extremely interesting insights into both the accretion process and neutron star structure. Existing tools in targeted

search pipelines should be adequate for tracking, if the signal is strong enough to be detected in the first place.

- The Stochastic group's Radiometer pipeline has been fully reviewed and a paper including a Sco X-1 search published in 2011, with another expected in 2014. The Sideband review is well under way, with a journal-ready paper planned for the end of 2014. The now-mature TwoSpect pipeline developed for an all-sky search (see below) can also be applied to directed binary searches with only modest code changes, with review expected to be complete in late 2014. All three pipelines should be available for use with advanced detector data. The cross-correlation and Polynomial algorithms (optimized for a directed search) are less mature, but reviews have begun, and vetted pipelines should be available by the time advanced detector data taking begins. All pipelines are participating in a dedicated mock data challenge, with results to be published in a methods article in late 2014.

Additional Planned Searches:

1. *All-sky search for unknown stars in binary systems.*

   - We plan to have at least one stable and proven all-sky search for unknown neutron stars in binary systems. The TwoSpect pipeline developed on S6/VSR2/VSR4 data provides a default, and the less mature Polynomial pipeline offers an alternative approach likely to be ready by the start of advanced detector data taking. Review of the TwoSpect pipeline completed in spring 2014. Review of the Polynomial pipeline has begun.
   - In the event of discovery, zooming in should be possible. Because significant sensitivity must be sacrificed (even more than for the all-sky isolated searches) to gain computational tractability, potential SNR gains are very large after initial detection. At the same time, zooming in on a binary may prove difficult and will need serious investigation between now and the advanced detector era. As for the discovery of an new isolated star, we favor strong interaction with electromagnetic astronomers once we have thoroughly explored the GW data. Again, statements in the discovery paper on neutron star population densities may be appropriate, depending on GW and EM observations.
   - The TwoSpect search completed in spring 2014 and an S6 / VSR2-3 paper approved for submission, and a reviewed, tagged pipeline is ready for use with advanced detector data (although some modest improvements will also be explored). The Polynomial search suffered a setback because of a graduated student's departure from the field, but another student has now taken on the project; it is hoped the pipeline will be complete and reviewed in time for advanced detector data taking. Software injections will be used to evaluate and compare these two pipelines (and any other that may arise).

2. *Search for transients from known stars.*

   - We hope to have at least one pipeline to search for long-lived transients associated with observed electromagnetic glitches in neutron stars, including magnetars. (It should be noted that both the Burst and Stochastic search groups have developed or are developing pipelines for post-glitch searches, using more generic approaches.) A transient CW pipeline has been developed recently for use in S5-S6 / VSR2-4 searches, but has not yet been reviewed.

- In the event of discovery, it may not be possible to zoom in to increase SNR, as the source may go quiet again for years. In that respect the transient search is similar to a burst search. Signal interpretation will depend on reconstructed signal strength, frequency and spindown. One could observe f modes, r modes or torsional modes, for example. Depending on how quickly a signal is detected, it may be appropriate to issue an ATel before publishing the discovery paper.

In addition, because many of the above pipelines are already mature or will reach maturity soon, we expect to have time between now and the advanced-detector era to expand our list of search targets and to develop new search algorithms.

New astrophysical targets include:

- Unknown isolated stars with unstable frequency or phase evolution. *(Search algorithms now under development, based on "loose coherence" and cross-correlation with frequency demodulation provide starting points for such searches.)*
- Newborn neutrons stars in our galaxy. Searching for continuous waves from a newborn star (should one be detected via a nearby supernova during the advanced-detector era) will be difficult because the star will likely spin down rapidly and with considerable phase instability. The radiometer directional pipeline of the Stochastic group may provide the best starting point for such a search.
- Extreme unknown binaries. All-sky binary searches currently under development may not perform well for binaries characterized by extreme eccentricity or large mass asymmetry.
- Very nearby stars with non-negligible proper motion. The first isolated star to be found in an all-sky search may well be very close to us. If so and if the proper motion of the star is high, then present coherent follow-up techniques will not be optimum, perhaps requiring a new search parameter.

Depending on simulation studies of performance of current search pipelines on the sources above, new algorithms may require development. It should be kept in mind that Moore's Law will automatically allow us to expand the parameter space searched. For example, more CPU time allows larger spindown ranges and/or more frequency derivatives to be explored in all-sky and directed searches. Similarly, the discrete number of sky locations explored by directed searches can be expanded beyond the list used in the initial-detector era.

But we can already identify some future algorithmic needs:

- Two pioneering all-sky search pipelines for unknown binaries, using very different approaches, are nearing maturity, but it's unlikely that these two approaches will fully explore what is feasible for an extremely computationally limited search. It is prudent to explore other approaches.
- Present all-sky upper limits on unknown CW sources do not take into account the spatial distribution believed to characterize neutron stars in our galaxy. More astronomically informative upper limits should be derivable by using population simulations.
- There are potentially large gains in computing from the use of Graphical Processor Units (GPUs), but such gains will, in most cases, require substantial restructuring of search pipelines to ensure that I/O limitations do not make the gains from arithmetic efficiency largely irrelevant.

Finally, we plan to build further on existing detector characterization infrastructure used currently to find spectral lines and identify their sources via coherence and/or spectral coincidence with auxiliary channels. The real-time monitoring will be enhanced with increasing sophistication and automation. Cataloguing of lines will also be improved to be more comprehensive and informative.

In regard to upcoming engineering runs, we are most interested in hardware runs with partial or full interferometers, such as the summer 2012 H2 One Arm Test and the 2013-2014 half-interferometer tests, allowing an early cataloguing of new spectral lines likely to plague advanced detector data. Purely software engineering runs are less relevant to our searches, since we don't strive for low latency and have implemented our own software injections. We note that single-interferometer engineering runs with sensitivity better than S6 or VSR4 data are astrophysically interesting, since one interferometer is sufficient to establish definitive discovery.

4.4 Stochastic Group

This section describes stochastic group work which is not described in search plans. This includes: initial LIGO-Virgo analyses, work that may eventually be formulated as a search plan, work that is in collaboration with search plans outside the stochastic group, work that does not fit into a search plan (e.g., detector characterization in support of stochastic analyses), and research and development for new ideas and techniques.

4.4.1 Initial LIGO-Virgo Searches

Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009–2010 LIGO and Virgo Data (in review). This search [325] implements the stochastic group's flagship isotropic analysis with data from S6 and VSR2-3. It improves on the S5 analyses by analyzing the complete frequency range from 41.5–1726 Hz as well as by including Virgo data at all frequencies. The strain sensitivity of the LIGO $4\,\mathrm{km}$ interferometers (H1 and L1) is improved compared to S5 at frequencies above 300 Hz. However, the coincidence time is limited by the relatively low duty cycle, especially early in the run. We find no evidence for an isotropic signal and set upper limits on the SGWB energy density. In the frequency band of 41.5–169.25 Hz, we set a 95% CL upper limit of 5.6×10^{-6}, a 38% improvement over the previous limit. For 600–1000 Hz, we set a limit of 0.14, a factor of 2.5 improvement over the earlier limit. In the bands of 170–600 Hz and 1000–1726 Hz, we set the first direct limits with values of 1.8×10^{-4} and 1.0 respectively.

Searching for stochastic gravitational waves using data from the two co-located LIGO Hanford interferometers (in review): The isotropic searches performed up to date have used non-collocated interferometer pairs because of the reduced chances of environmental correlations. The LHO interferometer pair, however, could potentially be $\sim 10\times$ more sensitive to stochastic GW background due to the more favorable overlap reduction function. However, the collocated interferometer pair also suffers from instrumental correlations, because the two interferometers share the same environment and the same sources of instrumental noise. The stochastic group developed two methods to estimate and suppress the instrumental correlations, as discussed above in more detail. The group has applied these methods to the S5 data, and the results [326] indicate that the PEM-coherence and the time-shift approaches identify well the grossly contaminated frequency bands, which are then removed from the analysis. Moreover, the PEM-coherence approach can be used to estimate the residual contamination in the ”good” frequency bands. The co-located interferometer (H1H2) analysis was performed in two frequency bands: low-frequency (80-$160\,\mathrm{Hz}$) and high-frequency (460-$1000\,\mathrm{Hz}$). The low-frequency band was found to contain residual correlated noise at a level that rendered the resulting upper limit noncompetitive. In the high-frequency band, however, we obtain an upper limit of $\Omega(f) < 7.7 \times 10^{-4}(f/900\,\mathrm{Hz})^3$, which represents a factor of $> 450\times$ improvement over the previous best limit. In addition to this important result, we expect the paper to serve as a useful reference for future analyses dealing with correlated noise.

Directional search for persistent GWs with S6/VSR2/VSR3 data: The stochastic grouphas conducted a directional search using S4 and S5 data [342, 343], producing upper limit maps of the gravitational wave sky and producing constraints on some of the most promissing point sources on the sky, including Sco-X1, the Galactic Center, and the SN1987. This search aims to improve on the past directional searches, taking advantage of the improved strain sensitivity of LIGO and Virgo detectors, and including the Virgo data into the directional search for the first time. New upper limits have already been produced for the SN1987 and the Galactic Center. The group plans to repeat the search in the Sco-X1 direction using an improved radiometer pipeline, which will choose the frequency binning so as to improve the sensitivity of the search by up to a factor of 2. The group will also produce new maps of the gravitational wave sky using the radiometer and spherical harmonics algorithms. We expect the search to be completed by the end of 2014, with the internal review commencing in early 2015.

An all-sky (untriggered) search for long-lived gravitational waves. STAMP (Stochastic Transient Analysis Multi-detector Pipeline) is an extension of the stochastic pipeline. STAMP analyzes the same cross-correlation data products used in a stochastic search in order to identify long-lived transient signals lasting $\gtrsim 10\,\mathrm{s}$. The STAMP All-Sky (STAMP-AS) analysis is an untriggered search looking for long duration $\sim 10 - 600\,\mathrm{s}$ GW transients [569]. Possible astrophysical sources for such signals are accretion disk instabilities, instabilities in nascent neutron stars, eccentric black hole binary systems, and neutron stars quasi-periodic oscillations. To carry out this search, several enhancements have been made with respect to the original STAMP code such as a faster clustering algorithm, allowing the code to look for a signal everywhere in the sky even without investigating every sky direction, and a more efficient mechanism to cross-correlate data. A dedicated STAMP-AS pipeline has been run on $\approx 1\,\mathrm{yr}$ of coincident Hanford-Livingston data from S5, including 200 time slides for background estimation and a thorough sensitivity study for a dozen test waveforms. (The STAMP-AS team has also performed extensive data-quality investigations.) At present, the analysis team is ready to open the box and to begin a review.

4.4.2 Advanced detector era

Searches for Long-Duration Transients. The STAMP code suite (see above) has reached a mature stage. The first analysis targeting GRBs has been published [490]. The more computationally challenging STAMP all-sky analysis is nearly complete; see above. The improved sensitivity of aLIGO, coupled with recent improvements to STAMP clustering algorithms, should allow us to probe long-duration transient models out to astrophysically interesting distances: $\approx 600\,\mathrm{Mpc}$ for extreme models emitting $E_{\mathrm{GW}} = 0.1 M_\odot$ of energy in GWs [564]. Recent work has demonstrated that seedless clustering can improve the sensitivity of both targeted and all-sky searches [564, 565] by a factor of up to two in distance. Other work has investigated new techniques for parameter estimation [574] and specialized searches targeting compact binary systems [568] (envisioned as a follow-up to matched filter triggers). The ADE triggered STAMP search will be applied to interesting electromagnetic triggers including gamma-ray bursts and supernovae. The ADE all-sky STAMP search will be run on all high-sensitivity data.

Galactic Neutron Star Search. Directional searches for the anisotropic SGWB have been conducted in the past, both in the spherical harmonic and in the pixel basis. These searches can be naturally extended to search for a stochastic signal localized in the Milky Way galactic plane [333]. The primary science target for this search would be the galactic millisecond pulsars, whose individual contributions can be summed up to a stochastic signal localized to the galactic plane. Other GW production mechanisms based on neutron stars (e.g. models based on r-mode or bar-mode instabilities) could also contribute to this signal. We are developing the new techniques needed to perform parameter estimation for such spatially extended models. While the past parameter estimation efforts were focused on the amplitude and shape of the SGWB frequency spectrum, the new techniques will also estimate parameters defining the spatial extent of the signal. The new search will be tested using data from initial detectors, as a part of the preparation for the advanced detector era.

4.4.3 Work in support of stochastic analyses

Modeling the stochastic background. Astrophysical SGWBs may result from the superposition of a large number of unresolved sources since the beginning of stellar activity. Among the most promising models are mergers of compact binary neutron stars and/or black holes, instabilities in rotating neutron stars such as r-modes, bar-modes, or quasinormal modes, magnetars, and core collapse supernovae to neutron stars or black holes including hypothetical population-III stars. Detection of an astrophysical GW background would not only provide information about the physical properties of the respective astrophysical objects (such as compact binaries of neutron stars), complementing individual GW detections and electromagnetic observations of these objects such as gamma ray bursts, but it would also elucidate the evolution of these objects with redshift and trace the star formation history or the metallicity. Recent studies suggest that the astrophysical SGWB from compact binary coalescences may be detectable with Advanced LIGO-Virgo [315]. Characterizing the SGWB from compact binary coalescences is a major focus of the group as we prepare for the advanced detector era.

The stochastic group is developing a catalog of cosmological and astrophysical sources of the SGWB. While many SGWB models have been proposed in the literature, most of them have not been systematically studied in terms of their accessibility to future GW detectors. It is therefore critical and timely to perform a systematic study of these models and identify specific science questions in the framework of these models that could be addressed by the second-generation GW detector network. Furthermore, as the GW community is beginning to conceive the design of the third-generation detectors, it is also important to identify the SGWB science targets that could be pursued by these detectors, thereby influencing the design requirements for such detectors. This effort is closely linked with the development of the parameter estimation techniques and the mock data project (see below).

This effort has already yielded some very interesting results: for example, if one assumes realistic rates of coalescences of binary neutron stars and/or black holes (such rates would imply tens of individual coalescences observed per year by the second-generation detectors), then the GW background generated by summing signals from all binary coalescences in the Universe may also be detectable by the second-generation detectors [315]. In a separate study [341], our group have shown that a large part of the parameter space of the model for stochastic GW background due to magnetars is also accessible to the second-generation detectors. Consequently, a measurement of the stochastic GW background by second-generation detectors would constrain the ellipticity and the equation of state of these neutron stars in strong magnetic fields. Other studies have shown that similar constraints could be placed on models of cosmic (super)strings, pre-Big-Bang models, and polarized stochastic backgrounds from parity-violating inflation models [672].

