LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	**LIGO–T1000476**

The LSC-Virgo white paper on gravitational wave data analysis

Science goals, status and plans, priorities

(2010-2011 public edition)

The LSC-Virgo Data Analysis Working 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 2010/08/15

Contents

1 Introduction 5

2 The characterization of the data 7
2.1 LIGO Detector Characterization . . . 7
2.1.1 Introduction . . . 7
2.1.2 Data Run Support . . . 8
2.1.3 Software Infrastructure . . . 8
2.1.4 Glitch Investigations . . . 10
2.1.5 Spectral features Group . . . 11
2.1.6 Calibrations . . . 12
2.1.7 Timing . . . 13
2.2 GEO Detector Characterization . . . 14
2.2.1 Introduction . . . 14
2.2.2 Long-term Plan . . . 15
2.2.3 Short-term Plan . . . 18
2.2.4 Opportunities for LIGO-VIRGO Members . . . 18
2.3 Virgo Detector Characterization . . . 18

3 Searches for signals from compact binary coalescence 19
3.1 Science goals . . . 19
3.2 Gravitational waves from the coalescence of compact binary systems . . . 20
3.3 Rates and Astrophysical Populations . . . 22
3.4 Search Strategies . . . 22
3.4.1 Matched filtering and Signal Consistency Tests . . . 23
3.4.2 Low latency pipeline . . . 24
3.4.3 Batch Analysis pipeline . . . 25
3.4.4 Background estimation . . . 26
3.4.5 Instrumental and environmental vetoes . . . 26
3.4.6 Efficiency evaluation . . . 27
3.4.7 Pipeline of pipelines . . . 27
3.4.8 Externally triggered searches . . . 28
3.4.9 Blind analyses . . . 28
3.4.10 Follow-up of candidate events . . . 29
3.4.11 Interpretation of results . . . 29
3.5 CBC Searches and Results from LIGO S1 to S4 data . . . 30
3.6 CBC Search Results from LIGO S5 and Virgo VSR1 data . . . 30
3.7 Ongoing CBC Searches and Studies from LIGO S5 and Virgo VSR1 data . . . 31
3.7.1 S5/VSR1 GRB searches: . . . 31
3.7.2 S5 high mass search: . . . 31
3.7.3 S5/VSR1 ringdown search: . . . 32
3.7.4 S5/VSR1 IMR comparison: . . . 32
3.7.5 NINJA 2: . . . 32
3.8 CBC Searches with the LIGO S6 and Virgo VSR2 data . . . 33
3.8.1 Detector Characterization . . . 33
3.8.2 Low Mass CBC Search . . . 34
3.8.3 High mass Inspiral–Merger–Ringdown CBC: . . . 34

3.8.4 Externally Triggered Searches: . . . 35
3.8.5 Low-latency search and EM follow-up: . . . 35
3.9 CBC Group Priorities and Goals for FY2011 . . . 36
3.9.1 S5/VSR1 Analyses . . . 36
3.9.2 S6/VSR2-3 Analyses . . . 36
3.10 High priority analysis development . . . 37
3.11 Development prior to advanced detector era . . . 40
3.11.1 Searches in the "Before Advanced detector Era" . . . 40
3.11.2 Analysis development in the "Before Advanced detector Era" . . . 41
3.12 Advanced detector era . . . 42

4 Searches for general burst signals **45**
4.1 Recent observational results . . . 46
4.2 Basic methodologies used in burst searches . . . 46
4.2.1 Signal extraction methods . . . 47
4.2.2 Detector characterization . . . 48
4.2.3 Accidental background estimation . . . 49
4.2.4 Simulations . . . 50
4.2.5 Hardware signal injections . . . 50
4.3 Science goals . . . 50
4.3.1 Search as broadly as possible for GWs . . . 51
4.3.2 Use knowledge of astrophysically plausible signals . . . 52
4.3.3 Look for GW transients triggered by energetic astrophysical events . . . 53
4.3.4 Seek confirmation of GW candidates with follow-up observations . . . 59
4.3.5 Be open to alternative theories of gravity . . . 61
4.3.6 Be prepared to detect a GW signal with confidence . . . 62
4.4 First science targets for S6/VSR2+3 . . . 63
4.5 GW burst science in the advanced detector era . . . 64
4.6 Preparing for the advanced detector era . . . 68

5 Searches for continuous-wave signals **69**
5.1 Non-accreting pulsars . . . 69
5.1.1 Fully coherent targeted searches . . . 70
5.1.2 Wide Parameter Searches for Known Pulsars . . . 73
5.2 Non-pulsing non-accreting neutron stars and favorable directions . . . 73
5.2.1 Cassiopeia A . . . 74
5.2.2 Other coherent directed searches . . . 74
5.2.3 Search for Calvera . . . 74
5.2.4 Searches for sources near the galactic center . . . 75
5.2.5 Supernova 1987A using the cross-correlation technique . . . 75
5.2.6 Semi-targeted search using stroboscopic resampling . . . 76
5.2.7 Semi-targeted searches for "transient CW signals" . . . 77
5.2.8 Directed search using $\mathcal{F}$-statistic . . . 78
5.2.9 Other targets . . . 78
5.3 All-sky searches for isolated neutron stars . . . 78
5.3.1 PowerFlux method . . . 79
5.3.2 Hough transform method . . . 80
5.3.3 `Einstein@Home Searches` . . . 81

5.3.4 Followup-searches to confirm or veto CW signal candidates . . . 82
5.3.5 Hierarchical Hough searche . . . 83
5.3.6 $\mathcal{F}$-statistic all-sky search . . . 83
5.3.7 Loosely coherent search . . . 83
5.3.8 Cross-correlation search . . . 83
5.4 Accreting and unknown binary neutron stars . . . 84
5.4.1 Known binary systems . . . 84
5.4.2 Unknown binary systems . . . 86
5.5 Search for Non-Sinusoidal Periodic Waves . . . 87
5.6 Needs and plans for S6/VSR2 . . . 88
5.7 Scientific goals for the advanced-detector era . . . 88
5.8 Preparing for the advanced-detector era . . . 90

6 Searches for stochastic backgrounds **92**
6.1 Sources of Stochastic Gravitational-wave Background . . . 92
6.2 Stochastic Search Method . . . 92
6.2.1 All-Sky Search . . . 92
6.2.2 Directional Search . . . 94
6.2.3 Mapping . . . 94
6.2.4 Multi-baseline: LSC/VIRGO joint search . . . 95
6.2.5 H1H2 All-Sky Search . . . 95
6.2.6 Searching for Long-Lasting Transients . . . 96
6.3 Results and Plans . . . 97
6.3.1 Status of S5 Searches . . . 97
6.3.2 Plans for S6 Searches . . . 98
6.3.3 Plans for Advanced Detector Era . . . 99

7 Computing and Software **100**
7.1 Current status . . . 100
7.2 Activities in support of LDG Operations . . . 100
7.3 Data Analysis Software Development Activities . . . 102
7.4 Intermediate-term development activities . . . 104
7.5 Preparing for the advanced detector era . . . 109
7.5.1 Software . . . 109
7.5.2 Services . . . 115

References **119**

1 Introduction

As for the previous editions, the data analysis white paper annually describes the goals, status, and plans of the data analysis teams in the LSC and Virgo. The document is revised and updated every year in the summer. This is the document for 2010-2011. It is intended to

- explain the science that we are doing
- identify "holes" in our science plan and tasks that demand more manpower
- prioritize of our objectives
- identify areas when manpower should be shifted to and/or removed from
- exploit synergies among the work carried out in different search groups
- facilitate an harmonious exploitation of common resources

In July 2009 began the 6th science run of the of the LIGO detectors (S6) and the 2nd/3rd science run of the Virgo detector (VSR2/3). The run has been interrupted by extended periods of commissioning: 070709-240809 for LLO and 070709-010909 LHO (S6A), then 210909-080110 even though data was taken until 120110 (S6B) and 160110-201010 for both LLO and LHO. For Virgo VSR2 took place during 070709-080110 and VSR3 040810-041010. In January 2010, Virgo shut down for a major change to the mirror suspensions; Virgo was expected to resume running in July 2010 (now August), and since the detector will be significantly different, the run will be called VSR3 from that point. At the time of writing the search groups are still analyzing data from the S5 and the S5/VSR1 run and have begun to look at S6/VSR2 data with their lower latency analysis tools.

A chapter in this year's white paper is devoted to the science and data analysis in the era of the advanced detectors. Much of the work that we are doing now is inspired by the vision for the science in that era.

The data analysis activities are organized in four groups which, broadly speaking, map into four different search approaches, depending on the different signals: compact binary coalescence signals, burst signals, continuous wave signals and stochastic backgrounds. This classification is historical in origin and as the searches that we carry out evolve, becoming more ambitious and broader, the boundaries between the different signals and the boundaries between the different search techniques become somewhat blurred and this distinction is only indicative.

Since April 2007 the LSC and Virgo have been operating their instruments as a network and the analysis of the data of the two detectors is carried out jointly. This allows to increase the chances of a detection as well as the extraction of more information on the detected signal: source localization and polarization may be disentangled and derived from the data of three or more detectors. Localizing the source opens up the possibility to also leverage information from other observatories, radio, optical, X-ray and neutrinos, since most of the gravitational-wave emission scenarios also predict significant emission of other forms of radiation in ways that may be correlated with a gravitational wave observation. This would significantly increase the confidence of a gravitational wave detection, as well as unveiling important information on source populations, astrophysical distances, the validity of post-Newtonian approximations in the strong gravity regime, alternative theories of gravity as well as astrophysical phenomena such as gamma ray bursts, soft gamma repeaters, core-collapse supernovae and prompt radio pulses. In January we have seen the first prompt production of GW-alerts sent out to allow fast pointing of (SWIFT and optical) telescopes. The signature of the MOU with Antares (and soon with Ice Cube) will also allow the first a-posteriori GW-HEN analysis of triggers. In the sections of this paper devoted to the searches for short gravitational wave signals one can see how the focus of our research is shifting in this direction, as we analyze S6/VSR2/VSR3 data and think ahead at the era of the advanced detectors, in which we expect routine detections and this type of work constitutes the main science output of our research.

Since the understanding of artifacts in our data is an essential part of the analysis work (allowing us to reduce the false alarm rate and increase our confidence that we have the tools to reliably interpret the output of the detector), we begin with a section on Detector Characterization. 'Detector characterization' is a term that indicates a variety of activities at the boundary between the detector and the data analysis. These activities are aimed at supporting the experimental effort by understanding the detector sensitivity performance to various types of signals and spotting critical artifacts that degrade it. They are also aimed at supporting the data analysis efforts by providing lists of "safe" vetoes for times and for triggers produced by the search pipelines which can be correlated with malfunctioning of known origin in the instrument and hence that can be discarded as not being of astrophysical origin. This section also includes information on the calibration procedures and expected calibration accuracy in the upcoming runs.

Finally, since data analysis work both drives and is constrained by the computing environment and facilities where it develops, the last section of this document describes the development and maintenance of software tools and the management of software and computing resources.

2 The characterization of the data

2.1 LIGO Detector Characterization

2.1.1 Introduction

Analysis of LIGO data requires a systematic understanding and characterization of the detector: its response function, timing stability, noise behavior and sensitivity to the environment, including correlated noise between interferometers. The confidence associated with source detection or upper limits for detection depends on detector performance characteristics, including: power spectra, the probability distribution of the detector output, stationarity of the noise, line noise sources, and the statistics of transients.

Commissioning too is part of detector characterization. In particular, understanding which instrumental or environmental sources define the current noise floor at any given frequency is critical to eliminating or ameliorating those sources.

In practice, detector characterization is carried out at several different levels within the LSC and by a variety of scientists focused on different problems. Commissioners operate at the cutting edge of detector characterization, evaluating and updating interferometer noise budgets, as improvements are made between data runs. By the nature of commissioning, long-term stability and statistics of transient artifacts is difficult to evaluate when such work is most intense. In the past, data runs have served as testing grounds for that stability, and there have been some unpleasant surprises. As experience has accumulated, as background monitoring tools have improved, and as more data have been collected in science mode, the rapidity of diagnostic feedback has improved dramatically. Feedback useful for mitigation and commissioning happens very often, if not routinely. Some investigations focus on interferometer-based detector characterization, such as investigation of line artifacts or environmental disturbances, while others focus on astrophysics-search-targeted artifacts, such as coherent or accidentally coincident glitches that could pollute inspiral and burst searches, or wandering line features that could mimic a pulsar.

The LSC Detector Characterization (DetChar) community has a broad membership, including full-time commissioners, full-time data analysts, and many in between. In practice, the DetChar working groups have concentrated most on providing characterization tools, and on providing characterization 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. There are four working groups within Detchar, dedicated to software infrastructure, timing diagnostics, "glitch" (short transients) investigations, and spectral features in the gravitational wave data. There is a separate but closely interacting working group dedicated to the calibration of the instrument.

In 2009-2010, the group's activities have been mostly directed towards diagnosing and improving the data quality of the LIGO data gathered in the Sixth Science Run (S6), started in July 2009 and expected to end in October 2010. In the 2010-2011 cycle, the group expects to produce final results and documentation for the work done on S6 data, as well as, and perhaps more importantly, begin working on the characterization of the advanced LIGO detector systems, as they begin to be installed in the LIGO Observatories in 2011.

As mentioned, a high priority of the Detchar group is to document the work done in S6, the methods used and results obtained, and a general summary of data quality of the LIGO detectors during the science run. The (ambitious) goal is to produce a publishable paper, which will serve as guide and reference for the work in Advanced LIGO, and for the understanding (or lack thereof) of instrumental problems that may affect the discovery of gravitational waves.

Also in 2010-2011, but related now to Advanced LIGO detectors, we will investigate the data quality of individual channels and subsystems as they are built in order to pick off individual problems as they arise rather than waiting for the first engineering runs and trying to sort the myriad of problems that arise in the detector output back to their individual sources. This includes using our tested tools to diagnose and fix when possible glitches, lines, stationarity, and noise budget for each system. This is a significantly different approach than we have used in previous runs, the logic being that if each individual subsystem can be monitored and checked off as well-behaved and understood, we will be much more likely to have a clean overall system when a complete detector system is first made operational. This will also serve the role of making detector characterization group members familiar with the individual subsystems, which should again pay dividends in the first science runs, and more importantly, in the first several detections of gravitational waves.

In the following subsections, we first describe how we envision data run support for Advanced LIGO, then describe the software infrastructure and storage of data quality information. We then turn to the activities and priorities of the different working groups working in DetChar.

2.1.2 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 operator per interferometer and at least one scientist (scimon) per observatory. The scimons have been responsible for monitoring the quality of the data, carrying out investigations (including the causes of lock loss), 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 approved in 2009 a policy for a 7-day minimum stay for people working as "scientific monitors"; this has resulted in more efficient and expert data monitoring in S6. We will critically evaluate the performance of this new system, and use the experience for the system to be used in future Advanced LIGO runs (although no scientific run is expected in 2010-2011). It is apparent that the longer stays have proven beneficial, and we will investigate ways to foster ways to encourage even longer stays of scimons and other LSC scientists at the Observatories. Also, the experience of a DetChar member having extended stays at the Hanford Observatory during many years proved critical to the understanding of many artifacts, especially those related to the coupling of the physical environment. We want to find ways to have more such LSC members closely involved at the Observatories.

Another important aspect of data run support is injection of simulated astrophysical signals into the detector hardware [28], to validate data analysis pipelines with high confidence. LIGO Laboratory and LSC scientists have provided the manpower to set up the injection infrastructure and carry out the injections during S6. Again, we will use the experience in S1-S6, including the results of "blind" injection exercise, to plan for future runs.

Although there will be no science runs in 2010-2011, we will propose short time intervals (one to several days) dedicated to study the data quality of the different Advanced LIGO subsystems as they are installed. We expect this activity will start in 2011, and want to take full advantage of the possibility to exercise data monitoring tools, as well as get an early understanding of the artifacts in each of the building blocks of what will be a complex instrument for routine detections of gravitational waves. This experience will need a close communication between commissioning teams and detector characterization groups, which will be an added advantage for the early stages of detector operations.

2.1.3 Software Infrastructure

Over the years, many tools have been developed for online and offline monitoring of detector status and data quality. Many of these software tools (EPICS[2], DTT[3] 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[3], is mostly used as a background-process environment for continuous monitoring. The DMT system provides many critical functions in S6, including the online production of calibrated strain files that are used for online analysis, online production of data quality information that is used for vetoing results of online analysis, 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.

The monitoring of the data quality is done reviewing the results produced by the DMT programs and results of data analysis programs and data collecting scripts running in the LIGO Data Grid clusters at the Observatories. This review is done continuously during science runs by scientific monitors ("scimons") at a basic level, and periodically at a deeper level by members of the Detchar Glitch and Spectral features groups. These investigations result in the diagnosing, identification ("flagging") and sometimes fixing of artifacts in the gravitational wave channel, which reduce the background for the searches of astrophysical signals in the data.

The DetChar group compiles information from its investigations to create a repository of data quality (DQ) information for each engineering and science 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 analyzable, or in "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 has been significantly revamped in S6 by the LSC Software Working group, as described in 7, resulting in a reliable system where flags are introduced online by well tested DMT monitoring programs, and offline by the DetChar group and by scimons.

Many scientists who do not routinely spend their time in the Observatories' control rooms did not use to have easy access to the detector data, especially for interactive graphical tools like DTT or DataViewer. The live data sharing was limited to few people, since heavy volume requests could cause problems with the online data archiving. The interest fore remote data access was high enough that very useful and popular data retrieving and graphical Matlab tools developed for non-live data access ("ligoDV") were adapted to the online data, which increased the demand for data access outside the Observatories. Also, Matlab tools in the control room that had access to live or recent data could not access older data archived outside the control rooms. The situation has qualitatively improved in 2009-2010 due the LIGO Laboratory's work on network data servers (NDS2) that serve live data, but are independent from the live archiving processes, and thus are safe for satisfying the greater demand. Notice that for diagnosing purposes, it is not just the gravitational wave channel what is of interest, but it is many other data streams (and their archived trends) that are needed. This need is probably going to be only greater and more complex in Advanced LIGO, since the systems themselves are going to be significantly more complex than for current detectors (multiple suspensions, active seismic isolation, signal recycling readout with output mode cleaners, etc), and have more channels being recorded.

Despite the critical importance of Detector Characterization software, there is currently only a single member with full time dedication to DetChar software (mostly dedicated to DMT and NDS2), and only a few other members partially dedicated to develop and maintain software tools. More complex Advanced LIGO detectors, expected to detect gravitational waves often will need a much more robust infrastructure and human support, with more personnel from the LIGO Scientific Collaboration (as opposed to just the LIGO Observatory).

The high priority areas we have identified for the 2010-2011 cycle about DetChar software infrastructure are listed below. Most of these are coordinated with the LSC Software working group, as listed in 7,

- identify the list of DMT tools critical for operations, so their performance can be monitored and guaranteed even in the absence of detector gravitational data.

- recruit more LSC and LIGO Lab members of the Detchar Software working group to guarantee long term reliability of software tools;
- expanded features of the segment database, including:
  - ability for creation of DQ flags for non-gravitational wave detector channels,
  - improvement of fast, safe and friendly ways to query the database from the control room,
  - allowing intervals with fractions of seconds.
- improve algorithms to detect qualitative changes in the character of data streams of non-gravitational wave channels;
- improve DMT hardware and software infrastructure to facilitate development of new online monitoring software;
- perform a critical review of existing and needed monitor programs, in light of use during science runs S1-S6;
- create a thorough documentation of channels available from each Advanced LIGO detector subsystem, as well as a list of critical channels to be included in "reduced data sets";

2.1.4 Glitch Investigations

The largest DetChar working group[17] carries out studies of interferometer noise transients, or "glitches". Composed of experimentalists and analysts, the working group has broad expertise, with its work closely coupled to the burst and inspiral searches.

The short-term goal of the Glitch Working Group is to identify the times of brief transients in the data taken during engineering and science runs that will affect the astrophysical searches. These transients make the data very non-stationary and non-Gaussian. Its long-term goal is to provide the information for experimentalists and builders of future detectors needed to achieve interferometer noise that *is* stationary and Gaussian.

More specifically, the goals of the Glitch group is to provide a statistical description of transients in the gravitational wave channel and in relevant auxiliary data channels; and identify possible correlations between transients in the auxiliary channels and in the gravitational wave channel, collaborating with the detector commissioners in the search for their cause.

These goals are pursued both online and offline. During the S6 science run, the Glitch Working Group reported weekly on found anomalies and investigations of them[19].

This rapid-feedback analysis is based on transients found in the gravitational wave channel and in auxiliary channels (e.g. KleineWelle and Omega triggers) and of the output of DMT monitors such as SenseMon. This was accomplished during S6 via weekly shifts of volunteers, weekly teleconferences, and through participation in scimon shifts at the observatories. In the offline analysis, as new data quality flags and event candidates are produced, the working group explores their correlation in order to establish which data quality flags and veto strategies are appropriate for burst and inspiral searches, taking into account the different needs of each search, but aiming at a consistent usage of vetos and data quality flags.

As mentioned in the introduction, although there will be no science runs in 2010-2011, we envision short time intervals (one to several days) dedicated to study the data quality of the different Advanced LIGO subsystems as they are installed. We expect this activity will start in 2011 with pre-stabilized laser, input mode cleaner, small chamber seismic isolation, small suspensions in L1 and large suspensions and arm cavities in H2. 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 artifacts in each of the building blocks of the very complex gravitational wave detectors.

The list of high priority activities for the Glitch group in the 2010-2011 are:

- Continue the monitoring of transients affecting data quality in S6 until it ends in October 2010, including
  - staffing of ”glitch shifts”
  - investigations of current and arising data quality problems
  - insertion of DQ flags for identified bad times
  - summarizing transients issues in S6 for a data quality paper
- Test current tools and develop tools as needed for finding glitch rates and outliers in each subsystem channel: this includes characterization of subsystems (e.g. trigger rates over a day) and check for good signals (above ADC background) of channels.
- Test proposed S6 strategies for automatic grouping of glitches on subsystem glitches.
- Investigate and decide which of DMT KleineWelle or LDAS Omega should be the official trigger generator used for auxiliary channels, adapt algorithms and infrastructure as needed.
- Improve current veto methods utilizing more information (i.e. expand “used percentage veto” [29] and hveto strategies for known bilinear couplings).
- Devise improved ways to diagnose problems arising from data acquisition, data sampling and/or imperfect timing in digital control systems.
- Choose methods and graphics from S6 glitch shifts/reports and create a hierarchy to make content easier to follow.

2.1.5 Spectral features Group

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 seismicity, high wind, acoustic noise, and electromagnetic interference. Some sources are natural, but many are 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, but unusual artifacts typically require detailed on-site investigations and eventually mitigation, work carried out by scientists from the Observatories and from several 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. The list of high priority activities related to characterizing spectral features in 2010-2011 are::

- Continue the monitoring of lines affecting data quality in S6 until it ends in October 2010, including summarizing frequency line issues in S6 for a data quality paper.
- *Noise budget for subsystems*: Measure the environment about subsystems to identify periodic signals so as to develop a catalog of potential noise lines that could enter the gravitational wave signal channel. 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 sub-systems. Use seismometers, accelerometers, microphones, magnetometers, voltage line monitors and other devices to map out noise, and how it couples into subsystems. Use existing line finding tools, such as Fscan (a pulsar search code, applied to auxiliary channels), and coherence (which calculates the coherence between the gravity wave channel and auxiliary channels).
- *Investigate coherence of environmental channels with the different subsystems*: Use the coherence tool to monitor the coherence between various signals. The intermediate pipeline (IM) 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; the IM pipeline, as well as bicoherence code, will be used to study up-conversion of noise in sub-systems.

2.1.6 Calibrations

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 DetChar group, although there are still many common members and activities.

Calibration of a detector is a complex task that involves instrumental hardware measurements, detector modeling, computer programs, and extensive validation and review. The Calibration Committee responsible for this essential work includes LIGO Laboratory and other LSC scientists, 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[8] available to the Collaboration, as well as recorded in the electronic logs, software repositories, and LIGO documents[9].

The calibration procedure has evolved in sophistication since the S1 run, most notably in automation, modeling, and redundant validation methods, with calibration provided both in the frequency domain (a frequency-indexed response function to be applied to the Fourier transform of the gravitational wave channel) and in the time-domain (a derived digital time series, "h(t)", representing strain as a function of time, which LIGO started to generate in S4)[10, 11]. Starting with the S6 run, the time domain data has been the main calibration product. The generation of the time domain data is complex enough a job that needed a dedicated team for calibration and another one for the review. The Calibration Committee is therefore co-chaired by a time-domain chair and an experimental chair. There are also some efforts to calibrate the detector data at higher frequencies, near the 4-km cavities' free spectral range at 37 kHz, where the detectors are, in principle, comparable in sensitivity to gravitational waves as in the baseband near 100 Hz.

Estimation and reduction of the errors in the calibration data products has been a major effort in recent years, and these investigations will continue. An alternative method of calibration using auxiliary laser pressure actuation ("photon calibrator") and interferometer laser frequency modulation developed and implemented in the S5 run, have also been used during S6. In S5 the various methods agreed to within 10%. In the S6 run, we have had calibrations by the coil calibration, the photon calibration, and other methods, with agreement at 10%-20% level. Understanding the origin of these differences is essential for the development of more accurate calibrations.

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 §2.1.7.

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.

As the necessity of the analysis using the data from multiple gravitational wave projects increases, so does the urgency to share the information about the calibration of various gravitational wave detectors transparently. There has been a very fruitful exchange of ideas and methods with the scientists performing the calibration from Virgo and GEO. Also important is an exchange of ideas about the review process. Though there hasn't been much communication between the calibration reviewers from different projects, it is desired that more communication channel is established by the end of the S6 run. 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.

The Calibration Committee's membership has been augmented in recent years by the graduate students and the scientists alike from several LSC institutions. The work load neccesary for the calibration of LIGO instruments increased drastically both in hardware- and software-related tasks since S1, and the participation of motivated persons from broad backgrounds proved highly successful and indeed indispensable for satisfying the goal of the Calibration and the Calibration Review Committee, i.e. the timely deliverly of the vetted calibration. In addition, for students this provide valuable instrumental training. It would be highly desirable to sustain this broad participation.

2.1.7 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 collaboration with the Calibration team (also see 2.1.6), the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator, 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 eliminate 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 remaining short term high priorities for S6 are:

- Organization of TSWG (coordination, training, supervision, publication, etc.)

- Timing verification studies and documentation of timing stability for the August-October section of S6
- Post S6 timing system snapshot and photon calibration based timing measurements
- Support of potential discoveries in the remaining of S6

Some of the longer term tasks in preparation of the advanced detector era:

- Further develop and test injection techniques to determine accurate timing through direct test mass excitations
- If necessary address emergencies and/or expand the capabilities of presently running data monitoring tools related to timing;
- Support S6 data analysis publications relying on timing performance
- Enhance the availability of timing diagnostics provided for various subsystems (e.g., LSC, OMC, etc.)
- Measure and document the timing performance of mission critical digital subsystems such as LSC and OMC DAQs
- 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

2.2 GEO Detector Characterization

2.2.1 Introduction

The GEO600 detector is currently undergoing an upgrade program called GEO-HF (LIGO-L1000195). This commissioning effort is mainly focused on improving the high frequency sensitivity of GEO600 and in the next year we will surpass all previous and current gravitational wave detectors in the region above 1 kHz. In two years, time we will have increased the high frequency sensitivity of GEO600 by a factor of 10. This still beats than the plans for Advanced LIGO without a signal recycling mirror. In not much longer than a year, we will be the only kilometer scale detector in the world that is in operation while the LIGO and VIRGO detectors decommission to install the second generation instruments.

The GEO-HF upgrade program consists of three major steps that will each have a significant effect on the shot noise limited region of the GEO600 sensitivity. The first step which began commissioning in April of this year is to inject a squeezed vacuum state into the output port of the interferometer. When we have obtained 6 dB of squeezing, this will have increased the high frequency sensitivity by a factor of two. The next step which will begin commissioning in November of this year is to replace the signal recycling mirror to increase the bandwidth. This will increase the high frequency sensitivity by another factor of two. The final step in the GEO-HF plan is to upgrade the laser and change the mode cleaners to supply GEO600 with 6.25 time more power than is currently available. This will increase the high frequency sensitivity by another factor of 2.5 bringing the total increase to a factor of 10 as stated. We expect this last step to take the longest. The power increase commissioning with begin with the new laser in January 2011 but we do not expect to be running at full power until somewhere in the middle of 2012.

Currently, even though we are at the heart of the commissioning time, we still operate GEO600 in a night and weekend astrowatch style mode. This will continue throughout the entire upgrade program with only a few months of breaks scattered over this period when large installations are happening. In two years, these upgrades and optimizations will start to slow down and we will be prepared to offer the community some extended science runs should the desire exist. These preparations call for an increased effort in the characterization of GEO600. In addition to this reason, it has become evident to the commissioning team that an increased detector characterization effort can greatly facilitate the effort to reach the designed sensitivity of the GEO-HF upgrade program.

2.2.2 Long-term Plan

Here we describe the long-term plan of detector characterization work at GEO600. This covers work that needs to be done through the end of the GEO-HF upgrade program and subsequent work in the data taking period that follows.

We view detector characterization work as having two different end goals: aiding commissioning and accessing data quality. The characterization work that is needed to accomplish these goals often has a large but not complete overlap. While active commissioning is going on at GEO600 we will be focusing most of our efforts on aiding commissioning. One of the strong benifits of working at GEO600 is that we have no restrictions about the boundries between commissioning activities and detector characterization activities. As such, we would like the detector characterization team to work intimately with the commissioners and even do some commissioning work themselves! As the commissioning work starts to die down and the GEO-HF upgrade nears completion we plan to switch the focus over to the data quality side of characterization. In the subsections that follow, we describe various tasks that we will be investigating and comment on how they can contribute to the two major goals of detector characterization.

State Space Model

Interferometric gravitational wave detectors require extremely complex and interconnected control schemes in order to keep the detector locked on its operating point and provide a quiet enough environment to search for possible gravitational wave signals. Because of this connectivity, unwanted signals often show up in many subsystems in close temporal proximity making it often difficult to locate the source of the problem. In addition to this, the control architecture can also create unexpected couplings in the detector that hurt the validity of analyzes that do not take them into account. In order to shed some light on these issues we will be working on creating a state space model of GEO600.

State space models are a way of representing physical systems by inputs, outputs, and state variables in phase space, connected by a set of first order equations of motion. For linear systems the equations of motion can be simplified into a matrix representation that allows easy time evolution modeling of the system. The inputs and outputs can be used to connect many different models and create control loops in a very complex way. With a completed model one can use the inputs to inject arbitrary signals and observe the time-domain response of the entire system. Such a model of GEO600 would be a very powerful tool that can help in a lot of the characterization tasks described below. A summer student has already started working on a very simple model containing only the Michelson loop for the purposes of verifying the relative calibration as described in Section **??**. We will continue work in the future by adding subsystems to create a more realistic model.

Lock Loss Studies

The usefulness of commissioning improvements that are aimed at improving the long term stability of the detector are often difficult to measure. One way of accomplishing this task is to have detailed statistics on many lock losses which pinpoint a possible cause for each loss of lock. Because of the complexity of the detector this is not easy to carry out and often times quick investigations aimed at a single lock loss prove to be unsuccessful at determining its cause. However, some figure of merit for stability improvements is needed to direct future commissioning efforts in a more efficient manner. Therefore, we deem deep investigations of lock losses a high priority to the commissioning of the GEO-HF upgrade to GEO600. This work will have the ultimate goal of providing an automatic tool for investigating all lock losses, providing flags for suspicious activity in various channels, and suggesting possible causes. Lock loss studies will be greatly facilitated by a state space model of GEO600 because it will give us the capability to verify these causes.

This tool will become essential when we start to commission the various power increases planned in GEO-HF. During these times the interferometer will become more unstable and the main goal following the hardware installations will be to counteract these instabilities with improvements in other systems of the interferometer. We will also gain from having an autonomous lock loss analyzer in the production of data quality flags for gravitational wave data analysis. It would enable characterization of the different types of lock losses that happen in a given interferometer configuration and then allow flags to be written which designate how much faulty data should be removed prior to a lock loss.

Noise Projections

Another very important aspect of detector characterization is understanding the various noise contributions to the sensitivity. This has been traditionally done at GEO600 and the other gravitational wave detectors around the world by injecting enough noise into a particular subsystem channel to dominate the sensitivity. This enables a transfer function to be measured from that channel into the gravitational wave channel $h(t)$. Assuming linearity, we take this to be the noise coupling when the injection is not done and use it to project the noise level of this channel during normal operation of the detector into $h(t)$. In the past we have also gone further and injected constant single lines in various subsystems, large enough to be seen with good SNR in $h(t)$, to provide a monitor for the noise coupling.

These studies are essential to commissioning as it allows us to pinpoint which subsystems are limiting the sensitivity of our detector telling us where to focus our efforts. The monitors also direct us in the process of debugging problems as they enable us to keep watch over these couplings. For data characterization the noise projections and monitors will play a large role in identifying which channels are important to watch and use for creating vetoes due to their influence on $h(t)$.

Although we have done a lot of this type of work in the past, GEO600 is currently going through rapid changes and we need to reconsider which channels to project from. We will also be doing more environmental injections in the future to measure their coupling into $h(t)$. One other area where we can improve on in the noise projections is an investigation of correlations between different subsystem channels that cause their combined coupling to be different from each channel separately. This study can make use of a mature state space model of GEO600 and possibly shed light on the reason behind these correlations.

Transient Studies

In addition to understanding and trying to improve the static sensitivity of GEO600, it is very important to pay attention to transient noise. For searches of gravitational wave bursts that do not have a defined template for the signal and only rely on coincidences between different detectors, these glitches ultimately provide the limit to the sensitivity of the search. In GEO600, we have mostly been focused on providing good high frequency sensitivity which is further emphasized with the current high frequency upgrade GEO-HF. Because of this, it turns out that GEO600 can only reasonably expect to have a chance at burst signals and signals originating from a small subset of continuous wave sources. Thus we have put, and will continue to put a very high priority on characterizing the transient noise in GEO600.

We currently have an on-line tool running called the Hierarchical Algorithm for Clusters and Ridges (HACR) that searches spectrograms of many channels for transient events. Coincidence analysis is routinely performed between many instrumental/environmental signals and $h(t)$. The results of this analysis are then automatically posted on detector summary pages which tell us which subsystems need to be studied further for repair, improvement, or data flagging.