Parameter estimation. Until recently, stochastic isotropic searches were limited to estimating the amplitude of the stochastic background (for a specific assumed spectral shape). The group has recently developed techniques that can be used to estimate additional model parameters [311]. Parameter estimation techniques can be used to constrain parameters in generic models (such as the power index in the power-law spectral form), or to constrain fundamental parameters in SGWB models, as identified in the systematic SGWB model studies. Past work on stochastic parameter estimation has focused on estimating the parameters of simple (two-parameter) models including generic power-law models and models of the SGWB from binary neutron stars and binary black holes [311]. In the case of an SGWB from compact binary coalescences, stochastic measurements can be combined with transient measurements (of individual coalescences) to gain more information about the source of an observed stochastic background [311].

Ongoing and future work will explore a wider variety of models including those characterized by more than two parameters. Exploring the higher-dimensional likelihood functions requires the application of numerical integration tools such as nested sampling [673] and/or Markov Chain Monte Carlo. The group has developed the infrastructure needed for application of such tools, and is currently using it for a systematic study of the SGWB model due to the stellar collapse of population-III stars. We expect this infrastructure to enable future work on combining information from other sources (such as transient GW detections) in order to disentangle different sources of the SGWB. We have also started development of parameter estimation techniques that would estimate not only frequency spectrum parameters, but also the parameters defining the spatial distribution of the stochastic signal. These techniques are being applied to the SGWB model due to the galactic neutron stars, which is limited to the Milky Way galactic plane.

Mock Data Project. In order to prepare for the advanced detector era, the stochastic group has created infrastructure for generating mock data (in *.gwf frame files) containing various SGWB signatures, such as from a population of BNS/BBH coalescences. The group is using this mock dataset to validate existing analysis tools and as a testbed to study new ones. For example, preliminary tests have been carried out to as-certain if there are any complications that might arise when we try to measure a non-Gaussian astrophysical stochastic background [547]. We plan to use this dataset to test new parameter estimation schemes, identify and remove astrophysical foreground, and to benchmark new non-Gaussian pipelines. In this way, the mock data project will facilitate a number of other group projects.

In order to create a realistic mock dataset, we draw on the experience gained from our projects to model sources of the SGWB (see above). The current mock dataset, populated with compact binary coalescences, utilizes observations and population synthesis results (for example the evolutionary code StarTrack) to generate sources realistically distributed in the parameter space (mass, spin, merging time, orientation, sky position) but also in time and redshift, accounting for the star formation history of the universe. The mock data generation code [674, 675] based on the LAL library was first developed in the context of the Einstein Telescope and then extended to the network of advanced detectors. It includes the possibility to simulate Gaussian SGWB (isotropic or from a specific direction in the sky) , popcorn SGWB, and future work will include recolored (glitchy) noise, and the use of galaxy catalogs to model anisotropies in the SGWB.

Sco-X1 Mock Data Challenge. The stochastic groupis participating in the Sco-X1 Mock Data Challenge organized by the CW group, whose goal is to compare different search pipelines that could be used to target accreting (and uknown) binary neutron stars. In particular, our group is applying the radiometer pipeline to analyze the simulated data containing potential GW signals coming from the direction of Sco-X1. The group is currently assessing the results of the simulated search, and a paper summarizing these results is planned for 2014.

Detector characterization for stochastic searches *Correlated noise.* The greatest detector characterization challenge facing the stochastic group in the advanced detector era is the possibility of correlated noise at widely separated detectors. Recent studies suggest that correlated noise from global Schumann resonances may occur at a level that is problematic for advanced detector SGWB analyses [332]. Correlated noise is problematic for SGWB analyses because it can bias our estimator for $\Omega(f)$.

The group has undertaken a multi-pronged approach to address correlated noise. First, we are working with instrumental experts and commissioners to determine if the coupling of magnetic fields to test mass motion can be reduced to a level that is acceptable for aLIGO SGWB analyses. Second, we are investigating the possibility of a Wiener filter subtraction scheme in order to measure and subtract correlated noise from Schumann resonances [344]. As part of this effort, we are studying whether the efficacy of subtraction can be improved by relocating magnetometers to magnetically quiet locations. We are looking into the performance of different magnetometers. Various schemes and proposals are evaluated using a Monte Carlo model. Finally, we are working to develop frameworks in which we can make astrophysical statements in the presence of correlated noise. This last effort overlaps with recent developments in the S5 H1H2 analysis.

Noise lines. The stochastic group continues to investigate noise lines. 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 timing in the data acquisition systems. A strong line in one interferometer, along with a large random fluctuation in the other, can also produce a signal at a specific frequency in the stochastic search pipeline. In the advanced detector era the coherence between pairs interferometers' output will be calculated in real time. 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 even be addressed and fixed in the laboratory.

STAMP-PEM. The STAMP pipeline has been demonstrated to be a good tool to search for long-duration noise transients in interferometer data [573, 572, 676]. The group has developed a stand-alone package, STAMP-PEM, that can be run by people not familiar with the STAMP pipeline. This code generates STAMP results and summary web pages that document results. The pipeline is useful as a real-time monitor of coherence between the strain channel (or some proxy strain channel) and physical environmental monitor channels. Currently the pipeline is used to produce hourly summary pages during engineering runs documenting coherence between main interferometer channel and other auxiliary and PEM channels Hourly summary pages are produced during science/engineering runs that document coherence between the strain channel and a variety of other channels over the previous hour. STAMP-PEM can also be used as an effective follow-up tool for characterizing noise in specific subsystems, investigating coherence between various degrees of freedom within the detector, and following up on short-duration transients caused by environmental factors [677]. This noise search effort will be done in collaboration with the detector characterization group.

Stochmon: realtime monitor for stochastic searches. The group is developing a web-based monitoring tool called *stochmon*. The *stochmon* webpage will include a variety of standard diagnostic plots and metrics, which will give the group an idea of the sensitivity for stochastic searches achievable with the data collected so far (without unblinding the final results). It will also serve to alert the group to detector characterization problems such as coherent electronic lines and non-stationary noise. By providing this information in near real time, *stochmon* will facilitate rapid feedback to commissioners. It will help us to plan our publication strategy by showing when our sensitivity is projected to reach interesting levels. Finally, it will set the stage for real-time data analysis, helping the group to obtain final results shortly after the completion of a science run.

4.4.4 Work in development

Non-Gaussian pipelines Astrophysical sources of the SGWB tend to be non-Gaussian whereas cosmological sources tend to be Gaussian. Measuring the non-Gaussianity of the SGWB may therefore provide a tool for disentangling astrophysical and cosmological sources. It may also be possible to improve our search sensitivity by including information about the Gaussianity of the signal. The group is exploring the possibility of one or more non-Gaussian pipelines. A framework for measuring the Gaussianity of the SGWB using realistic (non-colocated) detectors has been proposed [678], as well as an independent approach, relying on the fourth-order Edgeworth expansion [679]. Future work will investigate the application of this technique (and possibly others) to the measurement of astrophysical SGWBs. This project will build on the mock data project described above.

Stochastic backgrounds with non-standard polarizations with VIRGO and LIGO. General relativity predicts gravitational waves with two independent transverse polarizations, which transform as tensors for rotations around the propagation direction. The purpose of this project is to search for a stochastic backgrounds associated with non-standard polarizations, of vector and scalar character. These are a generic prediction of extended theories of gravitation. For an isotropic search, a minimal modification of the standard pipeline is required, namely the correct overlap reduction functions must be introduced. Preliminary investigations are underway to assess the possibility of an analysis with S5/VSR1 and S6/VSR2 data. The expected output are a set of upper limits which are expected to be the best ones in the relevant frequency range.

Fast time-domain cross correlation. The purpose of this project is to develop a very compact real–time pipeline for cross correlating the outputs of multiple gravitational wave detectors and generating event triggers from this cross correlation data. Key elements of this pipeline are real-time preprocessing of frame data including recolouring of raw frame data to optimise sensitivity and characterisation and removal of sinusoidal backgrounds (lines), an efficient cross correlation algorithm; an event trigger generator from the cross correlation output. The nature of the pipeline makes it a promising approach to analysis of broadband signals, such as those probed by white noise burst injections. It is suitable for probing unmodelled signals of any time duration, so we are interested in probing the same long time duration burst signal regime as is targeted by STAMP-PEM. The pipeline is implemented in C and C++ and is currently being applied to mock data challenge injections on S6 frame data to determine receiver operating characteristics. In addition, elements of this pipeline, particularly the recolouring and line removal codes, are applicable to detector characterisation tasks, and have been used for investigations of sinusoidal backgrounds as part of the work of this group. The short–term goal of this project is to optimise the pipeline for sensitivity and computational efficiency and to prepare for review of the codes and algorithms in preparation for use in the advanced detector era.

STAMP investigation of r-modes. Over the decades, r-modes from neutron stars have been very controversial. Slowly spinning stars may cool below some threshold before they have a chance to be spun up to the point where thermal run-away can occur. However, newborn neutron stars may present a window of opportunity for the observation of r-modes (especially those with rapid spin-down rates), since newborn neutron stars can be expected to be rapidly rotating, relatively hot and highly volatile. Moreover observations of both the r-mode frequency and the spin-down rate could begin to put constraints on neutron star equations of state since, with a modest amount of astrophysical data, those observations are enough to pin down the moment of inertia of a star which, in turn constrains the allowable mass and radius for the star. This possibility is an active area of investigation.

Members of the group have been exploring the sensitivity of Advanced LIGO-like instruments to possible detection of r-mode emission from newborn neutron stars, as a function of initial frequency and the spin-down rate. Along with similar trends elsewhere in the LSC, members are also exploring the novel potential of machine learning algorithms to improve the sensitivity from that of clustering algorithms toward that of match filtering. The aim is to have all the relevant search machinery in place before the high sensitivity data from aLIGO becomes available.

5 Characterization of the Detectors and Their Data

5.1 LSC Detector Characterization

5.1.1 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 completed six periods of scientific running, S1-S6, over the past decade, amassing several years of coincident observation with its detectors at design sensitivity, and often in coincidence with its international partners GEO600 and Virgo. LIGO is in the midst 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. Advanced LIGO is expected to make the first gravitational-wave detections and begin an era of gravitational-wave astronomy.

The LSC detector characterization group [680] directly supports LIGO's mission because a thorough characterization of the detectors is required to confidently detect gravitational waves. Gravitational-wave searches require accurately calibrated data with precise timing. 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.

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.

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 improved 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 problems. For initial LIGO, the online monitoring systems included the Data Monitoring Tool (DMT)[681] with a number of targeted monitors, the controls system software (EPICS)[682], and search-oriented monitors such as the trigger generators Omega [683] and KleineWelle [684] for the burst search, and daily iHope [685] 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 plan to build 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.The DetChar working group concentrated most of its effort in the initial detector era on 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. Every search working group has active DetChar members, and information flows in both directions, to mutual benefit. *The Det Char group has few members exclusively or even mostly dedicated to the group. This brings a beneficial diversity of ideas, but reduces efficiency with respect to full-time members. A goal of our group is to recruit more members that commit a substantial amount of their time to the group. The current LSC-Virgo MOU includes a less-restrictive publication policy for detector characterization work that supports this goal by allowing short author list publications on limited sets of interferometer data.*

In the following subsections, we describe the requirements on detector characterization for our collaboration to achieve its scientific goals during the advanced detector era (5.1.2), the LIGO-specific priorities for detector characterization (5.1.3), status and plans for data run support (5.1.4), commissioning support (5.1.5), and software infrastructure (5.1.6), as well as the activities and priorities of the different working groups, noise transients (5.1.7), spectral features (5.1.8), calibration (5.1.9) and timing (5.1.10).

5.1.2 Overview Requirements to Achieve aLIGO Science

Advanced LIGO is expected to begin scientific data collection in 2015 and perform observations of increasing sensitivity and duration in the years that follow [114]. The overarching goal of the LIGO detector characterization group over the next few years is to ensure that Advanced LIGO data is 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 get into the specifics of how they will be addressed until Subsection 5.1.3.

1. **Provide accurate calibration, timing, and state information for the gravitational-wave channels.** It is imperative for confident detections and deep searches that the gravitational-wave channels for Advanced LIGO are accurately calibrated into physical units (strain) and have precise timing, and that the state of the detectors is accurately recorded (e.g. to signal what data should be analyzed and the expected quality of that data).

2. **Remove egregious data quality problems in the detectors.** Neutron star binary coalescence signals are the most likely source to be detected by Advanced LIGO. Owing to their minutes-long waveforms, binary neutron star (BNS) searches are expected to be relatively robust against short and quiet noise transients. However, loud transients or dramatic non-stationarity could complicate or blind a potential detection. To maximize detections, BNS searches require that egregious data quality problems are 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.** The initial LIGO detectors exhibited a high rate of non-Gaussian noise transients (glitches). These acted 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. If Advanced LIGO has a similar glitch rate, these searches will require a significant reduction (possibly $> 90\%$ of single-detector noise transients above SNR 8). 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.

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.

5.1.3 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 the requirements listed above will be met by the first aLIGO science run in 2015.

1. **Characterize the Advanced LIGO subsystems as they are brought online.** [*Supports requirements 2,3,5 above.*] We will investigate the data quality of Advanced LIGO subsystems as they are deployed. Each subsystem will have 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.
    - Performing deeper and more specific investigations on the subsystems such as those listed in Section 5.1.5.

This work is aimed at identifying and fixing problems early - which is preferable to waiting for the first locks of the full interferometers in ca. 2014 and then trying to sort the myriad of artefacts present in the detector output back to their individual sources. This “ground up” approach is, we believe, an improvement over the “top down" approach that was typically used in Initial LIGO. It allows each subsystem to be carefully studied and “checked off" as well-behaved and understood during commissioning, which should result in a much cleaner output once the complete detectors are made operational. It will also train a larger number of detector characterization group members who are familiar with the individual subsystems. It should be noted however, that many data quality issues, such as complicated noise couplings, will only become apparent when the full interferometers are operating so we expect characterization in the later phases of commissioning and running (2014 and beyond) to also be very important.