In the future we will continue to use HACR as a tool for directing further commissioning or data characterization work but we are also looking to implement on-line analysis using the Ω-Pipeline. This pipeline is one of the standard analysis tools for creating burst search triggers and will provide us with a figure of merit that is familiar to analysts in the search groups.

Glitch studies will also greatly benefit by having a mature state space model of GEO600. The fact that the state space model enables simulation in the time-domain makes it particularly fitting for studies of transients in the control loops. This will allow us to both pinpoint and verify the source of the glitch by injection at a given point into the model and observing the response at different places.

Calibration

Calibration of an interferometric gravitational wave detector consists of two very different measurements. A relative calibration is carried out at GEO600 every second and measures the amplitude and time response of the detector as a function of frequency relative to one reference frequency. This calibration allows us to determine the shape of the $h(t)$ curve from the electrical output of the detector. Once we have this, we must pin down both the absolute amplitude the $h(t)$ timestream requiring a separate measurement.

Relative Calibration: The relative calibration of GEO600 is carried out by measuring the optical response of the detector every second with continuously injected lines. A script computes digital filters that invert this response to get back to test mass motion. To correct for the lower frequency region where the loop gain is high and the test masses are held fixed, we calculate the induced displacements arising from feedback loops and subtract them from the test mass motion. After applying the absolute calibration this signal is the calibrated strain output $h(t)$ of the detector.

This process was very heavily validated for the previous configuration of GEO600 but as we start to reconfigure throughout the GEO-HF upgrade, we will need to further validate the relative calibration process. In addition to this we have never carried out a the validation process to frequencies greater than 2 kHz which will be important for high frequency searches. Since the feedback signal is subtracted from the detector output to create $h(t)$, it is important to validate the relative calibration process with an injection input that is not seen by the main Michelson loop. We have recently added an input to do this with the standard test mass actuators and we also have a photon pressure actuator to use for this purpose. Future studies will investigate both of these methods for this validation. We also intend to use the state space model as an aid to understanding already visible systematic errors in the calibration at high frequencies.

Absolute Calibration: Since GEO600 will be mostly operating in a single detector era after most of the GEO-HF upgrades are finished, the absolute calibration does not play a huge role in the ability to find gravitational waves in the data. Thus it should be sufficient to provide a calibrated strain output to within 10% in amplitude and 10° in phase over the measurement band.

We currently have three methods of absolute calibration at GEO600. The traditional method uses the length of the first mode cleaner as a length reference. However, this method requires a chain of many separate measurements that assume a wide range of linear behavior at each step and is difficult to reproduce.

Recently we have investigated a new method of calibrating which uses a very similar technique to the free-swinging Michelson calibration used in the LIGO detectors. This method ultimately uses the laser wavelength as the length standard and requires only a single measurement. We have shown that the free-swinging Michelson calibration results are very stable and decided that this will be the fiducial method for GEO600 absolute calibration. This method is so simple that it can be done every time the detector looses lock but this will take some work on automating the process.

The last type of calibration is to use the photon pressure actuator to put a known motion on the test mass. This method has been employed at GEO600 for some time now and some of the important details about such actuators were worked out by GEO commissioners. However, the hardware was ill maintained over the last few years and we have also discovered a problem with the way we were sensing the actuator laser power. We have started implementing a solution to this but there is still work to be done to finish the setup and use it to provide an absolute calibration result.

2.2.3 Short-term Plan

We now discuss the plan for detector characterization of GEO600 over the next year. Currently we do not have the staff to carry out most of the activities described above. The team that works on GEO600 consists mainly of the handful of commissioners on site and a very small group of people doing detector characterization to aid commissioning at the University of Glasgow. We are planning on building a larger group of people, both on site and distributed at the Albert Einstein Institute in Hannover and at Cardiff University, and have the funds to do so. In this regard this paper should be seen as a call to action! With no new members to the GEO team for detector characterization we will be able to continue calibration work by manually carrying out free swinging Michelson measurements (Section **??**) and validations of the relative calibration using the out-of-loop input to the standard test mass actuators mentioned in Section **??**. We will also be able to continue doing noise projections (Section **??**) in an improved manner with a better channel selection but no added environmental injections. Transient studies as described in Section **??** will be constrained to keeping HACR running and monitoring its output for the purpose of commissioning.

With only a small number of added members on site, we can start working on the other areas mentioned above. The state space model (Section **??**) of GEO600 will continue slow progess in parallel to other more pressing tasks. With the ability to increase the input laser power coming soon, lock loss studies will become more and more pressing. This will be one of the first priorities of added man power to the detector characterization effort at GEO600. Minimal effort will be put on analyzing small handfuls of glitches until the state space model becomes more mature or we have a larger number of added members.

2.2.4 Opportunities for LIGO-VIRGO Members

Due to the upcoming period where the LIGO and VIRGO interferometers are being decommissioned, we would like to invite members of the communities involved in their characterization to play some role in the characterization of GEO600! This will be particularly applicable to the advanced detector era in two areas. First and most exciting is that GEO600 will be operating with an injected squeezed vacuum soon. This is an upgrade that is currently being discussed for Advanced LIGO. This decision to some effect rests on the sucess of the implementation at GEO600 and its characteristics. We are in a pivotal time now for detector characterization teams where this new enhancement is for the first time being added to one of the large gravitational wave detectors in the world. All of the characterization tasks described in Section 2.2.2 we need to be studied for the new squeezed light source. The second area where GEO600 characterization can play an important role in the advanced detectors is in working with a dual recycled interferometer. We have been operating with a signal recycling mirror almost from the very beginning. Studies invovling looking into noise couplings and glitches that result in a deeper understanding of the relationship between the signal recycling mirror and $h(t)$ will be valuable information for the advanced detectors when they start to implement dual recycling.

On the logistical side of carrying out detector characterization work far off site of GEO600 we would like organize this in a more compartmentalized manner. We feel it would best work out if most of the global interferometer work is done on site or by people who have a strong connection to the on site staff and are willing to spend a lot of time on site. However, characterization work involving individual subsystems is something that we think can easily be tackled by far off site participants. We envision a short two week long visit should be enough to give a quick introduction to the subsystem and then people can carry out investigations remotely. The on site characterization team will be dedicated to interacting with these subsystem teams as questions arise.

2.3 Virgo Detector Characterization

3 Searches for signals from compact binary coalescence

3.1 Science goals

The inspiral and merger of a compact binary system generates gravitational waves which sweep upward in frequency and amplitude through the sensitive band of the Earth-based detectors [30]. The detection of gravitational waves from these astrophysical sources will provide a great deal of information about strong field gravity, dense matter, and the populations of neutron stars and black holes in the Universe. The scientific program of the LSC/Virgo Compact Binary Coalescence (CBC) Group is designed to identify GW signals from compact binary sources in the detector data, estimate the waveform parameters with confidence in the correctness and validity of the results, and use these signals for the study of gravity and the astrophysics of the sources [30].

The immediate goal of the CBC group is to make the first detections of gravitational waves from compact binary systems with data from the LIGO and Virgo detectors, through computational methods that are as close as possible to optimal (that is, making full use of the detectors' sensitivities), and to develop robust and unbiased methods to gain confidence that the detections are not false alarms due to detector noise fluctuations. Our primary tools are event triggers generated by matched filtering, and coincidence between triggers from different detectors.

The detection of gravitational waves from compact binary systems will bring a great deal of information about the component objects and populations of neutron stars and black holes in the Universe. The CBC group has identified the following topics that will be important group tasks:

- Estimate the rate of compact binary coalescences in the Universe by direct observation of gravitational waves. In the event of a detection, this will take the form of a rate interval; in the absence of a detection, it can provide rate upper bounds [31, 32, 33] as a function of system parameters (masses, etc.).

- Associate gravitational waves from binary coalescence with coincident observations using other astronomical detectors, including radio, optical, x-ray and gamma ray telescopes, and low-energy and high energy neutrino detectors. Implement prompt methods for the flow of external triggers both to and from gravitational wave observatories to extract the most science out of these associations. In the absence of a gravitational waves signal at the time of an electromagnetic or neutrino trigger, use this information to exclude certain progenitor models, as was done for GRB 070201 [34].

- Establish the relationship between gamma ray bursts and compact binary coalescences involving neutron stars [34, 93].

- Measure the masses and spins of detected binary signals and develop a catalog of binaries from which further understanding of populations can be discerned.

- Measure the inclination, polarization, sky location, and distance as allowed by the use of multiple observatories, higher harmonic content and/or spin modulation.

- Determine the precise waveform and energy content of gravitational waves during the merger of binary black hole systems.

- Use consistency of parameters determined from different phases (inspiral, merger and ringdown) to test strong field predictions of analytical and numerical relativity.

- Probe the disruption of neutron stars during binary merger and thereby determine the equation of state of neutron stars.

- Test post-Newtonian theory [35] and alternative theories of gravity, such as scalar-tensor theories which can result in modified phasing of the gravitational waves from binary inspiral [36].
- Bound the mass of the graviton by direct observation of the gravitational waves from binary inspiral [36].
- In the case of high-mass ratio binaries, develop and implement methods to map the spacetime structure of the more massive object by observing the gravitational waves [37, 38].

In the remainder of this section, we will lay out in detail the strategies which are being pursued in order to achieve these goals. We begin with a discussion of the gravitational waves emitted during binary coalescence, then describe the search strategies employed by the group. We then provide a brief recap of the search results to date before outlining the future directions of the group.

3.2 Gravitational waves from the coalescence of compact binary systems

Compact binary systems which generate gravitational waves with frequencies above $\sim$ 10 Hz are prime candidates for gravitational-wave astronomy. These systems include binaries with masses as low as $1M_{\odot}$ up to several thousands of solar masses, composed of black holes, neutron stars, and perhaps other exotic compact objects with densities similar to or greater than neutron stars. Traditionally, the gravitational waveform from binary coalescence has been split into three parts: inspiral, merger and ringdown. Depending upon the masses of the binary's components, different parts of the waveform will lie within the detector's sensitive band. The beauty of these systems for data analysis is that the inspiral phase of the waveforms (and for binary black holes, the entire waveform) can be computed within the context of General Relativity and/or alternative theories of gravity (e.g., [39, 40]).

When the components of the binary are widely separated, the gravitational waves from these objects sweep upward in frequency and amplitude as loss of energy to gravitational waves causes the binary orbits to shrink, thus reducing the period of the orbit. This process is called the inspiral phase of the binary evolution. From an astrophysical perspective, these compact binary systems are expected to be clean with the inspiral dynamics controlled primarily by pure gravitational effects. For this reason, theoretical waveforms calculated within the post-Newtonian framework [39] should provide accurate representations of the gravitational waves if the calculations can be carried out to high enough accuracy. At present, there is good evidence that post-Newtonian waveforms provide an accurate representation at frequencies below $\approx 700(2.8M_{\odot}/M_{\rm total})$ Hz. Various estimates suggest that inspiral waves should extend to frequencies $\approx 1500(2.8M_{\odot}/M_{\rm total})$ Hz (the innermost stable circular orbit, or ISCO) or even higher. Since the band of optimal sensitivity for these sweeping signals lies in the range $40-800$ Hz for the current LIGO detectors, post-Newtonian waveforms should be adequate to detect inspirals of systems with total mass as high as $20-30M_{\odot}$.

Waveform models that are accurate to a fraction of a cycle over many cycles — more than 1600 cycles over 25 seconds for neutron star binaries between 40 and 1500 Hz — allow optimal integration of the signal and thus optimal detection sensitivity. Advanced detectors will extend to lower frequencies and thus much longer signal durations, in the regime where post-Newtonian waveforms are quite reliable. As the total mass of the binary system increases, the merger phase moves to lower frequencies, and the waveforms spend less time and fewer cycles in the most sensitive frequency band of the LIGO and Virgo detectors.

When a binary system with total mass above 1.4–$2.0M_{\odot}$ merges, it is likely that the end product is a single, perturbed black hole which rings down, by emitting gravitational radiation, to a stationary configuration. The ringdown waves are primarily emitted in the quadrupolar quasi-normal mode of the black hole; the frequency and quality factor of this mode depend on the mass and spin of the final black hole (for a recent review see, [41]). The ringdown waves will be in the detector's sensitive band for black holes of masses above $\sim 100 M_{\odot}$, up to several thousand solar masses for advanced detectors [42]. Observation of these waves will enhance our ability to measure parameters of the binary and to test numerical and analytical models of the merger phase. Moreover, ringdown waves can be detectable even if the inspiral waves do not enter the LIGO/Virgo band (for systems with high total mass). Perturbation theory provides waveforms emitted during this settling phase, the ringdown phase. These waveforms depend on both the mass and spin of the final black hole [41]. Higher harmonics will be present in the ringdown, which can be exploited to improve detection sensitivity and confidence, and as a sensitive test of black hole perturbation theory.

Recently, compact binary coalescence waveforms, including a number of cycles from the inspiral phase and the merger and ringdown phases, have been obtained in numerical relativity for a subset of astrophysically interesting mass-ratios, spin and orbital angular momentum values [43, 44, 45, 46]. Analytical models, including the effective one body (EOB) [47, 48, 143, 50] framework, and hybrid or phenomenological approaches [51], match the numerical results extremely well for the case of non-spinning binaries with comparable component masses and provide, for the first time, a template which covers the entire coalescence waveform. There is much work in progress to extend these models to include spin, at least in restricted regions of the parameter space. Phenomenological models that include the effects of components with spins aligned with the orbital angular momentum (and thus no precession of the orbital plane) have recently become available [91, 92]. These full waveforms allow for a coherent search over the complete coalescence signal, thereby enhancing the ability to both detect these signals and to accurately extract the parameter values. There have also been successes modelling the merger of neutron star binaries and neutron star black hole binaries [52, 53, 54]. These results will also allow for improved sensitivity of our searches.

The inspiral-merger-ringdown waveform models tuned the numerical relativity simulations are still under active development as numerical simulations improve and improved techniques of matching numerical and post-Newtonian waveforms are developed. This is true even in the best-understood case of spinless systems with mass ratios less than 10:1. For example, between the first and second generation of phenomenological and EOBNR waveforms, the differences in the predicted SNR can be as large as 50% or more (resulting in uncertainties of order 1.5^3 in the achievable limits). It seems likely that the first generation of waveforms covering any new region of parameter space will have similar uncertainties which will decrease over time as simulations and our understanding of them improve.

The gravitational waveforms received at a given detector depend on the location of the source, the orientation of its orbital plane, the polarization angle of the gravitational waves, the masses and spins of the component objects. For a single detector, many of these parameters are degenerate either with the constant phase offset of the signal or the distance to the source. The spins, however, can produce significant differences in the waveforms allowing their direct measurement. Most notable is the amplitude modulation of the inspiral waveform due to spin-orbit-coupling-induced precession of the orbital plane, interacting with the detectors' antenna pattern. There is also spin-orbit-coupling-induced phase modulation. These effects allow the spins to be extracted from the waveforms, but they also complicate the detection problem, since it is not yet computationally feasible to cover the full parameter space with a template bank [55, 56, 57]. This is an area of ongoing development and research [58, 59].

Eccentric orbits will also modify the gravitational waveform [60, 61]. By the time the waves are high enough in frequency to enter the sensitive band of ground-based detectors, the orbit is expected to be effectively circularized through radiation back-reaction, with negligible eccentricity. This prediction can be tested with observed waveforms, assuming that anomalous eccentricity does not significantly reduce the detectability. Similar considerations apply for anomalous waveforms due to soft neutron star equations of state, alternative theories of gravity such as scalar-tensor theories, massive gravitons, etc.

3.3 Rates and Astrophysical Populations

Astrophysical estimates for compact-binary coalescence rates depend on a number of assumptions and unknown model parameters, and are still uncertain. New papers in the field appear regularly, as better theoretical understanding allows more sophisticated models to be built, while additional electromagnetic observations of binaries with compact objects (pulsars, X-ray binaries, and short GRBs) provide tighter constraints on those models. The CBC group has published a document that summarizes the predictions in the literature as of 2009 [80]. However, this is meant to be a living document, which the authors will maintain in order to keep the information current.

There are two distinct methods for estimating NS-NS merger rates. The most confident among these estimates are the rate predictions for coalescing binary neutron stars which are based on extrapolations from observed binary pulsars in our Galaxy. The second method is based on population-synthesis codes, in which some of the unknown model parameters are constrained by observations and others are constrained by theoretical considerations. These estimates yield a likely coalescence rate of 100 Myr^{-1} per Milky Way Equivalent Galaxy (MWEG), although the rate could plausibly range from 1 Myr^{-1} MWEG^{-1} to 1000 Myr^{-1} MWEG^{-1}.

In the simplest models, the coalescence rates are assumed to be proportional to the stellar birth rate in nearby spiral galaxies, which can be estimated from their blue luminosity. We therefore express the coalescence rates per unit L_{10} (i.e., 10^{10} times the Solar blue-light luminosity), using the conversion factor of 1.7 L_{10}/MWEG [82]. Blue-light luminosity is a reasonable indication of the current star-formation rate in spiral galaxies, but it does not accurately track star-formation rates in the past; in particular, it ignores the contribution of older populations in elliptical galaxies. As the sensitivity reach of the detectors increases, it is also natural to express these results per Mpc^3, with a conversion factor of $0.02 L_{10}/\mathrm{Mpc}$. This yields a likely rate of 1 Myr^{-1} Mpc^{-3}, plausibly ranging from 0.01 to 10 Myr^{-1} Mpc^{-3}.

Due of the lack of observations of coalescing compact-object binaries containing black holes, NS-BH and BH-Bh rates can only be based on predictions from population-synthesis models. There are two distinct scenarios for the formation of double black-hole binaries close enough to coalesce through gravitational-wave emission. The first is the isolated binary-evolution scenario, which is expected to be the dominant scenario for NS-NS and NS-BH systems described above. The second scenario, which can be significant for BH-BH systems because of their higher mass, is the dynamical-formation scenario, in which dynamical interactions in dense stellar environments play a significant role in forming and/or hardening the black-hole binary before coalescence driven by radiation reaction.

The distributions of masses, mass ratios, and spin in NS-BH and BH-BH systems close to merger are unknown, and can only be estimated using population-synthesis codes which have large uncertainties. Consequently, in searching for gravitational waves from these systems, we endeavor to avoid bias by covering the entire parameter space to which we are sensitive, regardless of whether current models disfavor some regions. We evaluate our sensitivity and observed constraints (upper limits or confidence intervals) in terms of the rate per unit volume (in units of Mpc^3), as a function of the masses and spins of the binary system. Such sensitivity statements are computationally limited by the number of simulations (software injections) we can perform in each region of parameter (mass and spin) space.

3.4 Search Strategies

We aim to develop and employ search pipelines that

- effectively search over the full parameter space of binary systems to which the LIGO and Virgo detectors are sensitive, including all phases (inspiral, merger and ringdown) and component spins;

- perform a prompt search of the data, in order to quickly identify potential detections and to enable follow-ups using ground- and space-based telescopes;
- use all available data quality information for deciding what data to analyze and what data needs to be vetoed;
- promptly identify periods when the data has poor quality for analysis purposes, for causes not previously identified;
- promptly evaluate our confidence in candidate detections by understanding the detector properties, data quality and background trigger rate near the time of a candidate;
- promptly estimate the parameters of a candidate detection, in order to improve the detection confidence and to locate the source in the sky;
- make optimal use of input from other astrophysical observations, such as GRBs, to improve the sensitivity of our searches and to extract as much information as possible about the astrophysical source;
- employ robust methods to constrain the astrophysical rate of compact binary coalescences in the Universe.

The CBC group currently makes use of two different pipelines to perform the tasks listed above. The Virgo developed MBTA analysis [62] is designed to be an online, low latency trigger generator, and it is used as such. A second analysis, based upon the LSC developed pipeline [63, 33], was developed as a batch mode analysis and is run offline to analyze blocks of data (currently of a week's length), and perform a detailed estimate of the noise background, the astrophysical sensitivity and significance of any event candidates. In the remainder of this section, we describe the components of the analysis procedure in greater detail.

3.4.1 Matched filtering and Signal Consistency Tests

There is a well developed theory and practice to search for signals with a known waveform buried inside a noisy time series [64]. For Gaussian noise with a known additive signal, this theory leads to the matched filter. In gravitational-wave astronomy, as in many other fields which use matched filtering, the signal is not known exactly. For compact binaries, the inspiral signal depends on many unknown parameters as follows:

1. The ending time t_0 and the ending phase Φ_0 of the inspiral waves are unknown in advance. Physically the first can be thought of as the time when gradual inspiral ends and the merger begins; similarly the phase is the angle around the orbit when this transition occurs.
2. The gravitational waves also depend on the masses m_1 and m_2 of the compact objects and their spins $\vec{s}_1$ and $\vec{s}_2$. These parameters have strong effects on the evolution of the system's orbital frequency (and hence the phase of the signal) with time. They also appear in a variety of combinations in the amplitude part of the signal.
3. The amplitude of the waveform measured in a given detector also depends on a combination of the right ascension α and declination δ of the source, the inclination ι, the polarization angle of the waves, and the distance to the source.

We generate *triggers* by filtering the data from each detector with matched filters designed to detect the expected signals. At the single interferometer level, for non-spinning templates the angles can all be absorbed into the amplitude of the source giving an effective distance which is larger than the physical distance. Each trigger has an associated *signal-to-noise ratio* (SNR), coalescence time t_0, ending phase Φ_0 and parameters from the template that matched the data, such as the masses of the individual stars.

Since compact binaries with slightly different masses and spins would produce slightly different waveforms, we construct a *bank* of templates with different parameters such that the loss of SNR due to the mismatch of the true waveform from that of the best fitting waveform in the bank is less than 3–5% [65]. The template banks for the mass-space are well in hand [66, 67] for searches using Post-Newtonian approximations; work is ongoing to include spins in the most efficient manner [68].

Although a threshold on the matched filter output ρ would be the optimal detection criterion for an accurately-known inspiral waveform in the case of stationary, Gaussian noise, the character of the real data is known to be neither stationary nor Gaussian. Indeed, many classes of transient instrumental artifacts have been categorized, some of which produce copious numbers of spurious, large SNR events. When the origin of the instrumental artifacts is known, through understood coupling mechanisms or correlations of auxiliary data and the gravitational wave channel, the times are *vetoed*. However, many data transients remain that are not understood and produce an excess of false alarms. In order to reduce the number of spurious event triggers, we adopted a now-standard χ^2 requirement when using physical waveforms [69]. Instrumental artifacts tend to produce very large χ^2 values and can be rejected by requiring χ^2 to be less than some reasonable threshold. This test has proved to be one of the most powerful methods of dealing with noise glitches for the CBC group. Another very powerful tool for discriminating signals from noise glitches is the requirement of a coincident signal in more than one detector in the network. We require the signal to be consistent in both time (accounting for the light travel time between detectors) and physical parameters, such as mass [70]. Other signal dependent discriminators are being used in current analyses [71, 72], and many others are being explored. Fully coherent methods that combine information from multiple detectors, extracting or making use of source sky position [94], can be used to discriminate signal and background. General methods that employ multivariate statistical classification are under development to make more efficient use of information information about each event to discriminate signal and background.

3.4.2 Low latency pipeline

We are also using *low-latency* pipelines to generate triggers in close to real time. These are used for detector characterization, including near-realtime feedback to the control room via the InspiralMon figure of merit. Low-latency coincident triggers will be used, in conjunction with fast source localization algorithms, to generate external triggers for pointing EM telescopes. The low latency search will focus on the low mass range: total mass between 2 $M_\odot$ and 35 $M_\odot$ and component masses greater than 1 $M_\odot$. The baseline low-latency pipeline is based on the MBTA (Multi-Band Template Analysis) package[62] developed by the Virgo Collaboration.

MBTA splits the matched filtering of the data into two frequency bands for signal processing efficiency and then combines coherently the results to extract the full signal to noise ratio. Second order post-Newtonian templates in the time domain are used. MBTA includes adaptive mechanisms to follow detector non-stationarities, which are necessary to run online. Other features of the search include:

- Using a reduced and fast version of signal-based vetoes to eliminate some of the background.
- Making use of the data quality flags and veto information produced with low latency to further reduce instrumental artifacts.
- Extracting triggers detected in coincidence in several detectors.
- Monitoring the current background level of the detectors to estimate the false alarm rate of coincident triggers.

- Interfacing the pipeline output with source localization algorithms, an event archiving system and the alert procedure.
- Establishing parallel instances of the pipeline, some with software injections in order to monitor the efficiency of the search.

Another pipeline, called gstlal, is under development. This is based on the *gstreamer* open source multimedia framework [95], with CBC-specific plug-ins based on code in the LIGO Algorithm Library (LAL [96]). It features efficient multi-banding, an efficient numerical SVD-based decomposition of the template bank, fast and conditional computation of signal-based χ^2 quantities, and the ability to run continuously in "streaming" mode.

3.4.3 Batch Analysis pipeline

While matched filtering is the core analysis method used by the CBC group to search for these signals, it is only part of the complete detection pipeline which has been developed over the past five years [63, 33]. The current pipeline, implemented in the publicly available LIGO Algorithm Library suite [96], has the following steps:

1. Determine which data satisfies a minimal set of data quality cuts determined by operating characteristics of the instrument and bulk properties of the data. For example, we require (obviously) that the instrument function be flagged as nominal by the operators and scientific monitors and require that there are no flagged malfunctions in data acquisition ("Science mode"). All data satisfying these minimal data quality criteria are analyzed for signals.
2. Perform a matched filtering analysis on these data. The SNR $\rho(t)$ is computed for each template in the bank. Whenever $\rho(t)$ exceeds a threshold ρ^*, the local maximum of $\rho(t)$ is recorded as a *trigger*. Each trigger is represented by a vector of values: the masses and spins which define the template, the maximum value of ρ, the inferred coalescence time, the effective distance D_{eff} (derived from the trigger SNR), and the coalescence phase. These are inspiral-level-1 triggers.
3. Triggers generated at the single interferometer level are then compared between all instruments that were operating nominally. Any triggers identified as coincident, in time and other parameters [70], between at least two instruments are kept and recorded as coincidence-level-1 triggers.
4. Surviving coincident triggers are then used to define a smaller template bank for a second stage of single detector data filtering, wherein more computationally expensive quantities, such as the χ^2 and other signal-based vetoes, are computed. When a new trigger is found satisfying the signal-based vetoes, an inspiral-level-2 trigger is generated.
5. The inspiral-level-2 triggers are again subjected to the coincidence requirement. In addition, triggers which occurred in times of known poor data quality are flagged and removed. The coincidence-level-2 triggers are then recorded.
6. These surviving coincident triggers are used to compute a combined detection statistic from the signal-to-noise ratios and χ^2's of the single detector triggers comprising the coincident trigger. All the relevant quantities for each coincident trigger are recorded for further analysis.
7. The pipeline incorporates methods for estimating background from accidental coincidence of noise triggers (Sec. 3.4.4) and for measuring the pipeline efficiency across the parameter space including distance to the source (Sec. 3.4.6).

This pipeline was developed and refined over the S3 and S4 analyses [33] and has been used in the S5 analyses [73, 74]. The pipeline is designed and implemented to be both flexible and extensible. This pipeline has been described in the context of only the inspiral phase of compact binary evolution. It has been designed, however, to allow the easy inclusion of filtering techniques for the merger and ringdown phases. The first end-to-end ringdown search [42] (in S4 data) has been published, and searches in S5 data for the combined inspiral-merger phase and for the ringdown phase using the full pipeline are currently in progress.

3.4.4 Background estimation

The nature of gravitational-wave detectors makes it impossible to go off source to estimate the background in a single instrument. The CBC group requires coincidence between triggers from two or more detectors in time and template parameters (masses, etc); the dominant background is thus accidental coincidences of noise triggers from detectors with uncorrelated noise. The rate of such background coincident triggers can be estimated by time-sliding the triggers from one detector with respect to the other, by amounts that are long compared to the timescales on which autocorrelations between triggers are significant, but short in comparison to the detector non-stationarity. This method applies an artificial time slide to triggers from each detector and carries out all of the later stages of the analysis pipeline in exactly the same way as for the original data. This provides an estimate of the rate of coincident triggers satisfying all criteria used in the pipeline, but known to be false alarms. This method fails to account for noise triggers that are correlated between detectors due to some external terrestrial disturbance; eg, the H1 and H2 detectors exhibit such noise trigger correlations and it is thus difficult to estimate the background for events that are coincident in H1H2 only. The automated CBC pipeline typically uses 100 time-slides to estimate the background. In order to obtain false alarm rates of less than 1% for the loudest observed events, more time slides will need to be performed, and/or different methods must be applied, such as estimating accidental coincidence from single detector trigger rates (requiring significant modifications to our current pipeline); work along these lines is in progress. For externally triggered searches (eg, for short-hard GRBs [34, 93]) in which the time of the external trigger, and thus of the expected GW signal, is known in advance, the background can be estimated by looking at the data in nearby time intervals (“off-source”), sufficiently well separated from the external trigger time. The time-slide method can also be used in this and other searches for transients.

3.4.5 Instrumental and environmental vetoes

The complicated nature of interferometric gravitational-wave detectors means that instrumental and environmental effects produce non-gravitational-wave triggers in the detector output. To combat this problem, the CBC group uses vetoes based on a large number of different approaches to the data. Working closely with the DetChar, Glitch and Burst groups, the CBC group has adopted the convention to divide these vetoes into different categories depending on the degree to which the instrumental or environmental disturbance is understood. For example, if data quality information indicates that large transients could be expected in the data due to an instrumental malfunction or strong environmental disturbances, this would strictly veto a trigger. Similarly, triggers which are associated with a subsytstem malfunction as identified by analysis of auxiliary channels would provide a strict veto if the path from the sub-system to the gravitational wave channel is understood. Other categories of vetoes provide weaker evidence that something was wrong with the instrument or the environment and are used to flag triggers as less likely to be of gravitational-wave origin. For blind analyses, vetoes are identified by reference to the set of level-1 single detector inspiral triggers, to coincident triggers identified as false alarms (from the time-slide analysis), and to coincident triggers found in a subset of about 10% of the data distributed uniformly over the run, called the *playground*. Also, the triggers studied are limited to large signal-noise triggers clustered over a window much broader than the expected resolution of a real signal. Identification and confident use of instrumental and environmental vetoes based on analysis of auxiliary data channels is the subject of continued investigation. In coordination with the DetChar and Burst groups, many of these veto identification procedures have been automated for S6/VSR2. However, S6/VSR2 has brought a variety of new challenges to the detector characterization effort (see the DetChar chapter). As a result, during S6 it still requires some weeks of study after the data are collected before we can arrive at veto definitions with confidence.

3.4.6 Efficiency evaluation

The efficiency of the analysis pipeline described above for detecting gravitational waves from binary inspirals in LIGO and Virgo data can be evaluated by injecting many thousands of theoretically predicted waveforms into the data streams and identifying whether they are found by the pipeline with signal-to-noise above some relevant threshold. We aim to cover the full parameter space of such systems, including the broadest range of component masses and physical distances to which the detectors can be sensitive. The injected waveforms are usually generated "on the fly", although it is possible to read in waveforms from external files. They can be used for (a) testing the analysis pipeline code; (b) tuning various pipeline parameters and thresholds; (c) studying the effect of various systematic errors such as calibration uncertainties; and most importantly, (d) evaluating the efficiency and the cumulative source luminosity to which the search is sensitive, to establish astrophysical rate upper limits or confidence intervals (as described below). In addition, waveforms are injected directly into the detector test mass positions via the control system (*hardware injections*). These are used as an additional test of the pipeline response, as a test of the safety of instrumental vetoes, and as a diagnostic on our understanding of the calibrated detector response.

Theoretical waveforms such as post-Newtonian approximations have limitations in their domains of validity. This does not invalidate the use of these waveforms in searches, it simply reduces the sensitivity of the search relative to filtering with the exact physical waveform. The CBC group continues to follow the theoretical literature on computing waveforms from compact binary inspiral, merger and ringdown, including the new waveforms produced by the numerical relativity community. As new waveform approximations become available, they are coded to allow simulated injections into the data stream. These simulations help identify weaknesses in the current search techniques and determine the urgency with which the new information should be incorporated into the search pipeline.

For higher mass systems the primary gravitational-wave signal accessible to ground-based detectors is from the last stages of inspiral, merger, and ringdown. In the past year, fully phase coherent inspiral-merger-ringdown waveform models have been developed through detailed comparisons with numerical simulations, that continuously span the paramater space for non-spinning systems with mass ratios up to 10:1 [51] or higher [143]. More recently, waveform families have been developed which also incorporate aligned spins [91, 92]. These waveform families are suitable as both search templates and as software injections to test the pipeline and evaluate the efficiency. The S5 high mass search [97] has used the waveforms from [48, 143] as search templates, and waveforms from [143, 51], and [91] in order to evaluate the efficiency across the parameter space. The S6 high mass search is currently using the same waveforms. Currently, there is an effort to implement IMR waveforms with systems for generic spins.

3.4.7 Pipeline of pipelines

The analysis described above consists of many steps: selecting the data; filtering through the template bank and generating coincident triggers; evaluating the background with time slides; evaluating the efficiency with many sets of software injections; following up on detection candidates; evaluating the efficiency and computing upper limits; generating a rather large number of plots and tables that summarize the results and diagnose problems; and characterizing the detector data to establish data quality and identify vetoes.

This entire process consists of many sub-pipelines, and the entire process is repeated for every large data-taking interval. To make it easier for different people to repeat this process reliably and reproducibly, the CBC group has assembled a "pipeline of pielines" which we refer to as *ihope*. The *ihope* infrastructure has now become the standard way to run all batch mode CBC analyses. Many group members are involved in running this program, thereby developing the knowledge and expertise to further improve and develop its capabilities.

3.4.8 Externally triggered searches

Compact binary coalescence is expected to produce other observable signatures in addition to gravitational radiation. Most notably, short GRBs are thought to arise from binary neutron star coalescence [76], although binary coalescence may also be accompanied by optical, radio, x-ray and neutrino signals. By making use of the astronomical information derived from other "messengers", we perform deep searches for gravitational waves from these sources. The search shares many features with the all-time, all sky, full parameter space searches described above. However, limiting the time, sky location and parameter space of the search based on external information improves the sensitivity of the search.

In order to take full advantage of the external trigger, it is necessary to translate the trigger information into expectations for the gravitational wave signal. It is straightforward to incorporate a known sky location by requiring the observed time-delay of the signal between instruments to be consistent with the known sky location. The observed time of the event, as produced by other astronomical instruments, must be translated into the time at which we expect an associated gravitational wave would reach earth. There are uncertainties in this time due to instrumental effects and unknown astrophysics in the engines which might be generating the gravitational waves and other detectable signatures. These systematics govern the choice of the time window around the external trigger which is searched. For short GRBs during S5 [93], we have chosen to analyze a six second window, five seconds before and one second after the observed GRB time.