2. **Upgrade the LIGO Physical Environmental Monitoring Systems for Advanced LIGO** [*Primarily supports requirement 2, but supports all goals above.*] The LIGO PEM system must be taken apart for aLIGO installation, affording an opportunity for redistribution of the sensors, system upgrades, and channel renaming based on lessons learned in initial LIGO. In addition, changes associated with Advanced LIGO will require redeployment of sensors to new coupling sites, installation of new sensors in new rooms and on new vacuum chambers, as well as redeployment of sensors made redundant by seismic sensors in the active isolation system. Plans for this PEM upgrade are laid out in the aLIGO PEM Upgrade document [686]. Associated with the upgrade are a number of hardware and software projects for LVC members, detailed in the upgrade document and listed here.

   Hardware

   - power meters for roof radio monitors
   - RF monitors at the main modulation frequencies for inside the LVEA
   - an RF spectrum monitoring system that sweeps from a few kHz to a couple of GHz
   - 1 Hz to 10,000 Hz RF monitor
   - an electrostatic field monitor
   - several coil magnetometers
   - a sky observation system
   - an upgraded cosmic-ray detection system
   - infrasound monitors

   Software

   - updated dead channel monitor
   - channel snapshots
   - statistical channel monitor
   - channel location and calibration web page
   - direction to source finder using propagation delays
   - code to search for “pulsars" in selected auxiliary channels using modified all-sky and/or specific pulsar search code
   - modified stochastic code to search for signal between aux channels
   - significance figure of merit for Carleton DARM-aux coherence line monitor
   - 1Hz (and other) comb monitor

3. **Participate actively in the Advanced LIGO engineering runs.** [*Supports all requirements above.*] In the upcoming years the LIGO Scientific Collaboration plans to have a 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. The duration of these runs and the role played by the real interferometer data is expected to increase steadily from 2012 through 2014. In the detector characterization group we will work toward having key investigations completed and critical software (calibration, timing, state, data quality monitoring, etc.) implemented and tested in these engineering runs. We expect these periods will provide excellent opportunities to observe the longer term stability of the interferometer and its subsystems than is often possible during heavy commissioning.

4. **Develop improved methods to uncover the causes of and veto noise transients.** [*Supports requirement 2,4,5 above.*] During S6 we had some success using burst and CBC search algorithms [684] to parameterize glitches in the detector outputs and a large number of auxiliary channels and then using automated tools such as UPV [687] and hveto [688] 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 should be followed are:

    - Improved glitch parameterization that works well over a broad parameter space in frequency, duration, and SNR, and runs on the detector strain channels and all high-sample-rate auxiliary channels.
    - 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 5.1.7.

5. **Validate and test new tools and data prior to using them in production mode.** The development of new tools is important for achieving the goals of detector characterization. To make the most efficient use of these tools they should be demonstrated to work effectively on test data sets before being used broadly by the collaboration. Two key examples are, i) glitch parameterization tools should be shown to effectively identify and accurately parameterize known or injected glitches and ii) 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 remaining S6 analyses and engineering run analysis.** It is important that the detector characterization group provide support for data quality issues in all remaining Initial LIGO analyses, and support analyses that are performed on the Advanced LIGO engineering 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.

5.1.4 Data Run Support

For most of the engineering runs and all of the science runs, LSC policy has been to staff the control rooms around the clock with one LIGO Lab operator and at least one LSC scientific monitor per observatory. The scientific monitors are responsible for monitoring the quality of the data, carrying out investigations, and making decisions on when to take science data *vs.* when to make adjustments / repairs, in consultation with the operator on duty and the local run coordinator, when appropriate.

The LSC is currently in the process of critically evaluating the performance of the scientific monitoring system and is develop a proposed plan to be tested in upcoming engineering runs and implemented in the Advanced LIGO science runs [689]. Two key goals that are emerging 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.

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 use the experience with initial LIGO, including the results of “blind" injection exercise, to plan for future runs.

Although there will be no science runs in the upcoming year, there will be engineering runs. We will use these periods to test data monitoring systems and interactions with the searches, and to gain an early understanding of the artefacts in each of the subsystems. This will require a close communication between commissioning teams and detector characterization groups, a relationship we would like to foster.

5.1.5 Commissioning Support

Heavy commissioning of Advanced LIGO is expected to begin and 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 needs to work more 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 Commissioners have spelled out a list of high priority projects that should be done by LSC detector characterization folks over the coming year. An abbreviated list of projects is below, and the full list is available on the DetChar wiki pages [690]. These projects are aimed at using the resources in the detector characterization group to solve a problem with LIGO that might not get solved otherwise due to limited human resources in the commissioning group.

- Set up daily summary pages that display the key figures of merit of the instrument and are iterated with commissioners.
- Measure suspension crackle noise by demodulating strain energy at bounce/roll frequencies.
- Create a monitor for mechanical modes of the optics or suspensions that can be excited and saturate some of the electronics, corrupting data.
- Create an early warning system for laser death and understand what the timescales involved are.
- Help identify the reasons for losses of interferometer control (lock). This will help improve stability and give longer locked periods for analysis.
- Track the test mass scattering and absorption to check whether the optics or their cleanliness are degrading over time.
- Create an estimator of the test mass thermal state to improve the accuracy of the thermal compensation system and reduce the down time caused by warming up after lock acquisition.
- Create an early warning system for trains, earthquakes, and traffic to trigger changes in the control state of the interferometer that will be more likely to ride out these disturbances, decreasing down time.
- Investigate the anomalously high coupling of acoustic noise to hardware in the vacuum chamber that houses the output modecleaner.

Subsystem characterization Characterization of the Advanced LIGO subsystems is Priority 1 from Section 5.1.3 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” [691] 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.

5.1.6 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[682], DTT[692] 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[692], is used as a background-process environment for continuous monitoring. The DMT system provided many critical functions in S6, 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 done 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 [688] and UPV [687], 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 science runs by scientific monitors 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 are working to bring 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. The database infrastructure was first used in S5, and was significantly revamped for S6 by the LSC Software Working group (**??**). This resulted in a reliable system where flags are introduced online by well tested DMT monitoring programs, and offline by the DetChar group and by scimons. For aLIGO we plan to move 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 [693], a script-based Matlab interface, mDV [694], and a python interface to the NDS server, pynds. Starting in 2011 a new web-based data viewer, 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.

Despite the critical importance of Detector Characterization, there are only two collaboration members with near full time dedication to detector characterization software (one dedicated to NDS2 and DMT, and another dedicated to ligoDV and ligoDV-web), and only a few other members partially dedicated to developing and maintaining software tools. To thoroughly characterize the more complex Advanced LIGO detectors and enable the first detections, the detector characterization group requires a more robust software infrastructure and human resources to support it.

Software priorities This section describes priorities for detector characterization software work over the next few years. In general we want to build on the successful software infrastructure from S6 and expand it to meet the demands of searches in the Advanced detector era. These activities will be coordinated with the Software Working Group, as described in Section **??**.

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, adopt the common trigger handling format defined in T1300468 [695] for the upcoming engineering runs.
2. Implement a LIGO daily report webpage monitoring system inspired by the GEO summary pages.
3. 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.
4. Continue development of new and improved Channel Information System (CIS) `cis.ligo.org` containing channel names, sample frequencies, editable descriptions, links to appropriate subsystem models, and other information.
5. 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.
6. 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.
7. Produce software to monitor the first subsystems of Advanced LIGO that will form the foundation for data quality flags in the first runs.
8. Continue development of the LIGO segment database to increase input and output speed, robustness and to improved user interface tools.

9. 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.
10. 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.
11. Continue refinement of veto production algorithms and test these improvements on aLIGO subsystem data.
12. Migrate data quality and veto flags that proved useful in S6 and are likely to be useful in Advanced LIGO to on-line production.
13. Maintain appropriate reduced data sets for Advanced LIGO to be used for detector characterization and for data analysis. This includes data from engineering runs.

5.1.7 Noise Transients

The largest detector characterization subgroup, the Glitch Group[696], carries out studies of interferometer noise transients, or “glitches”. Composed of experimentalists and analysts, the working group has broad expertise, and its work is closely coupled to the burst and CBC searches.

The goals of the Glitch Working Group 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.

In the years before the first advanced-era science runs, there will be short engineering runs dedicated to studying the data quality of the different Advanced LIGO subsystems as they are installed. This activity began in 2012 with the pre-stabilized laser, which was characterized during Engineering Run 3. Future runs will involve the input mode cleaner, suspensions, seismic isolation, and the Dual-Recycled Michelson Interferometer test at L1. We expect, with the LIGO Laboratory’s help, to take full advantage of the possibility to exercise data monitoring tools, as well as get an early understanding of the artefacts in each of the building blocks of the very complex gravitational wave detectors.

The priorities for the coming year are:

- Participate in commissioning of the aLIGO subsystems. Identify glitches in subsystem channels that would/will affect the interferometer in future science runs, with special focus on rare glitches.
- Participate in engineering runs, including staffing glitch shifts during times designated for intensive characterization of the instrument. When instrumental data is used as a fake gravitational wave channel, provide data quality information for the astrophysical searches on the fake data.

- 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.
- Run burst-like searches on many auxiliary instrumental and environmental channels in real time. Improve these searches, their tuning, and their interoperability with data quality tools to provide the most useful information to glitch hunters.
- 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. Test this on S6 data as well as the new data that is coming in.

One of the goals above is to carry out glitch shifts during the engineering runs, similar to the way that glitch shifts were carried out in S6. In that case, a wiki page of useful plots was generated automatically. For each detector site and week of data, a person was assigned to review this page and the electronic logs for data quality / glitch issues and to report on them to the glitch call. Going forward, we would like use the daily summary pages to provide the automated plots which are used for the glitch shifts. The glitch shifts are not anticipated to be run regularly at first, but to occur opportunistically when commissioners request particular concentration on a subsystem or when there is a period of especially interesting data, or running that is relatively undisturbed by configuration changes, during the engineering runs.

Two month test

Glitch identification challenge

The first and foremost goal for the advanced detector eta 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 facilite 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 occurence of glitches.

5.1.8 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]. The retrofitting of LLO vacuum chambers with feed-forward, hydraulic pre-actuators led to dramatic improvement in L1 duty cycle, allowing the interferometer to ride out the passage of trains without lock loss. Nonetheless, significant increase in gravitational wave channel noise is seen during such a passage and in general 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. There have been extensive studies during S5 and S6 to understand better the sources of steady-state environmental couplings, particularly lines; these studies are now extending to advanced LIGO subsystems.

The list of high priority activities related to characterizing spectral features in 2012-2013 are::

- Continue to analyze and investigate noise lines affecting data quality in S6 data, including summarizing frequency line issues in S6 for a data quality paper.
- *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 sub-systems. 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.

As various advanced LIGO subsystems come on-line the software for spectral line identifications and interchannel correlations can be applied; this will serve as a means to identify noise in the subsystems, and prepare the routines for application on advanced LIGO h(t) data when it becomes available.

5.1.9 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[697] available to the LSC, and as with previous science runs, will continue to be recorded in the electronic logs, software repositories, and LIGO documents[698].

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 5.1.10.

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. Work on the aLIGO photon calibrator subsystem design, fabrication, installation, and commissioning is ongoing, dealing with subtle performance issues, such as the elastic deformation of the test masses resulting from radiation pressure from the photon calibrator beams and forces from electrostatic actuators.

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 independent techniques are being developed 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 would be 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. Initially the plan is utilize this capability purely to help the commissioners in the control room as opposed to providing final calibration. However, as it evolves and if it proves to be accurate and robust enough, the final calibration could be moved over from the more traditional calibration methods to utilizing this scheme.

The calibration team is developing a low-latency (sub-second latency) gstreamer-based pipeline for time domain calibration in aLIGO. This will be 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 is under development and was successfully tested during aLIGO's third engineering run. The offline infrastructure for the pipeline is also currently under development and was successfully tested during offline data reproduction after the third engineering run. The calibration team is continuing to develop both the infrastructure required for and the inner-workings of the low-latency aLIGO time domain calibration pipeline, and it will be further tested during the upcoming fourth engineering run and future engineering runs.

5.1.10 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 5.1.9), the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator[699], 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] J. Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Physical Review D*, 88:122004, 2013.

[2] J. Aasi et al. First Searches for Optical Counterparts to Gravitational-wave Candidate Events. *Astrophys. J. Suppl.*, 211:7, 2014.

[3] J. Aasi et al. Constraints on cosmic strings from the LIGO-Virgo gravitational-wave detectors. *Physical Review Letters*, 112:131101, 2014.

[4] J. Aasi 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*, 89:122003, 2014.

[5] 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.

[6] 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.

[7] J. Aasi et al. The NINJA-2 project: Detecting and characterizing gravitational waveforms modelled using numerical binary black hole simulations. *Classical & Quantum Gravity*, 31:115004, 2014.

[8] J. Aasi et al. Search for gravitational wave ringdowns from perturbed intermediate mass black holes in LIGO-Virgo data from 2005-2010. *Physical Review D*, 89:102006, 2014.

[9] 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.

[10] J. Aasi et al. Gravitational waves from known pulsars: results from the initial detector era. *Astrophys.J.*, 785:119, 2014.

[11] J. Aasi et al. A directed search for continuous Gravitational Waves from the Galactic Center. *Phys.Rev.*, D88:102002, 2013.

[12] 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.

[13] J. Aasi et al. Implementation of an F-statistic all-sky search for continuous gravitational waves in Virgo VSR1 data. *Submitted to Class. Quant. Grav.*, 2014.

[14] 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.

[15] 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.

[16] 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.

[17] S. Klimenko et al. A coherent method for detection of gravitational wave bursts. *Class. Quantum Grav.*, 25(11):114029–+, June 2008.

[18] 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.

[19] 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.

[20] E. Thrane and M. Coughlin. Seedless clustering in all-sky searches for gravitational-wave transients. *ArXiv e-prints*, January 2014. arXiv:1401.8060.

[21] S. Chamberlin et al. Applications of the ExcessPower Data Analysis Pipeline to Gravitational Wave Detection Efforts, 2013.

[22] 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.

[23] Tyson B. Littenberg and Neil J. Cornish. Separating gravitational wave signals from instrument artifacts. *Phys. Rev. D*, 82:103007, Nov 2010.

[24] 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.

[25] 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.

[26] Mark Edwards and Patrick J. Sutton. A new glitch-rejection algorithm forged in the spherical harmonic basis. *J.Phys.Conf.Ser.*, 363:012025, 2012.

[27] 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.

[28] 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.

[29] Bulent Kiziltan, Athanasios Kottas, and Stephen E. Thorsett. The Neutron Star Mass Distribution, 2010.

[30] 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.

[31] 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.

[32] S.E. Woosley, A. Heger, and T.A. Weaver. The evolution and explosion of massive stars. *Rev.Mod.Phys.*, 74:1015–1071, 2002.

[33] M. Kramer and N. Wex. The double pulsar system: a unique laboratory for gravity. *Classical and Quantum Gravity*, 26:073001, 2009.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] 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.

[39] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 17:2, 2014.

[40] 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.

[41] S. Fairhurst. Triangulation of gravitational wave sources with a network of detectors. *New Journal of Physics*, 11:123006, December 2009.

[42] 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.

[43] F. Beauville et al. Detailed comparison of LIGO and Virgo inspiral pipelines in preparation for a joint search. *Class. Quantum Grav.*, 25:045001, 2008.

[44] 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.

[45] Kipp Cannon et al. Toward early-warning detection of gravitational waves from compact binary coalescence. *The Astrophysical Journal*, 748(2):136, 2012.

[46] 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.

[47] 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.

[48] 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.

[49] 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.

[50] J. M. Lattimer. The Nuclear Equation of State and Neutron Star Masses. *Annual Review of Nuclear and Particle Science*, 62:485–515, November 2012.

[51] 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.

[52] 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.

[53] M. Favata. Systematic Parameter Errors in Inspiraling Neutron Star Binaries. *Physical Review Letters*, 112(10):101101, March 2014.

[54] 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.

[55] 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.

[56] 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.

[57] 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.

[58] 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.

[59] L. Blanchet and B. S. Sathyaprakash. Signal analysis of gravitational wave tails. *Classical and Quantum Gravity*, 11:2807–2831, November 1994.

[60] L. Blanchet and B. S. Sathyaprakash. Detecting a Tail Effect in Gravitational-Wave Experiments. *Physical Review Letters*, 74:1067–1070, February 1995.

[61] 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.

[62] 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.

[63] 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.

[64] Bruce Allen. A χ^2 time-frequency discriminator for gravitational wave detection. *Physical Review D*, 71:062001, 2005.

[65] Andres Rodríguez. Reducing false alarms in searches for gravitational waves from coalescing binary systems. Master's thesis, Louisiana State University, 2007.

[66] 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.

[67] 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.

[68] B. Abbott et al. Analysis of LIGO data for gravitational waves from binary neutron stars. *Physical Review D*, 69:122001, 2004.

[69] B. Abbott et al. Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Physical Review D*, 72:082001, 2005.

[70] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D*, 77:062002, 2008.

[71] 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.

[72] 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.

[73] J. Abadie et al. Search for compact binary coalescences in LIGO and Virgo data from S5 and VSR1. *Physical Review D*, 82:102001, 2010.

[74] 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.

[75] 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.

[76] 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.

[77] F. Raab. Working Timeline to First Observing Run of the Advanced Detector Era. Technical Report LIGO-G1301309, LIGO Scientific Collaboration and Virgo Collaboration, 2014.

[78] 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.

[79] Feryal Ozel, Dimitrios Psaltis, Ramesh Narayan, and Jeffrey E. McClintock. The Black Hole Mass Distribution in the Galaxy. *Astrophys.J.*, 725:1918–1927, 2010.

[80] 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.

[81] A. Merloni. Cosmological evolution of supermassive black holes and AGN: a synthesis model for accretion and feedback. Mem. Soc. Astron. Italiana, 79:1310, 2008.

[82] 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.

[83] K. Belczynski, G. Wiktorowicz, C. Fryer, D. Holz, and V. Kalogera. Missing Black Holes Unveil The Supernova Explosion Mechanism. *Astrophys.J.*, 757:91, 2012.

[84] 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.

[85] 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.

[86] 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.

[87] 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.

[88] 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.

[89] 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.

[90] Krzysztof Belczynski et al. On The Maximum Mass of Stellar Black Holes. *Astrophysical Journal*, 714:1217–1226, 2010.

[91] 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.

[92] Michal Dominik and Krzysztof Belczynski. Private communication, 2012.

[93] Krzysztof Belczynski, Alessandra Buonanno, Matteo Cantiello, Chris L. Fryer, Daniel E. Holz, et al. The Formation and Gravitational-Wave Detection of Massive Stellar Black-Hole Binaries. *Astrophys.J.*, 789:120, 2014.

[94] 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.

[95] 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.

[96] Jeffrey E. McClintock et al. The Spin of the Near-Extreme Kerr Black Hole GRS 1915+105. *Astrophysical Journal*, 652:518–539, 2006.

[97] 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.

[98] 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.

[99] 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.

[100] 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.

[101] 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.

[102] 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.

[103] 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.

[104] 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.

[105] Krzysztof Belczynski, Matthew Benacquista, and Tomasz Bulik. Double Compact Objects as Low-frequency Gravitational Wave Sources. *Astrophys.J.*, 725:816–823, 2010.

[106] V. Kalogera. Spin-orbit misalignment in close binaries with two compact objects. *Astrophysical Journal*, 541:042003, 2000.

[107] 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.

[108] T. Bulik and K. Belczyński. Constraints on the Binary Evolution from Chirp Mass Measurements. *Astrophysical Journal Letters*, 589:L37–L40, May 2003.

[109] Ilya Mandel. Parameter estimation on gravitational waves from multiple coalescing binaries. *Physical Review D*, D81:084029, 2010.

[110] R. OÕShaughnessy. Data-driven methods to explore a large space of computationally costly compact binary progenitor models. *Phys.Rev.*, D88(8):084061, 2013.

[111] J. Abadie et al. Search for gravitational waves from binary black hole inspiral, merger and ringdown. *Physical Review D*, 83:122005, 2011.

[112] 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.

[113] J. Abadie et al. All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run. *Phys.Rev.*, D81:102001, 2010.

[114] 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`.

[115] 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.

[116] 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.

[117] Michael P§rrer. Frequency domain reduced order models for gravitational waves from aligned-spin compact binaries. *Class.Quant.Grav.*, 31(19):195010, 2014.

[118] Thomas Cokelaer. Gravitational wave from inspiralling compact binaries: hexagonal template placement and its efficiency in detecting physical signals. *Physical Review D*, 76:102004, 2007.

[119] 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.

[120] 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.

[121] 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.

[122] 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.

[123] 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.

[124] 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.

[125] 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.

[126] 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.

[127] J. Aasi et al. Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network. *Phys.Rev.*, D88:062001, 2013.

[128] 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.

[129] 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.

[130] Harbir Antil, Scott E. Field, Frank Herrmann, Ricardo H. Nochetto, and Manuel Tiglio. Two-step greedy algorithm for reduced order quadratures, 2012.

[131] 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.

[132] 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.

[133] Rory Smith, Chad Hanna, Ilya Mandel, and Alberto Vecchio. Rapidly evaluating the compact binary likelihood function via interpolation. *Phys.Rev.*, D90:044074, 2014.

[134] 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.

[135] 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.

[136] 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.

[137] Katerina Chatziioannou, Neil Cornish, Antoine Klein, and Nicolas Yunes. Spin-Precession: Breaking the Black Hole–Neutron Star Degeneracy, 2014.

[138] 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.

[139] Vassiliki Kalogera. Spin orbit misalignment in close binaries with two compact objects. *Astrophys.J.*, 541:319–328, 2000.

[140] 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.

[141] 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.

[142] Z. Arzoumanian, D.F. Chernoffs, and J.M. Cordes. The Velocity distribution of isolated radio pulsars. *Astrophys.J.*, 568:289–301, 2002.

[143] Brad M.S. Hansen and E. Sterl Phinney. The Pulsar kick velocity distribution. *Mon.Not.Roy.Astron.Soc.*, 291:569, 1997.

[144] 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.

[145] 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.

[146] 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.

[147] Matthew J. Benacquista and Jonathan M.B. Downing. Relativistic Binaries in Globular Clusters. *Living Rev.Rel.*, 16:4, 2013.

[148] 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.

[149] 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.

[150] 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.

[151] 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.

[152] 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.

[153] 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.

[154] 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.

[155] 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.

[156] Francois Foucart. Black Hole-Neutron Star Mergers: Disk Mass Predictions. *Phys.Rev.*, D86:124007, 2012.

[157] 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.

[158] 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.

[159] 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.

[160] E. Berger. The Host Galaxies of Short-Duration Gamma-Ray Bursts: Luminosities, Metallicities, and Star Formation Rates. *Astrophys.J.*, 690:231–237, 2009.

[161] 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.

[162] 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.

[163] 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.

[164] 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.

[165] 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.

[166] 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.

[167] 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.

[168] 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.

[169] Andrea Merloni. Cosmological evolution of supermassive black holes and AGN: a synthesis model for accretion and feedback. *Mem. Soc. Astron. Italiana*, 79:1310, 2008.

[170] 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.

[171] 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.

[172] M. Kramer and N. Wex. The double pulsar system: A unique laboratory for gravity. *Class.Quant.Grav.*, 26:073001, 2009.

[173] 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.

[174] Bruce Allen. A chi**2 time-frequency discriminator for gravitational wave detection. *Phys.Rev.*, D71:062001, 2005.

[175] 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.

[176] 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.

[177] 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).

[178] 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.

[179] 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.

[180] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D*, 77:062002, 2008.

[181] 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.

[182] 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.

[183] 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.

[184] 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.

[185] 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.

[186] 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.

[187] Diego Fazi. *Development of a Physical-Template Search for Gravitational Waves from Spinning Compact-Object Binaries with LIGO*. PhD thesis, Bologna University, 2009.

[188] 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.

[189] 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.

[190] J. Abadie et al. First low-latency LIGO+Virgo search for binary inspirals and their electromagnetic counterparts. *Astron. Astrophys.*, 541:A155, May 2012.

[191] Marc van der Sluys et al. Parameter estimation of spinning binary inspirals using Markov-chain Monte Carlo. *Class. Quant. Grav.*, 25:184011, 2008.

[192] M. V. van der Sluys et al. Gravitational-Wave Astronomy with Inspiral Signals of Spinning Compact-Object Binaries, 2007.

[193] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys.Rev.*, D81:062003, 2010.

[194] 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.

[195] 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.

[196] 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.

[197] Detector characterization group. Plan for Providing Data Quality Information for Advanced LIGO Searches. *LIGO Document T1400002*, 2014.

[198] L. Blackburn et al. Equinox Event. Technical Report LIGO-G080559, LIGO Scientific Collaboration and Virgo Collaboration, 2008.

[199] 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).

[200] 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.

[201] 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.

[202] 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.

[203] 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.

[204] 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.

[205] E. Berger, W. Fong, and R. Chornock. An r-process Kilonova Associated with the Short-hard GRB 130603B. *Astrophys.J.*, 774:L23, 2013.

[206] S. Desai, E.O. Kahya, and R.P. Woodard. Reduced time delay for gravitational waves with dark matter emulators. *Phys.Rev.*, D77:124041, 2008.

[207] B. F. Schutz. Determining the Hubble constant from gravitational wave observations. *Nature*, 323:310, September 1986.

[208] 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.

[209] 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.

[210] IceCube Collaboration. `http://gcn.gsfc.nasa.gov/gcn3/14520.gcn3`.

[211] J. Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Phys.Rev.*, D88:122004, 2013.

[212] E. Thrane and M. Coughlin. Searching for gravitational-wave transients with a qualitative signal model: seedless clustering strategies. Preprint arXiv:1308.5292, 2013.

[213] M. C. Miller and E. J. M. Colbert. Intermediate-Mass Black Holes. *International Journal of Modern Physics D*, 13:1–64, January 2004.

[214] 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.

[215] 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.

[216] 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.

[217] K. Belczynski, A. Buonanno, M. Cantiello, D. E. Holz, C. L. Fryer, I. Mandel, M. C. Miller, and M. Walczak. The Formation and Gravitational-Wave Detection of Massive Stellar Black-Hole Binaries. *ArXiv e-prints*, March 2014. arXiv:1403.0677.

[218] 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, and et al. Search for gravitational radiation from intermediate mass black hole binaries in data from the second LIGO-Virgo joint science run. *ArXiv e-prints*, April 2014. arXiv:1404.2199.

[219] 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.

[220] 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.

[221] 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.

[222] Sergey Klimenko and Chris Pankow. Detection statistic for multiple algorithms, networks and data sets. Technical Report LIGO-T1300869, LIGO Project, 2013.

[223] L. Bildsten. Gravitational radiation and rotation of accreting neutron stars. *Astrophys.J.*, 501:L89, 1998.

[224] 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.

[225] B. J. Owen. Maximum elastic deformations of compact stars with exotic equations of state. *Phys.Rev.Lett.*, 95:211101, 2005.

[226] C. Cutler. Gravitational waves from neutron stars with large toroidal B fields. *Phys.Rev.*, D66:084025, 2002.

[227] 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.

[228] 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.

[229] 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.

[230] D.I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon.Not.Roy.Astron.Soc.*, 331:203, 2002.

[231] 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.

[232] R. Narayan. The birthrate and initial spin period of single radio pulsars. *Astrophysical Journal*, 319:162–179, August 1987.

[233] D. R. Lorimer. Binary and millisecond pulsars. *Living Reviews in Relativity*, 8(7), 2005.

[234] 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.

[235] 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/`.

[236] B. Knispel and B. Allen. Blandford's Argument: The Strongest Continuous Gravitational Wave Signal. *Phys.Rev.*, D78:044031, 2008.

[237] B. Abbott et al. All-sky search for periodic gravitational waves in LIGO S4 data. *Phys.Rev.*, D77:022001, 2008.

[238] V. Dergachev. Description of PowerFlux Algorithms and Implementation, 2005. LIGO-T050186, available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[239] P. R. Brady and T. Creighton. Searching for periodic sources with LIGO. 2. Hierarchical searches. *Phys.Rev.*, D61:082001, 2000.

[240] V. Dergachev. On blind searches for noise dominated signals: a loosely coherent approach. *Class.Quant.Grav.*, 27:205017, 2010.

[241] V. Dergachev. Loosely coherent searches for sets of well-modeled signals. *Phys.Rev.*, D85:062003, 2012.

[242] B. Abbott et al. The Einstein@Home search for periodic gravitational waves in LIGO S4 data. *Phys.Rev.*, D79:022001, 2009.

[243] 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.