For a time-restricted search, we can make use of off-source times near the trigger, in addition to time-slides, to determine a background estimate at the time of the trigger. This allows us to account for correlated noise that might be present near the time of the trigger, something that is more difficult in the case of all-time searches. Furthermore, the short analysis time allows for a lowering of the SNR thresholds in the analysis, thereby making the search more sensitive.

3.4.9 Blind analyses

The CBC group pursues "blind" searches for GW signals in the LIGO-Virgo data. We wish to avoid injecting bias into our analysis procedures, which could unduly "elevate" the significance of a candidate event, or "kill" an event which we believe to be background. We wish to quantitatively evaluate the significance of a potential detection by making use of all of the information available about the event in question, the expected background, and the anticipated signals. To avoid bias, we want to fully establish the criteria by which we make that quantitative evaluation *before* we look at a specific detection candidate event. At present, this is difficult, because of the non-stationarity and glitchiness of the LIGO and Virgo detectors, and because we have not yet determined what all the best tools are for making a quantitative evaluation. So, for example, we have a detection candidate "follow-up" procedure that still involves subjective steps like looking at Qscans.

Our analysis pipeline, as described in the preceding sections, consists of a sequence of steps which are governed by a small number of tunable pipeline parameters (thresholds, windows, cuts). These parameters are chosen to suppress detector background and maximize detection efficiency. Keeping the parameters general and their number small helps to avoid the over-tuning that can unduly elevate or kill individual events. To keep our analysis pipeline "blind", we choose our tunable pipeline parameters based only on the background estimation (from time-slide or off-source coincident events) and the detection efficiency (from software injections). Additionally, to verify the correctness of the analysis, we produce a large number of sanity checks — plots and tables of numbers that give us confidence in all steps of our analysis procedure, especially the proper values of the tunable parameters.

Only when all sanity checks are made and found to be satisfactory do we "open the box" by looking at the in-time or on-source coincident events, and rank-order them by detection statistic. Detection candidates consist of those events whose detection statistic value is inconsistent with the background distribution; for example, the false alarm probability is much less than one.

3.4.10 Follow-up of candidate events

The most significant detection candidate events are flagged for greater scrutiny by the follow-up procedure, which still has subjective elements and thus is still subject to bias. However, we endeavor to minimize bias as much as possible in our analysis pipeline before the "follow-up" procedure, and to automate the present follow-up steps so that they can become part of the standard pipeline, thereby reducing the role of subjectivity. Note that the follow-up procedure cannot increase the *statistical* significance of a candidate; we don't run the risk of creating a candidate during the followup stage.

The CBC group uses a detection checklist that is applied to all statistically significant candidate-events. As a sanity check, the CBC group also applies the detection checklist to the loudest candidates of the search even if these candidates have a high false alarm probability. The detection checklist consists of a series of tests that aims to corroborate a detection or to eliminate a false alarm. Any manual tests which are found to be particularly useful are subsequently incorporated into the automated analysis pipeline. Here we outline the main tests which comprise the current detection checklist.

- *Status of the interferometers*: The state of the interferometers, their sensitivity and data quality near the time of the candidate are checked. This is intended to supplement the data quality studies which have already been performed, and possibly highlight previously unknown issues in the detectors.
- *Environmental or intrumental causes*: As with the status of the interferometers, we look more closely at the environmental and auxiliary instrumental channels at the time of the candidate. In order to characterize the statistical significance of instrumental transients found at the time of inspiral candidates, the noise properties of auxiliary channels throughout the science run is also estimated. Finally, an analysis of the auxiliary channels at the time of hardware injections is performed to determine the safety of auxiliary channels in ruling out gravitational-wave candidates.
- *Candidates' appearance*: The candidates' appearance is examined through analysis tools which have not yet been implemented in the automated pipeline. These include time-frequency spectrograms of the data, time-series of the multi-detector coherent and null SNR, full parameter estimation. The interpretation of these tests is based on comparative studies for simulated hardware and software injections and time shifted coincidences.

3.4.11 Interpretation of results

The primary goal of our analysis pipeline is to detect gravitational waves from compact binary coalescences and measure the physical parameters of the sources. Another important goal, reachable even without any detections, is to constrain the rate of such events in the universe. To do this, we must understand the source population, despite the fact that little is known about compact binary systems. The group maintains and regularly updates a document that summarizes the range of predictions for CBC rates accessible by LIGO and Virgo [80].

Given the reach of the detectors to low mass binary coalescences during the S3 through S5 searches, such systems are expected to largely follow the blue light luminosity of galaxies [81]. Therefore, we have chosen to quote upper limits on rates in terms of events per year per L_{10}, where L_{10} corresponds to the blue-light luminosity of 10^{10} suns (the Milky Way is approximately $1.7\ L_{10}$). In the absence of a detection, we evaluate the *cumulative luminosity* to which the search was sensitive in units of L_{10}. This is done by convolving the pipeline efficiency as a function of distance (and other parameters) with a list of source galaxies. The efficiency is determined with the same pipeline as is used for the search, so the only systematic errors on the efficiency are associated with Monte Carlo statistics, calibration, and waveform uncertainty. Rate upper limits are established using a Bayesian procedure based on the loudest observed events [31, 32], allowing for a straightforward marginalization over systematic uncertainties, and combining of rate probability distribution functions from independent observations, yielding confidence intervals and/or upper limits.

For the S6 searches and beyond, the reach of the detectors will be large enough that it is reasonable to assume a uniform distribution of sources, which motivates setting upper limits in units of events per unit time per unit volume (in Mpc^3). Additionally, the early star formation in elliptical galaxies is expected to make additional contributions, requiring other tracers of CBC sources beyond blue light luminosity (e.g., galaxy mass). Furthermore, searches for higher mass systems (eg, the S5 high mass search and the S4 and S5 ringdown searches) suffer from near-complete uncertainty about the population of astrophysical sources. In all of these cases, upper limits are quoted as a function of system mass(es) in units of events per unit time per unit volume.

Interpretation of externally triggered searches is, in most ways, similar to the interpretation of other searches for gravitational waves, except that search sensitivity is quantified in terms of linear distance rather than volume. In this case, additionally, the results of searches can be used to say something about the association of the external trigger with a particular coalescing compact binary progenitor (as was done with GRB 070201 [34]). We can also make statistical statements about the presence of detectable gravitational waves from the entire population of GRBs examined by the searches [93].

3.5 CBC Searches and Results from LIGO S1 to S4 data

The CBC searches of data from the first four LIGO science runs has been completed and the results reviewed and papers written. These searches comprised:

- A search for binary neutron stars using post-Newtonian templates in S1 [83], S2 [84, 85], S3 and S4 [33] data
- A search for primordial black holes using post-Newtonian templates in S2 [86], S3 and S4 [33] data
- A search for stellar mass black holes (with component mass greater than $3M_\odot$) performed using phenomenological templates [87] in S2 [88], S3 and S4 [33] data
- A search for spinning binary black hole systems in the S3 data [58]
- A search for the ringdown of black holes formed after coalescence of binary systems, in S4 data [42].

3.6 CBC Search Results from LIGO S5 and Virgo VSR1 data

The LIGO S5 run lasted for two years, the last six months of which coincided with the first Virgo science run (VSR1). The search of this data for CBC signals, using the methods described in the previous sections, is nearing completion. The search parameter space has been broken up somewhat differently than in the previous four science runs. The "low mass" CBC search, performed with post-Newtonian templates, covers binaries with a total mass between $1M_{\odot}$ and $\sim 35M_{\odot}$ and component masses not less than $1M_{\odot}$. For these systems, the waveform is predominantly in the inspiral phase in the sensitive band of the LIGO and Virgo detectors, and post-Newtonian Stationary Phase Approximation [89] frequency domain templates are sufficiently accurate. The "high mass" search covers binaries with total mass between 25 and 100 $M_{\odot}$ (overlapping with the low-mass search in the mass range between 25 and 35 $M_{\odot}$). At these higher masses, the merger and ringdown of the signal are in the sensitive band and contribute a significant fraction of the signal-to-noise ratio. Therefore, developments in understanding the full waveform have a more significant impact on these higher mass signals. At even higher masses, a search for black hole ringdowns is being performed. In addition, a search for gravitational waves associated to short GRBs which occurred during S5-VSR1 has been done.

These searches required the development of many new analysis tools, including: the automated *ihope* pipeline infrastructure; the development of improved detection statistics to improve the separation of signal and background in the presence of glitchy data and its quantification in terms of inverse false alarm rate; the development of improved and automated data quality flags and vetoes; the development of improved and automated detection candidate followup procedures; and the development of richer search summaries (*ihope_page*), the implementation of databases for efficient handling of coincident triggers and associated information; and the implementation of two different families of NR-inspired IMR waveforms into LAL.

The CBC group has written the following papers describing astrophysical results of searches from S5/VSR1:

- Gravitational waves associated to Gamma Ray Burst 070201 whose allowable sky location range ovelapped the Andromeda galaxy [34].
- Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO's fifth and Virgo's first science run [93].
- Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data [73].
- Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run [74].
- Search of S5/VSR1

3.7 Ongoing CBC Searches and Studies from LIGO S5 and Virgo VSR1 data

3.7.1 S5/VSR1 GRB searches:

A search of a gravitational wave signal in data around GRB 051103, which happened just one day before the official start of S5 and whose sky location overlaps M81, is ongoing.

3.7.2 S5 high mass search:

The S5 high mass search focuses on systems with total mass between 25 and 100 $M_{\odot}$, where the merger and ringdown of the signal are in the sensitive band and contribute a significant fraction of the SNR. The S5 high mass search has used time-domain waveform templates known as EOB-NR, which make use of the Effective One Body (EOB) formalism [47, 143, 50], and incorporate information from numerical relativity simulations in order to reliably model the inspiral, merger and ringdown (IMR) phases for binary coalescence with non-spinning components. These waveforms, and the phenomenological waveform family described in [51, 142, 91] have been used to tune, test, and evaluate the efficiency of the search pipeline. The high mass search is being performed using the same *ihope* pipeline as the low mass search. Because of the emphasis on low frequencies, the search is LIGO-only (ie, does not include Virgo data), covering all 24 months of S5 (including the overlap with VSR1). Much effort has gone into handling discontinuities in the templates and injected waveforms, re-tuning the signal-based vetoes, understanding the parameter estimation and coincidence windows, and developing appropriate data quality vetoes. The paper describing this search is complete and currently under review in preparation for publication. Despite the great advances in modeling inspiral, merger and ringdown waveforms, the uncertainties in the waveform amplitdes can be as large as 50% or more at high masses ($\sim 75 M_{\odot}$ and above), which has a significant effect upon the interpretation of the results. Evaluation of the detection sensitivity for spinning waveforms is an ongoing challenge.

3.7.3 S5/VSR1 ringdown search:

The S5 ringdown search, focusing on black holes with masses between 75 and 750 $M_{\odot}$, is in progress. For high mass binary coalescences ($M \gtrsim 100 M_{\odot}$) , the majority of the power received by the LIGO and Virgo detectors will be from the ringdown part of the coalescence. Therefore, the search for coalescencing binary systems can be done looking for the final ringdown which has known waveforms parametrized by mass and spin (determining frequency and damping time of the ringdown). The uncertainty in the theoretical predictions for how the inspiral, merger and ringdown phases couple into a single waveform governed by a single set of source parameters (masses and spins) leads us to pursue such a ringdown-only search. This search has been completed on the S4 data, and is now being run on the S5 data using the same *ihope* pipeline as the low mass search. Because of the emphasis on low frequencies, the search is LIGO-only (ie, does not include Virgo data), covering all 24 months of S5 (including the overlap with VSR1).

3.7.4 S5/VSR1 IMR comparison:

The expected waveforms from the inspiral, merger and ringdown of compact binaries shifts to lower frequencies for systems of higher total mass ($f_{merger} \sim 1/M_{tot}$). For total masses below around 35 $M_{\odot}$, there are many cycles of inspiral waveform in the (Initial) LIGO band, and matched filtering techniques are optimal for identifying a signal in noisy data. For higher masses, only a small number cycles from the inspiral, merger, and/or ringdown phases are in the LIGO band, and the in-band signal waveform begins to resemble a low-Q wavelet or sine-Gaussian; the kind of waveform that burst searches developed by the LSC are designed to identify (Section 4). In this regime, burst searches may be as good as matched-filter searches such as the high mass and ringdown searches pursued by the CBC group, and we can expect similar efficiencies (for fixed fake rate) and considerable overlap in the parameter space. It is important to compare the CBC and burst search strategies on an "apples-to-apples" level; ie, evaluating the detection efficiency for high-mass CBC sources at the same false-alarm rate. A joint CBC-Burst working group is pursuing this goal, making use of the high mass and ringdown pipelines from the CBC group and the coherent Waveburst and Omega pipelines from the Burst group. These studies are complete for several, but not all, of the analyses. The work still needs to be written up and published

3.7.5 NINJA 2:

The Numerical Injection Analysis (NINJA) project [101] was a first attempt to bring together the numerical relativity and gravitational wave search communities to evaluate the sensitivities of existing search pipelines to the black hole waveforms generated by numerical simulations. While this first project was a success, it suffered from several shortcomings. In particular, due to the limited scope of the project, it proved difficult to draw quantitative conclusions from the analysis. This was further hampered by the use of simulated Gaussian data which does not present the same complications as real data.

To address these issues, a second NINJA project is being undertaken. This project aims to perform a systematic study of the sensitivity of existing gravitational wave searches to numerical waveforms emitted by non-precessing black hole binaries. NINJA 2 places more stringent requirements upon the quality of numerical relativity waveforms and their attachment to early post-Newtonian inspiral. Additionally, it will make use of real data from the LIGO S5 and Virgo VSR1 runs.

3.8 CBC Searches with the LIGO S6 and Virgo VSR2 data

The LIGO S6 and Virgo VSR2 runs began in July, 2009. VSR2 ended in mid-January 2010 when the Virgo detector was taken down for upgrade. The S6 run is ongoing and is scheduled to end in October 2010 to prepare for Advanced LIGO. The Virgo upgrade is due to be completed in July and Virgo expects to start VSR3 in July, 2010, thereby giving a few months of overlap between VSR3 and S6.

The S6 and VSR2 runs have brought significantly improved broadband noise, but also new classes of detector glitches (associated with the LIGO OMC) and increased sensitivity to low-frequency seismic disturbances such as ocean storms. These directly impact our ability to make reliable detections near the noise threshold.

The main CBC group activities in S6/VSR2 primarily follow those in S5/VSR1, namely: low mass CBC, high mass CBC and externally triggered GRB searches. In the S6/VSR2 run there has been a great push towards reducing the latency with which the analysis is completed and detailed results are obtained and examined. This has greatly improved the CBC group's ability to contribute to detector characterization efforts. Additionally, a very low latency analysis has opened up the possibility of producing alerts for followup by electromagnetic observers. In the following, we outline each of these efforts in turn

3.8.1 Detector Characterization

There is greatly enhanced leadership and involvement of CBC group members in the effort to produce timely and reliable data quality vetoes, analysis segment definitions, and reliable calibration during S6. In particular:

- CBC group members have developed comprehensive weekly data quality summary pages to supplement the work of the glitch group.

- Known glitch-producing detector anomalies are rapidly and automatically identified and inserted into the database. New anomalies must be discovered, investigated and vetoes developed through careful analysis. Their safety (against false dismissal, tested using hardware injections) must be carefully evaluated. A significant ongoing effort (through the end of S6/VSR3) is required to: (i) identify auxiliary channels containing noise events which are correlated with inspiral triggers in the gravitational-wave channels; (ii) investigate the couplings between these channels and the gravitational-wave channel; (iii) understand and identify particular artifacts by the signature across all appropriate channels; (iv) develop methods to clean the data of such artifacts, e.g., by vetoing noisy times; and (v) automate these methods within the pipeline and the on-line analysis.

- Correlations between glitches in the auxiliary channels and gravitational wave channel are studied using UPV and hveto. These tools have been used to identify instrumental artefacts which affect the CBC searches, as well as producing veto times to be excluded from the analysis.

- A new class of "bilinear" vetoes have been developed and implemented, which are very effective at identifying glitches in the h(t) channel while incurring very small deadtime and being completely safe against false dismissal.

- CBC low mass triggers and diagnostic plots are now produced daily (*daily ihope*). This information plays a crucial role in early discovery of detector anomalies and glitches that can cause loud CBC triggers, as well as rapid identification of hardware injections.
- The MBTA pipeline, running continuously in Cascina, produces single-detector and triple coincident lowmass triggers with latencies of a few minutes or less. In addition to its tremendous value for rapid data quality studies and identification of hardware injections, these triggers are in preparation for use for rapid EM follow-up in the upcoming S6/VSR3 run (see below).

3.8.2 Low Mass CBC Search

A search for binaries with masses between $1M_\odot$ and $35M_\odot$ and component masses greater than $1M_\odot$ which covers the binary neutron star, neutron star-black hole and lower mass binary black hole space.

- The CBC group has been running the *ihope*-based low mass and pipeline on a weekly or bi-weekly basis, largely unchanged from the ones employed for S5.
- We continue to follow "blind analysis" protocols, in which all analysis parameters are tuned and finalized before the in-time coincident triggers are examined, to avoid bias in the tuning. The latency for "opening the box" on in-time coincident triggers has decreased during the course of the run and now stands at around 1 month.
- Coherent (H1L1V1) low mass and high mass CBC hardware injections are performed regularly according to a predetermined schedule, in close coordination with the hardware injection teams at all sites (so far, all harware injections in S6 have been for systems with total mass less than 35 $M_\odot$). These injections are searched for with the *MBTA*, *daily ihope*, *low mass weekly ihope* and *high mass weekly ihope* pipelines, and deviations from expected performance are promptly investigated.
- The CBC low mass search participates in the Blind Injection Challenge in S6/VSR2. To this end, the group endeavours to finalise the search within 3 months of the data being taken, and prepares "quarterly reports" [102] on the results. In some cases, results have been available in less than 3 months. The up to $\sim$ 6 month delay reflects the need to understand and feel confident about the rapidly-changing data quality conditions during S6/VSR2, and the need to have a sufficiently large time-slide sample to estimate the background and software injection sample to estimate the pipeline sensitivity.
- Results from the analysis of the entire S6, VSR2/3 data will be combined (incorporating previous results) to produce a detection or upper limit statement.

3.8.3 High mass Inspiral–Merger–Ringdown CBC:

A search for binaries with total mass M between $25 - 100M_\odot$, using EOB-NR templates. There have been numerous innovations in the high mass search, in an attempt to improve its sensitivity and to keep up with the ever improving knowledge of the waveforms emitted by these systems.

- We continue to follow "blind analysis" protocols, in which all analysis parameters are tuned and finalized before the in-time coincident triggers are examined, to avoid bias in the tuning.
- the CBC *ihope* highmass search has recently incorporated a new "multivariate" detection statistic (*MVSC, pronounced "music"*) which makes use of more than a dozen numerical quantities characterizing each coincident trigger candidate event, in order to optimally separate signal from background

in that multi-dimensional space. This new detection statistic will significantly improve our sensitivity for signals with false alarm rate below a fixed and low value.

- New waveform families are being developed and implemented. When available these are used to produce simulations which are subsequently used to evaluate the efficiency of the search.
- Results from the analysis of the entire S6, VSR2/3 data will be combined (incorporating previous results) to produce a detection or upper limit statement.

3.8.4 Externally Triggered Searches:

A deeper (lower threshold), targeted search for CBC signals arriving within 6 seconds of a GRB event. The search uses inspiral templates for systems in which one component is a neutron star with mass between 1 and 3 $M_\odot$ and the companion mass is between 1 and 30 $M_\odot$.

- This search uses the same pipeline as was employed for the S5/VSR1 search whereby 40 minutes of data around the time of the GRB is analyzed. Six seconds are classified as the on source data, while the remaining time is off source data used for background estimation.
- It is launched automatically as soon as a GRB alert is posted on the Gamma-ray burst Coordinates Network [174] during LIGO's S6 run [103], regardless of whether the GRB is identified as short-hard or long-soft (since such an identification can come much later than the first alert, and the classification is not univocal).
- For GRBs subsequently identified as short-hard, a more powerful but also more compute-intensive detection statistic based on a likelihood ratio for signal/background is employed. It is a blind analysis with periodic box openings.
- The background estimation for this search is being improved by performing time-shifts of the off-source data. This should improve our ability to estimate the significance of candidate events by two orders of magnitude.

3.8.5 Low-latency search and EM follow-up:

The CBC group is preparing a low-latency search during the triple-coincident data collection of S6/VSR3. These will be included in the fast EM-followup programs developed by the Burst group.

- A low-latency search for low mass CBC signals. The search is executed with the *MBTA* pipeline at Cascina, within minutes of data acquisition.
- Triple coincident (H1L1V1) triggers are identified and a sky localization map is generated using timing and amplitude information.
- Candidate detections passing to-be-determined selection criteria will be submitted to the GraceDB event database [100], appropriately scrutinized, and then potentially passed to the Lumin [98] and Swift [99] follow-up processing pipeline for rapid observational follow-up by ground-based robotic telescopes and the space-based Swift telescope, respectively.

The above activities will result in observational results in the searches for low mass CBCs, high mass CBCs, and GRBs. There will be technical papers, if not observational results, associated with the low-latency EM-followup activity. There will be opportunities for joint publications with the burst group on

GRB searches, high-mass CBC searches, and low-latency EM follow-ups. We expect several technical papers on data quality procedures, the use of IMR waveforms, and many of the ongoing analysis activities described below. Specific plans for most of these publications are still under discussion.

As of June 2010, the LIGO H1 and L1 detectors are taking the best data we have ever acquired, with SenseMon BNS ranges approaching 20 Mpc, and low glitch rates. We are employing the most rapid, well-understood and sensitive detection pipelines that we've ever had. We have a well developed detection confidence follow-up pipeline and increasingly accurate and useful parameter estimation tools. With the addition of Virgo this summer, we will have the most powerful and sensitive gravitational wave detection network ever assembled and will be well prepared for detection and for the setting of the most stringent event rate confidence bands.

However, many important analysis tools remain to be developed, only some of which may be in place for the full exploitation of S6 data in the coming year. These are discussed in section 3.10.

3.9 CBC Group Priorities and Goals for FY2011

The *highest priorities* for the CBC group during FY2011 are listed below. The main focus is the analysis of S5-6 and VSR1-3 data. The "CBC" code will aid in referring to these tasks in group MOUs to the LSC.

3.9.1 S5/VSR1 Analyses

- **GRB 051103 Search - CBC-S5-1** Complete of the search of (pre-)S5 data at the time of GRB 051103. Preparate and review the results paper.
- **S5 high mass CBC search - CBC-S5-2** Complete the review of the S5 high mass CBC search and publish results paper.
- **S5 ringdown search - CBC-S5-3** Perform the S5 ringdown search. Write, review and publish the results paper.
- **S5 IMR comparision - CBC-S5-4** Completion of the comparison of IMR methods used on the S5 data. Publication of the results.
- **NINJA 2 - CBC-S5-5** Perform NINJA 2 analysis on S5/VSR1 data, review results, publish findings.

3.9.2 S6/VSR2-3 Analyses

- **Data infrastructure - CBC-S6-1**. Primarily under computing infrastructure.
- **Data Characterization - CBC-S6-2**. Continue the rapid identification and development of data quality characterization and vetoes through to the end of the S6/VSR2/VSR3 run.
- **Hardware Injections - CBC-S6-3**. Continue the maintenance, development and prompt follow up of coherent CBC hardware injections at the three sites.
- **Low Latency Analysis and EM follow-up - CBC-S6-4**. The CBC group is preparing to produce low-latency triple-coincident, sky-localized detection candidates for rapid EM follow-up. We expect to demonstrate this process in the coming S6/VSR3 run.
- **Low mass ihope Analysis - CBC-S6-5**. Complete the S6/VSR2 CBC low mass search, identify and evaluate detection candidates, evaluate detection sensitivity and compute rate limits, prepare publishable results and paper(s), and conduct a thorough review.

- **High mass ihope Analysis - CBC-S6-6**. Complete the S6/VSR2 CBC high mass search, identify and evaluate detection candidates, evaluate detection sensitivity and compute rate limits, prepare publishable results and paper(s), and conduct a thorough review.
- **GRB Externally Triggered Search - CBC-S6-7**. Complete the S6/VSR2 CBC GRB/external trigger search, identify and evaluate detection candidates, evaluate detection sensitivity and compute rate limits, prepare publishable results and paper(s), and conduct a thorough review.
- **Followup of Candidates - CBC-S6-8**. Careful examination of any GW candidates to emerge from the searches. This will necessitate close collaboration with both those performing the searches and the detector characterization effort.

3.10 High priority analysis development

In order to achieve these science goals, the CBC group will continue the following ongoing analysis development tasks. The "CBC" code will aid in referring to these tasks in group MOUs to the LSC.

- **Improved Background and FAR estimation - CBC-S6D-1**: Improved estimation of the false alarm rate (FAR) for the most significant coincident event candidates. We currently estimate the FAR using 100 time-slides, limiting the FAR estimate to no less than 1%. The long auto-correlation time for CBC triggers may make it difficult to substantially improve this. Work is in progress to make reliable FAR estimates using single-detector trigger rates, but this may require a substantial modification of the current two-stage analysis pipeline. A significant effort is required to: (i) study and understand the limitations of the time-slide method dues to trigger auto-correlations, especially in the context of large parameter space searches; (ii) study alternative methods of estimating the background; (iii) implement necessary changes to the existing pipeline in order to make use of these alternative methods; (iv) investigate possible sources of bias in the background estimate from environmental or other correlations.
- **Extension of externally triggered search - CBC-S6D-2**: Currently, the externally triggered search is triggered by announcements of GRBs. There are arguments that binary coalescences will also produce signals in other bands, such as radio, x-ray and neutrinos. The externally triggered search will be extended to include transients which are observed in these wavelengths. In many cases, these external triggers will provide precise sky localization information but rather loose time windows (eg, days). The incorporation of an improved background estimation using timeslide methods is also one of the high priority goals.
- **Coherent Analysis - CBC-S6D-3**: Incorporation of coherent multi-detector network information to more effectively distinguish signal from background. The existing CBC pipelines incorporate network detection information by requiring single-detector triggers to be near-coincident both in time and in waveform template. We employ additional software tools that perform a "coherent" analysis around a coincident trigger event candidate, generating coherent and null-stream SNR time series. However, these tools have not yet resulted in significant additional power to separate signal from background. Work is ongoing.
- **Searches incorporating spin - CBC-S6D-4**: Work is underway to identify regions of parameter space where spin waveforms differ significantly from the non-spinning templates. In parallel, search using a physical template family of spinning waveforms [57, 68] is being developed. These waveforms have been implemented into the CBC software base, and can be used for both template bank construction and simulated signal injection for efficiency evaluation. Evaluation of the efficacy of

a PTF-based search is ongoing. Depending upon the results of these studies, we may mount an S6 search using spinning templates — either as an independent search or a followup to candidates produced by the non-spinning waveforms. The issue of spin is equally significant, if not more so, in the high mass regime. For these masses, the post-Newtonian approximation does not adequately describe the revelant part of the waveform. Numerical results for spinning systems are being obtained, but a full exploration of the spin parameter space will take some time.

- **Improved signal based vetoes - CBC-S6D-5:** In non-stationary data, a high signal to noise ratio does not necessarily correspond to a signal. A number of additional signal consistency tests have been developed to help discriminate signals from noise transients in the data. Two of these, the so called χ^2 and r^2 tests have been incorporated as standard parts of the ihope analysis. Other signal-based tests have been proposed and developed. They have the potential to further assist in separating signal and background, and have the benefit of being less computationally costly than the χ^2 test. Work is needed to complete the development of these vetoes and to deploy them into the analysis pipeline.

- **Improved Classification of Signal and Background - CBC-S6D-9**: The SNR alone is not an optimal way to distinguish signal from background due to non-gaussian detector glitches. Combining SNR with χ^2 ("effective SNR" or "newSNR") and carefully tuned coincidence criteria helps the low mass search achieve sensitivity close to the Gaussian limit, but it is still difficult to assess the false alarm rate for different combinations of coincident triggers. In the course of the analysis of the LIGO data from S1 to S5, the detection statistic used to separate signal from background has evolved to incorporate more information from the analysis — SNR; effective SNR; IFAR; effective likelihood. However, the separation of signal and background still makes use of only a fraction of the information which is known about a given candidate (signal to noise ratio in each detector, χ^2 value, r^2 value, mass and spin parameters, coincidence between detectors, estimated effective distance in each detector, sky location, data quality information, coherent SNR, null stream, information from auxiliary detector channels or environmental monitoring channels, . . .). By making use of more of this information, a better separation between signal and background is possible. A better detection statistic would provide a ranking of the "signal-like" nature of in-time coincident triggers, going beyond the simple SNR "loudness". This would improve our rejection of noise glitches and immunity to non-Gaussian noise behavior, and thereby increase the sensitivity of the search. There are several projects currently underway which are attempts to do this, including Multivariate Classification techniques, Likelihood Ratio techniques, and Bayesian Likelihood techniques.

- **Overlap with Burst searches for high-mass systems - CBC-S6D-10**: This remains an important question. There are two concrete avenues through which this being investigated: the S5/VSR1 IMR comparison effort and NINJA 2.

- **Parameter estimation - CBC-S6D-11**: A good understanding of parameter estimation abilities of the detector network is critical for both the design of our analyses, and in extracting astrophysical parameters from gravitational wave observations. There are several areas of active development:

  - Understand the systematic effects of calibration errors on parameter estimation. This study must be carried out using real and simulated detector noise.
  - Use Bayesian parameter estimation techniques to determine posteriors on the parameters of a detected signal. Use MCMC codes and analytical techniques to systematically study approximate symmetries and degeneracies in the parameter space. These include degeneracies in the extrinsic parameter space, such as sky location, which can be removed if multiple detectors are online.

  - Implement tools to include higher-order post-Newtonian corrections to the waveform amplitudes, and hence other harmonic structure. It is known that the other harmonics can significantly enhance our ability to detect gravitational waves from binary systems as the mass ratio decreases. These harmonics also break degeneracies between many of the parameters. By including this information in an observation, it may be possible to better measure the binary parameters.
  - We expect astrophysical coalescing binaries to have spinning components, and a direct observation of this through unambiguous parameter estimation will provide the first *direct* measurement of black hole spin. Various parameter degeneracies can be broken by including spin and can much improve sky localization.
  - Make a quantitative comparison of tradeoffs between the full parameter estimation codes and rapid follow-up codes regarding the accuracy to which the location of a binary signal can be localised on the sky.

  All these tools need development in order to be better understood and more useful in the event of the first detections.

- **Improved detection confidence procedures - CBC-S6D-12**: Coincident events in the CBC pipelines that are significantly inconsistent with the background estimates are subjected to a series of checks and tests (the "follow-up pipeline and checklist") to look for signs that the events were caused by instrumental glitches. These include the examination of coincident triggers satisfying different data quality category requirements or different coincidence criteria; incorporating information from auxiliary channels; incorporating information from the coherent and null-stream analyses; information from Qscans; etc. We endeavor to make these tools as quantitative as possible. However, most of these detection confidence checks are at best semi-quantiative, and in many cases, are wholly qualitative and subjective, requiring experience and good judgment in their application. We still consider them to be very valuable in the development of confidence in our first detections, and indeed they may be the deciding factors. There is much room to incorporate much of this information into automated, unbiased, statistically significant detection criteria (eg, via MVSC).

- **Next Generation Low Latency Pipelines - CBC-S6D-13**: The Virgo MBTA pipeline is in place during S6 for the generation of CBC triggers for detector characterization, and for rapid identification of coincident triggers that could be GW detections. Work is in progress to rapidly locate sources in the sky associated with such triggers, for EM follow-up. In addition, work is in progress to develop new, stream-based low latency search and sky localization pipeline infrastructure for both S6 and for the long-duration chirps expected in Advanced LIGO.

- **GPU-enhanced CBC pipelines - CBC-S6D-14**: The optimal filtering and chisq computations that are at the heart of the CBC pipeline are very cpu-intensive. Modern Graphical Processor Units (GPUs), optimized for these kinds of computations, can greatly speed up, and/or greatly reduce the cost of, conducting a wide-parameter search for CBC (and also continuous wave) signals. The CUDA library greatly facilitates the conversion of our existing code to make use of GPUs. We are developing the necessary software modifications, and designing next-generation computing clusters to implement GPU-enhanced CBC search pipelines.

- **Development of new waveform families - CBC-S6D-15**: CBC pipelines use binary coalescence waveforms for: tuning and testing the efficiency of our search pipelines; as templates for matched filtering at the heart of the search pipeline; and as templates for parameter estimation after detection. Much work has gone into developing waveforms of increasing accuracy, covering all phases of coalescence (inspiral, merger, ringdown: IMR), and spanning broader parameter spaces; especially mass and spin, but also matter effects, ellipticity, and other novel phenomena close to merger. Numerical relativity (NR) groups continue to generate waveforms that explore these effects; but they sparsely sample the parameter space, making them unsuitable for matched filter template banks or parameter estimation. Three families of parametrized waveforms tuned to NR have been developed, implemented in LAL, and used in the S5 high mass search: EOB-NR, Phenom-A and Phenom-B (spin-aligned). As numerical results continue to improve and probe an ever greater region of the parameter space, new phenomenological waveforms will be developed to better reflect our understanding of the waveform. These will also be incorporated into our analysis pipelines. Additionally, the ideal template spacing for these waveforms will differ from the post-Newtonian one, so a template space metric, stochastic template placement, or other new technique is required.

- **Accurate evaluation of search sensitivity - CBC-S6D-16**: We employ many thousands of software injections to evaluate our search efficiency as a function of source parameters (distance, masses, and spins) for each data taking epoch. These simulations are compute-intensive and can be problematic (defining whether an injection is found is not unambiguous). We never have enough injections, and simulation ("Monte Carlo") statistics are often a dominant systematic error when setting rate confidence bands. This is especially true for large search parameter spaces, as in the high mass search. Since source populations are essentially unknown, we cannot integrate over assumed parameter distributions, and must evaluate our efficiency in bins of the source parameters. Large-scale simulation efforts are required, and alternative methods for evaluation of search sensitivity can be developed.
- **Development of improved detector characterization and data quality techniques - CBC-S6D-17**: Enhance and develop the data characterization and data quality methods pioneered during the initial detector era to produce robust, low latency characterization of the advanced LIGO and advanced Virgo data, ready when these detectors are first being commissioned.