[244] Llucia Sancho de la Jordana and Alicia M. Sintes. A chi**2 veto for continuous wave searches. *Class.Quant.Grav.*, 25:184014, 2008.

[245] 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.

[246] 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.

[247] C. Palomba, P. Astone, and S. Frasca. Adaptive Hough transform for the search of periodic sources. *Class.Quant.Grav.*, 22:S1255–S1264, 2005.

[248] 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.

[249] 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.

[250] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan. Cross-correlation search for periodic gravitational waves. *Phys.Rev.*, D77:082001, 2008.

[251] J. Abadie et al. All-sky Search for Periodic Gravitational Waves in the Full S5 LIGO Data. *Phys.Rev.*, D85:022001, 2012.

[252] M. Shaltev and R. Prix. Fully coherent follow-up of continuous gravitational-wave candidates. *Phys.Rev.*, D87(8):084057, 2013.

[253] J. Aasi et al. Einstein@Home all-sky search for periodic gravitational waves in LIGO S5 data. *Phys.Rev.*, D87(4):042001, 2013.

[254] J. Aasi et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. Preprint arXiv:1304.0670, 2013.

[255] B. Abbott et al. All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data. *Phys.Rev.Lett.*, 102:111102, 2009.

[256] V. Dergachev. Novel universal statistic for computing upper limits in an ill-behaved background. *Phys.Rev.*, D87(6):062001, 2013.

[257] B.P. Abbott et al. Einstein and Home search for periodic gravitational waves in early S5 LIGO data. *Phys.Rev.*, D80:042003, 2009.

[258] 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.

[259] J. Aasi et al. Gravitational waves from known pulsars: results from the initial detector era. *Astrophys.J.*, 785:119, 2014.

[260] B.J. Owen. Detectability of periodic gravitational waves by initial interferometers. *Class.Quant.Grav.*, 23:S1–S8, 2006.

[261] 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.

[262] B. Abbott et al. Limits on gravitational wave emission from selected pulsars using LIGO data. *Phys.Rev.Lett.*, 94:181103, 2005.

[263] B. Abbott et al. Upper limits on gravitational wave emission from 78 radio pulsars. *Phys.Rev.*, D76:042001, 2007.

[264] B. Abbott et al. Beating the spin-down limit on gravitational wave emission from the Crab pulsar. *Astrophys.J.*, 683:L45–L50, 2008.

[265] B.P. Abbott et al. Searches for gravitational waves from known pulsars with S5 LIGO data. *Astrophys.J.*, 713:671–685, 2010.

[266] J. Abadie et al. Beating the spin-down limit on gravitational wave emission from the Vela pulsar. *Astrophys. J.*, 737:93, 2011.

[267] Dupuis, R.J. *Bayesian Searches for Gravitational Waves from Pulsars*. PhD thesis, University of Glasgow, 2004. LIGO-P050014-00-R.

[268] M. Pitkin and G. Woan. Binary system delays and timing noise in searches for gravitational waves from known pulsars. *Phys.Rev.*, D76:042006, 2007.

[269] 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.

[270] 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.

[271] 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.

[272] Gill, C. D. *Searching for gravitational waves from pulsars*. PhD thesis, University of Glasgow, 2012.

[273] 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.

[274] P. Astone, A. Colla, S. D'Antonio, S. Frasca, and C. Palomba, 2012. poster at GWPAW12.

[275] J. Aasi et al. Improved Upper Limits on the Stochastic Gravitational-Wave Background from 2009-2010 LIGO and Virgo Data, 2014.

[276] E. W. Kolb & M. S. Turner. *The Early Universe*. Westview Press, 1994.

[277] A. A. Starobinskii. Spectrum of relic gravitational radiation and the early state of the universe. *JETP Lett.*, 30, 1979.

[278] R. Bar-Kana. Limits on Direct Detection of Gravitational Waves. *Phys. Rev. D*, 50, 1994.

[279] T. W. B. Kibble. Topology of cosmic domains and strings. *J. Phys.*, A9, 1976.

[280] T. Damour & A. Vilenkin. Gravitational radiation from cosmic (super)strings: bursts, stochastic background, and observational windows. *Phys. Rev. D*, 71, 2005.

[281] S. Olmez, V. Mandic, and X. Siemens. Gravitational-Wave Stochastic Background from Kinks and Cusps on Cosmic Strings. *Phys. Rev. D*, 81:104028, 2010.

[282] S. Olmez, V. Mandic, and X. Siemens. Anisotropies in the Gravitational-Wave Stochastic Background. *J. Cosmol. Astropart. Phys.*, 2012:009, 2011.

[283] A. Buonanno. Spectrum of relic gravitational waves in string cosmology. *Phys. Rev. D*, 55, 1997.

[284] V. Mandic & A. Buonanno. Accessibility of the Pre-Big-Bang Models to LIGO. *Phys. Rev. D*, 73, 2006.

[285] T. Regimbau. The astrophysical gravitational wave stochastic background. *Research in Astronomy and Astrophysics*, 11:369–390, April 2011.

[286] 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.

[287] P. Sandick, K. A. Olive, F. Daigne, and E. Vangioni. Gravitational waves from the first stars. *Physical Review D*, 73(10):104024, May 2006.

[288] 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.

[289] 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.

[290] 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.

[291] P. A. Rosado. Gravitational wave background from rotating neutron stars. *Physical Review D*, 86(10):104007, November 2012.

[292] T. Regimbau and J. A. de Freitas Pacheco. Gravitational wave background from magnetars. *Astronomy and Astrophysics*, 447:1–7, February 2006.

[293] 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.

[294] 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.

[295] 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.

[296] G Sigl. Cosmological gravitational wave background from phase transitions in neutron stars. *J. Cosmol. Astropart. Phys.*, JCAP04:002, 2006.

[297] 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.

[298] 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.

[299] 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.

[300] 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.

[301] A J Farmer & E S Phinney. The gravitational wave background from cosmological compact binaries. *Mon. Not. R. Ast. Soc.*, 346:1197, 2003.

[302] 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.

[303] P. A. Rosado. Gravitational wave background from binary systems. *Physical Review D*, 84(8):084004, October 2011.

[304] 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.

[305] 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.

[306] 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.

[307] L Barack & C Cutler. Confusion noise from LISA capture sources. *Phys. Rev. D*, 70:122002, 2004.

[308] G Sigl & J Schnittman & A Buonanno. Gravitational-wave background from compact objects embedded in AGN accretion disks. *Phys. Rev. D*, 75:024034, 2007.

[309] P.A.R. Ade et al. BICEP2 I: Detection Of B-mode Polarization at Degree Angular Scales, 2014. `http://arxiv.org/abs/1403.3985`.

[310] M Maggiore. Gravitational Wave Experiments and Early Universe Cosmology. *Phys. Rept.*, 331:283, 2000.

[311] 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.

[312] 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.

[313] J L Cook and L Sorbo. Particle production during inflation and gravitational waves detectable by ground-based interferometers. *Phys. Rev. D*, 85:023534, 2012.

[314] B. Abbott et al. Searching for Stochastic Gravitational Waves with LIGO. *Nature*, 460:990, 2009.

[315] 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.

[316] 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.

[317] C. Brans and R.H. Dicke. Mach's principle and a relativistic theory of gravitation. *Phys.Rev.*, 124:925–935, 1961.

[318] Y. Fujii and K. Maeda. *The Scalar-Tensor Theory of Gravitation*. Cambridge Monograph on Mathematical Physics. Cambridge University Press, Cambridge, 2002.

[319] Thomas P. Sotiriou and Valerio Faraoni. f(R) Theories Of Gravity. *Rev.Mod.Phys.*, 82:451–497, 2010.

[320] Antonio De Felice and Shinji Tsujikawa. f(R) Theories. *Living Rev. Relativity*, 13(3), 2010.

[321] Matt Visser. Mass for the graviton. *General Relativity and Gravitation*, 30(12):1717–1728, 1998.

[322] Claudia de Rham, Gregory Gabadadze, and Andrew J. Tolley. Resummation of massive gravity. *Phys. Rev. Lett.*, 106:231101, Jun 2011.

[323] B Abbott et al. Searching for a Stochastic Background of Gravitational Waves with the Laser Interferometer Gravitational-Wave Observatory. *Astrophys. J.*, 659:918, 2007.

[324] 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.

[325] J Aasi et al. Improved upper limits on the stochastic gravitational-wave background from the LIGO S6 and Virgo VSR2-3 science runs. *In Preparation*, 2014. `https://dcc.ligo.org/LIGO-P1300154`.

[326] The LIGO Scientific Collaboration and Virgo Collaboration. Searching for stochastic gravitational waves using data from the two co-located LIGO Hanford interferometers. *In Preparation*, 2014. `https://dcc.ligo.org/LIGO-P1000112`.

[327] Bruce Allen and Joseph D. Romano. Detecting a stochastic background of gravitational radiation: Signal processing strategies and sensitivities. *Phys.Rev.*, D59:102001, 1999.

[328] N Christensen. Measuring the Stochastic Gravitational Radiation Background with Laser Interferometric Antennas. *Phys. Rev. D*, 46:5250, 1992.

[329] B Abbott et al. Analysis of first LIGO science data for stochastic gravitational waves. *Phys. Rev. D*, 69:122004, 2004.

[330] 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.

[331] E. Thrane for the The Stochastic Group. Calibration requirements for stochastic searches in the advanced detector era, 2014.
`https://dcc.ligo.org/LIGO-T1200562`.

[332] 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.

[333] 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`.

[334] T Regimbau & B Chauvineaux. A stochastic background from extra-galactic double neutron stars. *Class. Quantum Grav.*, 24:627, 2007.

[335] E Howell et al. The gravitational wave background from neutron star birth throughout the cosmos. *Mon. Not. R. Ast. Soc.*, 351:1237, 2004.

[336] 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.

[337] 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.

[338] T Regimbau & J A de Freitas Pacheco. Cosmic background of gravitational waves from rotating neutron stars. *Astron. Astrophys.*, 376:381, 2001.

[339] T Regimbau & J A de Freitas Pacheco. Gravitational wave background from magnetars. *Astron. Astrophys.*, 447:1, 2006.

[340] J. Aasi et al. Constraints on cosmic strings from the LIGO-Virgo gravitational-wave detectors. *Phys.Rev.Lett.*, 112:131101, 2014.

[341] 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.

[342] B Abbott et al. Upper limit map of a background of gravitational waves. *Phys. Rev. D*, 76:082003, 2007.

[343] B Abbott et al. Directional limits on persistent gravitational waves using LIGO S5 science data. *Phys. Rev. Lett.*, 107:271102, 2011.

[344] 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.

[345] K. Wette et al. Searching for gravitational waves from Cassiopeia A with LIGO. *Class.Quant.Grav.*, 25:235011, 2008.

[346] J. Abadie et al. First search for gravitational waves from the youngest known neutron star. *Astrophys.J.*, 722:1504–1513, 2010.

[347] 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.

[348] 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.

[349] R.V. Wagoner. Gravitational radiation from accreting neutron stars. *Astrophys. J.*, 278:345, 1984.

[350] J. Papaloizou and J.E. Pringle. Gravitational radiation and the stability of rotating stars. *Mon. Not. R. Ast. Soc.*, 184:501, 1978.

[351] Nils Andersson. A New class of unstable modes of rotating relativistic stars. *Astrophys.J.*, 502:708–713, 1998.

[352] J.L. Friedman and S.M. Morsink. Axial instability of rotating relativistic stars. *Astrophys. J.*, 502:714, 1998.

[353] Y. Levin. Runaway Heating by r-Modes of Neutron Stars in Low-Mass X-Ray Binaries. *Astrophys. J.*, 517:328, 1999.

[354] 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.

[355] D. Chakrabarty et al. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. *Nature*, 424:42–44, 2003.

[356] 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.

[357] P. Arras and et al. Saturation of the r-Mode Instability. *Astrophys. J.*, 591:1129, 2003.

[358] 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.

[359] 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.

[360] 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.

[361] A. Arvanitaki, S. Dimopoulos, S. Dubovsky, N. Kaloper, and J. March-Russell. String axiverse. *Physical Review D*, 81(12):123530, June 2010.

[362] A. Arvanitaki and S. Dubovsky. Exploring the string axiverse with precision black hole physics. *Physical Review D*, 83(4):044026, February 2011.

[363] Hirotaka Yoshino and Hideo Kodama. Probing the string axiverse by gravitational waves from Cygnus X-1, 2014.

[364] H.J. Bulten et al. Gravitational waves from Sco X-1 : A comparison of search methods. in preparation.

[365] J. Aasi et al. A directed search for gravitational waves from scorpius x-1 with initial ligo. in preparation.

[366] 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.

[367] 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.

[368] 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.

[369] 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.

[370] J.T. Whelan, S. Sundaresan, Y. Zhang, and P. Peiris. A Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1. in preparation, LIGO-P1200142.

[371] 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.

[372] T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J.Phys.A*, A9:1387–1398, 1976.

[373] A. Vilenkin and E. Shellard. *Cosmic strings and other Topological Defects*. Cambridge University Press, 1994.

[374] I. B. Zeldovich. Cosmological fluctuations produced near a singularity. *Mon. Not. R. Astron. Soc.*, 192:663–667, sep 1980.

[375] Mark Wyman, Levon Pogosian, and Ira Wasserman. Bounds on cosmic strings from WMAP and SDSS. *Phys.Rev.*, D72:023513, 2005.

[376] Neil Bevis, Mark Hindmarsh, Martin Kunz, and Jon Urrestilla. Fitting CMB data with cosmic strings and inflation. *Phys.Rev.Lett.*, 100:021301, 2008.

[377] Richard Battye, Bjorn Garbrecht, and Adam Moss. Tight constraints on F- and D-term hybrid inflation scenarios. *Phys.Rev.*, D81:123512, 2010.

[378] Richard Battye and Adam Moss. Updated constraints on the cosmic string tension. *Phys.Rev.*, D82:023521, 2010.

[379] P.A.R. Ade et al. Planck 2013 results. XXV. Searches for cosmic strings and other topological defects, 2013.

[380] Andrei D. Linde. Hybrid inflation. *Phys.Rev.*, D49:748–754, 1994.

[381] 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.

[382] 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.

[383] Saswat Sarangi and S.H. Henry Tye. Cosmic string production towards the end of brane inflation. *Phys.Lett.*, B536:185–192, 2002.