3.11 Development prior to advanced detector era

The Enhanced LIGO run (S6) will end in the fall of 2010, when the installation of the Advanced LIGO detectors will begin [105]. The Virgo VSR3 run will continue; the current plan is for VSR3 to end in May 2011 when Advanced Virgo installation begins [105]. The GEO600 detector is being upgraded to GEO-HF, concentrating on the detection of high frequency signals and the demonstration of advanced technologies [106]. They plan to run in "Astrowatch" mode through at least 2011. The advanced LIGO and Virgo detectors will begin commissioning activities as early as 2013, with the goal of first scientific observations in 2014 or 2015.

3.11.1 Searches in the "Before Advanced detector Era"

During the period when Virgo and/or GEO-HF are running and before the advanced LIGO and Virgo detectors turn on, we will have the opportunity to continue searching for CBC signals. Detection confidence will be greatly strengthened if signals are observed in coincidence with electromagnetic signals from extragalactic short GRBs. (GW burst signals can also be found in coincidence with GRBs, galactic SGRs, and core-collapse supernovas.) We therefore plan on continuing the search for CBC signals during this "Before Advanced detector Era" (BAE).

- **GRB-triggered searches - CBC-BAE-1**: We will employ the GRB-triggered CBC search on short GRB external triggers, using data from Virgo and/or GEO-HF, if available.

- **CBC coincident searches - CBC-BAE-2**: We may decide to employ the CBC low mass search on coincident data from Virgo and GEO-HF, if available. This depends on the performance and duty cycle of these detectors. The prospects for such a search will be studied in the coming months.

3.11.2 Analysis development in the "Before Advanced detector Era"

The CBC analysis pipelines are built on a solid base of well developed, tested and reviewed code in lalsuite. The pipelines themselves have grown in complexity in the last 10 years in order to handle the many requirements of a complete analysis of observational data. The end of S6 provides us with an opportunity to revisit our large code base, with the goals of making our analysis pipelines more robust, sensitive to GW signals, computationally efficient, flexible, powerful, scalable, maintainable, understandable and reviewable.

As discussed above in section 3.10 and below in section 3.12, many significant enhancements (or major architectural changes) will be needed in order for the CBC analysis codes to be used for accomplishing all of our scientific goals in the advanced detector era. These include handling the long-duration chirps expected in Advanced detectors, and performing coherent network analysis and incorporating additional features in order to improve detection sensitivity and parameter estimation.

We expect that there will be at least three ground-based advanced detectors in the coming years: LIGO Hanford, LIGO Livingston, and Virgo. Another LIGO detector may be sited at Hanford or in Western Australia. It is increasingly likely that a detector will be build in at Kamioka, Japan. We need to understand the impact of each detector on the overall sensitivity of the detector network, detection confidence, and especially, sky localization and parameter estimation. Multiple detections of the same wave can provide more powerful tests of the fundamental properties of gravitational waves as predicted by General Relativity. We will prepare the tools required to fully exploit the full detector network for detection, parameter estimation, and perform fundamental tests of GR.

A design study of thrid generation gravitational wave detectors, funded by the European Framework Programme - 7 for three years, began in May 2008. The goal of this study is to explore the scientific benefits of a detector that is some ten times more sensitive than advanced detectors and capable of observing down to 1 Hz in frequency and to identify technical and data analysis challenges associated with a study. A science vision document for such a detector has already been developed by the design study team [108] and the team is now looking at tradeoffs between different detector geometries, sensitivities and site locations. Additionally, the team is involved in the development of a mock data challenge to understand the complexity of analysis of data from a detector in which millions of compact binary signals can be expected each year.

The commissioning of the advanced detectors will begin in 2013. As we have during S6, the CBC group intends to play a major role in detector characterization and data quality / glitch studies. In order to be ready for the challenges and surprises that we are bound to encounter, we intend to develop software tools and expertise, and apply these as early as possible during commissioning.

We expect the transition from commissioning to astrophysical observation mode to occur in stages during 2014. The first science run will be a significant milestone, which should occur when there is a significant chance of confident detection. For the case of CBC signals, estimating the chance of detection brings together our best estimates of the rate of binary mergers per Mpc^3 with our best estimates of our detection reach in Mpc^3. The former requires us to continually update our knowledge of the astrophysical predictions, as described in section 3.3. The latter requires us to understand our detection thresholds for confident detection, in terms of individual detector SNR, network SNR, the detectors' noise spectra and their glitchiness, our ability to estimate the background in the regime of very small false alarm rates, the required false alarm probabilities (for "evidence", "first detection", "detection with associated EM observation(s)", etc), and any appropriate astrophysical priors. Tools for understanding and quickly evaluating these quantities need to be fully developed when the first science quality data are avaliable. Some thoughts on the threshold for the first advanced detector science runs are in [107].

A second significant milestone will be the first detections. Tools for robustly evaluating detection significance / false alarm rate, and extracting binary parameters, need to be fully developed and understood. Another significant milestone will be when we are making, or capable of making, routine detections. Many science goals associated with these milestones are listed below in section 3.12.

At some point, the LIGO and Virgo Collaborations will no longer be able to lay claim to exclusive possession of our data and the extraction of scientific results from them. We will need to make our strain data publicly available (in addition to the data associated with detections that will be published). This will be necessary even in the absence of detections by our Collaborations, if the detectors are deemed sufficiently sensitive that detections are "to be expected". We must clearly define the threshold of sensitivity that should trigger the public release of the strain data. We must also consider the science that we can do in the absence of detection after months of observation at "sufficient" sensitivity. What rate limits will we be able to obtain? Will they significantly constrain astrophysical models and expectations? What level of certainty is required in the strain calibration and the signal waveforms, and will those levels of certainty be possible / practical? We intend to explore and study various such scenarios before the advanced detector era begins.

Specific R&D goals for the CBC group during the "before advanced detector era" (BAE) are listed below.

- **Advanced detector analysis pipeline development - CBC-BAE-3**.
- **Detector Characterization and data quality preparation for the advanced detectors - CBC-BAE-4**.
- **Detection sensitivity and parameter estimation in the advanced detector era - CBC-BAE-5**. Understanding our detection SNR thresholds versus false alarm probability as a function of detector glitchiness, etc.
- **Establishing thresholds for science running and for data release - CBC-BAE-6**. This makes use of our best estimates of astrophysical rates.
- **Studies of the advanced and third generation detector network - CBC-BAE-7**. This includes understanding the impact of detectors in Japan, Australia, and elsewhere on the detection sensitivity, detection confidence, sky localization, etc. It also includes, contribution to the design study of the Einstein Telescope, highlighting the science goals and data analysis challenges.

3.12 Advanced detector era

As the Earth-based gravitational-wave detectors move into a phase where they are making routine astronomical observations, the nature of the analysis efforts will change. We intend to continue to perform searches for binary coalescences, with: minimum delay (ideally, near-real-time); maximum coverage of the plausible space of component masses and spins; prompt and reliable follow-up of candidate events to establish detection confidence; and prompt source localization and parameter estimation. We look forward to pursuing the following innovations:

1. **Multi-band and Hierarchical Search Methods**: Advanced detectors will have greatly reduced noise at low frequencies, allowing for detection of gravitational waves down to 10 Hz. Inspiral waveforms can spend many minutes in the low frequency band. It will be necessary to restructure the existing pipeline, perhaps along the lines of the Virgo MBTA pipeline, in order to handle such long signals coherently. Hierarchical methods need to be implemented and tested. Several different hierarchical schemes exist which search over coarse grained parameters in an early part of a search, following up on triggers from this coarse grained search with finer grained parameters. An effort is required to: (i)

develop code to translate triggers into hierarchical banks for injection and use by the filtering code; (ii) test and develop intuition for hierarchical searches using the current pipeline; (iii) develop tools to tune this hierarchical pipeline.

2. **New approaches to filtering**: The core of the analysis pipeline is the filtering of the data through bank(s) of templates, in the frequency domain (sections 3.4.1 and 3.4.3). This process is time consuming and (for large bank parameter spaces) computationally limited. This results in a variety of problems: analysis turn-around time is slow; securing adequate and sufficiently reliable computing resources is difficult; it is difficult to develop new (or tune existing) analysis features such as signal-based vetoes; and there is an inherent latency in the pipeline, making it unsuitable for real-time searches or searches using very long-duration templates (necessary for Advanced LIGO data). There is exploratory work in progress to redesign the filtering stage to make it more modular, make use of multi-band hierarchical filtering, employ low-latency time-domain filtering, speeding up the processing (or reducing the number of computers) using GPUs, and other innovations.

3. **Astronomical alerts** will be provided with low latency to allow observations of various events by other astronomical observatories if possible. Observing a source/event simultaneously in different windows has recently provided rich information about transient γ-ray events. Future gravitational-wave observations will benefit from such multi-messenger astronomy, which requires these alerts to be provided quickly but also reliably.

4. **Parameter estimation**: The continued development of the multi-project analysis efforts and joint analysis projects is of critical importance to doing astronomy with Earth-based gravitational wave detectors. Analysis using multiple detectors makes it possible to accurately locate the source in the sky, determine the parameters of the binary system, estimate the errors, and feed this information into catalogs and models, as described below.

5. **Catalog of events**: Database catalogs containing information about all detected events will be developed; tools to visualize, mine and interpret the results will be provided. Such a database could be used to study cosmological models and predictions and address important issues such as the compact binary population of the Universe, evolution of the star formation rate, equation of state of dark energy, etc.

6. **Constraining astrophysical source population:** A growing database of detected events will directly constrain source population models, and eventually produce model-independent measurements of the CBC source population in terms of all relevant parameters (distance, mass, mass ratio, spin, etc).

7. **Standard candle**: A binary inspiral is an astronomers ideal standard candle: by measuring the frequency evolution one can determine the mass parameter(s) and thereby measure the luminosity distance. Inspiral events can therefore be used to build new astrophysical distance ladders and to confirm the current ones. However, for sources at cosmological distances the mass is "blue-shifted" thereby requiring observations of the host galaxies to break the mass-redshift degeneracy.

8. **IMRI search**: Implement a search for intermediate mass ratio inspirals (IMRI) using waveforms from perturbation theory. Typically, the commissioning of a new search takes about two years, so this is a long term project which has potentially great interest.

9. **Validity of post-Newtonian expansion**: Using parametrized post-Newtonian signal model determine the validity of the post-Newtonian theory and determine how far into the strong gravity regime post-Newtonian theory is valid. This should help in developing analytical insights into and performing more accurate numerical relativity of the merger phase.

10. **Alternative gravity theories**: Determine a useful way to constrain alternative theories of gravity using the results from our core search pipelines. It is important to understand to what extent the existing searches can be used to achieve this, or if it is necessary to perform searches using theoretical waveform templates from these other theories.

4 Searches for general burst signals

The mission of the LSC-Virgo Burst Analysis Working Group (also known as the "Burst Group") is a broad search for gravitational-wave "bursts", *i.e.*, short-duration gravitational-wave signals. A variety of astrophysical sources are expected to produce such signals; sophisticated models are available in some cases, while in other cases the amplitude and waveform of the GW signal are highly uncertain. Thus, the Burst Group must utilize robust signal detection methods. Initially, the group focused on signals much shorter than 1 second in duration, with little or no assumption about their morphology [109]. Currently, the Burst Group is also pursuing longer signals and the incorporation of available signal waveform knowledge, to seek specific targets as well as unknown phenomena and mechanisms of emission and propagation. Section 4.1 summarizes recent GW burst search results, while section 4.2 describes many of the basic methodologies used in burst searches.

Sections 4.3 and 4.4 describe the goals and plans for GW burst searches with the LIGO, GEO and Virgo instruments over the next ∼3 years. During this time we will complete the analysis of data from S6/VSR2+3. Some specialized searches may also use S5/VSR1 as well as "AstroWatch" data. This document mentions Several analyses and investigations that are ongoing or planned activities of Burst Group members. In addition, the group is open to new, scientifically motivated activities that will support its goals: the analysis techniques and tools that we continue to refine over the next few years will form the foundation of burst data analysis in the Advanced LIGO/Virgo era. GEO-HF data collected before that time will also provide a testing ground for evolving techniques, as well as good results in certain cases, such as searches for high-frequency signals with external triggers.

The Burst Group's ultimate goal in S6/VSR2+3 is the unequivocal detection and astrophysical interpretation of gravitational-wave burst signals. In the absence of any detection, our goal is to set physically significant upper limits on the gravitational-wave emission and astrophysical or cosmological mechanisms responsible for that. There are several approaches to the pursuit of these goals. The broadest net is cast by the so-called **untriggered burst search**, which "listens" to the whole sky, at all times allowed by the instruments operation, for statistically significant excursions of signal power within their sensitive frequency band (sec. 4.3.1). In some cases, astrophysically motivated assumptions about sources—GW waveforms or the source population—point to **customized untriggered searches** which are better than the general untriggered search in terms of sensitivity and ability to interpret the results (sec. 4.3.2). Another direction of interest is to look for gravitational-wave bursts starting with energetic transient astrophysical events observed in the electromagnetic or neutrino spectrum. The known sky position and time of these events, combined with the source phenomenology, offer sensitivity improvements and rich potential for interpretation within an astrophysical context. These are the so-called **externally triggered burst searches** (sec. 4.3.3). We also plan to symmetrize this relation by **seeking electromagnetic counterparts** for gravitational-wave burst candidates. For this we expect to form collaborations with observer groups and pursue Target of Opportunity (ToO) observations that will provide the possibily for gravitational-wave burst candidates to trigger non-gravitational-wave observatories (sec. 4.3.4).

In some cases, an overlap in astrophysical targets and methods exists with other LSC-Virgo groups. Such cases include short-duration gamma-ray bursts and the merger and ringdown of binary compact objects, which are pursued by both Burst and CBC Groups. Longer duration transients (∼minutes or longer) may also be pursued using methods developed by the stochastic and continuous waves groups for their corresponding searches. The Burst Group will coordinate with other working groups in areas of common interest to ensure that the best possible scientific results and corresponding publications are brought forward.

4.1 Recent observational results

The LSC-Virgo Burst Group has published a number of scientific results in the past year. A brief overview of these results is provided in this section.

During the S5/VSR1 run, over 130 gamma-ray bursts (GRBs) were detected by satellite-based gamma-ray experiments during times when two or more GW detectors were operating; a joint LIGO-Virgo search for GW bursts associated with those GRBs was published in 2010 [110]. This was a coherent network analysis excess-power type search for gravitational-wave bursts associated with 137 GRB events which occurred during the fifth LIGO science run S5 and the first Virgo science run VSR1. The gravitational-wave data were collected from 2005 November 4 to 2007 October 1, and most of the GRB triggers were from the Swift satellite. The search found no evidence for gravitational-wave burst signals associated with this sample of GRBs. Using simulated short-duration (< 1 s) waveforms, upper limits were set on the amplitude of gravitational waves associated with each GRB. Lower bounds on the distance to each GRB were also estimated under the assumption of a fixed energy emission in gravitational waves, with a median limit of $D \sim 12$ Mpc $(E_{\rm GW}^{\rm iso}/0.01 {\rm M}_{\odot} {\rm c}^2)^{1/2}$ for emission at frequencies around 150 Hz, where the LIGO-Virgo detector network has best sensitivity.

Untriggered all-sky searches over the course of LIGO's fifth science run have resulted in several publications over the last year. The results obtained from the first calendar year of the LIGO S5 run are described in two back-to-back papers detailing analysis below 2 kHz [111] and between 2 and 6 kHz [112] respectively. S5 is the first time untriggered burst analyses have extended above 3 kHz. All-sky results from the second calendar year of S5 and Virgo's first science run (VSR1) were published in a single paper several months later [113]. The combined search for untriggered bursts over the entire S5-VSR1 science run allows us to place a limit of 2.0 events per year on detectable bursts in the 64-2048 Hz band at 90% confidence and place limits on the root-sum-squared strain amplitude in the range $6 \times 10^{-22}\,{\rm Hz}^{-1/2}$ to $2 \times 10^{-22}\,{\rm Hz}^{-1/2}$. In terms of astrophysical reach, it is estimated that the sensitivity near the optimal frequency (around 153 Hz) is sufficient to detect a mass conversion of $1.8 \times 10^{-8}\,{\rm M}_{\odot}$ at a distance of 10 kpc or $0.046 {\rm M}_{\odot}$ at 16 Mpc with 50% efficiency. Using models produced by Ott et al [114], supernova core collapse gravitational wave emission was estimated to be detectable at a range up to 30 kpc, for one of the more promising models (s25WW). These combined S5-VSR1 results represent the most sensitive untriggered GW burst search performed so far and includes the first untriggered burst search during the long-term operation of a worldwide network of interferometers of similar performance at three different sites.

The first targeted search for cosmic strings was published in 2009 [115]. This was a matched-filter search for predicted cosmic string cusp gravitational-wave bursts using LIGO data from the fourth science run S4 (February and March 2005). No gravitational waves were detected in 14.9 days of data from times when all three LIGO detectors were operating. The search set frequentist loudest event upper limits on the rate of gravitational wave bursts from cosmic string cusps, thereby constraining the parameter space (string tension, reconnection probability, and loop sizes) of cosmic string models. While the sensitivity of this search did not allow constraints as tight as the indirect bounds from Big Bang Nucleosynthesis, analyses using data from future LIGO runs are expected to surpass these limits for large areas of the cosmic string parameter space.

4.2 Basic methodologies used in burst searches

This section summarizes some of the basic approaches that are used to search for GW bursts. We explain how we can search for signals without knowing their form, how we select good-quality data for burst searches, how we evaluate the remaining "background" rate of false triggers from detector noise fluctuations, and how we estimate the sensitivity of our searches for plausible GW signals.

4.2.1 Signal extraction methods

Burst searches target detection of gravitational waves from violent events in the Universe: supernovae, gamma ray bursts, mergers of binary systems and other sources. In most cases little is known about the anticipated GW waveforms, therefore burst searches use a variety of methods to find any transient signature in the data which is inconsistent with the baseline noise level. Also burst searches rely heavily on coincidence between multiple gravitational wave detectors. A "search pipeline" generally consists of one or more major signal processing algorithms along with post-processing and diagnostics. A pipeline produces a set of "triggers" which, if they pass the full set of significance tests and consistency checks, are considered to be candidate GW events.

In a few special cases when an accurate signal model is available, such as for cosmic string cusps, a search can be done using matched filtering with a bank of templates. Otherwise, un-modeled bursts can be identified in the detector output data as excess-power events localized in time and frequency. Therefore, for better localization of burst events usually the analysis is performed in the time-frequency (TF) domain. To obtain a TF representation of data a number of transformations are used, including windowed Fourier transforms, discrete wavelet decompositions [116] (Symlets, Meyer wavelets) and continuous wavelet transforms [117] (Q-transform). These transformations have been actively used in GW burst search algorithms. At the same time, the Burst Group is exploring other interesting approaches such as the Hibert-Huang Transform (HHT) [118], an adaptive time-frequency decomposition which can be used for more detailed study of the TF content of gravitational wave events.

A handful other burst search methods have been implemented which do not start with a TF representation. These include a change-point analysis of bandpassed data [119] (BlockNormal) and a cross-correlation analysis using pairs of detectors [120] (CorrPower).

Networks of GW detectors are very important for the first detection of gravitational waves. Detection of a consistent signal by multiple instruments will increase confidence in an event, especially for a burst signal which otherwise may not be distinguished from the instrumental and environmental artifacts produced in the detectors. Also, with multiple detectors it is possible to reconstruct the two gravitational wave polarizations and determine the direction to the source. For these reasons, the Burst Group has developed multi-detector search algorithms which can be approximately devided into two groups: incoherent and coherent. In the incoherent algorithms [119, 117, 121, 122], excess-power events in individual detectors are identified and then a time coincidence is required between the events in different detectors. The incoherent algorithms are particularly useful for characterization of individual detectors and they are actively used for studies of enviromental and instrumental artifacts (see Section 4.2.2). Coherent algorithms are based either on cross-correlating pairs of detectors [123, 124] or on a more general approach called coherent network analysis (CNA).

Coherent network analysis addresses the problem of detection and reconstraction of gravitational waves with networks of GW detectors. In coherent methods, a statistic is built as a coherent sum over the detector responses and, in general, it is more optimal (better sensitivity at the same false alarm rate) than the detection statistics of individual detectors. It is a high priority of the group to continue development of more advanced CNA algorithms and apply them to burst searches. There are several CNA approaches employed by the Burst Group. The constrained likelihood method (coherent WaveBurst [125, 126, 127]) has been developed and used for the S4 and S5/VSR1 all-sky searches. Also a likelihood method for triggered searches (X-pipeline [128]) was used during the S5/VSR1 run. These methods have being upgraded and now useed for analysis of the S6/VSR2+3 data. The group is working on a Bayesian formulation [129] of coherent network analysis and on maximum entropy methods (MaxEnt [130]). The first version of the Bayesian CNA tool is being used for the S6/VSR2+3 analysis as part of the Omega [131] analysis pipeline. Also dedicated CNA algorithms which may use partial information about burst sources and incomplete source models are being investigated by the Burst Group (see Section 4.3.2).

Coherent algorithms enable not only the detection of gravitational waves, but also astrophysical analyses of GW signals and measurements of source properties, including reconstruction of the GW waveforms and source coordinates. These are necessary tools for the nascent field of GW astronomy. Prompt detection of GW signals and estimation of source coordinates enables coincident observations with other astronomical instruments, which can significantly increase the confidence of detection. Such measurements may not only aid the first detection of gravitational waves but also they will give us fundamentally new information about the GW sources and their population distribution. The Burst Group has made significant progress in the development of source localization methods. The coherent WaveBurst and the Omega Bayesian tool were extensively tested during the Position Reconstruction Challenge conducted by the burst group and used during the on-line run in December 2009 - January 2010. The group will continue to work on the source localization problem and apply the coordinate reconstruction methods to the advanced detector networks.

LIGO, Virgo and other ground-based gravitational wave detectors have a linear response to the gravitational wave strain at the detector sites. The inferred gravitational wave strain at each detector site thus amounts to the greatest information that gravitational wave detector observations can provide for the purpose of astrophysical interpretation. Even in the absence of an electromagnetic counterpart an inferred waveform can provide basic information about a source and its dynamics. With an electromagnetic counterpart the accessible physics and astrophysics expands exponentially. Waveform inference is thus a basic desiradatum of gravitational wave detection and the pursuit of robust and reliable reconstruction algorithms that provide the collaboration this capability is one of the Burst Group priorities.

4.2.2 Detector characterization

Data quality plays a key role in burst searches, where the false alarm rate is dominated by noise transients, or "glitches", which can happen with similar morphology in multiple detectors and pass the coherence tests developed for the identification of gravitational wave candidates. Glitches represent the ultimate limit to the sensitivity of the burst search and to the confidence in a possible detection.

During the S6/VSR2+3 run, we are following the strategy of data quality and event-by-event vetoes originally developed in S5/VSR1. Many people contribute to this work, though it is still generally manpower-limited. In a coordinated effort of Burst and CBC Group members (the Glitch Group, a part of the Detector Characterization effort described elsewhere in this document), we study the correlation of the rate and strength of single detector transients to trends and transients in the sensing and control systems that maintain the interferometers in their operating condition, as well as monitors of the physical environment in which the interferometers sit: vibration, sound, magnetic and electric fields, power line voltages, and others. These studies led to the identification of times likely to be contaminated by non-GW effects, which the Burst Group uses to veto event candidates found by the GW channel analyses. Based on their duration and character, we distinguish between *data quality* vetoes and *event-by-event* vetoes.

Data Quality (DQ) vetoes are long time intervals (typically several seconds), during which auxiliary signals indicate that an interferometer was out of its proper operating condition. The vetoes are constructed from the DQ flags identified by the Detector Characterization group. Different flags have different correlation with transients in the gravitational-wave channel, thus we developed a categorization system for DQ flags to be used as vetoes:

- Category 1 vetoes define which data can be safely analyzed by the search algorithms; they remove features that could affect the power spectrum (for instance, detector instability prior to loss of lock or severe saturations).
- Category 2 vetoes define the *full* data set, where to search for detection candidates. They remove times when the detector is unambiguously misbehaving (examples are auxiliary channel saturations,

and power main glitches that couple magnetically in the detector) with a well-understood physical coupling. They introduce a small dead time (fraction of a percent) and have high efficiency for removing single-detector outliers.

- Category 3 vetoes define the *clean* data set to be used to set an upper limit in the case of no detection. They identify times with an excess of single-detector triggers in the gravitational-wave channel, but the correlation is not unambiguous and they may introduce a large dead time, up to 10%. Typically, these vetoes are associated with high seismic activity. If a detection candidate is found at these times, the category 3 flag is taken into account in the event followup, as described in section 4.3.6.

- Category 4 data quality are advisory flags: there is no obvious correlation with single detector transients, but they are known detector or environmental features, so they are to be taken into account in the followup of a candidate detection.

The classification of DQ flags into vetoes is based on their correlation with single-detector triggers from the *KleineWelle* algorithm and the *Omega* search. They are not tuned on multi-detector outputs, but once the tuning is complete, their effectiveness is tested on *coherent WaveBurst* triggers.

Event-by-event vetoes are short intervals (typically 100 ms or shorter) that mark individual transients in an interferometer's output with a coincident transient in one or more diagnostic signals. They are identified with extensive statistical studies of coincidence between *KleineWelle* triggers in auxiliary channels and event candidates, with a software package known as *Hierarchical Veto*, or *hveto*. *Hveto* considers many possible veto conditions in a hierarchical classification process in which the vetoes are ranked on the basis of the significance of their correlation with gravitational wave triggers, and their significance is re-evaluated in subsequent iterations, after each condition is applied. Their safety is tested against hardware injections.

To ensure the success of the online analysis, the Burst Group is collaborating with the Detector Characterization group for a prompt and documented definition of DQ flags and veto criteria. In particular, Omega single-interferometer triggers are produced online and time-frequency plots, as well as histograms and correlograms, will be available with a latency of 2-3 minutes. These plots are available for commissioning and science monitoring, and used for the identification of DQ flags. A standardized set of data quality flags are produced with low latency and are available for online analysis; additional studies and a periodic revision of vetoes and their categorization is performed for the offline burst analysis, as the understanding of each detector improves. The offline vetos are tested against coherent WaveBurst multi-detector triggers, especially relevant for understanding the background in the burst analysis.

4.2.3 Accidental background estimation

The key to any burst search pipeline is the discrimination between real GW signals and noise fluctuations which satisfy the analysis selection criteria, referred to as *false alarms* or *background*. The False Alarm Rate (FAR) or, alternatively, the False Alarm Probability (for the total observation time), depends on the properties of the detector noise and on the various selection criteria in the pipeline, including coincidence or coherence tests among the detectors. The FAR is normally assessed by constructing an *off-source* resampling of the observation data, i.e. switching off the possible effects of GW signals.

In a network of detectors with independent noise, the resampling is suitably performed by running the same signal processing on data sets built by *time-shifting* the data of the detectors relative to each other by more than the maximum light travel time between the sites. This destroys the coincidence for any real GW signals which may be present in the data and provides a reliable estimate of the accidental coincident background if the noise properties vary over time scales longer than the time-shifts. The procedure is repeated a number of times for different time shifts, so that each resample can be considered independent from the others. The sum of the resampled observation times should exceed the inverse of the target FAR by at least a factor of a few. In cases when the trigger time of the possible GW source is known, the background estimation procedure can take advantage of this information by defining off-source samples within a single detector.

The significance of event candidates found in the *on-source* (unshifted) analysis can be evaluated by comparison with the distribution of the accidental background, typically quantified with one or a few statistics describing the strength or quality of the signal. To have an objective, non controversial assessment of the FAR, we follow a *blind* statistical procedure, e.g. using the time-shifted data (without examining the unshifted data) to select the exact procedures to compute the test statistics and their thresholds.

This method cannot discriminate between GW signals, other foreground signals or correlated noise sources at different detectors. If the accidental events are not compliant with a Poisson point process, a problematic issue is the assessment of the uncertainty of the empirical off-source distribution, which propagates on the uncertainty of the False Alarm statements.

4.2.4 Simulations

While tuning an analysis pipeline, and after finalizing all of the event selection criteria, the *detection efficiency* of the observation is assessed to interpret the results in terms of different signal models. This is done using the *Mock Data Challenge* (MDC) technique, i.e. special runs of the analysis are done with simulated signals added to the actual detector data at pseudo-random times. The simulated signals are chosen to span the expected range of signal properties (frequency, duration, etc.) but not to exhaustively test all plausible signals; the robustness of the signal extraction methods allows to extrapolate to other signals. The detection efficiency is evaluated as a function of waveform and amplitude, either averaged over random sky positions (for all-sky burst searches) or at the fixed sky position of an astrophysical event used to trigger the search.

Systematic effects of calibration uncertainties on the detection efficiency are measured by performing MDC injections of signals with suitably mis-calibrated amplitude and phase (or time). These tests can be performed on subsets of the observation time to limit the computational load, since it is usual that a few-days subset is representative enough of the overall detection efficiency.

While the basic machinery of MDC injections is well-established, there is still much work to be done to expand the set of waveforms to include modeled or astrophysically motivated signals.

4.2.5 Hardware signal injections

We inject simulated signals into the interferometer hardware from time to time as an end-to-end test of the detector, data acquisition system and data analysis pipelines. By comparing the reconstructed signal against the injected one, we check the detector calibration. Hardware signal injections are also useful for establishing limits on the cross-coupling of loud GW signals into auxiliary data channels that might be used to define vetoes.

4.3 Science goals

Our fundamental goals are to detect gravitational-wave signals and, using them, test the general theory of relativity and learn new things about astrophysical objects. Looking into the future, we expect the advanced detectors to capture GW signals on a regular basis, enabling statistical analyses of source populations and emission mechanisms. But in the near term, our focus is on discovery—the identification and validation of that first signal (or, if we are lucky, first few signals) that will begin that era. The Burst Group's special role, of searching the whole space of possible transient GW signals, compels us to utilize every handle we can to distinguish real signals (even those whose form is unknown) from detector noise fluctuations.

In this section we describe six broad goals which, taken together, form our vision of how to search for the full spectrum of possible GW burst signals over the next few years. They drive (or will drive) many aspects of the actual searches that we implement and carry out now and in the near future using data from S5, VSR1, S6, VSR2 and VSR3. These six goals are not ordered by importance, and in fact many search efforts will draw on more than one of them. We have pursued many of these goals with S5/VSR1 data and intend to pursue all of them with S6/VSR2+3 data. The data analysis we have done as a collaboration over the past several years has provided us with valuable experience and sophisticated tools and procedures in many of these areas, while in other areas there is still much to be developed and learned.

4.3.1 Search as broadly as possible for GWs

There is strong astrophysical motivation [132] for searching for burst-like gravitational-wave signals using ground-based laser interferometers. The emphasis has historically been on astrophysical systems for which the resulting burst waveforms are either poorly modeled or remain unknown, including (but not limited to) binary compact star mergers and core-collapse supernovae. In recent years numerical relativity calculations have offered significant information on the waveform accompanying binary mergers [133] as well as important new information on the features of signals accompanying core-collapse [134]. Burst sources with well-modeled waveforms include emission from neutron star ringdowns following a pulsar glitch, black hole ringdowns or cosmic string cusps (sec. 4.3.2).

Typical signal durations predicted by these models are from less than a millisecond to hundreds of milliseconds and with signal power in the frequency range from 50 Hz and up to few kHz. Various models of GW emission from core-collapse supernovae [134] and gamma-ray burst progenitors [135] may result in signals lasting up to several seconds. Although the best sensitivity of ground-based interferometers is achieved around $\sim$150 Hz, the instruments remain sensitive to within a factor of $\sim$10 over most of their frequency band (60–6000 Hz). Given the broad spectrum of astrophysical sources, signal morphologies and their uncertainties and the possibility of unanticipated sources, it is of paramount importance for the burst search to provide an eyes-wide-open approach capable of detecting the widest range of signals.

Our approach relies on generic search methods and on minimal assumptions about the signal morphology to provide detection capability for the widest possible range of sources. The search analyzes as much of the available data as possible from times when two or more detectors are running well, and no assumption is made on the direction and the source(s) of gravitational-wave bursts. This kind of untriggered burst search is often referred to as *all-times, all-sky*. The search is traditionally tuned in such a way that a very low number ($\ll 1$) of background events is expected over the duration of the observation. Thus, any signal events (foreground) resulting from this analysis constitute detection candidates for which further and exhaustive investigation is performed (sec. 4.3.6). The immediate result of this search is a statement on the rate and strength of gravitational-wave bursts at the instruments. We can also give an astrophysical interpretation to the results through order-of-magnitude estimates of the reach (distance to which a signal could be detected) for certain astrophysical models.

Care should be given in the number of methods formally invoked in pursuing this, as well as any, search. Multiple methods are generally expected to increase sensitivity to some part of the signal phase-space covered by gravitational-wave bursts but at the cost of increased false alarm rate. Efforts toward new method development should address this question quantitatively, before being considered as part of the group-wide search for bursts.

Due to practical considerations, the all-times, all-sky burst search has been split into "low" (∼50–2000 Hz) and "high" (∼2000–6000 Hz) frequency bursts.

A continued all-times, all-sky for GW bursts over the full frequency spectrum allowed by the instrument and motivated by source phenomenology will remain of the highest priority in the S6/VSR2+3 science run. The analysis of the S6/VSR2+3 data set (Summer 2009-Fall 2010, expected) has been split into few-months (calendar) intervals. This has mostly been dictated by changes in the individual detectors' and the global network's operational characteristics. The search in S6a-VSR2 (Summer 2009) has been completed while the one in S6b-VSR2 (Fall 2009-Winter 2010) is in the final stages and remains the highest priority of the group. A single paper describing the search over the entire frequency regime and over the entire S6/VSR2+3 data set is expected be submitted in 2011.

In addition to this high priority search, the group will give expand analysis coverage to signals that several seconds. There are also questions surrounding the search for bursts that go beyond the all-sky "counting experiment" approach used so far. For example, can the distribution of weaker sources be revealed through the statistical analysis of the amplitude distribution of signal and background gravitational-wave burst candidates? Is the gravitational-wave burst sky uniform? How can the ability of the detector network to reconstruct event sky positions be used to perform sky map comparisons between foreground and background? Such questions can in principle improve sensitivity and provide hints for possible astrophysical models behind candidate signals. Moreover, they have been used routinely by cosmic ray experiments in the search for astrophysical sources of high energy photons and neutrinos [136, 137, 138, 139]. The availability of S5/VSR1 data provides an opportunity for the exploration and benchmarking of these search methods prior to applying them to data from the S6/VSR2+3 run.

4.3.2 Use knowledge of astrophysically plausible signals

Burst sources are un-modeled by definition. However, recent progress in astrophysics and numerical relativity provide knowledge about the intensity and morphology of expected gravitational wave signals from some sources; the most plausible being the coalescence of compact objects, over a broad parameter space, and core collapse supernovae. Predictions for binary black hole (BBH) coalescence rates are based on population synthesis models constrained by astronomical observations. Due to a number of assumptions and unknown model parameters, the rate estimates are still uncertain at present. According to estimates compiled from the literature [140], assuming that an Initial LIGO detector has a horizon distance of 173 Mpc for a merger of two 10-solar-mass BHs, it is likely to detect 0.01 BBH coalescences per year, with a plausible range between 0.0009 and 2 BBH coalescences per year. Assuming that an Advanced LIGO detector has a horizon distance of 2186 Mpc for such a system, it is likely to detect 30 BBH coalescences per year, with a plausible range between 2 and 4000 BBH coalescences per year. The expected rate of core collapse supernovae is 1/40-50 years in the Milky Way, about twice that in the local group, but at ∼5 Mpc the integrated rate may be as large as 1 every other year and reaches one a year at 10 Mpc. These are upper limits, supported by electromagnetic observations [141]. The actual range of burst searches for supernovae in LIGO/Virgo depends on how much gravitational wave energy is emitted, which is subject of a wide range of predictions [134]. The Burst Group aims to exploit such knowledge for an astrophysical interpretation of its results.

Burst searches do not normally use these calculated waveforms as templates for matched filtering. However, available waveforms are used to test the efficiency of the detection pipelines and quantify the reach of the search. In particular, the coalescence of black holes, whose waveform only has a few cycles in the interferometers' sensitive band, is of particular interest for the Burst Group. The IMR (*Inspiral-Merger-Ringdown*) working group is a collaborative effort with the CBC Group, where the same set of phenomenological and EOBNR waveforms [142, 143, 144, 145] is being analyzed by burst search methods, inspiral and ring-down matched filtering to compare detactability across the parameter space. This includes an event-by-event comparison of events found by each search (simulated and background) and a study of how to combine inspiral, burst and ringdown triggers in a single result. This study currently includes systems with 25-350 $M_\odot$ total mass and no spin. Preliminary studies based on numerical relativity waveforms show that the magnitude and orientation of the spin of the coalescing black holes has a significant impact on the range of burst searches. Also binary black hole objects can be formed in the galactic nuclei or globular clusters and enter the frequency band of detectors at high eccentricities [146, 147]. Studies of corresponding waveforms will help to devise new detection techniques and provide important guidance for the interpretation of search results. Similar studies are being pursued using current models for the gravitational wave emission from core-collapse supernovae.

In preparation for a detection, the group needs to finalize parameter estimation techniques, which allow to reconstruct the waveform, compare it to known models and extract source parameters. Some progress towards waveform reconstruction has already been made in this direction, via coherent techniques, but more progress is needed to compare a candidate to merger waveform parameters. Bayesian [148] and MCMC [149] techniques are currently being explored by the CBC Group as well as members of the Burst Group. New techniques such as the Hilbert-Huang transform [150, 118] may also prove to be useful for reconstructing waveforms.

For some targets, a customized search may be more appropriate than the all-sky flagship analysis of the Burst Group. Such searches may be pursued after preliminary studies show them to be significantly more sensitive than the flagship analysis or they provide scientifically more meaningful interpretation of the search results. In certain instances, the existing all-sky pipelines can be tuned to target a specific signature. For example, elliptical constraints can be applied to *coherent WaveBurst* to search for elliptically polarized signals such as black hole mergers [151]. Or the configuration file for *Omega* could be modified to target high frequency, long waveforms from certain models for core collapse supernovae [152, 153, 154, 155, 156]. Matched filtering can be used for some classes of GW burst signals, such as cosmic string cusps [115] or specific models of supernovae. In other cases, customized algorithms may be needed to search for certain extended signals, such as trains of quasi-periodic pulses from eccentric binaries.

Customized burst analyses are also used to exploit the source sky locations. The group developed and implemented techniques to search for gravitational waves in coincidence with electromagnetic events such as GRBs, SGRs and optical supernovae, for which the sky location is known; see section 4.3.3 for a full discussion. Dedicated searches have been designed and will be used to target specific sky locations, including but not restricted to the galactic center or M31. In addition, the Burst Group will use directional information from burst candidate triggers to form sky map probabilities of their astrophysical origin. A comparison of such maps with background could unveil directional correlations with known or new point sources, as well as correlations with ensembles of galactic and extragalactic sources (as has been done by the Auger Collabration [139], for instance). This directional anisotropy search, in preparation for S6/VSR2+3, relies on triggers and direction information from the all-sky searches. It may ultimately be optimized for each possible arrival direction, like a *radiometer*.

4.3.3 Look for GW transients triggered by energetic astrophysical events

Many gravitational-wave sources should be observable in more traditional channels. Directly relevant astrophysical observations are abundant, ranging 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 event increases the sensitivity of the GWB search compared to an untriggered all-sky search. The association with a known astrophysical event may be critical in increasing our confidence associated with a candidate GWB detection. Perhaps most importantly, joint studies can enable scientific insight that cannot be accessed through gravitational waves or other messengers alone.

A high priority for S6/VSR2 has been to invoke in near-real-time the baseline GRB search [157, 158] used in S5/VSR1. The prompt availability of results from reviewed and published analysis methods such as the GRB and SGR [159, 160] pipelines will allow relatively fast publications on exceptional astrophysical triggers.

ExtTrig members are also working towards increased interaction with external astronomers, including through workshops (such as the 2009 GW+HEN workshop in Paris, the 2010 GRB multimessenger meeting in Cardiff, and the 2010 SNe workshop in Caltech) and by regularly hosting presentations by non-LVC astronomers at weekly ExtTrig telecons. One example of the benefits of these interactions has been that, in response to requests from non-LVC GRB astronomers, a public webpage listing GRBs for which we could potentially perform GW analysis is under construction. This webpage will contribute to the GRB astronomers's decision of whether to perform follow-up observations for a given GRB.

ExtTrig sources and science mentioned in this section will also be relevant for the aLIGO era, as discussed in Section 4.5.

Triggers from Gamma-ray bursts (GRBs)

GRBs are intense flashes of gamma rays that are observed approximately once per day, isotropically distributed across the sky. GRBs are divided into two classes by their durations [161, 162]. Long ($\gtrsim$2 s) GRBs are associated with star-forming galaxies of redshifts of $z \lesssim 8$ and core-collapse supernovae [163, 164, 165, 166]. Short GRBs ($\lesssim$2 s) have been observed from distant galaxies of different types. Most short GRBs are believed to be due to the merger of neutron star or neutron star – black hole binaries [167, 168], while up to 15% may be due to SGRs [169, 170] (see below).

The central engine of both long GRBs and short GRBs produced by binary mergers (non-SGR GRBs) is likely a rapidly rotating black hole with accretion disk. Gravitational waves could be produced in long GRBs, for example, from fragmentation in the disk [171], or suspended accretion due to interactions between the BH spin and magnetic fields [172]. Binary mergers should also emit strong gravitational wave signals with relatively well-modeled amplitude and frequency evolution. The current LIGO-Virgo network is sensitive to GWB emission from these models to distances of $O(10)$ Mpc [157]. This compares favorably with estimates of the local rate density of low-luminosity GRBs (up to $10^3\,\mathrm{yr}^{-1}\mathrm{Gpc}^{-3}$ [173]), so that the enhanced or advanced detectors may detect gravitational waves associated with GRBs or constrain these progenitor models.

ExtTrig searches analyze all GRBs, both long and short, for which we have data from two or more detectors. The search algorithm looks for a GWB coincident in time and sky direction with the external trigger, and is designed to be sensitive to generic GWBs. The burst analysis of GRBs is complementary to the compact binary coalescence search, providing the ability to detect or constrain any GW emission that does not precisely match the inspiral waveform.

We have completed the analysis of all S5/VSR1 GRBs reported by the GCN [174] with coincident LIGO/Virgo data [157]. These were GRBs which had reported localizations, mostly from Swift, and for which there was good science data from at least two LIGO/Virgo interferometers. Since the completion of this analysis, we have started a collaboration with IPN (Interplanetary Network) to localize additional GRBs which were detected by IPN satellites during the time of the LIGO S5/VSR1 run. We estimate that up to ~ 300 additional GRBs will be available for analysis using S5/VSR1 data as a result of this collaboration. Given the good data quality and IFO sensitivity during the LIGO S5/VSR1 run, there is good scientific motivation for pursuing the analysis of these additional IPN GRBs. Further, pursuing this analysis will strengthen the collaborative relationship between LIGO/Virgo and IPN, and will help move it forward into the era of advanced detectors.

The analysis of AstroWatch GRBs using G1, H2, and V1 data is underway and will be completed by the end of 2010. A separate search is being conducted for GRB 051103 [175], a short GRB whose sky-position error box overlaps with the nearby galaxy group M81 (3.6 Mpc). Initiated by contact with non-LVC GRB astronomers, the 051103 search is of particular interest since this GRB it has properties indicative of an extra-galactic SGR giant flare. A joint search involving the burst, inspiral, and SGR pipelines is nearing completion and will contribute to the discussion of this event's progenitor.

Our primary goal for S6/VSR2 is to perform a fast search for all available GRB triggers. Currently all GCN-reported GRBs with coincident data are being processed autonomously by the online search. A paper draft covering the full S6/VSR2 GRB set will be written within 3 months after the end of the S6 run.

Low threshold gamma-ray triggers

Gamma-ray satellites such as Swift typically set thresholds for GRB detection high enough to ensure that the rate for false events is very small. Hence, there are likely to be good GRBs which are below the standard threshold. And there is good scientific motivation to conclude that GRBs which are not bright in gamma rays are not necessarily associated with low GW brightness. A clear example is provided by short-duration GRBs, which typically have much lower gamma-ray luminosity than long GRBs, but, given the current progenitor hypotheses, are expected to be much brighter in GWs. The other example is given above—GRBs for which the Earth is not centered on the gamma-ray beam. In either case, such low brightness GRBs could easily by relatively nearby. By examining GRB candidates which are below the standard threshold, one might hope that this expanded sample would include promising GRB-GWB candidates.

Currently a joint analysis of such an expanded-sample of gamma-ray bursts from Swift during the S5/VSR1 run is in development.

Triggers from Soft Gamma Repeaters (SGRs)

SGRs emit short-duration X-ray and gamma-ray bursts at irregular intervals, and (rarely) giant gamma-ray flares with luminosities up to 10^{47}erg/s [176]. Crustal deformations and catastrophic cracking [177, 178] and excitation of the star's nonradial modes [179, 180, 181] might lead to the emission of gravitational waves [180, 181, 182, 183, 184]. The pulsating X-ray tail of SGR1806−20 also revealed the presence of Quasi-Periodic Oscillations (QPOs) (e.g. see [185, 186, 187]) that were plausibly attributed to seismic modes of the neutron star, thus bringing up the possibility of associated gravitational wave emission [160, 188].

Less than two dozen Galactic SGRs and AXPs have been identified, at distances from $\sim$ 1.5kpc to $\sim$ 50kpc. SGR 1806-20 and two others (SGR 1900+14 and SGR 1627-41) are located in our Galaxy near the galactic plane, between 6 and 15 kpc distance. Another, SGR 0526-66, is located in the Large Magellanic Cloud, about 50 kpc away. Two others, SGR 0501+4516 and SGR 0418+5729, may be an order of magnitude closer to us than the other known Galactic SGRs located towards the galactic anticenter. Advanced LIGO and Virgo will probe the energetics of these Galactic sources several orders of magnitude below the typical electromagnetic energy output of giant SGR flares, constraining the unexplored and interesting region [181, 182, 183, 184].

For SGRs, we also perform two specialized searches. "Stacking" searches look for the collective signature of weak GWBs from repeated flares from a single SGR source. "QPO" searches look for long quasi monochromatic GWBs, with special emphasis on frequencies and durations matching those seen in the electromagnetic emission.

Two SGR searches on S5/VSR1 data have been published: the analysis of the first year, and a stacking analysis of the 2006 SGR 1900+14 storm. A paper is in preparation reporting the results of the analysis of flares from the second year of S5/VSR2 and Astrowatch, including from the recently discovered SGR 0501+4516, thought to be at a distance of only 1-2 kpc. The next priority will be the QPO search, triggered searches for GWBs associated with SGR flares during S6/VSR2-3, and a stacking analysis of isolated bursts from the nearby magnetars SGR 0501+4516 and AXP 1E 1547.0-5408.

Neutrinos as Astrophysical Triggers

Gravitational waves and high-energy neutrinos can originate from a common, very energetic source that exhibits burst activity. For example, GRBs are thought to be high energy neutrino emitters as implied by the internal shock model [189, 190, 191, 192, 193, 194, 195, 196, 197, 198, 199]. In the prompt and afterglow phases of long GRBs, high-energy neutrinos are expected to be generated by accelerated protons in relativistic shocks [189, 200, 201]. Short GRBs [168, 167, 202] are also thought to be associated with a detectable high-energy neutrino flux. "Low-luminosity" GRBs are associated with a particularly energetic population of core-collapse supernovae [166, 203, 164, 204, 205, 206, 207]. Typically discovered at relatively small distances (e.g. SN 1998bw at about 40 Mpc), the expected rate of these events in the local volume can be an order of magnitude larger than that of conventional long GRBs [208, 207] and strong neutrino emission is also expected from these sources [199, 209, 210]. A detectable number of neutrino events are additionally expected from the so-called "choked" GRBs [211, 212, 213, 214, 215]. They are particularly interesting as they can be hidden from conventional observational methods, therefore only neutrinos and gravitational waves can reveal their properties. SGRs may also be considered as common sources for high-energy neutrinos and gravitational waves [181, 216, 217].

The IceCube [218] and ANTARES [219] neutrino observatories have extragalactic reach in the high energy regime and they can determine both the arrival time and the direction (with complementary sky coverages) of individual neutrino events. These detectors can also reconstruct the energy of individual neutrino events [220], which can enhance background rejection as the spectrum of astrophysical neutrinos differs from that of atmospheric events [220]. The LIGO-Virgo detector network also provides extragalactic reach and relevant pointing information. A gravitational-wave – high-energy neutrino search therefore may significantly expand the scientific reach of all experiments. The past, present, and projected overlaps in the observation schedule of the collaborations ensures concurrent gravitational-wave and high-energy neutrino observation time.

The feasibility and expected performance of a hypothetical IceCube-LIGO-Virgo network was demonstrated [221, 222], indicating that the false alarm rate for the combined detector network is expected to be very low, $O(10^{-2})$ yr^{-1}. Closed-box analyses using fake HEN triggers and real LIGO-Virgo S5/VSR1 data have also been performed as the first step in the implementation of a full joint search.

Supernova triggers from detectors sensitive to low energy (10s of MeV) neutrinos can be used to initiate a search for associated gravitational waves. 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 [223] with a pointing accuracy of 4°. Unlike the electromagnetic signatures of supernova, neutrino bursts and gravitational waves mark the moment of core collapse, and are expected to be coincident in time to < 1 s. In fact, neutrinos and gravitational waves provide the only direct probes of the dynamics of core-collapse.

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. In general, a joint neutrino detection would boost confidence in a low-SNR GW detection.

We have studied various scenarios for a joint GW-neutrino supernova search, particulary with Super-K. The most recent supernova search by Super-K required two neutrinos in coincidence within a 20-second window, and an energy threshold of 17 MeV, at which the background rate is ~ 1 per day. A core-collapse supernova at Andromeda would have a $\sim 9\%$ probability of satisfying this search criterion. Requiring only one neutrino, e.g. in coincidence with a GW detection, would increase this probability to $\sim 35\%$ [224]. Decreasing the energy threshold to 7 MeV would result in detection probabilities of $\sim 15\%$ and $\sim 50\%$ for a two-neutrino and single-neutrino requirement, respectively.

At this lower energy threshold, the background rate is ~ 100 per day, but requiring coincidence with a GW event in a joint GW-neutrino search would keep the effective false alarm rates low [224]. For example, if we consider the expected background rates at Super-K which would result from energy thresholds at which the expected neutrino count at Super-K from SN sources at various distances (corresponding to the energy thresholds) is one neutrino event, and if we use the LIGO S5 first-year background rate threshold of $\sim 6.2 \times 10^{-4}$ per day, then the probability for a GW and a neutrino background event to be coincident, using a coincidence window of 1 second, would be $\mathcal{O}(10^{-6})$. If the GW background rate threshold is increased to ~ 1 per day, corresponding to an improvement in GW energy sensitivity of a factor ~ 2 at 150 Hz, then the coincidence probability would be $\mathcal{O}(10^{-3})$ for a single-neutrino requirement. Using a coincidence window of 20 seconds, and requiring two neutrinos in coincidence with a GW event, would result in a coincidence probability of $\mathcal{O}(10^{-5})$ if the GW background rate threshold is increased to ~ 10 per day, an improvement of a factor of ~ 3 in GW energy sensitivity.

Optical Supernovae as Astrophysical Triggers

Core collapse supernovae are interesting candidates as gravitational-wave sources and can also be associated with neutrino bursts/triggers. For extragalactic supernovae the external trigger must be derived from data from an early optical detection, leading to a large uncertainty on the derived event time (order of several hours), making the data analysis task challenging. The information on distance and direction is often well-known thus we can use directional search analysis methods. The large number of extragalactic supernova discoveries makes this line of analysis a possibility, even though theoretical motivation is still evolving.

Relatively close-by core collapse supernova events were observed during S5 and S6 in an early state of their evolution. Therefore reconstructing supernova lightcurves can place a bound on the core collapse time (trigger time) and therefore on the analysis time window. The analysis method for S5/VSR1 supernova is in development, and is expected to continue during S6/VSR2/VSR3.

Pulsar Glitches as Astrophysical Triggers

Pulsars emit periodic pulses of radio-frequency radiation. The periodicity of these pulses is, in general, extremely stable but, on many occassions, a sudden change in the periodicity is observed by radio astronomers. These sudden changes, often referred to as *glitches*, are most likely due to a increase in coupling between the solid outer crust and the liquid inner core. Put more simply, the outer crust "sticks" to the core momentarily and begins to rotate faster for a short period of time.

It is possible that a pulsar glitch will be accompanied by a burst of gravitational waves strong enough to be observed by Advanced LIGO and Virgo and possibly by the current "enhanced" detectors. The most likely signal waveform is a quasi-sinusoidal ring-down from the excited neutron star. A precise prediction of the time at which this signal should be seen is important for maximizing sensitivity and reducing the search parameter space. There are several pulsar timing programmes worldwide and, with varying degrees of precision, the times of glitches from pulsars are readily available.

A search for a GWB associated with the August 2006 Vela pulsar glitch has been performed, using a new Bayesian model selection algorithm [225], and a paper is in the advanced stages of internal review. This method will be applied to pulsar glitches of interest during S6/VSR2, and possibly to S5/VSR1 glitches from the X-ray pulsar J0537−6910.

Population study of astrophysical sources using data from other detectors

The objective of this analysis is to infer properties that characterize a population of astrophysical sources as a whole rather than any one individual member. The sources could have an electromagnetic or particle signature (i.e., external triggers) or could be candidate GWB events obtained from detectors that work in a different frequency band than LIGO and Virgo.

The first version of a population study method for GRB triggers was investigated in [226] and an improved one [227] was demonstrated using data from the S2, S3, and S4 runs of LIGO [228]. The population study method is complementary to the search for GW signals in coincidence with individual triggers [229]. An important feature of the method is the accumulation of signal-to-noise ratio as the sample size of triggers increases and a corresponding monotonic improvement in the astrophysical constraints obtained [226, 229].

It is foreseen that the cumulative use of GRB triggers from S5/VSR1 and the forthcoming S6/VSR2 runs will be significant. Among the improvements will be the use of coherent network analysis and a Bayesian scheme for assigning optimal weights to the triggers according to their sky location, measured redshifts (if any) and other relevant parameters obtained from the non-GW observations.

RXTE-LIGO coincidence search for GWs from Sco X-1

The Rossi X-ray Timing Explorer (RXTE) mission has been extended to the end of 2010 and observations therefore overlap with the S6/VSR2/VSR3 run. Multi-messenger observations of Sco X-1 allow us the opportunity to search for correlations between the X-ray and potential GW emission. Sco X-1 was the first extra-solar X-ray source to be discovered and there are still many unanswered questions regarding this object: What is the spin frequency? What mechanism is at work during low- and high-frequency quasi-periodic oscillations (QPOs) [230]? What is the limiting mechanism on the increase in spin frequency due to the accretion torque [231]? Recent papers have suggested that some of these issues may be related to gravitational radiation [232, 233, 234, 235]. In addition, gravitational-wave bursts from f-mode excitation may be induced by r-modes [236] or crust collapse due to mass accretion [237]. An exploratory cross-correlation analysis between RXTE and LIGO has been performed using S5 coincidence data. In addition to this ongoing analysis a second approach is under way where both RXTE and LIGO-VIRGO time-series data are analyzed to generate trigger lists using existing search pipelines developed for gravitational-wave burst detection. LIGO-VIRGO data is analyzed by a coherent network analysis pipeline and RXTE data by a single detector pipeline based on the excess power method. Masking the sky around Sco X-1, using a coincidence analysis can improve the detection confidence, and therefore enable us to set upper limits on GWs correlated with astrophysical events encoded in RXTE data.

Radio triggered searches

The existence of theoretical models in which various mechanisms give rise to a prompt pulse of radio emission from some putative gravitational wave sources, particularly coalescing compact binaries, motivates joint searches with radio telescopes.

Models for prompt radio emission from compact binaries generally require that one of the compact objects is a magnetar. In the simplest of these model classes [238, 239], the orbital motion of the binary generates time dependent magnetic fields and consequently induced electric and magnetic fields. These fields then lead to the emission of radiation, which is predicted to be in the radio band.

A second, larger class of models similarly require a high magnetic field from one member of the binary; in these models the field either confines or otherwise interacts with a plasma. For example, the unmagnetized companion object can develop surface charge that can then be ejected from the surface of the star and subsequently undergo acceleration as it follows the magnetic field lines of the magnetar. Alternatively [240], gravitational waves emitted during the inspiral and merger of a binary neutron star system may excite magnetohydrodynamic waves in the plasma (see also [241] for details), which then interact with charged particles in the plasma, inducing coherent radio emission.

A progenitor confirmation is still awaited for short hard GRBs (SHB). Results have been published on radio afterglows for SHB but the data shows only weak signals hours or days after the burst. There are a series of proposed models for radio afterglows for SHB that may be observed seconds to minutes after the burst. One of the models put forward by Usov and Katz [242] predicts that immediately after merger the rotational energy of the binary system is transferred to a highly magnetized, highly relativistic particle wind that interacts with the ambient warm gas and as a result EM radiation is emitted. The main bulk of radiation is below 1 MHz but its tail can reach 1–30 MHz. As a main prediction, according to the model we should detect an incoherent pulse of 1–30 MHz, duration of 1–100 s and fluence of a few percent of the total gamma-ray flux. The expected delay for such a pulse would be of around 10^4 s for a source placed at 3.2 Gpc.

Follow-up searches of radio triggers in gravitational wave data will allow us to dig deeper into our noise by focusing on short astrophysically interesting times. We expect these searches will increase our chances of finding gravitational waves. We propose to perform follow up searches in our data, starting with existing radio transients detected during S4 and S5/VSR1. An interesting aspect of follow-up of radio triggers is that for each event we will have the dispersion measure. This will provide an independent measure of the distance, and allow us to better predict when the gravitational wave should have arrived at our detectors.

4.3.4 Seek confirmation of GW candidates with follow-up observations

The previous section described many scenarios in which astrophysical systems are expected to emit electromagnetic (EM) radiation along with GW bursts, and the substantial benefits of joint observations. However, 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 [243]; 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.

GRB progenitors provide an intriguing example. While Fermi and other satellites have gamma-ray detectors which can see a large fraction of the sky, the gamma-ray emission is expected to be tightly collimated, so that only a small fraction of GRB progenitors are detected by these instruments. The corresponding afterglow emission at longer EM wavelengths, on the other hand, is believed to be much less collimated [244, 245]. Thus there may be short-lived "orphan afterglows" from nearby energetic events which are going undetected [168]. 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.

Other possible sources of joint EM/GW emission include decaying neutron star matter ejected during merger [246, 247] and supernovas [248]. More details on these sources and their expected observational signatures can be found in [249, 243, 250].

In the event of the detection of a truly astrophysical signal, such observations will be greatly beneficial for two reasons. First, they will help establish the event as astrophysical in nature, effectively increasing the reach of the interferometer in the case of a low-signal-to-noise-ratio event that might not otherwise stand out from the background. Additionally, having an associated EM signal increases the astrophysical information that can be mined from a GW event. Several recent papers have stressed the progress that can be achieved only with a combination of EM and GW data [251, 252, 253]. This second motivation will remain strong even in the future advanced detector era, when GW signals are expected to be observed regularly.

During the current S6/VSR2+3 run, we are executing an ambitious program to seek confirmation of GW candidates with electromagnetic follow-up observations. This project is sometimes called "LoocUp", which originally stood for "Locating and Observing Optical Counterparts to Unmodeled Pulses" [249], but the scope has evolved over time to include EM bands other than the optical and to be a joint effort with the CBC Group, who wish to follow up promising candidates from (well-modeled) inspiral searches.

There are several challenges that are currently being addressed to make this possible. One difficulty is the relatively large uncertainty expected for interferometer signal localization. For the LIGO-Virgo network, uncertainties of order a few degrees are expected, which is large compared to the field-of-view of many astronomical instruments. For this reason, it is natural to plan to work with wide-field instruments, with fields of view of several square degrees or more. We have had detailed discussions with a number of groups of astronomers who operate robotic wide-field optical telescopes around the globe—many of which are already being used to look for afterglows to GRBs—and have formalized arrangements with some to collect images following up on LIGO-Virgo GW triggers. Smaller field-of-view instruments may also be useful for secondary observations, including spectra, or in conjunction with a catalog of nearby galaxies to limit the region to be imaged, as outlined in [249]. To this end, we have made arrangements with the *Swift* project to plan for a limited number of target-of-opportunity (ToO) observations using that satellite's X-ray and UV/optical telescopes, and we have applied for ToO observing time on a few large-aperture ground-based telescopes. Another approach being investigated is to "point" into normal survey data from Fermi, *Swift*, and/or RXTE afterwards to search for low-significance counterparts in these high-energy EM data sets.

An "online", low-latency, GW data analysis pipeline is a key component of any follow-up plans. The Burst Group has recently made this a reality, with a complete version demonstrated during periods of S6/VSR2+3, delivering triggers to robotic telescopes with roughly 30 minute latencies. This online pipeline will continue to develop through S6/VSR3, along with LUMIN and GEM, collections of specialized tools for identifying relevant events and passing them to automated telescopes for directed observation.

The identification and classification of EM transients in the collected images is a major challenge. Several group members are beginning to engage with our astronomer colleagues to understand and implement this complex process. Developing an image processing pipeline, evaluating its efficiency, and measuring the transient background in collected optical and X-ray images are goals that are closely linked with the EM follow-up plan.

Even in the absence of an S6/VSR2+3 detection, we expect this development effort to be useful. The first benefit of S6/VSR2+3 follow-up observations is the development itself; the huge pay-offs of measuring linked EM and GW signals merits developing software tools and protocols in anticipation of Advanced LIGO/Virgo. The second pay-off is the ability to set unique upper limits on astrophysical models predicting both GW and EM emission. The process for setting such an upper limit contains the optical transient search as an additional coincidence test, and so may prove more sensitive than a GW-only limit.

Another area currently being explored is the use of gravitational wave detectors as a trigger for follow-up radio searches, which could provide a method of detecting faint radio transients that would otherwise be missed. Furthermore, certain classes of radio telescope, such as the LOFAR array, employ aperture synthesis to allow a combination of a field of view of a large fraction of the sky and angular resolution as good as a fraction of an arc second. Such telescopes use digital signal processing to allow matching of the telescope field of view with the error circle of a gravitational wave trigger. The bandwidth of LOFAR is limited to the UHF (40-240 MHz) by the usable bandwidth of the antennae. For other conventional steerable dish telescopes, use of the gravitational wave channel as the search trigger makes optimal use of the wide field of view of gravitational wave interferometers to maximize the probability of transient detection, whilst conventional dish antennae have the advantage of covering a wide band of radio frequencies in a frequency band that is better understood. Detailed plans for extending LoocUp to the radio band are currently being discussed.

4.3.5 Be open to alternative theories of gravity

An intriguing possibility is that gravity may be better described by a theory other than General Relativity. The direct detection of gravitational waves could help determine the correct theory of gravity [254, 255, 256, 257]. This, in turn, could help unravel outstanding problems in astrophysics, including dark energy and dark matter, and even provide clues to reconciling gravity and quantum mechanics [258]. LIGO provides a unique opportunity to study the detailed properties of gravitational radiation.

Alternative theories of gravity may result in a difference between the speed of light and the speed of propagation of gravitational waves [259]. Coordinated electromagnetic and gravitational wave searches could place stringent constraints on the propagation speed of gravitational waves.

Alternative metric theories of gravity also predict extra polarization states, in addition to the +- and $\times$-polarization modes of Einstein gravity. Indeed, every other known viable metric theory of gravity predicts more than two polarization states [259], and the most general gravitational wave can have up to six polarization states [254, 255]. This will have an impact on multi-detector coherent gravitational wave searches because the linear combination of data streams that maximizes the gravitational wave content depends on the number and type of additional polarization states. If the direction of the gravitational wave is known in advance then disentangling which polarization states are present in a signal is straightforward [255], provided there are at least as many detector data streams as polarization modes. This is the case for externally triggered searches, and for searches where the gravitational wave triggers have sufficiently large signal-to-noise that triangulation can be used to pinpoint the sky-location. For all-sky searches new techniques need to be developed to separate out the expanded set of polarization states.

Evidence for the the existence of extra polarization states would be fatal for General Relativity. A non-detection, however, may not rule out alternative theories of gravity because the strength of the extra polarization states depends on the source and emission mechanism. The emission mechanism can also be different in alternative theories of gravity. In particular, different multipole moments (such as the dipole) contribute to the radiation.

One non-tensorial mode search pipeline based on a coherent network analysis is being developed for current operational detectors. Even in the absence of the detection the pipeline can be set upper limit on GR. This search is complementary to current ongoing searches: the standard search for the tensorial mode might miss some sources which produce non-tensorial modes such as highly spherical core collapse events.

4.3.6 Be prepared to detect a GW signal with confidence

The Burst Search uses a variety of tests to see whether a candidate signal should be considered a genuine gravitational wave detection, including:

- coherence tests among coincident signals from different interferometers (included near the front end of some pipelines, toward the back end of others),
- establishment of a low probability of false alarm, from extensive background studies, and
- various checks that the instrument outputs are unlikely to have been generated by the interferometers themselves or by spurious "pickup" from their immediate physical environment.

We build tests of these items into our pipelines. Coherence tests are a feature of most of our search methods. The Detector Characterization section, above, explains the vetoes that are used to handle the various sources of glitches that we have identified. These tests are tuned on and applied to background derived from time-shifted data before being used in any un-shifted analysis. In this way they are already reflected in our ranking for the quality of each event, and therefore they contribute to the measurement of false-alarm-rate.

Beyond what we have been able to automate, it is necessary to check on a case-by-case basis any detection candidate that appears at the end of our pipelines. The "detection checklist" outlines a follow-up procedure for these events. It is meant to be applied to an event which has any chance of being identified as a gravitational wave. The checklist as a whole is applied to gravitational-wave candidates and blind injections (which remain blind throughout the follow-up procedure). Due to the amount of work involved, it is not currently applied in full to known injections or time-shift background events, though many individual checklist items do involve such benchmarks. Currently the detection checklist consists of an 80-point list of items to be checked by hand covering the following categories:

- zero-level sanity: reports in detector logs, check hardware injections
- data integrity: frame file check-sum, undocumented injections, check against raw frames
- state of the instrument: obvious disturbances reflected in auxiliary channels, verify coupling for any proposed veto, check by hand against known disturbances: dust, cosmic rays, power fluctuations, acoustic, etc.
- event properties: construct detailed spectrogram, reconstructed waveform and direction, compare background from various methods, check signal consistency across interferometers
- astrophysical interpretation: check for external EM or neutrino events, catalog sources consistent with reconstructed direction, compare waveform against simulations

The checklist serves several functions:

- provides a central resource for basic information about the event
- serves as a careful secondary review of the specifics of the end-to-end analysis that relate to the event
- provides an opportunity to check for obvious reasons to dismiss a candidate event (e.g. clear environmental cause) or increase our interest in an event (e.g. optical counterpart)
- outlines event details for each individual detector which otherwise could be buried in multi-detector statistics

- provides a test-bed for new ideas which have not been able to make their way into the quantitative detection statistic

The checklist comes after the blind analysis is complete and does not make any quantitative or objective end statement about the baseline statistical significance of an event (which is predetermined). However, items which prove to be particularly useful at separating signal from background will have a strong motivation to be worked into the blind detection statistic for future analyses.

Successful application of the detection checklist builds confidence in our procedures and interpretation of their results. We continue to work with members of the CBC group and others in the LIGO Scientific Collaboration and Virgo Collaboration, to improve our ability to use the event follow-up procedure to hear the "ring of truth" in a gravitational wave detection.

4.4 First science targets for S6/VSR2+3

The S6 science run is currently providing the most sensitive LIGO data yet. Combined with the Virgo VSR2 run (up to January 2010), the network provided the best three-site data to date; the VSR3 run, expected to begin in July, should have even better sensitivity. We will carry out a full range of GW burst searches with this data set. Nevertheless, there are a few topics that the Burst Group has identified as our *first* science targets for S6/VSR2+3. These are analyses of broad interest for which we have well-developed methods ready to apply to the new data from the outset of the science run, specifically:

- Rapid search for GW burst signals associated with exceptional astrophysical events such as a nearby GRB, an SGR giant flare, or a nearby supernova.
- All-sky search for arbitrary GW bursts (up to $\sim$1 second in duration) with frequency content in the range 40–5000 Hz (or higher, if the high-frequency calibration is understood well enough).
- Search for GW bursts assocated with GRBs detected during the run.
- Search for GW bursts assocated with SGR flares observed during the run.

For the first of these, we wish to be able to prepare a paper reporting our analysis results within one month after the exceptional event. For each of the others, we intend to write a paper summarizing the results from the full S6/VSR2+3 run within three months after the end of the joint running period.