[384] Rachel Jeannerot, Jonathan Rocher, and Mairi Sakellariadou. How generic is cosmic string formation in SUSY GUTs. *Phys.Rev.*, D68:103514, 2003.

[385] Edward Witten. Cosmic superstrings. *Physics Letters B*, 153(4???5):243 – 246, 1985.

[386] Edmund J. Copeland, Levon Pogosian, and Tanmay Vachaspati. Seeking String Theory in the Cosmos. *Class.Quant.Grav.*, 28:204009, 2011.

[387] N. Kaiser and A. Stebbins. Microwave anisotropy due to cosmic strings. *Nature*, 310:391–393, 1984.

[388] Amy S. Lo and Edward L. Wright. Signatures of cosmic strings in the cosmic microwave background, 2005.

[389] 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.

[390] 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.

[391] 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.

[392] Thibault Damour and Alexander Vilenkin. Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows. *Phys.Rev.*, D71:063510, 2005.

[393] S. Olmez, V. Mandic, and X. Siemens. Gravitational-Wave Stochastic Background from Kinks and Cusps on Cosmic Strings. *Phys.Rev.*, D81:104028, 2010.

[394] Mark G. Jackson, Nicholas T. Jones, and Joseph Polchinski. Collisions of cosmic F and D-strings. *JHEP*, 0510:013, 2005.

[395] T. Kibble. *The formation and evolution of cosmic strings*, chapter Cosmic strings – An overview. Cambridge Univ. Press, gibbons, g. and hawking, s edition, 1990.

[396] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys.Rev.Lett.*, 85:3761–3764, 2000.

[397] David P. Bennett and Francois R. Bouchet. Evidence for a Scaling Solution in Cosmic String Evolution. *Phys.Rev.Lett.*, 60:257, 1988.

[398] Christophe Ringeval, Mairi Sakellariadou, and Francois Bouchet. Cosmological evolution of cosmic string loops. *JCAP*, 0702:023, 2007.

[399] Jose J. Blanco-Pillado, Ken D. Olum, and Benjamin Shlaer. The number of cosmic string loops. *Phys.Rev.*, D89:023512, 2014.

[400] V. Berezinsky, B. Hnatyk, and A. Vilenkin. Superconducting cosmic strings as gamma-ray burst engines, 2000.

[401] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cusps and kinks on cosmic strings. *Phys.Rev.*, D64:064008, 2001.

[402] J. Aasi et al. The characterization of Virgo data and its impact on gravitational-wave searches. *Class.Quant.Grav.*, 29:155002, 2012.

[403] T. Accadia et al. Virgo Document VIR-0128A-12, 2012.

[404] 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.

[405] 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.

[406] Kipp C. Cannon. A Bayesian coincidence test for noise rejection in a gravitational-wave burst search. *Class.Quant.Grav.*, 25:105024, 2008.

[407] Tristan L. Smith, Elena Pierpaoli, and Marc Kamionkowski. A new cosmic microwave background constraint to primordial gravitational waves. *Phys.Rev.Lett.*, 97:021301, 2006.

[408] B.P. Abbott et al. An Upper Limit on the Stochastic Gravitational-Wave Background of Cosmological Origin. *Nature*, 460:990, 2009.

[409] 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.

[410] 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`.

[411] A. Vicere and M. Yvert. First tests of a moment autocorrelation method to detect gw from pulsars in binary systems, October 2013. Virgo technical note VIR-0292C-13 (October 2013).

[412] 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.

[413] 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.

[414] 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.

[415] Bence Kocsis and Janna Levin. Repeated bursts from relativistic scattering of compact objects in galactic nuclei. *Phys. Rev. D*, 85:123005, Jun 2012.

[416] 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.

[417] 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.

[418] 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.

[419] 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.

[420] M. Coleman Miller and Douglas P. Hamilton. Four-body effects in globular cluster black hole coalescence. *Astrophys. J.*, 576(2):894, 2002.

[421] S. J. Aarseth. Mergers and ejections of black holes in globular clusters. *Mon. Not. Roy. Astron. Soc.*, 422:841–848, May 2012.

[422] 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.

[423] 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.

[424] 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.

[425] 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.

[426] P. C. Peters. Gravitational radiation and the motion of two point masses. *Phys. Rev.*, 136:B1224–B1232, 1964.

[427] Neil J. Cornish and Tyson B. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, 2014.

[428] Kai Sheng Tai, Sean T. McWilliams, and Frans Pretorius. Detecting gravitational waves from highly eccentric compact binaries. *Physical Review D*, 90:103001, 2014.

[429] M. Millhouse, N. Cornish, and T. Littenberg, 2014. In preparation.

[430] 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].

[431] M. Coughlin, E. Thrane, and N. Christensen. Detecting compact binary coalescences with seedless clustering. *Phys. Rev. D*, 90:083005, 2014.

[432] 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.

[433] 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.

[434] 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.

[435] 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.

[436] 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.

[437] 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.

[438] I. Mandel and J. R. Gair. Can we detect intermediate mass ratio inspirals? *Classical and Quantum Gravity*, 26(9):094036, May 2009.

[439] 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.

[440] 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.

[441] T. Callister. Improving models of eccentric compact binary inspirals, MPhil Thesis (Cambridge University), LIGO-DCC-P1400238, 2014.

[442] 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.

[443] 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.

[444] T. Pradier. Coincidences between Gravitational Wave Interferometers and High Energy Neutrino Telescopes. *ArXiv e-prints*, 807, July 2008. arXiv:0807.2562.

[445] 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.

[446] 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.

[447] 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.

[448] Shin'ichiro Ando et al. Multimessenger astronomy with gravitational waves and high-energy neutrinos. *Rev. Mod. Phys.*, 85(4):1401–1420, 2013.

[449] E. Waxman and J. Bahcall. High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physical Review Letters*, 78:2292–2295, March 1997.

[450] E. Waxman and J. N. Bahcall. Neutrino Afterglow from Gamma-Ray Bursts: 10^{18} eV. *Astrophys. Journal*, 541:707, October 2000.

[451] P. Mészáros and E. Waxman. TeV Neutrinos from Successful and Choked Gamma-Ray Bursts. *Phys. Rev. Lett.*, 87(17):171102–+, October 2001.

[452] 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.

[453] 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.

[454] 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.

[455] S. Ando and J. F. Beacom. Revealing the Supernova Gamma-Ray Burst Connection with TeV Neutrinos. *Phys. Rev. Lett.*, 95(6):061103, August 2005.

[456] 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.

[457] 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.

[458] 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.

[459] 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.

[460] 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.

[461] 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.

[462] J Ahrens et al. Sensitivity of the IceCube detector to astrophysical sources of high energy muon neutrinos. *Astroparticle Physics*, 20(5):507 – 532, 2004.

[463] 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.

[464] 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.

[465] O. Benhar, V. Ferrari, and L. Gualtieri. Gravitational wave asteroseismology reexamined. *Physical Review D*, 70(12):124015–+, December 2004.

[466] G. L. Israel, T. Belloni, L. Stella, Y. Rephaeli, D. E. Gruber, P. Casella, S. DallâĂŹOsso, N. Rea, M. Persic, and R. E. Rothschild. The discovery of rapid x-ray oscillations in the tail of the sgr 1806âĂŞ20 hyperflare. *The Astrophysical Journal Letters*, 628(1):L53, 2005.

[467] 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.

[468] 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.

[469] B. Abbott et al. Search for gravitational-wave bursts from soft gamma repeaters. *Physical Review Letters*, 101(21):211102, 2008.

[470] 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.

[471] J. Abadie et al. Search for Gravitational Wave Bursts from Six Magnetars. *The Astrophysical Journal Letters*, 734(2):L35, 2011.

[472] 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.

[473] B. Abbott and for the LIGO Scientific Collaboration. Implications for the Origin of GRB 070201 from LIGO Observations. *Astrophysical Journal*, 681:1419–1428, 2008.

[474] 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.

[475] 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.

[476] M. Shibata. Constraining Nuclear Equations of State Using Gravitational Waves from Hypermassive Neutron Stars. *Physical Review Letters*, 94(20):201101, May 2005.

[477] 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.

[478] 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.

[479] 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.

[480] 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.

[481] 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.

[482] 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.

[483] 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.

[484] 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.

[485] Michal Was, Patrick J. Sutton, Gareth Jones, and Isabel Leonor. Performance of an externally triggered gravitational-wave burst search. *Phys.Rev.*, D86:022003, 2012.

[486] 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.

[487] 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.

[488] 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.

[489] 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.

[490] J Aasi et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Phys. Rev. D*, 88:122004, 2013.

[491] C. Ott. Analytic CCSN Waveforms: 1. Barmode. *LIGO-T1000553*, 2013.

[492] K. Kotake. Multiple physical elements to determine the gravitational-wave signatures of core-collapse supernovae. *Comptes Rendus Physique*, 14:318, 2013.

[493] C. D. Ott. TOPICAL REVIEW: The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26:063001, 2009.

[494] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Class. Quant. Grav.*, 26(20):204015, October 2009.

[495] 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.

[496] 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.

[497] W. Engels, R. Frey, and C. D. Ott. A parametric multivariate statistical model for gravitational waveforms from rotating core collapse. *Submitted to Phys. Rev. D.*, 2014.

[498] 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.

[499] 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.

[500] 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.

[501] 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.

[502] 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.

[503] S. E. Gossan, C. D. Ott, et al. Observing Gravitational Waves from Core-Collapse Supernovae. *In preparation*, 2014.

[504] 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.

[505] 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.

[506] J. F. Beacom and P. Vogel. Can a supernova be located by its neutrinos? *Physical Review D*, 60(3):033007, 1999.

[507] 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.

[508] S. Klimenko. Coherent WaveBurst 2G Update. *LIGO-G1400993*, 2014.

[509] 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.

[510] C.L. Fryer, D.E. Holz, and S.A. Hughes. Gravitational Wave Emission from Core Collapse of Massive Stars. *Astrophysical Journal*, 565:430, 2002.

[511] A. L. Piro and E. Pfahl. Fragmentation of Collapsar Disks and the Production of Gravitational Waves. *Astrophysical Journal*, 658:1173, April 2007.

[512] S. Sutton, I. S. Heng, G. Prodi, and E. Chassande-Motin. Search Plan: All-Sky Burst Search. *LIGO-T1400192*, 2014.

[513] Thomas S. Adams. Detector Characterisation and Searches for Gravitationa l Waves Using GEO600. 2014. https://dcc.ligo.org/LIGO-P1400088.

[514] M. Ikeda et al. Search for Supernova Neutrino Bursts at Super-Kamiokande. *Astrophysical Journal*, 669:519, November 2007.

[515] Super-Kamiokande.

[516] Borexino.

[517] 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.

[518] Large Volume Detector.

[519] C. Vigorito [LVD Collaboration]. Galactic supernovae monitoring at LVD. *Nuc. Phys. B Proc. Suppl.*, 221:410, December 2011.

[520] 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.

[521] 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.

[522] LSC/Virgo Burst Analysis Group. LALInference_Burst. *https://wiki.ligo.org/Bursts/LIB/WebHome*, 2014.

[523] N. Cornish and T. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Submitted to Class. Quantum Grav., LIGO-P1400185*, 2014.

[524] T. Z. Summerscales, A. Burrows, L. S. Finn, and C. D. Ott. Maximum Entropy for Gravitational Wave Data Analysis: Inferring the Physical Parameters of Core-Collapse Supernovae. *Astrophysical Journal*, 678:1142, May 2008.

[525] 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.

[526] E. O'Connor and C. D. Ott. Black Hole Formation in Failing Core-Collapse Supernovae. *Astrophysical Journal*, 730:70, April 2011.

[527] 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.

[528] E.F. Keane. Fast Radio Bursts - I: Initial Cogitation, 2014.

[529] V. Berezinsky, B. Hnatyk, and A. Vilenkin. Superconducting cosmic strings as gamma-ray burst engines. *Baltic Astronomy*, 13:289, 2004.

[530] Yun-Wei Yu, Kwong-Sang Cheng, Gary Shiu, and Henry Tye. Implications of fast radio bursts for superconducting cosmic strings, 2014.

[531] Tanmay Vachaspati. Cosmic sparks from superconducting strings. *Physical Review Letters*, 101(14):141301, Sep 2008.

[532] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys. Rev. Lett.*, 85(18):3761–3764, Oct 2000.

[533] V. M. Lipunov and I. E. Panchenko. Pulsars revived by gravitational waves. *Astronomy & Astrophysics*, 312:937, August 1996.

[534] M.S. Pshirkov and K.A. Postnov. Radio precursors to neutron star binary mergings. *Astrophysics and Space Science*, 330(1):13–18, 2010.

[535] B. M. S. Hansen and M. Lyutikov. Radio and X-ray signatures of merging neutron stars. *MNRAS*, 322:695, April 2001.

[536] J. Moortgat and J. Kuijpers. Gravitational waves in magnetized relativistic plasmas. *Phys. Rev. D*, 70:023001, Jul 2004.

[537] 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.

[538] Lucia M. Franco, Bennett Link, and Richard I. Epstein. Quaking neutron stars. *The Astrophysical Journal*, 543(2):987, 2000.

[539] K.D. Kokkotas and B.G. Schmidt. Quasi-normal modes of stars and black holes. *Living Rev. Rel.*, 2:2, 1999.

[540] 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.

[541] Abbott, B. and others. Search for Gravitational Wave Bursts from Soft Gamma Repeaters. *Physical Review Letters*, 101:211102, 2008.

[542] 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.

[543] A.G. Lyne et al. Unusual glitch activity in the RRAT J1819?1458: an exhausted magnetar? *Mon. Not. R. Astron. Soc.*, 400:1439, 2009.

[544] P. Weltevrede, S. Johnston, and C.M. Espinoza. The glitch-induced identity changes of PSR J1119-6127. *AIP Conference Proceedings*, 1357:109, 2011.

[545] 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.

[546] 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.

[547] D Meacher, E Thrane, and T Regimbau. Statistical properties of astrophysical gravitational-wave background. *Phys. Rev. D*, 89:084063, 2014.

[548] 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.

[549] R Prix, S Giampanis, and C Messenger. Search method for long-duration gravitational-wave transients from neutron stars. *Phys. Rev. D*, 84:023007, 2011.

[550] C A van Eysden and A Melatos. Gravitational radiation from pulsar glitches. *Class. Quantum Grav.*, 25:225020, 2008.

[551] 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.