In order to meet these targets, we are actively engaged in analyzing the data *during* the run. We are performing initial analyses using the preliminary calibrated data, with the goal of only re-analyzing once (if possible) after the calibration is refined. The data is being analyzed with mature, well-understood analysis pipelines that we have already developed and reviewed, so that only incremental changes need to be reviewed to validate the new results. A major effort is ongoing to study data quality and develop veto conditions on a daily and weekly basis as the data is collected, tracking changes that inevitably occur. For the all-sky and GRB searches, preliminary partial results are being produced *during* the science run by freezing the data quality, vetoes, and event selection cuts and "opening the box" to check for event candidates.

While these first science targets will be our top priority as a group, we expect that Burst Group members will be working on several other analyses too during the run, and continuing afterward to thoroughly address all of our science goals. We expect most of the remaining S6/VSR2+3 analyses in our full scientific portfolio to be completed by the end of 2012. Even if no GW signal is detected in this data set, we will explore new territory in sensitive searches and astrophysical interpretation, to be carried even further in the advanced detector era.

4.5 GW burst science in the advanced detector era

The list of known potential astrophysical sources of gravitational wave bursts does not change from the initial to the advanced detector era. However, due to the ~ 10 times greater reach and ~ 100 times greater energy sensitivity of advanced detectors, the chance of detection and the potential for parameter estimation and the extraction of interesting fundamental physics and astrophysics increases. For this, gravitational wave burst data analysis strategies in the advanced detector era will need to take into account more information from theoretical/computational astrophysical source models. In turn, more and higher-fidelity source modeling input will be required. We must actively approach the modeling community to encourage the development of improved models for advanced detector data analysis. Most important for parameter estimation and physics extraction will be to have a theoretical understanding *for each source* of the mapping between signal characteristics (frequency content; time-frequency behavior) and physics parameters including knowledge of potential degeneracies that may be broken by complementary information from EM or neutrino observations. For this, multi-messenger modeling input and observations will be a necessity.

Most targeted searches in the advanced detector era will likely be externally triggered, hence will have EM and/or neutrino counterpart observations. Strategies must be developed for the extraction of physics at the post-detection stage combining EM, neutrino, and gravitational wave information. Hence, it will be important to work with our external partners to extend collaborations (and MOUs) from trigger exchange to full data sharing and joint analysis, parameter estimation and physics extraction.

While we can expect to learn more astrophysics (even by obtaining better upper limits) about the known set of sources and potentially constrain aspects of fundamental physics (i.e., the nuclear equation of state), we must be ready for the unexpected. The unexpected may come in two variants: (i) a detected signal from a known source (e.g, with an EM or neutrino counterpart) is completely different from model predictions and (ii) a high-significance event is detected with unexpected characteristics and no EM or neutrino counterpart. We must develop strategies of how to handle both (i) and (ii).

In the following, we discuss key burst sources that are likely to be focal points of data analysis efforts in the advanced detector era and briefly discuss science possiblities and potential.

Soft Gamma Repeaters (SGRs) and Anomalous X-ray Pulsars (AXPs)
SGR and AXP outbursts and giant flares are believed to be due to dynamical events in the crust and/or magnetosphere of very strongly-magnetized neutron stars (magnetars). Together with glitching radio pulsars, magnetars are the closest known potential sources of gravitational waves with multiple known sources within $10\,\mathrm{kpc}$ and at least ~ 7 within $5\,\mathrm{kpc}$ [260].

If magnetar outbursts/giant flares are due to magnetic-stress induced crust quakes and if crust and core are efficiently coupled in magnetars, SGRs/AXPs will emit gravitational waves. A network of advanced detectors is likely to be able to establish energy upper limits for gravitational wave emission by magnetar outbursts that are a factor of 10^2 better than the S5y2/A5 search [261]. Provided reliable GW emission models that go beyond analytic estimates (i.e., ring downs), advanced detectors may be able to constrain the outburst mechanism by putting astrophysically interesting upper limits on emitted energies.

Core-Collapse Supernovae
Massive star collapse does not immediately lead to a supernova explosion. Within half a second of collapse, the inner stellar core reaches nuclear density. There the nuclear EOS stiffens, leading to a rebound of the inner core (core bounce) into the still infalling outer core material. This results in the formation of the supernova shock that initially moves out quickly in mass and radius, but soon is decelerated, then stalls due to the dissociation of heavy nuclei and neutrino losses. The shock must be revived (by the core-collapse supernova mechanism), but how precisely this occurs is currently uncertain [262]. It has been argued [263] that gravitational waves from a galactic core-collapse supernova could help constrain the mechanism.

Core-collapse supernovae, no matter how precisely they explode, are expected to involve a wealth of processes leading to the emission of gravitational waves [134]. Based on S5y2/VSR1 reach and all-sky search sine-gaussian upper limits [113], advanced detectors (assuming a factor 100 improvement in energy sensitivity) can be expected to detect core-collapse supernovae emitting $10^{-10} - 10^{-9}\, M_{\odot}c^2$ ($10^{-8} - 10^{-7}\, M_{\odot}c^2$) at $100\,\mathrm{Hz}$ ($1000\,\mathrm{Hz}$) throughout the galaxy. This is well within what is predicted by conservative models [134]. More optimistic predictions suggest gravitational wave energies of up to $\sim 0.1\, M_{\odot}c^2$ [248]. Advanced detectors are likely to be able to constrain these models out to tens of Mpc. Improved modeling, in particular full 3D models that provide signal predictions for both polarizations, will be necessary to profit from the improved reach/sensitivity that advanced detectors will offer.

The physics that may be learned from a detection of gravitational waves from a core-collapse supernova goes far beyond constraining emission and explosion mechanisms. Supernovae are cosmic laboratories for high-energy-density gravitational, plasma, nuclear, and particle physics. In particular, it may be possible to extract information on the nuclear EOS directly from gravitational-wave observations [264]. Combining information carried by neutrinos with that carried by gravitational waves may allow more robust parameter estimation and physics extraction, but the details of how these data can be combined are unknown and must be worked out by theorists.

Current constraints on core-collapse supernova gravitational wave emission come from all-sky blind burst searches [113]. Sensitivity improvements (by factors of order unity) can be expected from a targetted all-sky search using information from models and, in particular, from a search that uses EM triggers for extragalactic supernovae. Such searches are in the process of being developed. Ways need to be found to handle the large uncertainties (and correspondingly large on-source regions) on the time of core collapse and onset and end of gravitational wave emission inherent to an EM triggered search. Neutrino triggers will provide much better timing and present collaboration initiatives with neutrino collaborations should be intensified. It is, however, unlikely that neutrino triggers for events at $D \gtrsim 1\,\mathrm{Mpc}$ will be available in the advanced detector era.

An interesting challenge to be ready for is the detection of gravitational waves and neutrinos from a galactic core collapse event with no EM counterpart. This could be either a unnova (an event that does not lead to a strong explosion—or no explosion at all) and in which a black hole is formed after $\sim 1 - 2\,\mathrm{s}$ of gravitational-wave and neutrino emission or it could be a EM-obscured supernova. Being able to determine which of the two possibilities is the case (possibly, by catching the high-frequency black hole ringdown emission) will be crucial.

Accretion-Induced Collapse (AIC) of Massive White Dwarfs

AIC occurs when a white dwarf (WD) is pushed over its Chandrasekhar mass limit and conditions (central density / temperature / composition) favor collapse rather than thermonuclear explosion. AIC may occur in binary WD merger events or by accretion from a companion star. Their occurrence rate is probably multiple orders of magnitude smaller than that of regular core collapse (e.g., [265]). AIC will proceed like a normal core-collapse event, but unlike ordinary massive stars, AIC progenitors are quite likely rapidly spinning. Hence, AIC is likely to give a strong gravitational wave signal from core bounce. In addition, postbounce long-lasting nonaxisymmetric rotational instabilities are plausible [265, 248].

AIC are expected to lead to EM-subluminous supernova explosions and we may be faced with a strong gravitational wave and neutrino signal with a weak EM counterpart. Being able to differentiate the AIC case from the black-hole forming regular core-collapse case will be important.

Long Gamma-Ray Bursts (long GRBs) and Engine-Driven Supernovae

The question of the long GRB central engine is one of the major unsolved problems in relativistic astrophysics. It appears clear from EM observations that long GRBs are related to massive star death and core-collapse supernovae [266], but the precise nature of this relationship is unclear. Central engine scenarios either involve a very rapidly spinning (millisecond period) magnetar or a stellar-mass black hole with an accretion disk. Relativistic GRB outflows may be powered by neutrino pair annihilation in polar regions or extraction of spin energy (or a combination of these).

The early gravitational wave and neutrino signals to be expected by a long GRB will be similar to that of a rapidly spinning core-collapse supernova before explosion [134]. Hence, GRBs should be approached with similar modeling input as supernova searches. During the GRB stage, GW emission may come from accretion disk instabilities (clumping, fragmentation) [248, 172, 171] or nonaxisymmetric magnetar deformation [267, 248]. The most extreme of these models predict emitted energies in gravitational waves of order $0.1\,M_{\odot}c^2$ which advanced detectors may be able to constrain to many tens to hundreds of Mpc (depending on the frequency of the gravitational-wave emission) [110]. For nearby GRBs ($D \lesssim$ few Mpc), engine scenarios may be constrainable, but much more theoretical modeling will be necessary to establish signal shapes characteristic for particular central engine models. These should be taken into account in future searches.

An interesting class of objects between long GRBs and regular core-collapse supernovae are hyper-energetic type-Ib/c supernova explosions that do not produce an observed GRB, but exhibit late time energy input into the ejecta by a central engine seen in radio observations (e.g., [268, 266]). Engine-driven supernovae are occurring considerably more frequently in the local universe than long GRBs and may be extreme core collapse events with plausibly strong GW emission. An engine-driven supernova within a few tens of Mpc would be an interesting science target for advanced detectors.

Phase-transition Induced Collapse of Neutron Stars

Recent work on core-collapse supernova and neutron star modeling suggests that a QCD hadron-quark phase transition in a (proto-)neutron star could lead to the emission of gravitational wave by ringdown oscillations following a "minicollapse" of the neutron star and its stabilization of the neutron star at a higher-density equilibrium [269, 270]. If no stabilization occurs, a black hole will form and black hole quasinormal modes will emit gravitational waves as the hole rings down to an axisymmetric equilibrium [134]. In the former case, typical GW frequencies will be $1-3\,\mathrm{KHz}$, while in the latter emission will be predominantly at frequencies $\gtrsim 4-6\,\mathrm{Khz}$.

Given the high-frequency emission of this class of sources, advanced detectors will still be limited to nearby ($D \lesssim$ few kpc) [113, 112] events, but there are a number of accreting neutron stars in X-ray binaries that should be monitored for such high-frequency gravitational wave emission. Provided high-SNR detection, information on the nuclear EOS and object mass could be gained.

Pulsar Glitches

Pulsar spin evolution is well modelled by assuming magnetic dipole braking is the dominant mechanism for the spin-down. However, some pulsars occassionally exhibit sudden step jumps, called a *glitch*, in the rotation rate, followed by an exponential decay back to the (almost) pre-glitch rate.

There exist two main candidates for the mechanism behind pulsar glitches. One suggested mechanism is that glitches may be due to starquakes where the equilibrium configuration of the solid crust of the neutron star deforms as the pulsar spins down. The energetics of the more frequent 'glitchers', however, are indicative of a superfluid core rotating more or less independently of a solid crust. The crust experiences a magnetic torque and spins down, leading to differential rotation between the superfluid and the crust. The superfluid rotates by means of vortices, which are normally 'pinned' to the outer crust. It has been suggested (eg.,[271]) that once the differential rotation reaches a critical value, an instability sets in causing a dramatic un-pinning of the superfluid vortices, transferring angular momentum from the superfluid component to the crust, causing the observed temporary spin-up. There are a variety of mechanisms for gravitational wave emission associated with pulsar glitches, the details of which depend strongly on the glitch mechanism. For superfluid-driven glitches, there may be an incoherent, band-limited stochastic burst of gravitational waves due to an 'avalanche' of vortex rearrangement [272] which will occur during the rise-time of the glitch ($\leq 40\,\mathrm{s}$ before the observed jump in frequency). A possible consequence of this vortex avalanche is the excitation of f-mode oscillations in the neutron star in the frequency range $1-3\,\mathrm{KHz}$, leading to quasi-sinusoidal gravitational wave emission over timescales of $50 - 500\,\mathrm{ms}$. In the case of starquakes-driven glitches, it also seems reasonable to expect that the neutron star f-modes may be excited.

The relatively low energy associated with pulsar glitches (the energy associated with the change of angular velocity in a typical glitch is $\sim 10^{42}\,\mathrm{erg}$) and the high-frequency of emission, advanced detector searches triggered by pulsar glitches will still be limited to nearby sources, such as the Vela pulsar ($\sim 300\,\mathrm{pc}$), and are somewhat opportunistic. However, as with other searches related to neutron star oscillations, the science benefit of a high-SNR detection is enormous and, in this case, comes with the additional benefit of potentially directly probing the highly uncertain mechanics of pulsar glitches.

NSBH, NSNS Coalescence and Postmerger Evolution
Burst searches generally have the capability of picking up also compact binary (inspiral-merger-ringdown) events (e.g., [110, 151]). Recent analytic and computational work suggests that constraints on the nuclear EOS are possible by matched filtering of advanced detector data from the intermediate to late inspiral of NSNS (and NSBH) binaries [273]. Moreover, the discovery of an inspiral signal associated with a short GRB would clarify the pressing question about the short-hard GRB progenitor scenario. On the basis of scaled S5 GRB results [110] this may be possible with burst methods alone to $\sim 100\,\mathrm{Mpc}$ with advanced detectors.

Interesting science potential lies not in the inspiral alone. In the NSNS case, the merger signal as well as the gravitational waves emitted in the postmerger evolution may tell volumes about the mass and spin of the system, whether a black hole forms promptly or in a delayed fashion, and may also allow to put constraints on MHD spindown and the nuclear EOS. The postmerger signal (that cannot be templated) may also provide information on whether a short GRB is powered by a millisecond hypermassive neutron star or by a black hole – accretion disk system. However, most of the postmerger gravitational-wave emission will occur at frequencies of $\sim 1 - 4\,\mathrm{KHz}$, may emit up to $\sim 0.001 - 0.01\,M_{\odot}c^2$, but will most likely be detectable only for nearby ($D \lesssim$ few $10\,\mathrm{Mpc}$) events. It will still be worthwile to perform a targeted (though untriggered) search on the postmerger evolution for nearby events (the majority of which are probably EM silent) using advanced detector data. This search should be informed by the next generation of computational model predictions from self-consistent models that include all the necessary physics.

Binary Black Hole (BBH) Coalescence
BBH coalescence is the single source whose gravitational-wave emission can be templated completely and with confidence (under the assumption that general relativity is the right theory of gravity). Given the large parameter space of potential BBH pre-merger configurations, thousands of templates must be computed (by combing post-Newtonian theory with numerical relativity) if the conventional matched-filtering approach is to be used. Burst searches are an alternative approach that can do without detailed templates and generally yield results within an order of magnitude of matched-filtering approaches [151]. Current results suggest for advanced detectors detectability of many-stellar-mass BBH coalescence out to at least hundreds of Mpc with burst methods alone. In case of highly eccentric binaries accurate templates may not be readily available and the burst searches may be the only way to detect and study these sources. The relative sensitivity of burst and matched-filter searches, as a function of source masses, spins, eccentricity and other properties, is a matter of ongoing study.

Cosmic String Cusps

The advanced detectors will enable much more sensitive searches for cosmic string cusps, especially since those signals are predominantly at low frequencies where the sensitivity improvement will be the greatest.

Unknown Unknowns
Blind burst searches allowing for any kind of signal time-frequency content are the only way in which completely unexpected signals that do not have EM/neutrino triggers can be found. The dramatic increase

of the volumetric reach of advanced detectors will lead to a corresponding increase of the possiblity of discovering an unexpected event. A strategy and/or protocol of how to handle such an event and how to communicate with outside theorists should be established for the advanced detector era.

4.6 Preparing for the advanced detector era

The Burst group aims to be ready to extract science from early advanced detectors data, building on the online and offline analysis experience of the initial detectors. We expect the group will devote most of its attention to preparations for the advanced detector era once the S6/VSR2+3 first science target analyses are completed, but early preparation will be needed in some cases, especially if they will benefit, or benefit from, the use of GEO-HF data during the transition years. In this section, we highlight some general areas where preparations will be needed; a more detailed action plan will be formulated in the coming year.

- This is an appropriate time for the burst group to prioritize of its science goals, on the basis of the discussion in section 4.5. Simulations and astrophysical input will help decide whether these sources require targeted analyses or a different tuning and interpretation of the standard all-sky and externally triggered searches. Over the next 3–5 years, the group will invest manpower in simulation studies and in the integration of astrophysical knowledge in data analysis. This will involve standardizing the simulation code used by different pipelines and use, where available, knowledge on waveforms, as it has started for binary black hole mergers and core-collapse supernovae.
- The Burst group has established, over the past two years, several external collaborations, which remain key for the group's science goals. This applies both to the multi-messenger approach, where the external collaborators are other experiments or observatories, and for collaborations with numerical relativists and theoretical astrophysicists, who can provide guidance for our searches. The group will need, in the post S6/VSR3 era, to decide whether to maintain all the existing collaborations in their current form, or revise their scope, and define goals for these collaborations for the next 3-5 years.
- Understanding data quality and vetoes will be particularly important at the beginning of the advanced detector era, when the new detectors will likely exhibit a new zoo of transient noise. The Burst group will need to be more closely coupled with the detector characterization effort, with at least one, or more, liason members focused on vetoes for each search. During the transition era, the group needs to also decide whether to re-think its approach to vetoes, and give priority to the development and implementation of signal-based vetoes.
- Realistic calibration needs, both in terms of accuracy and latency, will be evaluated during the transition years.
- Future needs for hardware computing infrastructure will also need to be evaluated during the transition years. The Burst group will benefit from an effort to make the software user-friendly for all group members, with search codes packaged as common tools for all to run and configure for customized analyses, with reduced overhead for learning how to use the code, and a standardized output from all pipelines.

5 Searches for continuous-wave signals

Rapidly rotating neutron stars are the most promising sources of continuous-wave (CW) gravitational signals in the LIGO and Virgo frequency band.[1] These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [278, 279, 281], magnetic deformations [280, 284], unstable r-mode oscillations [282, 278, 283], and free precession [288], all of which operate differently in accreting and non-accreting stars. We present a review of these emission mechanisms in [277]. 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, non-accreting unknown stars, and accreting stars in known or 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. That said, however, we do adhere 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 spindown limit is achievable – see below) and on all-sky searches for unknown isolated pulsars.

The recent merging of LSC and Virgo CW efforts has revealed strong and well-developed programmes from both groups. An ongoing task is to combine these effectively, maximising both the return on time already invested in developing codes and the science delivered by new joint ventures. Our state right now is one of transition, where we are evaluating the scientific justification and available manpower for these searches, while balancing the importance of redundancy against current resources.

An important new element in this evaluation 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 are expected to lead to the abandonment (or at least de-prioritization) of some search pipelines.

5.1 Non-accreting 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 (see section 5.1.1 below).

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

Photon astronomy also sets an upper limit on the gravitational wave strain we could see from a known pulsar, assuming that all of the observed spin-down is due to gravitational waves. In terms of the distance D, gravitational wave frequency $f_{\rm gw}$ and its time derivative $\dot{f}_{\rm gw}$, this indirect limit is [277]

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

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 full S5 data has beaten this indirect "spin-down limit" by a factor of 7 for the Crab pulsar (59.45 Hz). As discussed below, analysis of the VSR2 data should lead to beating the spindown limit on the Vela pulsar (22.38 Hz). Other pulsars for which the spindown limit may be reached in S6/VSR2/VSR3 include PSRs J1913+1011 (50.59 Hz), J1952+3252 (55.69 Hz), J0737−3039A (88.11 Hz) and J0537−6910 (123.95 Hz) [285].

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 (5.1.2), 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.

5.1.1 Fully coherent 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.

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 below.

5.1.1.1 Time domain Bayesian method

The time-domain Bayesian method has been applied successfully to data from the first five LSC science runs [274, 275, 294, 292, 293] and is currently being used to produce results from the full data set from the sixth science run and from the Virgo VSR2 run. A detailed discussion of the method can be found in [291], with the implementation of the inclusion of binary system parameters in [295].

The method is designed to carry out robust signal extraction and optimal parameter estimation, rather than perform well in a large parameter space search. Its primary purposes are

- to perform searches for signals from known pulsars and,
- to determine the astrophysical parameters of candidate sources.

The method comprises a heterodyne and filtering stage to extract interferometer data in a tracked 1/60 Hz band centered on the expected gravitational wave signal, followed by a Bayesian parameter estimation stage. This second stage delivers an upper limit to the strain amplitude of any signal and an estimation of the signal's parameters, should one be detected, in the form of marginal posterior probability distributions for the signal parameters. The method has successfully determined the parameters of the injected signals in all our science runs. The most computationally intensive part of the search is the heterodyning and down-sampling of the raw data. Currently this takes about 25 min per pulsar per detector per day of data.

We have strong links with the radio groups at Jodrell Bank, Parkes, Effelsberg and Green Bank who have generated timing solutions for our pulsar candidates over the LIGO observing runs, and checked that no glitches have occurred in these pulsars. Newer collaborations with HartRAO and Hobart have given us timing solutions for the Vela pulsar. These collaborations have provided data for the S5 targeted searches and will enable an even wider range of targets for S6/VSR2 and beyond. We have initiated a collaboration with X-ray groups to be able to target promising X-ray objects. This collaboration has provided useful discussions and timing information for the young X-ray pulsar PSR J0537−6910, which is another pulsar for which we could soon beat the spin-down limit.

There are just over 200 pulsars within the sensitive band of the LIGO interferometers, that is with spin frequencies greater than 20 Hz, and nearly another 40 between 10 and 20 Hz. For all these we are able to perform the first stage of the search heterodyne process, but for S5 we had 116 with radio/X-ray observations overlapping the times of the run. For all pulsars, the radio timing data give uncertainties in the pulsars' parameters. In the past we have used these uncertainties, without taking account of covariances between source parameters, as an estimate on whether a single template search is valid for the pulsar. This has meant discarding some pulsars from our analysis. For the majority of pulsars we now have covariances for the parameters and can use these to make a better estimate of whether a single template search is sufficient. This has lead to the inclusion of 12 pulsars in the current search that would have been vetoed previously. We are also no longer restricted to a single template as we have now extended our parameter estimation techniques to use Markov Chain Monte Carlo techniques to search over the parameter uncertainties, whilst also taking into account the covariances between the phase parameters. In the cases where no signal is seen this will marginalise over the uncertainties and fold them into our upper limit in a natural way. For the majority of pulsars a signal template (i.e. not added extra phase parameters into the MCMC search) would still be sufficient, but for at least one pulsar the additional search space would be required to properly recover the signal. In addition, we now account for glitches in pulsars by adapting the timing model to allow for step changes in rotational phase at these points. For the S5 search this has been applied for three pulsars that were seen to glitch during the run: the Crab pulsar, J0537−6910 and B1951+32.

Of the pulsars for which we have accurate timing, the Crab pulsar is both the youngest, and the most rapidly spinning-down, candidate within our sensitivity range. The relatively large amount of intrinsic timing noise for this pulsar is tracked and corrected for within our search method [294, 295]. We have published the results of a search for gravitational waves from the Crab pulsar using data from the first nine months of S5 until the time that a large timing glitch was observed in the Crab pulsar [292] (also see §5.1.2.)

A follow-on search for 116 pulsars in the full S5 data set included the Crab and enabled us to beat the spin-down limit by a factor of 5.6 using uniform priors over the unknown parameters. Astrophysically motivated priors on the inclination and polarisation angles allowed us to further beat the spin-down limit by an extra factor of 1.3. This result has allowed us to constrain the amount of the pulsar spin-down power budget released by gravitational radiation to be less than about 2%.

Preliminary results from a search for the Vela pulsar in VSR2 data have been obtained for both uniform and restricted priors on its orientation. In the coming year these results will be submitted for a stand-alone publication on Vela, along with results from the other two targeted pipelines described below.

In addition, the time-domain Bayesian method will be applied in the coming year to search the full S6 and VSR2 data sets for all accessible known pulsars with available precise timing, and VSR3 data will also be searched.

In the longer term, the time-domain method will be used in a full VSR3 search for known pulsars, searching not only at the nominal 2*f frequency, but also at 1*f and in the vicinity of (4/3)*f for *r*-modes, as part of a broader effort to widen the parameter space in targeted searches. In addition, more use will be made of Fscans (discussed below) and other diagnostics to optimize searches for the Crab and Vela pulsars.

5.1.1.2 Signal Fourier 5 components method

The signal Fourier 5 components targeted search starts from the short FFT database (SFDB) and uses time domain data cleaning, based on an auto-regressive method which allows to efficiently remove short time domain disturbances[300]. The method is based on the computation of the analytical signal [296]. In this method we start from the short FFT database and extract the frequency band of interest. Then, the data are down-sampled by creating the analytical signal. The correction for the Doppler effect, spin-down or any other frequency variation is done by a over-resampling procedure. Further cleaning steps are done (to eliminate bad periods and big events) and finally the data are passed through a Wiener filtering stage in order to weight less the more noisy periods. The final stage consists in applying to the data a matched filter at the frequencies corresponding to the five signal Fourier components.

Preliminary results from a search for the Vela pulsar in VSR2 data have been obtained both with and without constraints on orientation. In the coming year these results will be submitted for a stand-alone publication on Vela,along with results from the other two targeted pipelines.

In addition, this method will be applied in the coming year to search the full S6 and VSR2 data sets for the Crab pulsar. The VSR3 data will be searched for all pulsars with reachable spindown limits, including searches of band surrounding 2*f, with publication expected in the following year.

5.1.1.3 Time domain matched-filter method using the $\mathcal{F}$-statistic

Assuming that other parameters of the gravitational wave signal i.e. the amplitude, the phase, polarization and inclination angles are unknown, matched filtering is realized by computing the $\mathcal{F}$-statistic [324]. If the computed value of the $\mathcal{F}$-statistic is not significant we derive an upper limit on the gravitational wave signal. From current observations of the Vela nebula the polarization and inclination angles can be estimated to a very good accuracy [325]. We use this knowledge to improve our search. This required derivation of a new optimum statistic called the G-statistic. In our analysis we take into account non-Gaussianity and non-stationarity of the noise.

We can apply these methods to target various interesting known pulsars. In particular they have been applied to the analysis of VSR1 and VSR2 data for the Vela pulsar, for which the spindown limit should be beaten in VSR2. For this analysis we are using updated ephemeris data provided by radio-astronomers (Aidan Hotan, Jim Palfreyman - Hobart Radiotelescope). Another obvious target is the Crab pulsar, for which about 2.5 months of data would allow to go below the current upper limit at $2f_{\rm rot}$, at the Virgo design sensitivity. The search for emission at $f_{\rm rot}$ will be also performed.

The status of the search for Vela in VSR1 data was described at the 2010 GWDAW meeting [306]. Final results of the VSR2 analyses are expected to be published in the coming year, along with those of the other two targeted pipelines described above.

5.1.2 Wide Parameter Searches for Known Pulsars

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. We also know of emission mechanisms which can cause gravitational wave emission at other harmonics, such as $(4/3)f_{\rm rot}$ for r-modes. In our search we consider the cases of free precession and a two-component model of the neutron star's crust and core, and work out how much difference might occur between the true gravitational frequency and twice the rotation frequency. We also consider the uncertainity in frequency associated with r-mode emission when searching around $(4/3)f_{\rm rot}$.

These calculations, along with considerations of the computational costs, will be applied to all the known pulsars within the LIGO band to produce a parameter space it is reasonable to search within, to compliment the exact searches carried out with the time domain search method. These searches will use the full S5 data set with an improved search code utilizing resampling of the data. This new code achieves a speed up in computation time proportional to the number of SFTs used in the search, effectively an improvement of 3 to 4 orders of magnitude in computation time. This speed up makes these searches possible, where as previously we had restricted the wide parameter search to just the Crab pulsar.

We have searched a smaller parameter space for just the Crab pulsar with nine months of data from the S5 run, up until a timing glitch in the pulsar [292]. By using known information on the orientation of the Crab pulsar, the search placed 95% confidence strain upper limit of 1.2×10^{-24}, beating the indirect spin-down limit of 1.4×10^{-24} across the entire parameter space searched. It is a less stringent upper limit than placed by the time domain search because of the use of 3×10^7 more templates and the statistics of the noise. This search will continued using the larger amount of data available after the glitch, which means the search should see over a 20% improvement in sensitivity for the Crab pulsar.

A paper describing these wider searches in S5 data is expected to be submitted for publication in late 2010 or early 2011.

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

This type 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. 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. (1) as

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

in terms of the age a, 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 ($\tau_{\rm sd} = 326\,{\rm yr}$, $h_{\rm IL} = 1.2 \times 10^{-24}$) – and the closest, Vela Junior ($D > 200\,{\rm pc}$, though the precise value is uncertain). Several more objects have indirect limits attainable with advanced LIGO, including the remnant of Supernova 1987A ($h_{\rm IL} = 3.2 \times 10^{-25}$). However this putative neutron star is only 20 years old and would require a search over a large parameter space including six or seven frequency derivatives. Searches over small sky areas (single "pixels") are computationally the same as searches for known point sources, and for several of these (such as the galactic center) even initial LIGO could beat indirect limits. We recently completed a search for Cassiopeia A (see section 5.2.1) and shortly will start searching the central parsec or so of the galactic center (see section 5.2.4) and at least one globular cluster (see section 5.2.6). 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 5.2.9).

The first search for a source with no timing (Cas A below) used the $\mathcal{F}$-statistic code with a single integration of $\mathcal{O}(10)$ d. Our estimate of computational cost and sensitivity [327] shows 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. In the near future we will try hierarchical searches (see other searches below) which will require algorithm and code development to adapt to the needs of this search. We are also evaluating the potential of resampling methods to reduce not only the computational cost for a given search, but also the cost's scaling with integration time. This, combined with hierarchical methods, will allow us to search a significant fraction of the S5 data set (and future sets of comparable length) rather than $\mathcal{O}(10)$ d.

5.2.1 Cassiopeia A

This search was recently completed and submitted for publication [326] after substantial development work on template generation using optimal lattices and careful treatment of the boundaries of the parameter space. The search pipeline developed for Cassiopeia A will also be applied to similar targets in S6/VSR2 data that require only one sky position. The sensitivity, computational cost, and indirect limits to aim at are described in [327].

5.2.2 Other coherent directed searches

The coherent search for Cas A will be extended to top targets from each category of directed search (see Sec. 5.2.9), with the aim of beating the indirect limits for of order ten sources. In order to "industralize" the Cas A prototype search, the template bank will be made more efficient by taking advantage of long-range correlations in the parameter space. Preliminary results are expected in the coming year with publication in the following year.

5.2.3 Search for Calvera

Calvera is an X-ray source originally detected by ROSAT and confirmed with Swift and Chandra measurements. It has no detected optical or radio counterpart and is thought to be an isolated neutron star, possibly a millisecond pulsar beaming away from us, but relatively close – O(100 pc) away.

A fully coherent search using a barycentric resampling method based on interpolation in the time domain has been carried out for Calvera over the 90-360 Hz band [328], assuming an age of O(10^6) years which severely restricts the spindown range to be searched. Because the sky location is known, the search benefits dramatically in reduced computational cost from the resampling step during preprocessing. Preliminary upper limits have been obtained, and final results will be submitted for publication in the coming year.

5.2.4 Searches for sources near the galactic center

The galactic center is a location where one might expect to find a large number of unknown, young neutron stars. Standard electromagnetic observations have identified only a small fraction of all pulsars thought to be present in the galactic center. The dispersion and scattering of the signal by interstellar material between potential sources and the Earth significantly reduces the depth of such observations. The current estimate of the total number of pulsars present in the galactic center (the inner $2\,\deg \times 0.8\,\deg$ of the galaxy) is 10^6 (Muno et al). Some of those objects could be promising sources of CW gravitational wave signals. Searching in the galactic center involves searching over a small sky-region but over large frequency and spin-down ranges. An important parameter is the minimum spin- down age we wish to search for, which will determine the number of spin-down parameters required and also the maximum amplitude of such a signal.

We are carrying out a search for traditional CW sources that last for the duration of the entire S5 run using the present `Einstein@Home` search code based on global correlatons (see below). While `Einstein@Home` is carrying out an all sky-search using the same code, this search is meant to be more sensitive over a more limited region of parameter space. Thus, it should search a larger range of spin-downs and also have a longer coherent integration time and reach younger spindown ages. In addition, the galactic center search will use barycentric sampling like that used for the Calvera search.

Additional work has focused on determining the parameter space region and template grid layout. Important parameters of the search are the sky position, a single template pointed directly to the galactic center; the frequency band, covering a sensitive region of the LIGO band, 77-496 Hz, for which the search can beat the indirect spindown limit for an age assumed to be greater than 200 years. The data searched is from S5, using only H1 and L1 detectors. The range of spindowns is specified by the spindown age via $\tau = f/\dot{f}$, where f is the gravitational wave frequency of a neutron star. The search uses 630 11.5-hour segments semi-coherently. Preliminary results have been obtained from a subset of the total search band. The search will be completed and final results derived in the coming year, with submission for publication expected in the following year.

5.2.5 Supernova 1987A using the cross-correlation technique

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 [334], can be found in [314]. 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 $T_{\mathrm{coh-min}}$ 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 weghts 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 $T_{\rm obs} \gg T_{\rm coh-min} \gg T_{\rm sft}$. This is however a computationally inefficient way of calculating the coherent statistic, for which it is better to use the existing $\mathcal{F}$-statistic, so we expect that the cross-correlation is useful only with $T_{\rm coh-min}$ either comparable or lesser than $T_{\rm sft}$.

The current plan for S5 is to use the cross-correlation method to search for periodic gravitational waves from Supernova 1987A and possibly also the galactic center using data from all three LIGO interferometers in a broad frequency range from 50 Hz to 1 kHz. The software for computing the cross-correlation statistic for isolated pulsars has been implemented in LAL/LALapps and is currently being validated, reviewed and tuned for an S5 search. The parameter $T_{\rm coh-min}$ will be tuned so that the search can be completed on a time scale of 1-2 weeks which leads to $T_{\rm coh-min} \approx 1$ hr. In searching for such a young object, searching over frequency derivatives can be prohibitive because one would need to search over higher derivatives as well. It turns out that the search space can be narrowed by using a physical model for the frequency evolution: $\dot{n u} = Q_1\nu^5 + Q_2\nu^n$. 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 this is not observed in practice). With this model, and using $T_{\rm coh-min} \approx 1$ hr, it turns out that the computational cost becomes manageable.