[552] A G Lyne, S L Shemar, and F G Smith. Statistical studies of pulsar glitches. *Mon. Not. R. Ast. Soc.*, 315:534, 2000.

[553] K Kashiyama and K Ioka. Magnetar asteroseismology with long-term gravitational waves. *Phys. Rev. D*, 83:081302, 2011.

[554] 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.

[555] 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.

[556] D I Jones and N Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Ast. Soc.*, 331:203, 2002.

[557] 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.

[558] 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.

[559] N Andersson and K D Kokkotas. The r-mode instability in rotating neutron stars. *Int. J. Mod. Phys. D*, 10:381, 2001.

[560] 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.

[561] A. L. Piro and E. Thrane. Gravitational Waves from Fallback Accretion onto Neutron Stars. *Astrophys. J.*, 761:63, 2012.

[562] A Melatos and M Priymak. Gravitational radiation from magnetically funnelled supernova fallback onto a magnetar. *Astrophys. J.*, 794:170, 2014.

[563] T. Prestegard, E. Thrane, et al. Identification of noise artifacts in searches for long-duration gravitational-wave transients. *Class. Quantum Grav.*, 29:095018, 2012.

[564] E Thrane and M Coughlin. Searching for gravitational-wave transients with a qualitative signal model: seedless clustering strategies. *Phys. Rev. D*, 88:083010, 2013.

[565] E. Thrane and M. Coughlin. Seedless clustering in all-sky searches for gravitational-wave transients. *Phys. Rev. D*, 89:063012, 2014.

[566] `http://ligo-vcs.phys.uwm.edu/wsvn/MatApps/packages/stochastic/trunk/stamp2/`.

[567] 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`.

[568] M. Coughlin, E. Thrane, and C. Christensen. Detecting compact binary coalescences with seedless clustering. *In preparation*, 2014.

[569] 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`.

[570] The LSC and Virgo. All-sky search for long-duration gravitational-wave transients with ligo, 2014. `https://dcc.ligo.org/LIGO-P1400138`.

[571] E Thrane, V Mandic, and N Christensen. Identification of very long transients in a stochastic search, 2014. `https://dcc.ligo.org/T1400559`.

[572] 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`.

[573] 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.

[574] M Coughlin, N Christensen, J Gair, S Kandhasamy, and E Thrane. Method for estimation of gravitational-wave transient model parameters in frequency-time maps. *Submitted to CQG*, 2014. `http://arxiv.org/abs/1404.4642`.

[575] 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.

[576] S. Ando, B. Baret, B. Bouhou, E. Chassande-Mottin, A. Kouchner, L. Moscoso, V. Van Elewyck, I. Bartos, S. Márka, Z. Márka, A. Corsi, I. Di Palma, M. A. Papa, A. Dietz, C. Donzaud, D. Eichler, C. Finley, D. Guetta, F. Halzen, G. Jones, P. J. Sutton, S. Kandhasamy, V. Mandic, E. Thrane, K. Kotake, T. Piran, T. Pradier, G. E. Romero, and E. Waxman. Multimessenger astronomy with gravitational waves and high-energy neutrinos. *ArXiv e-prints*, March 2012. arXiv:1203.5192.

[577] E. Waxman and J. Bahcall. High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physical Review Letters*, 78:2292–2295, March 1997.

[578] E. Waxman and J. N. Bahcall. Neutrino Afterglow from Gamma-Ray Bursts: 10^{18} eV. *Astrophys. Journal*, 541:707–711, October 2000.

[579] P. Mészáros and E. Waxman. TeV Neutrinos from Successful and Choked Gamma-Ray Bursts. *Physical Review Letters*, 87(17):171102, October 2001.

[580] S. Ando and J. F. Beacom. Revealing the Supernova Gamma-Ray Burst Connection with TeV Neutrinos. *Physical Review Letters*, 95(6):061103, August 2005.

[581] T. J. Galama, P. M. Vreeswijk, J. van Paradijs, C. Kouveliotou, T. Augusteijn, H. Böhnhardt, J. P. Brewer, V. Doublier, J.-F. Gonzalez, B. Leibundgut, C. Lidman, O. R. Hainaut, F. Patat, J. Heise, J. in't Zand, K. Hurley, P. J. Groot, R. G. Strom, P. A. Mazzali, K. Iwamoto, K. Nomoto, H. Umeda, T. Nakamura, T. R. Young, T. Suzuki, T. Shigeyama, T. Koshut, M. Kippen, C. Robinson, P. de Wildt, R. A. M. J. Wijers, N. Tanvir, J. Greiner, E. Pian, E. Palazzi, F. Frontera, N. Masetti, L. Nicastro, M. Feroci, E. Costa, L. Piro, B. A. Peterson, C. Tinney, B. Boyle, R. Cannon, R. Stathakis, E. Sadler, M. C. Begam, and P. Ianna. An unusual supernova in the error box of the γ-ray burst of 25 April 1998. *Nature*, 395:670–672, October 1998.

[582] S. R. Kulkarni, D. A. Frail, M. H. Wieringa, R. D. Ekers, E. M. Sadler, R. M. Wark, J. L. Higdon, E. S. Phinney, and J. S. Bloom. Radio emission from the unusual supernova 1998bw and its association with the γ-ray burst of 25 April 1998. *Nature*, 395:663–669, October 1998.

[583] D. Malesani, G. Tagliaferri, G. Chincarini, S. Covino, M. Della Valle, D. Fugazza, P. A. Mazzali, F. M. Zerbi, P. D'Avanzo, S. Kalogerakos, A. Simoncelli, L. A. Antonelli, L. Burderi, S. Campana, A. Cucchiara, F. Fiore, G. Ghirlanda, P. Goldoni, D. Götz, S. Mereghetti, I. F. Mirabel, P. Romano, L. Stella, T. Minezaki, Y. Yoshii, and K. Nomoto. SN 2003lw and GRB 031203: A Bright Supernova for a Faint Gamma-Ray Burst. *Astrophysical Journal Letters*, 609:L5–L8, July 2004.

[584] A. M. Soderberg, S. R. Kulkarni, E. Berger, D. W. Fox, M. Sako, D. A. Frail, A. Gal-Yam, D. S. Moon, S. B. Cenko, S. A. Yost, M. M. Phillips, S. E. Persson, W. L. Freedman, P. Wyatt, R. Jayawardhana, and D. Paulson. The sub-energetic γ-ray burst GRB 031203 as a cosmic analogue to the nearby GRB 980425. *Nature*, 430:648–650, August 2004.

[585] B. E. Cobb, C. D. Bailyn, P. G. van Dokkum, and P. Natarajan. SN 2006aj and the Nature of Low-Luminosity Gamma-Ray Bursts. *Astrophys. Journal*, 645:L113–L116, July 2006.

[586] E. Pian, P. A. Mazzali, N. Masetti, P. Ferrero, S. Klose, E. Palazzi, E. Ramirez-Ruiz, S. E. Woosley, C. Kouveliotou, J. Deng, A. V. Filippenko, R. J. Foley, J. P. U. Fynbo, D. A. Kann, W. Li, J. Hjorth, K. Nomoto, F. Patat, D. N. Sauer, J. Sollerman, P. M. Vreeswijk, E. W. Guenther, A. Levan, P. O'Brien, N. R. Tanvir, R. A. M. J. Wijers, C. Dumas, O. Hainaut, D. S. Wong, D. Baade, L. Wang, L. Amati, E. Cappellaro, A. J. Castro-Tirado, S. Ellison, F. Frontera, A. S. Fruchter, J. Greiner, K. Kawabata, C. Ledoux, K. Maeda, P. Møller, L. Nicastro, E. Rol, and R. Starling. An optical supernova associated with the X-ray flash XRF 060218. *Nature*, 442:1011–1013, August 2006.

[587] A. M. Soderberg, S. R. Kulkarni, E. Nakar, E. Berger, P. B. Cameron, D. B. Fox, D. Frail, A. Gal-Yam, R. Sari, S. B. Cenko, M. Kasliwal, R. A. Chevalier, T. Piran, P. A. Price, B. P. Schmidt, G. Pooley, D.-S. Moon, B. E. Penprase, E. Ofek, A. Rau, N. Gehrels, J. A. Nousek, D. N. Burrows, S. E. Persson, and P. J. McCarthy. Relativistic ejecta from X-ray flash XRF 060218 and the rate of cosmic explosions. *Nature*, 442:1014–1017, August 2006.

[588] D. M. Coward. Simulating a faint gamma-ray burst population. *MNRAS*, 360:L77–L81, June 2005.

[589] M. Ageron, J. A. Aguilar, I. Al Samarai, A. Albert, F. Ameli, M. André, M. Anghinolfi, G. Anton, S. Anvar, M. Ardid, and et al. ANTARES: The first undersea neutrino telescope. *Nuclear Instruments and Methods in Physics Research A*, 656:11–38, November 2011.

[590] M. de Jong. The KM3NeT neutrino telescope. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 623(1):445 – 447, 2010.

[591] A. Avrorin, V. Aynutdinov, V. Balkanov, I. Belolaptikov, S. Berezhnev, D. Bogorodsky, N. Budnev, I. Danilchenko, G. Domogatsky, A. Doroshenko, A. Dyachok, Zh. Dzhilkibaev, G. Ermakov, S. Fialkovsky, O. Gaponenko, K. Golubkov, O. Gres', T. Gres', N. Grishin, O. Grishin, A. Klabukov, A. Klimov, A. Kochanov, K. Konischev, A. Korobchenko, A. Koshechkin, V. Kulepov, D. Kuleshov, L. Kuzmichev, E. Middell, S. Mikheyev, M. Milenin, R. Mirgazov, E. Osipova, A. Pan'kov, L. Pan'kov, A. Panfilov, D. Petukhov, E. Pliskovsky, P. Pokhil, V. Poleschuk, E. Popova, V. Prosin, M. Rozanov, V. Rubtzov, A. Sheifler, A. Shirokov, B. Shoibonov, Ch. Spiering, O. Suvorova, B. Tarashansky, R. Wischnewski, A. Zagorodnikov, V. Zhukov, and A. Yagunov. The baikal neutrino experiment. *Nucl. Instrum. Methods Phys. Res., Sect. A*, 626-627(Supplement 1):S13 – S18, 2011.

[592] 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.

[593] C. L. Fryer and K. C. B. New. Gravitational Waves from Gravitational Collapse. *Living Reviews in Relativity*, 14:1–+, January 2011.

[594] K. Scholberg. Supernova Neutrino Detection. *Ann. Rev. Nuc. Part. Sc.*, 62:81, November 2012.

[595] C. D. Ott, E. Abdikamalov, E. O'Connor, C. Reisswig, R. Haas, P. Kalmus, S. Drasco, A. Burrows, and E. Schnetter. Correlated gravitational wave and neutrino signals from general-relativistic rapidly rotating iron core collapse. *Physical Review D*, 86(2):024026, July 2012.

[596] M. Edwards, R. Meyer, and N. Christensen. Bayesian parameter estimation of core collapse supernovae using gravitational wave simulations. *LIGO Document P1400034*, 2014.

[597] M. Shibata, K. Taniguchi, and K. Uryū. Merger of binary neutron stars with realistic equations of state in full general relativity. *Physical Review D*, 71(8):084021, April 2005.

[598] R. Oechslin, H.-T. Janka, and A. Marek. Relativistic neutron star merger simulations with non-zero temperature equations of state. I. Variation of binary parameters and equation of state. *Astron. Astrophys.*, 467:395–409, May 2007.

[599] R. Oechslin and H.-T. Janka. Gravitational Waves from Relativistic Neutron-Star Mergers with Microphysical Equations of State. *Physical Review Letters*, 99(12):121102, September 2007.

[600] M. Anderson, E. W. Hirschmann, L. Lehner, S. L. Liebling, P. M. Motl, D. Neilsen, C. Palenzuela, and J. E. Tohline. Simulating binary neutron stars: Dynamics and gravitational waves. *Physical Review D*, 77(2):024006, January 2008.

[601] Y. T. Liu, S. L. Shapiro, Z. B. Etienne, and K. Taniguchi. General relativistic simulations of magnetized binary neutron star mergers. *Physical Review D*, 78(2):024012, July 2008.

[602] K. Kiuchi, Y. Sekiguchi, M. Shibata, and K. Taniguchi. Long-term general relativistic simulation of binary neutron stars collapsing to a black hole. *Physical Review D*, 80(6):064037, September 2009.

[603] B. Giacomazzo, L. Rezzolla, and L. Baiotti. Accurate evolutions of inspiralling and magnetized neutron stars: Equal-mass binaries. *Physical Review D*, 83(4):044014, February 2011.

[604] Y. Sekiguchi, K. Kiuchi, K. Kyutoku, and M. Shibata. Gravitational waves and neutrino emission from the merger of binary neutron stars. *Phys. Rev. Lett.*, 107, 2011.

[605] A. Bauswein and H.-T. Janka. Measuring Neutron-Star Properties via Gravitational Waves from Neutron-Star Mergers. *Physical Review Letters*, 108(1):011101, January 2012.

[606] A. Bauswein, T. W. Baumgarte, and H.-T. Janka. Prompt Merger Collapse and the Maximum Mass of Neutron Stars. *Physical Review Letters*, 111(13):131101, September 2013.

[607] S. Bernuzzi, T. Dietrich, W. Tichy, and B. Bruegmann. Mergers of binary neutron stars with realistic spin. *ArXiv e-prints*, November 2013.

[608] K. Takami, L. Rezzolla, and L. Baiotti. Constraining the equation of state of neutron stars from binary mergers. *ArXiv e-prints*, March 2014.

[609] D. R. Lorimer, M. Bailes, M. A. McLaughlin, D. J. Narkevic, and F. Crawford. A Bright Millisecond Radio Burst of Extragalactic Origin. *Science*, 318:777–, November 2007.

[610] E. F. Keane, B. W. Stappers, M. Kramer, and A. G. Lyne. On the origin of a highly dispersed coherent radio burst. *MNRAS*, 425:L71–L75, September 2012.

[611] L. G. Spitler, J. M. Cordes, J. W. T. Hessels, D. R. Lorimer, M. A. McLaughlin, S. Chatterjee, F. Crawford, J. S. Deneva, V. M. Kaspi, R. S. Wharton, B. Allen, S. Bogdanov, A. Brazier, F. Camilo, P. C. C. Freire, F. A. Jenet, C. Karako-Argaman, B. Knispel, P. Lazarus, K. J. Lee, J. van Leeuwen, R. Lynch, A. G. Lyne, S. M. Ransom, P. Scholz, X. Siemens, I. H. Stairs, K. Stovall, J. K. Swiggum, A. Venkataraman, W. W. Zhu, C. Aulbert, and H. Fehrmann. Fast Radio Burst Discovered in the Arecibo Pulsar ALFA Survey. *ArXiv e-prints*, April 2014. arXiv:1404.2934.