In the coming year the search will be finished and results submitted for publication. It is hoped that completion of this semi-targeted search will be useful in developing an all-sky search based on cross correlation.

5.2.6 Semi-targeted search using stroboscopic resampling

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) [298]. The removal (or the duplication) of the samples takes place typically each a 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 [298] 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 the Torre's graduation thesis [315], or (more in summary) in [316]. 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 the down-sampled data (making negligible the computing time for reading data) is in progress. The amplitude modulation will be taken into account using a matched filtering in the frequency domain, in a way similar to the one suggested by the Rome group. The method is now being tested on hardware injections in VSR2 data and will soon be applied to an exploratory search of the globular cluster Tucanae. A method paper was recently submitted to PRD. It is too early to make firm predictions on publishing a results paper.

5.2.7 Semi-targeted searches for "transient CW signals"

This project aims at developing an efficient search method and pipeline to scan long stretches of data (of length $\sim$1 year) for *transient* quasi-monochromatic signals that only last for timescales between about a day to a few weeks. The motivation for this study comes from glitching pulsars, which illustrate that neutron stars can be in non-equilibrium states that relax back to equilibrium on timescales of order weeks. This makes it plausible that on such timescales neutron stars could be more strongly deformed than suggested by equilibrium studies of the maximal deformation of neutron stars. During such episodes they might therefore emit stronger gravitational waves, which look like "continuous GWs" but last only for a few days to weeks. A study of the search method has been carried out in Matlab, to quantify issues of computing cost and sensitivity, and accounting for the additional parameter-space, which now includes start time and duration of the signal. The results are promising, and the algorithm has been ported to C for production running, with a search to begin in summer 2010. Preliminary results are expected in the coming year, with submission for publication in the following year. A methods paper is in preparation.

5.2.8 Directed search using $\mathcal{F}$-statistic

Another directed-search pipeline is under development, building upon already implemented all-sky search code that employs the $\mathcal{F}$-statistic method (Sect. 5.3.6). Novel usage of graphical processor unit (GPU) solutions, framework and dedicated libraries promise that a search for GWs from astrophysically motivated sources over a large frequency and frequency derivative will dramatically reduce computational costs, as well as allow for much greater flexibility in comparison to what is possible in all-sky searches. Additionally, the pipeline will serve as a testing ground for new theoretical methods of parameter space optimization and hardware/software implementation techniques.

5.2.9 Other targets

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 nearly ten central compact objects in supernova remnants with no observed pulsations and indirect limits on h_0 high enough to beat with S5 or S6/VSR2 coherent and semi-coherent searches. "Calvera" [329] is not associated with a remnant, but may be the closest observed neutron star. 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 beginning to find many point source neutron star candidates, and thus it is likely that the interesting target list for LIGO will expand substantially even before advanced LIGO. Examples include HESS TeV gamma-ray sources which are followed up with Chandra and XMM-Newton X-ray observations to yield pulsar wind nebulae and sometimes the neutron stars themselves.

In the coming year a paper describing an interesting list of targets will be submitted for publication. As the advanced detector era approaches, these studies will be extended.

5.3 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

$$h_{\rm IL} = 4 \times 10^{-24} \left(\frac{30\,{\rm yr}}{\tau}\right)^{1/2}, \tag{3}$$

where τ is the mean time between supernovae in the Galaxy. The latest and most detailed derivation of this upper limit is given in [277]. Note, however, that a more recent simulation analysis finds significantly lower expectations that depend on the assumed source frequency and ellipticity [312].

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 dramatically in some search pipelines.

Five all-sky pipelines are currently in use 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-interfereometer 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. It is likely that the use of one or more of these pipelines will be discontinued in the next 1-2 years, following the systematic comparisons using standardized simulated data sets, as described above.

In addition, two new methods are under development that offer greater robustness against uncertainty in the source model: 1) a "loosely coherent" method [313] using the PowerFlux infrastructure; and 2) a cross-correlation method [314] 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.

5.3.1 PowerFlux method

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.

Within the last few years, we have explored three related methods for incoherent strain power summing: StackSlide [286], Hough [307, 276], and PowerFlux [308]. 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 science run. An article based on applying all three methods to the S4 data was published in early 2008 [309] 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 $\sim$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 [310]. These results cover the frequency range 50-1000 Hz and negative spindown 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 within a half-year. A two-interferometer power sum is being used, together with coincidence between H1 and L1, to push deeper into the noise than before. Preliminary upper limits for the 50-800 Hz band have been produced and will be submitted for publication in the coming year, along with results of outlier follow-ups.

A fully automated coherent followup pipeline has been implemented for PowerFlux. This pipeline uses the ComputeFStatistic code to perform a coherent followup of each outlier. This followup is performed in steps, with each step zooming in on the loudest event(s) from the previous step, and using around 3 to 4 times more data than the previous step. The followup step is more sensitive than the PowerFlux, allowing it to better determine if the outlier was due to noise or an actual gravitational wave. It is expected that several thousand outliers can be investigated using the new infrastructure. The pipeline can also be applied to the outliers from other all-sky searches.

In parallel, an alternative "loosely coherent" follow-up has been added directly to PowerFlux, one that allows for slow drifts or modulations in phase from one SFT epoch to the next [313]. 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.

In the coming year full-S5 PowerFlux results, including those from the fully coherent and loosely coherent follow-ups, will be submitted for publication. In parallel, preliminary S6/VSR2 results will be obtained, with publication expected in the following year. Further PowerFlux refinements, including a coherent IFO-sum option for each SFT to gain further SNR [311], are also planned.

5.3.2 Hough transform method

As in the PowerFlux method, the weighted Hough transform method (used already to analyze the S4 data [309]) takes into account the detector antenna pattern functions and the non-stationarities in the noise. This algorithm allows to combine data from different interferometers, to perform a multi-interferometer search.

In preparation for analyzing the full S5 data, a set of new features have been included into the Hough search code, such as dynamical selection of data depending on SFT noise floors and sky-positions, splitting of sky patches with frequency dependent size, creation of a top list of candidates, internal follow-up using the full data, and a chi-square test [317] to reduce the number of candidates and consequently increase the sensitivity of the search. Preliminary, separate analyses of the first and second calendar years of S5 data have been carried out, with with coincidence of outliers between years 1 and 2 to be imposed as a filter on outliers to be followed up.

Full-S5 Hough results, including those from outlier follow-ups, will be submitted for publication in the coming year, perhaps together with the PowerFlux results.

5.3.3 Einstein@Home Searches

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 blind 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 $\sim$ 200 Tflops 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 and not on setting precise upper limits. So far it has analysed LIGO data from S3, S4 and S5.

The analyses on S3 and S4 have been completed, a report and final results from the S3 search were posted on the `Einstein@Home` web page, and a paper on the S4 results has been published in PRD [321]. A similar search was run on S5 data (S5R1), and the paper presenting these results published in PRD [322].

S5 R2/R3 postprocessing: The second S5 analysis was based on a greatly improved search pipeline, which eliminates the main problem limiting the sensitivity in previous `Einstein@Home` searches: this search was based on a Hierarchical search pipeline, consisting of individual $\mathcal{F}$-statistic searches on $N_{\text{stack}} = 84$ data "segments" spanning no more than $T_{\text{stack}} = 25$ h each. Each of these segments contains at least 40 h of data from H1 and L1 (the multi-IFO coherent analysis is another significant improvement). The results from these $\mathcal{F}$-statistic stacks are then combined in a second stage using the Hough algorithm. As both of these steps are performed on a participating host computer *before* sending back the results, the optimal threshold on the $\mathcal{F}$-statistic stacks can be used, avoiding the limiting sensitivity bottleneck and are expected to substantially improving sensitivity with respect to the previous search method. After an initial shorter "test" run ("S5R2") with this pipeline, lasting for about 3 months, a further improved setup was launched as the 3rd `Einstein@Home` run on S5 data (codename "S5R3"), designed to run for about a year. This run has finished and the work on postprocessing of these results has nearly completed. A paper describing this search of the first year of S5 data will be submitted for publication in the coming year.

S5R5 postprocessing: The S5R3 run analysed only data from roughly the first year of S5. A follow-on `Einstein@Home` run (S5R5) was launched in Jan 2009, covering a frequency range up to 1kHz. S5R5 uses a mostly identical pipeline to the previous (S5R3) run, but include 121 new segments of data from the second half of S5. The search setup also takes account of the speedup of the science application by nearly a factor of 2, and an increase of more than 30% in participating hosts. A paper describing this definitive E@H search of the full S5 data is expected to be submitted for publication in the coming year.

S5GC1 search: A significant improvement to the hierarchical search comes from exploiting global correlations in source parameter space [323]. An engineering run (S5GCE) to test this improvement was completed in spring 2010, and a new production run (S5GC1) over the full S5 data is now underway. Results from this search are expected in the coming year with submission for publication in the following year. A follow-on search will target S6 and VSR2/VSR3 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. 5.3.4).

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 a lot of 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".

GPU optimizations for E@H: An effort is underway (in collaboration with NVIDIA) to leverage the potentially large computing power gained from optimizing our E@H science codes to benefit from the massively parallel capabilities of modern graphic chips (GPUs), currently mostly aiming at NVIDIA cards using the CUDA software library.

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.

5.3.4 Followup-searches to confirm or veto CW signal candidates

Better theoretical understanding and the development of software tools is required to be able to efficiently deal with following up interesting candidates from incoherent search pipelines, in a systematic and mostly automated way. This involves questions of required integration times for coherent followups, and number of "zoom" stages in order to successivly trim down the parameter space and accumulate sufficient SNR to gain confidence in CW signal candidates. Analytical results have been obtained recently, and a methods paper is in preparation. This work will be completed in the coming year, including the implementation of a preliminary follow-up pipeline. In the longer term, stable production code will be produced that includes automatic feeding of candidates with minimal user intervention. In addition, other semi-coherent methods will be investigated.

5.3.5 Hierarchical Hough searche

A hierarchical pipeline with an incoherent adaptive Hough step [299, 300, 301] will be applied to VSR2 data. First, calibrated data at 4kHz are cleaned in time domain; then the "short" FFT database is built. From it, time-frequency peak maps are produced. They are cleaned, removing lines of likely instrumental origin, in order to reduce the final number of candidates. The peak maps are the input of the Hough transform stage, which produces a set of candidates. Coincidences among candidates obtained from the analysis of different data sets, belonging to the same or different detectors, are done in order to reduce the false alarm probability. On the surviving candidates the coherent follow-up is applied[303]). A new Hough procedure, based on the transformation between the time-frequency plane and the source frequency/spin-down plane, has been developed[305] and will be applied to VSR2 data. A new procedure for the coincidences among candidates will be developed. This year we are going to analyze data of the second Virgo scientific run VSR2. Work on this search has been slowed by parallel work on the Vela search described above. In the coming year the search pipeline will be augmented by a refined Hough step and by coincidence analysis of candidates. In the longer term the search will be applied to the full VSR3 and S6 data.

5.3.6 $\mathcal{F}$-statistic all-sky search

Another analysis method currently used for all-sky searches is based on the $\mathcal{F}$-statistic [324]. It consists of a coherent stage over 2-day data segments, each covering 1 Hz bandwidth plus a follow up analysis of candidates in a 4-day data segment. We shall assume the minimal spin-down time of 1000 yr. We shall use a constrained optimal grid with minimal match MM = $\sqrt{3}/2$. The constraints are such that we need to resample the data only once for each sky position and such that we can use the FFT to calculate the F-statistic. With an available time of half a year for such computation we expect to analyze coherently 1000 data segments. Our guiding principle in the choice of a segment to analyze will be the quality of the data. The data for analysis will be narrow-banded and cleaned using the procedures described in 4.1.3. We shall set a low threshold of 30 for twice the value of the $\mathcal{F}$-statistic above which we shall register the parameters of the candidates. We shall verify the candidates by coincidence test among the candidates from different segments and by the F-test for the $\mathcal{F}$-statistic value gain when we increase the observation time twice. We shall collaborate on the coincidence analysis with the LSC E@HEinstein@Home team. We also plan to do search using the above method in the subspace of the parameter space defined by the spin down parameter equal to 0. For this subspace we plan to analyze 25000 2-day 1Hz band sequences.

A paper describing the all-sky VSR1 searches is expected to be submitted for publication in the coming year. In the longer term the search pipeline will be applied to the full VSR2 and VSR3 data, with improvements based on the global correlations code used in `Einstein@Home`.

5.3.7 Loosely coherent 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, in principle, the method could be applied from scratch in a dedicated pipeline to identify those outliers. A stand-alone program is in development to permit "blob" searches for narrow regions in frequency, frequency derivative and sky location. Exploratory work will be completed in the coming year, with application to full S6 data in the following year.

5.3.8 Cross-correlation search

The cross-correlation method has been desribed previously. The plan for S5 is to use it in a directed search but it could in principle also be used for an all-sky search complementing the existing semi-coherent searches. The current plan is to use this technique for all-sky and binary searches only in S6.

5.4 Accreting and unknown binary neutron stars

For this class of source the gravitational radiation is thought to be powered by accretion onto the neutron star and not, as is the case for isolated neutron starts, by its own rotation. In this scenario, first proposed by [289], 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 [277]. The resulting indirect limit can be put in terms of X-ray flux F_x and spin frequency $f_{\rm rot}$ as

$$h_{\rm IL} = 5 \times 10^{-27} \left(\frac{300\,\rm Hz}{f_{\rm rot}}\right)^{1/2} \left(\frac{F_x}{10^{-8}\,\rm erg\,cm^{-2}\,s^{-1}}\right)^{1/2}. \quad (4)$$

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 $\sim$85 known LMXBs the spin frequency of the neutron star is unknown but thought to lie in the range $\sim$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 $\sim$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 completed a single coherent analysis on the accreting neutron star in the LMXB Sco X-1 using S2 data [277]. 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 [334].

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.

5.4.1 Known binary systems

5.4.1.1 Sideband search for known binary systems

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 [319, 320]. 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 will be completed soon and preliminary testing and tuning performed on S5 and S6/VSR2 data. However, the expected sensitivity of this search will only become astrophysically interesting (i.e., will start challenging accretion balance upper-limits) for advanced LIGO and specifically for Sco X-1.

As written above, the method assumes constant frequency over the observation. But it can be extended to the case of changing frequency, e.g. due to fluctuating accretion rate, with semi-coherent methods. A natural choice to investigate in this context is the stack-slide method, which could use coherent integration lengths of order two weeks [286].

A paper describing an S5 search for Sco X-1 using the simplest version of the method is expected to be submitted for publication in the coming year.

5.4.1.2 Cross-correlation searches for known binary systems

The cross-correlation search described in section 5.2.5 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 [277], and a more sensitive search than the incoherent cross-correlation search done in [334].

We will extend the cross-correlation code written for isolated neutron stars, described in section 5.2.5, and apply it in a search for radiation from Sco X-1. The search will initially be performed on S6/VSR2 data, with an eye towards having the pipeline fully developed by the time advanced detectors come on line.

5.4.2 Unknown binary systems

5.4.2.1 Polynomial search for unknown binary systems

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 $f_{\rm ssb}$ detected in the solar system barycenter may be modeled as

$$f_{\rm ssb} = f_{\rm gw}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right) \tag{5}$$

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

$$\frac{df_{\rm ssb}}{dt} = f_{\rm gw}\gamma\left(1 - \frac{d\vec{v}}{dt}\cdot\frac{\vec{n}}{c}\right) + \frac{df_{\rm gw}}{dt}\gamma\left(1 - \frac{\vec{v}\cdot\vec{n}}{c}\right). \tag{6}$$

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 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 amount of parameters 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. [286] 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 has been tested on simulated data with white noise. The documentation of the code is being prepared, as well as a document describing the search strategy and the results of the tests with simulated data. A methods paper will be submitted for publication in the coming year and preliminary S6/VSR2 results obtained. Publication of results is expected in the following year.

In parallel, the Polynomial code has been ported to a platform compatible with graphical processor unit (GPU) cards, in the hope of gaining significant speed in this computationally limited search. The software is undergoing optimization, to make the best use of GPU architecture.

5.4.2.2 TwoSpect search for unknown binary systems

The TwoSpect search is a hierarchical method under development 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. We plan to complete the search pipeline and begin production running in late summer 2010. It seems unlikely that the search will have the sensitivity to make a detection in S5 or S6/VSR2 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 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.

Preliminary results from an S6/VSR2 search for unknown binaries will be produced in the coming year, with publication expected in the following year. A methods paper will also be submitted for publication in the coming year.

5.5 Search for Non-Sinusoidal Periodic Waves

Our searches for continuous waves focus mainly on waveforms that are nearly sinusoidal, with smooth modulations in frequency and amplitude. But in the the spirit of keeping eyes wide open, it is reasonable to look for other periodic gravitational waveforms, such as periodic pulses similar to the radio pulses that led to the original electromagnetic discovery of pulsars. In the Fourier domain these non-sinusoidal waveforms could contain many frequency harmonics, no one of which has sufficient power to be detectable in a conventional CW search.

A number of algorithms can be applied to detect such waveforms, including incoherent harmonic summing [330] and the Gregory-Loredo method [331]. We have begun to explore the use of the Gregory-Loredo method, which has been used previously in radio, X-ray and gamma ray astronomy to search for non-harmonic periodicity in time-series. It is designed to be efficient in detecting pulsating signals in sparse-sampled data. We will study the tradeoffs in detection efficiency *vs.* computational cost for non-sinusoidal pulses when applied to the high-duty-factor LIGO and Virgo data. Initial exploratory studies are being carried out in the Matlab environment. If the method proves to be promising, it will likely be implemented in the offline DMT environment, to increase computational efficiency.

It should be noted that the Gregory-Loredo method may also prove useful in detector characterization to identify periodic instrumental glitches or periodic non-stationarities. A DMT implementation of the search code could be applied straightforwardly for such glitch searching in online, real-time detector monitoring.

5.6 Needs and plans for S6/VSR2

The CW Group has traditionally carried out its searches on the most sensitive data available, since it has focused on sources that should persist across data runs. For the most part, we expect to continue this practice going forward. For example, although the newest search pipelines described above will likely be developed using S5 data, for some of them the first use for publication may well be carried out on S6/VSR2 data, which we expect to be substantially more sensitive than S5 data. Hence for off-line analysis of S6/VSR2 data, what is written above for S5 analysis carries over with little modification.

The CW Group has a strong interest in monitoring (and mitigating) instrumental spectral lines (see the detector characterization section) 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 have proven 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 in the coming year will be dedicated F-scans for bands relevant to special pulsars, such as the Crab and Vela.

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 will be given to narrow spectral bands around a half dozen known pulsars for which the spindown limit may be reached in S6/VSR2. 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 transient events and stationary spectral lines. It has been run offline on VSR2 data, but will run in real-time on VSR3 data. Additional planned improvements for the coming year include a new user-friendly interface and automated computation of coherences among channels.

5.7 Scientific goals 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 plan to have in place on the first day of science run data taking the following mature search pipelines:

- At least two independent targeted-search pipelines ready to search for O(100) pulsars with known ephemerides from radio, X-ray and gamma-ray astronomers. *(Three targeted pipelines are in use* *now, including the Glasgow pipeline used to search for 116 known pulsars in S5 data.)*

- At least one stable and proven fast-track all-sky search for isolated neutron stars. *(The PowerFlux program used in S4, S5 and S6 searches provide a default pipeline.)*
- A stable and proven version of Einstein@Home able to carry out production running and to finalize post-processed results within three years of the start of data taking.
- At least three independent, stable and proven pipelines able to search for Sco X-1. *(The current Sideband pipeline and the radiometer search pipeline of the Stochastic group provide defaults. We expect several other pipelines (Cross-correlation, Polynomial and TwoSpect) to be available for Sco X-1 application, too.)*

We also plan to have in place on Day One somewhat lower-priority searches that are nonetheless already proven or well under development, from which we can confidently expect timely, publishable results in the advanced-detector era:

- At least one stable and proven directed-search pipeline for isolated neutron stars (known direction, but unknown frequency) based on long coherence times. *(Recent searches for Cas A, Calvera, and the galactic center, all based on the $\mathcal{F}$-statistic, provide prototype approaches.)*
- At least one stable and proven directed-search pipeline for isolated neutron stars based on a robust, phase-insensitive method. *(The ongoing search for Supernova 1987A using cross-correlation techniques provides a prototype.*
- At least one stable and proven all-sky search for unknown neutron stars in binary systems. *(The Polynomial and TwoSpect pipelines now approaching maturity on S6/VSR2 searches provide prototypes.)*

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.

5.8 Preparing for the advanced-detector era

In preparing to meet its 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.

To understand why this consolidation is needed, it is helpful to consider the situation as of July 2010:

- We have **3** targeted pipelines, all of which have produced preliminary results on the Vela pulsar using VSR2 data.
- We have **3** all-sky isolated-star searches with preliminary results from the S5 data; **2** others reaching maturity on the VSR2 data, and **2** nascent pipelines.
- We have **1** directed search with results submitted for publication; **2** others with preliminary results, and **2** pipelines near completion.
- We have **1** search for Sco X-1 with preliminary S5/VSR1 results and **1** other pipeline under development.
- We have **2** all-sky searches for unknown binaries nearing maturity.
- We have **1** "transient CW" search (known direction, approximate frequency, high spindown rate) under development.
- We have **1** search for "quasi-periodic" gravitational waves under development.

This proliferation of pipelines is excessive, despite our wish to have a broad suite of approaches. We plan to prune some of these search pipelines in the coming years by developing a standard set of performance metrics that allow us to identify unnecessarily redundant pipelines.

The goal is not to eliminate every pipeline but one in each category. But we do want to understand clearly the justification for each one that remains. The justifications can include:

- Best sensitivity in most interesting region of parameter space
- 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 for safety in a critical, flagship 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.

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 quickly reveal which pipelines can be safely abandoned and which should receive the efforts needed for stabilization in the advanced-detector era.

6 Searches for stochastic backgrounds

6.1 Sources of Stochastic Gravitational-wave Background

The stochastic background searches target a broadband and continuous background of gravitational waves, that could be produced by a large collection of incoherent sources. Sources of stochastic gravitational-wave background could be cosmological (such as inflationary models, cosmic strings models etc) or astrophysical (such as rotating neutron stars, low-mass X-ray binaries (LMXBs) etc).

One of the searches performed by the Stochastic Background Working Group targets an isotropic gravitational-wave background. The isotropic background is predicted by different models, and it can be completely described in terms of dimensionless $\Omega_{\rm GW}(f)$, the gravitational-wave energy density per unit logarithmic frequency (in units of the closure density of the Universe). Different models predict different spectral shapes in the LIGO frequency band (roughly $50-150$ Hz), although they typically follow a power-law form. Hence, the group performs the stochastic background search for different power-law forms of $\Omega_{\rm GW}(f)$. The increasing sensitivity of LIGO interferometers has allowed the group to start exploring the implications of the stochastic background searches for various models. In particular, the most recent result of the isotropic background search, based on the LIGO S5 science run, has started to explore cosmic strings and pre-big-bang models. In the case of cosmic strings models, a population of models has been ruled out, that was not accessible to other measurements and observations.

The group is also performing searches for non-isotropic stochastic background. This includes the radiometer search which targets localized foreground (astrophysical) point sources and the spherical harmonics search which targets stochastic sources spatially extended across the sky. The potential point sources include low-mass X-ray binaries, rotating neutron stars etc, and are expected to follows the local matter distribution in our galactic neighborhood. The potential extended sources include a number of cosmological models, as well as the galactic plane as a whole.

Finally the group started looking for long-lasting transients, i.e. signals that are present for a duration of minutes, hours, or longer. Such signals may for instance originate from pulsar glitches, which are narrow-band and relatively long in duration (possibly days), as well as sources such as long-lasting Gamma Ray Bursts (GRBs) or Soft Gamma Repeaters (SGRs), which may be broadband in frequency but relatively short in duration (seconds or possibly minutes). The duration of those signals is long enough to make the stochastic method described below an efficient way to search for them.

6.2 Stochastic Search Method

The stochastic search method has evolved form a specific search for an isotropic GW background (see section 6.2.1), to a directional seach for point-like sources (section 6.2.2), to an algorithm estimating the maximum likelihood strain power distribution across the sky (section 6.2.3), to a search for long-lasting transient signals correlated between different detectors (section 6.2.6). The first two have been used to analyse LIGO data in the past. The third one is capable of producing the same results as the other two as a special case output, and thus has the prospect of superseding them.

6.2.1 All-Sky Search

A stochastic background of gravitational waves (GWs) is expected to arise as a superposition of a large number of unresolved sources, from different directions in the sky, and with different polarizations. It is usually described in terms of the logarithmic spectrum:

$$\Omega_{\rm GW}(f) = \frac{1}{\rho_c}\,\frac{d\rho_{\rm GW}}{d\ln f}\,, \tag{7}$$

where $\rho_{\rm GW}$ is the energy density of gravitational waves, ρ_c is the critical density of the Universe, and f is frequency. The effect of a SGWB is to generate correlations in the outputs s_A, s_B of a pair of GW detectors, which can be described for an isotropic background in the Fourier domain by

$$\langle \widetilde{s}_A^*(f)\, \widetilde{s}_B(f')\rangle = \frac{1}{2}\delta(f-f')\,\gamma_{AB}(f)\,S_{\rm gw}(f) \tag{8}$$

where $\tilde{s}_A$ and $\tilde{s}_B$ are the Fourier transforms of the strain time-series of two interferometers ($A \neq B$).

The raw correlation depends on the (one-sided) power spectral density $S_{\rm gw}(f)$ the SGWB would generate in an IFO with perpendicular arms, as well as the observing geometry. The geometrical dependence manifests itself via the overlap reduction function (ORF)[335], which can be written as

$$\gamma_{AB}(f) = d_{Aab}\, d_B^{cd}\, \frac{5}{4\pi} \iint d^2\Omega_{\hat{n}}\, P^{\rm TT\hat{n}}{}^{ab}_{cd}\, e^{i2\pi f \hat{n}\cdot(\vec{r}_2-\vec{r}_1)/c} \tag{9}$$

where each IFO's geometry is described by a response tensor constructed from unit vectors $\hat{x}$ and $\hat{y}$ down the two arms

$$d^{ab} = \frac{1}{2}(\hat{x}^a\hat{x}^b - \hat{y}^a\hat{y}^b)\,, \tag{10}$$

$\vec{r}_{1,2}$ is the respective interferometer's location and $P^{\rm TT\hat{n}}{}^{ab}_{cd}$ is a projector onto traceless symmetric tensors transverse to the unit vector $\hat{n}$.

We deploy a cross-correlation method to search for the stochastic GW background, following [336]. In particular, we define the following cross-correlation estimator:

$$Y_{AB} = \int_{-\infty}^{+\infty} df \int_{-\infty}^{+\infty} df'\; \delta_T(f-f')\; \tilde{s}_A(f)^*\; \tilde{s}_B(f')\; \tilde{Q}_{AB}(f')\,, \tag{11}$$

where δ_T is a finite-time approximation to the Dirac delta function, and $\tilde{Q}_{AB}$ is a filter function. Assuming that the detector noise is Gaussian, stationary, uncorrelated between the two interferometers, and uncorrelated with and much larger than the GW signal, the variance of the estimator Y_{AB} is given by:

$$\sigma^2_{Y_{AB}} = \frac{T}{2}\int_0^{+\infty} df\, P_A(f)P_B(f)\mid \tilde{Q}(f)\mid^2\,, \tag{12}$$

where $P_i(f)$ are the one-sided power spectral densities (PSDs) of the two interferometers, and T is the measurement time. Optimization of the signal-to-noise ratio leads to the following form of the optimal filter [336]:

$$\tilde{Q}_{AB}(f) = N_{AB}\, \frac{\gamma_{AB}(f)S_{GW}(f)}{P_A(f)P_B(f)}\,, \text{ where } S_{GW}(f) = \frac{3H_0^2}{10\pi^2}\,\frac{\Omega_{GW}(f)}{f^3}\,. \tag{13}$$

$S_{GW}(f)$ is the strain power spectrum of the stochastic GW background to be searched. Assuming a power-law template spectrum with index α, $\Omega_{GW}(f) = \Omega_\alpha(f/100\text{ Hz})^\alpha$, the normalization constant N_{AB} is chosen such that $\langle Y_{AB}\rangle = \Omega_\alpha T$.

In order to handle gaps in the data, data non-stationarity, and for purposes of computational feasibility, the data for an interferometer pair are divided into many intervals of equal duration (typically 1-3 minutes), and Y_I and σ_{Y_I} are calculated for each interval I. The loss in duty-cycle due to the finite interval size is of order 1 minute for each analyzable data segment (which is typically several hours). The data in each interval are decimated from 16384 Hz to 1024 Hz and high-passed filtered with a 40 Hz cut-off. 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 intervals are overlapped by 50% to recover the original signal-to-noise ratio. The effects of windowing are taken into account as discussed in [338].

The PSDs for each interval (needed for the calculation of $Q_I(f)$ and of σ_{Y_I}) are calculated using the two neighboring intervals. 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_I calculated using the neighboring intervals with σ'_I calculated using the interval I, we identify intervals 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-sec 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 intervals that pass all the data-quality cuts are averaged with $1/\sigma_I^2$ as weights, yielding the final estimates of Y and σ_Y.

6.2.2 Directional Search

The analysis described above is designed to search for the signal integrated over the whole sky. It is also possible to search for anisotropies in the GW background. One way to approach the problem is to define a sky-position dependent optimal filter. As discussed in [339], one can write:

$$Q(t,f,\hat{\Omega}) = N(t,\hat{\Omega}) \, \frac{\int d\hat{\Omega}' \, \gamma(t,f,\hat{\Omega}') \, A(\hat{\Omega},\hat{\Omega}') \, H(f)}{P_1(f) \, P_2(f)}, \tag{14}$$

where $A(\hat{\Omega},\hat{\Omega}')$ reflects the anisotropy in the GW spectrum across the sky. For point sources, one chooses $A(\hat{\Omega},\hat{\Omega}') = \delta^2(\hat{\Omega},\hat{\Omega}')$. Note, also, that the overlap reduction function γ is now dependent on the sky-position, as well as on the sidereal time t. Following the procedure analogous to the one outlined in the previous Section leads to an estimate of Y and σ_Y for every direction on the sky - i.e. a map of the GW background. However, this map is "blurred" by the antenna patterns of the interferometers. The problem of deconvolving the antenna pattern from this map is non-trivial and is being actively pursued.

6.2.3 Mapping

The methods described in 6.2.1 and 6.2.2 are optimal under the assumption that the background is either isotropic or dominated by point sources, but neither addresses the question of estimating the actual spatial distribution of a stochastic background. A method that does this is described in this section.

The spatial distribution $\mathcal{P}(\hat{\Omega})$ of the strain power of stochastic background can be expanded with respect to any set of basis vectors on the sphere:

$$\mathcal{P}(\hat{\Omega}) = \mathcal{P}_\alpha \mathbf{e}_\alpha(\hat{\Omega}) \, , \tag{15}$$

Defining $C(f,t)$ as the cross-power between the output of the two detectors,

$$C(f,t) = \frac{2}{\tau} \, \tilde{s}_1^*(f,t)\tilde{s}_2(f,t) \, , \tag{16}$$

one can show that its expectation value is given by

$$\langle C(f,t)\rangle = H(f)\gamma_\alpha(f,t)\mathcal{P}_\alpha \, , \tag{17}$$

with $H(f)$ the strain power spectrum of the stochastic background. The $\gamma_\alpha(f,t)$ are basis dependent geometric factors that can be pre-calculated and play the role of the overlap reduction function in the isotropic analysis. The covariance matrix of $C(f,t)$ is given by

$$\begin{aligned} N_{ft,t't'} &= \langle C_{ft}C^*_{f't'}\rangle - \langle C_{ft}\rangle\langle C^*_{f't'}\rangle \qquad (18)\\ &\approx \delta_{tt'}\delta(f-f')P_1(f,t)P_2(f,t), \qquad (19) \end{aligned}$$

with P_1 and P_2 the strain noise power spectra of the two detectors.

Assuming Gaussian noise, the likelihood for measuring a specific cross-power $C(f,t)$ is

$$p(C_{ft}|\mathcal{P}_\alpha) \propto \exp\Bigg[-\frac{1}{2}((C^*_{ft} - \langle C^*_{ft}\rangle)N^{-1}_{ft,f't'} (C_{f't'} - \langle C_{f't'}\rangle)\Bigg] \quad (20)$$

where $\langle C_{ft}\rangle$ is given by 17 and repeated ft and $f't'$ indices are summed and integrated over—e.g., $\sum_t \int_{-\infty}^{\infty} df$..

Now one can ask for the $\mathcal{P}_\alpha$ that maximize this likelihood. They are given by

$$\hat{\mathcal{P}}_\alpha = (\Gamma^{-1})_{\alpha\beta} X_\beta \quad (21)$$

where

$$X_\beta = \sum_t \tau \int_{-\infty}^{\infty} df\, \gamma^*_\beta(f,t)\, \frac{H(f)}{P_1(f,t)P_2(f,t)}\, C(f,t)\,, \quad (22)$$

$$\Gamma_{\alpha\beta} = \sum_t \tau \int_{-\infty}^{\infty} df\, \gamma^*_\alpha(f,t)\, \frac{H^2(f)}{P_1(f,t)P_2(f,t)}\, \gamma_\beta(f,t)\,. \quad (23)$$

The matrix inversion in 21 in practise requires a regularization scheme because the interferometer pair can be insensitive to particular background distributions.

Note that if one restricts the basis set to either just an isotropic component or just a point source at a given location, one will get exactly the analysis described in 6.2.1 and 6.2.2 respectively.

While this algorithm in principle would work in any basis, a basis with a natural resolution cut-off will reduce the required number basis vectors and thus simplifies the required matrix inversion. One obvious such basis is formed by Spherical Harmonics.

6.2.4 Multi-baseline: LSC/VIRGO joint search

As shown in [336], the optimal method for combining more than two detectors is to make pairwise correlation measurements, and then combine these results in the same way measurements from different times are combined: average the point estimates Y with a relative weighting of σ^{-2}, or equivalently in the mapping formalism, sum up the X_β and the Fisher matrices $\Gamma_{\alpha\beta}$. As discussed in [337] the inclusion of the LIGO-Virgo pairs can enhance the sensitivity of the global GW detector network to an isotropic background of gravitational waves, particularly at frequencies above 200 Hz. Furthermore, the addition of a third instrument with comparable live time and sensitivity improves both the resolution and sensitivity of the mapping algorithm, effectively simplifying the regularization problem mentioned in section 6.2.3.