[612] S. Burke-Spolaor, M. Bailes, R. Ekers, J.-P. Macquart, and F. Crawford, III. Radio Bursts with Extragalactic Spectral Characteristics Show Terrestrial Origins. *Astrophysical Journal*, 727:18, January 2011.

[613] P. Saint-Hilaire, A. O. Benz, and C. Monstein. Short-duration Radio Bursts with Apparent Extragalactic Dispersion. *ArXiv e-prints*, February 2014. arXiv:1402.0664.

[614] D. R. Lorimer. Blind surveys for radio pulsars and transients. *AIP Conference Proceedings*, 1357, 2011.

[615] C. W. Stubbs. *Class. Quantum Grav.*, 25:184033, 2008.

[616] E. Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442:166–236, April 2007.

[617] L. Li and B. Pacyznski. *Astrophys. J.*, 507:L59–62, 1998.

[618] S. Kulkarni. Modeling supernova-like explosions associated with gamma-ray bursts with short durations. 2005.

[619] B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. *MNRAS*, 406:2650, August 2010.

[620] L. F. Roberts, D. Kasen, W. H. Lee, and E. Ramirez-Ruiz. Electromagnetic Transients Powered by Nuclear Decay in the Tidal Tails of Coalescing Compact Binaries. *Astrophysical Journal Letters*, 736:L21, July 2011.

[621] J. Sylvestre. *Astrophys. J.*, 591:1152–6, 2003.

[622] J. S. Bloom et al. Astro2010 decadal survey whitepaper: Coordinated science in the gravitational and electromagnetic skies. 2009.

[623] S. Desai, E. O. Kahya, and R. P. Woodard. *Phys. Rev. D*, 77:124041, 2008.

[624] N. Dalal, D. E. Holz, S. Hughes, and B. Jain. *Phys. Rev. D*, 74:063006, 2006.

[625] J. Kanner, T. L. Huard, S. Marka, D. C. Murphy, J. Piscionere, M. Reed, and P. Shawhan. Looc up: Locating and observing optical counterparts to gravitational wave bursts. *Class. Quantum Grav.*, 25:184034, 2008.

[626] S. Nissanke, M. Kasliwal, and A. Georgieva. Identifying Elusive Electromagnetic Counterparts to Gravitational Wave Mergers: An End-to-end Simulation. *Astrophysical Journal*, 767:124, April 2013.

[627] J. Abadie et al. Implementation and testing of the first prompt search for gravitational wave transients with electromagnetic counterparts. *Astron. Astrophys.*, 539:A124, 2012.

[628] C. Pankow et al. A burst search for gravitational waves from binary black holes. *Classical and Quantum Gravity*, 26:204004, 2009.

[629] G. Mazzolo, F. Salemi, M. Drago, V. Necula, C. Pankow, et al. Prospects for intermediate mass black hole binary searches with advanced gravitational-wave detectors. *Phys.Rev.*, D90:063002, 2014.

[630] R. M. O'Leary, B. Kocsis, and A. Loeb. Gravitational waves from scattering of stellar-mass black holes in galactic nuclei. *MNRAS*, 395:2127–2146, June 2009.

[631] B. Kocsis, M. E. Gáspár, and S. Márka. Detection Rate Estimates of Gravity Waves Emitted during Parabolic Encounters of Stellar Black Holes in Globular Clusters. *Astrophysical Journal*, 648:411–429, September 2006.

[632] P. Ajith et al. A template bank for gravitational waveforms from coalescing binary black holes: I. non-spinning binaries. *Phys. Rev.*, D77:104017, 2008.

[633] Yi Pan et al. A data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. *Phys. Rev. D*, 77:024014, 2008.

[634] P. Ajith, M. Hannam, S. Husa, Y. Chen, B. Bruegmann, N. Dorband, D. Mueller, F. Ohme, D. Pollney, C. Reisswig, L. Santamaria, and J. Seiler. "Complete" gravitational waveforms for black-hole binaries with non-precessing spins. *ArXiv e-prints*, September 2009. arXiv:0909.2867.

[635] W. E. East, S. T. McWilliams, J. Levin, and F. Pretorius. Observing complete gravitational wave signals from dynamical capture binaries. *Phys. Rev. D*, 87:043004–+, 2013.

[636] A. Gopukumar and G. Shafer. Gravitational wave phasing for spinning compact binaries in inspiraling eccentric orbits. *Phys. Rev. D*, 84:124007–+, 2011.

[637] John Veitch and Alberto Vecchio. Assigning confidence to inspiral gravitational wave candidates with Bayesian model selection. *Class. Quant. Grav.*, 25:184010, 2008.

[638] B. Abbott et al. First LIGO search for gravitational wave bursts from cosmic (super)strings. *Physical Review D*, 80(6):062002, 2009.

[639] P. Hořava, A. Mohd, C. M. Melby-Thompson, and P. Shawhan. GR 20 parallel session A3: modified gravity. *General Relativity and Gravitation*, 46:1720, April 2014.

[640] S. A. Farrell, N. A. Webb, D. Barret, O. Godet, and J. M. Rodrigues. An intermediate-mass black hole of over 500 solar masses in the galaxy ESO243-49. *Nature*, 460:73–75, July 2009.

[641] I. Mandel. Spin distribution following minor mergers and the effect of spin on the detection range for low-mass-ratio inspirals. *ArXiv e-prints*, July 2007. 0707.0711.

[642] H. Nakano, Y. Zlochower, C. O. Lousto, and M. Campanelli. Intermediate-mass-ratio black hole binaries. II. Modeling trajectories and gravitational waveforms. *Physical Review D*, 84(12):124006, December 2011.

[643] 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.

[644] 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.

[645] I. Kamaretsos, M. Hannam, S. Husa and B. S. Sathyaprakash. Black-hole hair loss: learning about binary progenitors from ringdown signals. *Physical Review D*, 85:024018, 2011.

[646] A. F. Siegel. Testing for Periodicity in a Time Series. *Journal of the American Statistical Association*, 75:345, 1980.

[647] B. Abbott et al. Upper limit map of a background of gravitational waves. *Phys.Rev.*, D76:082003, 2007.

[648] J. Abadie et al. Directional limits on persistent gravitational waves using LIGO S5 science data. *Phys.Rev.Lett.*, 107:271102, 2011.

[649] D. Passuello and S. Braccini. Virgo Note (VIR-046A-07).

[650] Torre, O. *Resampling method for semi-targeted continuous gravitational wave search in Virgo interferomete*. PhD thesis, University of Pisa, 2010. Virgo-LIGO note VIR-0013A-10, available at `https://pub3.ego-gw.it/itf/tds/index.php`.

[651] S. Braccini, G. Cella, I. Ferrante, D. Passuello, O. Torre. A resampling technique to correct for Doppler effect in gravitational wave search, 2010. Virgo-LIGO note VIR-0090A-10, available at `https://pub3.ego-gw.it/itf/tds/index.php`.

[652] S. Braccini, G. Cella, I. Ferrante, D. Passuello, and O. Torre. Resampling technique to correct for the Doppler effect in a search for gravitational waves. *Phys.Rev.*, D83:044033, 2011.

[653] Reinhard Prix, Stefanos Giampanis, and Chris Messenger. Search method for long-duration gravitational-wave transients from neutron stars. *Phys.Rev.*, D84:023007, 2011.

[654] I. Santiago-Prieto, I. S. Heng, D. I. Jones, and J. Clark. Prospects for transient gravitational waves at r-mode frequencies associated with pulsar glitches. *Journal of Physics Conference Series*, 363(1):012042, June 2012.

[655] Y. Levin and G. Ushomirsky. Crust core coupling and r mode damping in neutron stars: A Toy model. *Mon.Not.Roy.Astron.Soc.*, 324:917, 2001.

[656] T. Strohmayer and S. Mahmoodifar. A Non-radial Oscillation Mode in an Accreting Millisecond Pulsar? *Astrophysical Journal*, 784:72, March 2014.

[657] M. Bejger and A. Kr?lak. Searching for gravitational waves from known pulsars at once and twice the spin frequency. *Class.Quant.Grav.*, 31:105011, 2014.

[658] J. Zhang, C. Jarman, and K. Riles. Adding Coherent Summing to PowerFlux, 2008. *LIGO Technical Document* LIGO-T080061, available at `https://dcc.ligo.org`.

[659] Holger J. Pletsch and Bruce Allen. Exploiting global correlations to detect continuous gravitational waves. *Phys.Rev.Lett.*, 103:181102, 2009.

[660] D. Keitel, R. Prix, M. A. Papa, P. Leaci, and M. Siddiqi. Search for continuous gravitational waves: Improving robustness versus instrumental artifacts. *Phys. Rev. D*, 89(6):064023, March 2014.

[661] R.V. Wagoner. Gravitational Radiation from Accreting Neutron Stars. *Astrophys.J.*, 278:345–348, 1984.

[662] J. T. Whelan et al. A Model-Based Cross-Correlation Search for Gravitational Waves from Scorpius X-1, 2013. *LIGO Technical Document* LIGO-P1200142, available at `https://dcc.ligo.org`.

[663] S. Sundaresan and J. T. Whelan. Windowing and Leakage in the Cross-Correlation Search for Periodic Gravitational Waves, 2012. *LIGO Technical Document* LIGO-T1200431, available at `https://dcc.ligo.org`.

[664] Chris Messenger. A semi-coherent search strategy for known continuous wave sources in binary systems. *Phys.Rev.*, D84:083003, 2011.

[665] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters. *Physical Review D*, 53:3033–3055, 1996.

[666] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Physical Review D*, 53:6749–6761, 1996.

[667] Patrick R. Brady, Teviet Creighton, Curt Cutler, and Bernard F. Schutz. Searching for periodic sources with LIGO. *Phys.Rev.*, D57:2101–2116, 1998.

[668] Reinhard Prix. Search for continuous gravitational waves: Metric of the multi-detector F-statistic. *Phys.Rev.*, D75:023004, 2007.

[669] 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.

[670] `https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html`.

[671] T. Accadia, F. Acernese, M. Agathos, P. Astone, G. Ballardin, et al. The NoEMi (Noise Frequency Event Miner) framework. *J.Phys.Conf.Ser.*, 363:012037, 2012.

[672] S G Crowder, R Namba, V Mandic, S Mukohyama, and M Peloso. Measurement of parity violation in the early universe using gravitational-wave detectors. *Phys. Lett. B*, 726:66, 2013.

[673] J. Skilling. "bayesian inference and maximum entropy methods in science and engineering". In R. Fischer, R. Preuss, and U. von Toussaint, editors, *AIP Conf. Proc.*, volume 735, page 395. Melville: AIP, 2004.

[674] T. Regimbau, T. Dent, W. Del Pozzo, S. Giampanis, T. G. F. Li, C. Robinson, C. Van Den Broeck, D. Meacher, C. Rodriguez, B. S. Sathyaprakash, and K. Wójcik. Mock data challenge for the einstein gravitational-wave telescope. *Phys. Rev. D*, 86:122001, Dec 2012.

[675] T. Regimbau, D. Meacher, and M. Coughlin. Second Einstein Telescope mock science challenge: Detection of the gravitational-wave stochastic background from compact binary coalescences. *Physical Review D*, 89(8):084046, April 2014.

[676] R Schofield et al. DetChar line hunters remind us that I/O box fans and switcher are contaminating channels, 2014.
`https://alog.ligo-wa.caltech.edu/aLOG/index.php?callRep=11931`.

[677] P Meyers. Discussion of Loud Time Segment in H1L1 S6 Epoch C for Radiometer Analysis, 2014.
`http://www.ldas-sw.ligo.caltech.edu/ilog/pub/ilog.cgi?group=stochastic&task=view&date_to_view=02/18/2014&anchor_to_scroll_to=2014:02:18:07:50:38-meyers`.

[678] E Thrane. Measuring the non-Gaussian stochastic gravitational-wave background: a method for realistic interferometer data. *Phys. Rev. D*, 87:043009, 2013.

[679] L. Martellini and T. Regimbau. A Semi-Parametric Approach to the Detection of Non-Gaussian Gravitational Wave Stochastic Backgrounds. *ArXiv e-prints*, May 2014.

[680] LSC Detector Characterization Group home page.
`https://wiki.ligo.org/foswiki/bin/view/DetChar/WebHome`.

[681] J. Zweizig. Data Monitor Tool Project .
`http://www.ligo.caltech.edu/~jzweizig/dmt/DMTProject/`.

[682] EPICS home page. `http://www.aps.anl.gov/epics/index.php`.

[683] Home page for the omega pipeline:.
`https://trac.ligo.caltech.edu/omega/`.

[684] 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.

[685] 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.

[686] R. Schofield and A. Effler. aLIGO PEM System Upgrade, LIGO Document T1200221-v5. Technical report, LIGO, 2012.

[687] 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.

[688] 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.

[689] J. Smith, G. Gonzalez, J. Giaime, F. Raab, D. Brown, Z. Marka, G. Kang, T. Souradeep, B. Willke, K. Riles, B. Slagmolen, R. Adhikari, T. Summerscales, X. Siemens, S. Klimenko, and M. Cavaglia. Report of the LSC Committee for Scientific Monitoring in Advanced LIGO, LIGO Document L1200249-v12. Technical report, LIGO, 2012.
`https://dcc.ligo.org/LIGO-L1200249`.

[690] Detector Commissioning and Characterization Projects.
`https://wiki.ligo.org/DetChar/DetectorDetCharCommissioningProjects`.

[691] Advanced LIGO Subsystem Characterization Matrix.
`https://wiki.ligo.org/DetChar/AligoSubsystemMatrix`.

[692] 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.

[693] ligoDV home page.
`https://www.lsc-group.phys.uwm.edu/daswg/projects/ligodv.html`.

[694] mDV is available from the 40m svn.

[695] 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.

[696] S6 glitch group homepage. `https://wiki.ligo.org/DetChar/GlitchStudies`.

[697] LIGO calibration web page.
`http://blue.ligo-wa.caltech.edu/engrun/Calib_Home/`.

[698] 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).

[699] 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.