6.2.5 H1H2 All-Sky Search

The all-sky search outlined above is usually applied to the non-collocated interferometers (such as the two 4-km interferometers at Hanford and Livingston), in order to minimize the instrumental correlations. However, the overlap reduction for this interferometer pair is significant above 50 Hz. Hence, the collocated pair of Hanford interferometers could potentially lead to a $\sim 10\times$ more sensitive all-sky stochastic result, but it is also more susceptible to instrumental correlations. The stochastic group has developed two methods to handle this problem.

One approach relies on the coherence, defined as

$$\Gamma_{XY}(f) = \frac{|P_{XY}(f)|^2}{P_{XX}(f)P_{YY}(f)} \quad (24)$$

where P_{XY} is the cross-power spectrum between channels X and Y, and P_{XX} and P_{YY} are the two power spectra. As discussed in [340], it is possible to estimate the instrumental correlations between interferometers 1 and 2 by

$$\Gamma_{instr,12} \approx \max_i(\Gamma_{1Z_i} \times \Gamma_{2Z_i}) \tag{25}$$

where Z_i are the numerous environmental channels, including microphones, seismometers, accelerometers, power-line monitors etc. As discussed in [340], this method can be used to identify frequency bands in which the instrumental correlations between two interferometers are large. These bands could then be removed from the all-sky stochastic search. Moreover, the method can be used to estimate the residual contamination in the "good" frequency bands.

The second approach relies on time-shifting one GW channel with respect to the other. Since the stochastic GW background is expected to be broadband, its coherence time is much shorter than ∼1-sec, so the GW correlations between the two channels are expected to disappear at 1-sec time-shift. However, narrow-band features (of width ∼ 1 Hz) are expected to survive 1-sec time-shift. Hence, this method can also be used to identify narrow-banded instrumental correlations. The first tests indicate that the two methods agree well, but further studies of the systematic errors of the two methods are still required.

6.2.6 Searching for Long-Lasting Transients

The stochastic group has initiated a new effort targeting long-duration transient gravitational-wave signals, on the time-scale of minutes, hours, or longer. These searches are a natural extension of the cross-correlation techniques described above. In particular, the group started generating the Stochastic Intermediate Data (SID), containing the power and cross spectral densities for a pair of interferometers, calculated for every minute of data. The original motivation for this upgrade was the fact that different stochastic searches (isotropic, radiometer, and spherical harmonics) could directly use the SID instead of having to individually access and process the time-series data, thereby simplifying the analysis procedures and reducing the computational requirements of the group.

It is, of course, possible to parse the SID in search for correlated transients on the time-scale on which the SID is generated (minute or longer), or for purposes of characterization of the detector performance. These intermediate duration transients are not targeted by the traditional searches for transients (which tend to focus on second or sub-second bursts, or on several-second long inspiral signals).

Astrophysical studies performed to date have shown that the possible sources on these time-scales include pulsar glitches, which are narrow-band and relatively long in duration (possibly days), as well as sources such as long-lasting Gamma Ray Bursts (GRBs) or Soft Gamma Repeaters (SGRs), which may be broadband in frequency but relatively short in duration (seconds or possibly minutes). Consequently, the stochastic group started developing several search pipelines within the scope of this effort. Some of them are targeting broad-band signals (eg using the power excess algorithm), and some are targeting narrow-band signals (eg using Hough or Radon transforms).

Moreover, there are different possible modes in which these pipelines could be applied to the LIGO and Virgo data. In particular, the analyses could be performed in the "all-sky" mode, in which the entire sky is scanned at all frequencies. While the clear advantage of this mode is its breadth, the computational constraints imply that the parameter space can be searched relatively crudely. An alternative is to search for "externally triggered" sources - for example, focusing on the GRBs or pulsar glitches that are detected by electromagnetic observations would allow much finer parsing of the parameter space.

6.3 Results and Plans

6.3.1 Status of S5 Searches

Isotropic Search The stochastic group has repeating the isotropic search with LHO-LLO interferometer pairs using the S5 data. The final results is a new 95% upper limit on the gravitational-wave energy density $\Omega_0 < 6.9 \times 10^{-6}$ for a frequency independent spectrum ($\alpha = 0$) in the band 41-169 Hz. This result is 10 times more sensitive than the previous upper limit based on S4 data [332], and it is more sensitive than the Big-Bang Nucleosynthesis bound and the Cosmic Microwave Background bound in the LIGO frequency band. The result was published in Nature [341], including the implications of the new result for the models of early-universe cosmology, for cosmic (super)string models and for pre-big-bang models.

Radiometer Search The stochastic group is repeating the radiometer search using the S5 data of the LHO-LLO 4-km interferometer pair (and using the new pipeline). This analysis is expected to produce $\sim 10\times$ more sensitive maps of the GW sky than those produced using the S4 data [334], and it will apply an algorithm for deconvolution of the antenna pattern from the maps. This search will also produce a second isotropic measurement, with similar sensitivity improvement over the S4 all-sky result.

Spherical Harmonics Search In addition to the radiometer analysis, the stochastic group is pursuing a directional search based on spherical harmonics. The goal of the search is to estimate the spherical-harmonic decomposition of the gravitational-wave sky, similarly to what is done in the field of Cosmic Microwave Background. This method would allow searches for complex source distributions on the sky. The group has developed the formalism for this search, has performed studies of expected sensitivities and correlations between different spherical-harmonic components, and has completed a series of simulations and tests geared toward understanding the relationships between the three different searches (isotropic, radiometer, and spherical harmonics), their relative advantages, and their limitations. These studies have been summarized in a method paper which was published in Physical Review D [342]. The group has also applied the new algorithm to the S5 data, the results of which are currently being finalized.

Isotropic Search using Collocated Hanford Interferometers The isotropic searches performed up to date have preferred using the non-collocated interferometer pairs because of their insensitivity to instrumental or environmental correlations. The LHO interferometer pair, however, could potentially be $\sim 10\times$ more sensitive to stochastic GW background, because the antenna pattern overlap of collocated interferometers is optimal. However, the collocated interferometer pair also suffers from the instrumental correlations, because the two interferometers share the same environment and the same sources of instrumental noise.

The stochastic group is developing 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 preliminary results 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. More effort is needed to understand the systematic errors of the applied techniques, and to assess the possibility of some remaining (undetected) instrumental or environmental correlation contaminating the result.

All-sky Search at Free-Spectral-Range Frequencies The stochastic group has performed an isotropic stochastic search at the free-spectral-range (FSR) frequencies (37.5 kHz) using one month of the S5 data of the collocated LHO interferometers, and has produced a measurement of the stochastic GW background at 37.5 kHz. Efforts were made to understand the timing accuracy of the fast gravitational-wave channels of the interferometers, and the interferometers' calibration at these frequencies. The results are currently being internally reviewed by the LSC.

LIGO-VIRGO Searches The stochastic group is also conducting a joint LIGO-VIRGO stochastic search, using the shared S5/VSR1 data (data acquired between May 18, 2007 and October 1, 2007). Although the LIGO-VIRGO interferometer pairs are less sensitive than the LIGO 4-km interferometer pair to the isotropic stochastic background at frequencies below 800 Hz, above 800 Hz the LIGO-VIRGO pairs are similar or even more sensitive than the LIGO-LIGO pairs. Moreover, the LIGO-VIRGO pairs have different zeroes in the overlap reduction function, which could improve the overall network sensitivity even at lower frequencies. The analysis is mostly complete and the internal review process is underway.

Non-Gaussian Search The group is exploring the possibility of searching for non-Gaussian stochastic gravitational-wave background, also known as the "popcorn noise". The basic formalism for this search has been developed for co-located, white detectors. Preliminary tests of the formalism have been successfully performed, however the extension of this method to realistic interferometer pairs is still under investigation. The hope of the group is to perform the non-Gaussian search using all of S5 data, thereby improving on the sensitivity to non-Gaussian stochastic signals as compared to the standard isotropic search.

Pipeline Upgrade The stochastic group has completed an upgrade of the analysis pipeline. In particular, since several searches rely on similar quantities (such as strain cross and power spectral densities of different interferometers), the group has decided to produce Stochastic Intermediate Data (SID), stored in the frame format, and containing the commonly used quantities calculated for segments throughout the S5 run. In addition to simplifying the standard stochastic searches, the SID frames also finds use in detector characterization studies and in searches for GW transients on minute or hour time-scales. In particular, the SID frames combined with the new algorithms searching for long-lasting transients have led to a new S6 data quality flag identifying transients due to passing airplanes.

The first set of SID frames was produced for 52-sec long segments. This segment duration allows relatively simple averaging of the intermediate data to produce the cross and power spectral densities for segments of one sidereal day. The advantage of this approach is that the resulting data set is small enough that it could be stored on a personal computer, consequently simplifying different stochastic searches. However, shorter and longer segment durations are of interest to transient searches, so the group is preparing to produce multiple versions of SID, corresponding to different segment durations.

Searches for Long-Lasting Transients The stochastic group has initiated an effort to search for transient GW signals on the time scales of minutes, hours or longer. For this purpose, the SID frames are converted into time-frequency maps of the cross-correlation between two interferometers with time resolution of order 1 minute and with the frequency resolution of order 0.25 Hz. The time-frequency map is then parsed in search for different types of GW signals (broad-band or narrow-band). Several different algorithms are being implemented (power-excess, Radon transform, Hough transform etc), the group is currently assessing their performance, and preparing a method paper on the topic. The group is also drawing from experiences of other groups (such as the burst group) performing searches for short duration transients. The algorithms will be applied to S5 data starting in the summer or fall 2010, and are expected to produce first publications in 2011.

6.3.2 Plans for S6 Searches

Isotropic and Directional Searches The stochastic group plans to merge the isotropic, radiometer, and spherical harmonic decomposition (SHD) searches into one search, based on the SHD algorithm. The search will rely on the stochastic intermediate data (SID), likely in the collapsed form (to one sidereal day).

The strain sensitivity of the 4-km interferometers (H1 and L1) during S6 has improved as compared to S5. In addition, the sensitivity of the VIRGO interferometer during S6 is expected to be similar to those of LIGO interferometers at low frequencies (around 100 Hz). Hence, we expect the S6 stochastic searches to produce more sensitive measurements of the stochastic gravitational-wave background. Moreover, due to the longer baseline between the LIGO and VIRGO sites, adding VIRGO interferometer is expected to improve the angular resolution of the directional search.

Other Searches As noted above, the stochastic intermediate data (SID) will allow additional searches that have not been performed to date. In particular, the SID can be parsed in search for correlated transients on minute or hour time-scale. This search would benefit from closer collaboration with the burst group, as it will likely deploy an algorithm searching for power excess in the data.

The group will also pursue search for non-Gaussian stochastic gravitational-wave background. As noted above, this search is currently in developing stages, and will likely be applied to S5 LIGO data. The improved interferometer sensitivity during S6 and adding the VIRGO interferometer in the analysis could substantially increase the sensitivity of the search and will justify conducting it with the S6 data.

Other Activities The searches for long-lasting transients and the directed searches have one thing in common: they are performing several measurements (for different directions or times), averaging over less data than the previously performed isotropic search. Thus effects due to the non-Gaussian interferometer data will play a more important role for them. The group will investigate the influence of these effect on the analyses.

6.3.3 Plans for Advanced Detector Era

The network of advanced detectors is expected to be sensitive to the stochastic gravitational wave background at the level of $\Omega_{GW} \sim 10^{-9}$ or better. At this sensitivity, a number of astrophysical models (such as magnetars or binary neutron star coalescences) and of cosmological models (cosmic strings, alternative cosmology models such as pre-big-bang models) will become accessible. Moreover, the larger number of non-collocated detectors of similar sensitivity (up to five, if including LIGO-South option and LCGT) will also provide unprecedented angular resolution. These two moments define the most important preparatory efforts that the group will undertake:

- Consolidate the searches for anisotropic stochastic gravitational-wave background, and solidify our understanding of the algorithms and of possible systematics. This will likely require further development of the signal simulation infrastructure.

- Perform a systematic study of accessible models, and investigate possible ways of distinguishing between different models based on their frequency content and on the predicted anisotropy levels and angular scales.

- Improve the search algorithm to include estimating two different polarization components (current searches average over the two polarizations).

The superior sensitivity and angular resolution of the advanced detector network will also have implications for the development of the searches for intermediate duration signals. Similarly to above, the preparatory efforts of the group will include:

- Strengthen our understanding of the various search algorithms that are being developed, with attention on understanding possible systematics and biases, false-alarm rates etc. This includes further development of the signal simulation infrastructure.

- Investigate accessibility of different models of intermediate-duration transients. This includes models of magnetoturbulence in GRBs, protoneutron star convection, rotational instabilities in supermassive neutron stars etc.

- Investigate the possibilities of doing targeted searches, relying on the electromagnetic and neutrino-based observatories that will be available when the advanced detector network comes online. Develop the necessary infrastructure that would enable these searches to take place.

7 Computing and Software

The LIGO instruments deliver about 1TB/day of data. Even with only about 1% of this data in the gravitational-wave strain channel (the rest consists of detector and environment monitoring information) LIGO data analysis is a formidable computing challenge. Binary inspiral, burst and stochastic searches can utilize many Tflops of computing power to analyze the data at the rate it is acquired. *LIGO's scientific pay-off is therefore bounded by the ability to perform computations on this data.*

The LSC has adopted commodity computer clusters as the solution that meets its computational needs most cost effectively. Compute centers are geographically distributed at several sites around the world; this approach has the advantage that it puts resources close to the university researchers who are analyzing the data. Grid middleware allows for relatively easy access to data and computing power. If local resources are inadequate or a poor match, a researcher can access additional resources on the grid.

The LSC also developed the Einstein@Home project to leverage an alternative distributed computing paradigm for its most formidable computing challenge, the search for gravitational waves from isolated pulsars. The pulsar analysis puts reduced demand on quick turn-around and has low data flow, but requires PFlops of computing power. The analysis engine that underlies Einstein@Home utilizes much of the standard LSC software infrastructure described below; BOINC [2] is used to distribute work to thousands of volunteered personal computers world-wide.

7.1 Current status

The LIGO Data Grid (LDG) is the combination of computational and data storage resources, grid computing middleware and LSC services which, together, create a coherent data analysis environment for gravitational-wave science. With resources located at LIGO Laboratory centers (Caltech, MIT, LHO and LLO) and LSC institutions (UWM, Syracuse, and 3 sites in the EU managed by the GEO-600 collaboration), the LDG is a true distributed facility.

The LIGO Data Grid currently offers the minimal services required on a fully functional data grid. LIGO is in continuous science operation at unprecedented sensitivity, and the LDG continues to see growth in the number of users, higher demand for the resources, and construction of more sophisticated workflows. It is essential, therefore, to provide support of the LDG infrastructure, to provide user support and documentation, and to create the new services that gravitational-wave scientists will require. These services include: improved resource monitoring service and a resource brokering service to ensure that optimal use is made of LDG resources at all times; a metadata service to provide collation, distribution and access to the scientific results of searches; and a virtual organization management service to facilitate access control of LDG resources.

We anticipate evolution of the usage model as the community gains experience, so we are committed to a modular approach which allows us to remain light on our feet and to implement solutions which enable the best gravitational-wave science. A detailed description of the program of work on the LIGO Data Grid follows.

7.2 Activities in support of LDG Operations

1. **Hardware and Operating System Maintenance** The LDG clusters are all commodity clusters as this offers the most GFLOPs/dollar of capital investment. Using Linux requires an investment to track, and in some cases work around, this developing operating system. These are the traditional system-administration roles independent of grid activities.

[2] `http://boinc.berkeley.edu`

2. **Grid Middleware Administration** Each local cluster must maintain an evolving set of middleware in as stable a fashion as possible. The primary means to do this is the LDG Server software, discussed below. This software is rapidly evolving and requires effort to configure, support and maintain, independent of the effort required to create and maintain the LDG Server package itself.

3. **Data Distribution and Storage** The LDG currently uses the commercial SAM-QFS mass storage software from Sun Microsystems, commodity storage in the compute nodes, Linux based RAID servers, and the LIGO Data Replicator (LDR) to store and distribute data. Input data is common to the majority of analysis pipelines, and so is distributed to all LDG centers in advance of job scheduling.

4. **Reduced Data and Calibrated Data Products** The LIGO raw full-frame data files contain 13 to 15 thousand channels, including the uncalibrated gravitational-wave channel. Within the LIGO Data Analysis Systems (LDAS) at the observatories, a Level 1 Reduced Data Set (RDS) is generated that contains on the order of 300 channels (some of which are downsampled) that are the most useful for analysis and detector characterization. The Level 1 RDS files are about 20% the size of the raw frames, facilitating their distributed storage on cluster disk space for rapid I/O and for transfer to downstream LDG clusters. An even smaller Level 3 RDS is generated that contains just the gravitational-wave channel and instrument status information. Within the RDS infrastructure, LDAS also generates calibrated strain data using code from the LSC Algorithm Library, which it distributes as a separate set of frames files that are used for most offline analyses.

5. **Certificate Authority and User Accounts** LIGO uses X.509 certificates for authentication of users on the LDG. Several international Grid Certificate Authorities (CAs) supply user certificates, including DOEGrids CA in the USA. The LSC provides a simplified script interface for DOEGrids CA users within LIGO. RA agents are required to verify certificate requests to the CA and then approve them. LDG user accounts are requested via a web interface; these are also verified, and approvals are sent to each LDG site where local admins add the accounts.

6. **LIGO Data Grid Client/Server Bundles** LSC staff leveraged experience with the VDT and built upon the Virtual Data Toolkit (VDT) to create the LIGO Data Grid Client and Server packages. The server bundle enables LSC administrators to easily deploy standard grid services and middleware such as Globus GRAM and GridFTP across the LDG. The client bundle provides quick one-stop installation of all the software needed to gain access to the LDG resources by users in the LSC. Moreover, the LDG Client bundle provides scripts specific to the LDG to simplify certificate requests and other activities that users perform. Over the past year, the LSC has worked with the VDT team to migrate the LDG Client and Server to use native packaging for Linux platforms. The LSC now maintains these bundles in the LSCSOFT repositories for easy installation and configuration on the LDG. A Mac OS client suite is maintained in order to support the increasing number of scientists using this platform to access LDG resources. The LSC continues to collaborate with the VDT team to provide feedback on their software and distribution mechanisms.

7. **User Support** The LDG predominantly uses Condor for job queue management. As the analysis workflows for this new branch of astronomy are evolving rapidly, significant effort is required to work closely with the Condor development team to ensure efficient use of the LDG clusters. This feedback has been productive, with many timely bug fixes and feature enhancements being provided, however this requires significant effort from LDG administrators to isolate and troubleshoot issues that are particular to gravitational-wave data analysis. Compared with our High Energy Physics colleagues, the workflows that are being developed on the LDG are not yet as mature or stable, causing a significant burden on cluster administrative staff. Since the LDG users are generally scientists and not grid ex-

perts, staff are required to offer performance tuning in terms of GFLOP/s, job scheduling efficiencies, memory utilization, file management, and general debugging support for intermittent job failures.

8. **LIGO VO Support for OSG** Provide primary support for OSG usage of LIGO VO resources, continue to fulfill the responsibilities of OSG point of contact, security contact, and support center for LIGO, and handle any issues that arise for OSG users, OSG administrators and the OSG Grid Operations Center (GOC) while using LIGO facilities; regular participation in OSG Operations, OSG Integration, and OSG Support Center telecons. Maintain and administer the Virtual Organization Membership Service (VOMS) and LIGO Accounts Management System (LAMS) used to track users with Grid certicates approved to use LIGO Data Grid resources.

7.3 Data Analysis Software Development Activities

A suite of software tools are supported, developed and released by the LSC for the purpose of analyzing data from gravitational-wave experiments. These data analysis software projects are developed under the umbrella of the *Data Analysis Software Working Groups* (DASWG). Many of these projects have evolved into full scale software projects which enable most of the large scale analysis efforts within the LSC, thus requiring substantial effort to maintain them. Moreover, the LSC and the international community of gravitational-wave astronomers have embraced the grid-computing model and its associated technologies placing further demands on the software tools developed by DASWG.

1. **Data Monitoring Tools** The Data Monitoring Toolbox or DMT is a C++ software environment designed for use in developing instrumental and data quality monitors. About 50 such monitor programs have already been developed by members of the LIGO Scientific Community. DMT monitors are run continuously while LIGO is in operation, and displays produced by these monitors are relied on to give the operators immediate quantitative feedback on the data quality and interferometer state. In addition to their on-line use, the monitors and the software infrastructure they are based on have many offline applications including detector characterization, data quality determination and gravitational wave analysis. To facilitate the use of the DMT environment and monitors offline, the majority of the DMT package has been ported to the LSC offline processing clusters. Porting and packaging the DMT for offline use will continue to be supported.

2. **GLUE** The Grid LSC User Environment (GLUE) provides workflow creation tools and metadata services, written in Python, which allow LSC scientists to efficiently use grid computing resources within and external to the LIGO Data Grid. Analysis of data from gravitational-wave detectors is a complicated process typically involving many steps: filtering of the data from each individual detector, moving trigger data to a central location to apply multiple instrument coincidence tests, investigating auxiliary channels, and coherent combination of data from all detectors in the network. The description of these complicated workflows requires a flexible and easy to use toolkit to construct a virtual representation of the workflow and then execute it on a single cluster, across the entire LIGO Data Grid, or on external compute resources such as the OSG. The GLUE pipeline module provides this facility and is used by numerous LSC-Virgo data-analysis pipelines. GLUE is integrated with Pegasus workflow planner, allowing scientists to better manage the workflows generated using the pipeline module. Direct generation of Condor workflows is also supported. GLUE also provides an extensive suite of metadata management tools. The ligolw module provides a toolkit for generating and manipulating LIGO light-weight XML documents (LIGOLW XML). These documents are used to store many data products, from detector data quality and metadata information to the scientific products of searches. GLUE also probides the server tools to manage the extensive detector state metadata generated by the LIGO, Virgo and GEO detectors, as well as client tools used by LSC-Virco scientists to access these data.

3. **LSC Algorithm Library Suite** The LSC Algorithm Library Suite (LAL) is a collection of C language routine libraries that form the engine of the computationally-intensive data analysis programs. LALSuite routines are used in LAL Applications (collected in the LALApps package) which are programs that perform specific data analysis searches, and the LAL-Python interface (PyLAL) that provides access to LAL routines within the Python scripting environment. LALSuite contains (i) general purpose data analysis routines that provide common data analysis tools (e.g., routines to perform time-domain filtering, Fourier and spectral analysis, differential equation integrators), astrometric tools (e.g., routines for converting between sky coordinate systems and time systems), and gravitational-wave specific tools for signal simulation and data calibration; (ii) routines for reading and writing data in standard LIGO data formats; and (iii) implementations of search-specific gravitational data analysis algorithms. Enhancements are planned to improve the I/O routines to interface with LDR data catalogs directly and to leverage Grid tools to directly access data stored remotely. Also planned are significant improvements to the interface of the core analysis routines to make these routines easier to integrate into other software.

   C language applications for performing specific searches are contained in the LALApps package which is freely available under the GPL. This package provides a set of stand-alone programs that use LAL routines to perform specific pieces of a search pipeline. The programs can be strung together to form a data analysis workflow: a sequence of steps that transform the raw interferometer output into a set of candidate events. These applications continue to be enhanced and new ones developed.

   PyLAL is a Python module that includes extension modules that link against LAL, thereby making LAL routines available within the Python scripting environment. PyLAL thus provides a mechanism for rapid data analysis application development, for data exploration and graphing, and for performing quick follow-up analyses. As PyLAL matures, many more LAL routines will be incorporated so that significant aspects of the data analysis pipelines will be written in Python.

4. **MATLAB Applications** The MATLAB software suite is a commercial product which is widely used within the LIGO Scientific Collaboration (and the broader gravitational wave detection community beyond) for on-line and off-line data analysis, detector characterization, and operations. The MATLAB Applications package (MatApps) is a collection of gravitational-wave data analysis tools for use within the MATLAB environment that were written by the LSC members in support the analysis of LIGO, Virgo, and GEO data. This software is now maintained as part of the LSC MATLAB Applications (MatApps) project. Many of the contributions to MatApps are complete analysis tools developed by individual scientists and, as a result, there was considerable duplication within the repository. Recent initiatives seek to streamline the repository, better document its contents and share this knowledge with the MatApps community in order to minimize the duplication of efforts and increase ease of use. Streamlining has taken the form of migrating MatApps from a CVS to an SVN and flattening the repository structure for better intuitive use; this effort is ongoing. Improving the communication within MatApps includes the creation of a MatApps wiki, where users (including MatApps leadership) are continually developing the documentation content, and the creation of a MatApps discussion email list where users ask questions of the community at-large. A pilot newsletter has been issued and will be used in the future to communicate general information that may affect a user's interaction with MatApps (outages, new features in the latest MATLAB release, etc.). Better user support efforts are ongoing and include the creation of a dedicated RT tracking system for users to seek assistance with MatApps. Finally, MatApps intends to further reduce duplication of efforts by integrating more with other software projects within the LSC (e.g. LAL/LALApps, PyLAL, GLUE). Specifically, improvement to I/O routines can be made by interfacing with LDR and LDAS data catalogs. Through these streamlining and communication efforts, the collaboration will significantly increase the verifi-

ability and maintainability of this analysis software, while simultaneously reducing the barrier to the development of analysis software by individual researchers, educators and students.

5. **LIGO Data Analysis Systems Software** The LIGO Data Analysis Systems (LDAS) includes an important software component which provides (among other things) a frame API for interacting and reducing gravitational-wave frame data, a diskcache API for tracking the location of tens of millions of files mounted on hundreds of filesystems, a job management service for running frame and diskcache API jobs, and the maintenance of a C++ library for interacting with frame data. LDAS RDS and calibrated strain data generation, the data finding service provided by GLUE, and the data replication publishing service provided within LDR are among the software components that use LDAS software services.

6. **Support and Release of Software** The LSC now releases a unified build of the LSCSoft bundle for use by the LSC and other gravitational-wave scientists. This release method will be enhanced to include better support of platforms other than the cluster operating systems selected by the CompComm and DASWG.

   A well defined LSCSoft Software Release Protocol [3] has been developed over the past two years and is currently in use. This protocol requires that inclusion of new or modification/updating of existing packages in the LSCSoft bundle must be approved first by the Software Change Control Board (SCCB). These packages are then built, by the repository maintainers, for the officially supported operating systems [*CentOS 5.3, Debian 5.0 Lenny and MacOS X Leopard*].

   These packages [*rpm*'s for CentOS, *deb*'s for Debian] are maintained in `YUM` [for CentOS] and `APTITUDE` [for Debian] repositories at UWM. The external *MacPorts* repository is used for MacOS X. These comprise *Testing* and *Production*, 32 and 64 bit repositories. The CentOS build process leverages modern virtualization technologies, i.e., a testbed of Virtual Machines at UWM [*Xen, VMWare, VirtualBox*] which are used for building & testing the built the software before publishing it to the testing repositories and announcing their availability to the DASWG email list. For the Debian packages, a similar process [but without using virtualization technologies] is carried out by the Debian Team at Hannover Univ., which also maintains a mirror of the UWM repository. Once the testing phase ends, and if no errors are found in the packages, they are moved to the production repositories upon approval by the SCCB. The corresponding announcement of the official release is then made to DASWG email list.

   The next step in this project is to deliver fully functional virtual-machine images to downstream users. Initially, virtual machines will include the fill LSCSoft bundle and LDG client installed and configured to provide a fully integrated environment for analysis. It is further anticipated that users may wish to have custom configured virtual machines with selected software and applications installed. An interface will be developed to allow users to request such VMs which will be automatically built and delievered to them. In the long term, this approach will allow the LSC to maximally leverage Cloud Computing technioligies and may provide a route to reduce the total cost of computing for gravitational-wave astronomy.

7.4 Intermediate-term development activities

The distributed LDG relies on a number of grid services to allow robust, efficient operations. A minimal subset are currently deployed on the LDG. The full set is outlined here along with estimated personnel requirements to support, enhance and deploy them where appropriate.

[3] `https://www.lsc-group.phys.uwm.edu/daswg/wiki/SoftwareReleaseProtocol`

1. **Problem Tracking and Support** Robust operation of the LDG requires detailed problem tracking to insure that services are maintained and that security issues are quickly and efficiently addressed. There is already web based problem tracking facilities. This service needs to be extended and integrated with the LDG monitoring services. Over the next year, the informal knowledge base that exists in mailing lists and sprinkled throughout web pages and wikis will be harvested to develop a powerful and extensible help system. Furthermore, problem reporting and tracking will be simplified.

2. **Authentication and Authorization** The LSC relies on the Grid Security Infrastructure (GSI) from the Globus toolkit to authenticate users. GSI authenticates using X.509 certificates, which are currently obtained from a number of nationally operated grid certificate authorities from countries in which LSC member institutions reside. User access is provided at each site via hand-maintained grid map files. Users access standard unix accounts which are provisioned by hand by administrators. This approach does not scale sufficiently for the LDG. The OSG is using the VOMS-GUMS-PRIMA model for this purpose. The LSC has deployed these tools to share resources with OSG, but needs to explore all technologies that meet the collaboration's needs.

   Within the next year, this model will change substantially. Using the centralized authentication and authorization infrastructure currently being developed in the LSC and LIGO lab, short-lived X.509 certificates and proxy certificates will be supplied by MyProxy backed by LIGO.ORG CAs. MyProxy will leverage the existing centralized authentication infrastructure (in particular the LIGO.ORG and LIGO-GUEST.ORG kerberos realms) to link these certificates to user's identity in the LIGO.ORG LDAP. This will allow the capability for fine-grained access control and for automatic and uniform account provisioning on the LDG. Over the next several years, LIGO will be seeking TAG-PMA accreditation for the LIGO CAs to allow LIGO users to seamlessly interact with other scientific grids such as OSG.

   These developments are part of a larger effort, known as the Auth Project, which is described in more detail in 7.5.

3. **Monitoring Services** While the current LDG infrastructure is working well, it lacks of a fully deployed monitoring/information system. Having easy access to current information about the health of the LDG would allow us to prevent problems and/or troubleshooting issues much more effectively. Moreover, having access to historical data about usage and health of the LDG would facilitate decision making when the time comes to enhance or adjust the LDG. It is clear that aLIGO will require a considerable growth in the current computational infrastructure that will benefit from a fully functional monitoring service.

   One type of information inherent to grid computing models describes the status of clusters, their processes, their services, the status of jobs on the cluster, and the status of connectivity between clusters. In order to maximize the throughput, users and job submitting agents need to have access to this information. The LDG currently uses Ganglia to obtain snapshots of the status of clusters at different locations and then reports them to a central Ganglia metadata server. Enhancing monitoring services by including new tools to collate the information collected and to provide a consolidated Grid friendly interface is an essential step to improve efficiency.

   A prototype information service, the LSC Grid Information Service (LSCGIS), has been deployed which uses standard cluster monitoring tools and scripts to gather the required information and then exposes it via a RESTFul web service. This LDG-customized project can be enhanced by integrating it together with more general tools such as *Nagios*, for a finer metadata gathering. While this information is currently used to prevent/fix problems, it is clear that it can also be used to feed information into analysis pipelines or workflows to make them aware of available infrastructure and to make them more intelligent. The prototype LSCGIS should will continue to be evolved to address all of these possibilities.

4. **LIGO Data Replicator** The LIGO Data Replicator (LDR) replicates in bulk interferometer data files to LIGO Data Grid (LDG) computing sites, as well as the Virgo site in Cascina, Italy (CSC). LDR provides a metadata catalog for gravitational wave data files (typically with extensions .gwf and .sft) that in conjunction with other tools allows LIGO and Virgo scientists and their codes to discover data and other files within the LDG. Replication begins when data is *published* into the LDR network at a site. Publishing implies that relevant metadata about a file is entered into the local metadata catalog that is part of LDR and that a mapping from the logical filename (LFN) to an access path (typically a URL) or physical filename (PFN) is created in the local replica catalog (LRC). By the end of the LIGO S6 science run the LDR metadata catalog is expected to contain metadata information on more then 35 million files and each RLS replica catalog is expected to hold between 1 and 50 million mappings, depending on the data sets replicated to each site. Currently LDR is deployed at the LIGO Hanford site (LHO), LIGO Livingston site (LLO), Caltech (CIT), Massachusetts Institute of Technology (MIT), Syracuse University (SYR), University of Wisconsin-Milwaukee (UWM), Albert Einstein Institute Hannover (HAN), Cardiff University (CDF), and Birmingham University (BHM), as well as the Virgo site CSC. The CDF and BHM deployments leverage the “LDR as a service” model where only a GridFTP server is deployed at the site and the rest of the LDR logic and tools are hosted at UWM and provided as a service to the site. Investigations and testing continue to ensure scalability and performance meet the demands for the post enhanced LIGO era, especially since data from both the Virgo and GEO instruments will continue to be published, replicated, and discovered using LDR even as the LIGO instruments turn off after S6. Specific directions include tightly integrated web based monitoring to further ease the administrative burden, as well as migrating the LRC to a web services and more robust server platform.

5. **Data Quality and Segment Database** The lightweight database daemon (LDBD) provides a client and server framework for scientific meta-data services. LDBD is built on top of the existing LIGO authentication and authorization services, with a relational database back-end (DB2). This framework is designed to be extensible; the first application using it is the interferometer data quality service. Tools have been developed for low latency discovery and archival of Science and DQ segments for the S6 online and offline analysis. A production server at Caltech and a development server at Syracuse are currently providing critical metadata services for the LSC and Virgo collaborations. There are several tasks remaining to be completed in the short term: (i) the segment database lacks robust monitoring and fail-over solutions. The production database is backed up, but no hot spare exists. An automated fail-over solution must be developed, along with replication of segment information to redundant off-site systems; (ii) the existing LSC-developed infrastructure needs to be integrated with the Virgo segment services (This is currently performed by hand by LSC and Virgo scientists). In the intermediate term, production-level support will be provided for the LSC and Virgo collaborations through the end of the S6 run and during the era that VSR3 and GEO-HF will be operating.

6. **Event Database and Archival Project** The gravitational-wave candidate event database (GraCEDb) is a prototype system to organize candidate events from gravitational-wave searches and to provide an environment to record information about follow-ups. A simple client tool is provided in Glue to submit a candidate event to the database.

   An entity submits an event to Gracedb using the client tool in Glue or via the web interface. At the time of submission, the following things happen: 1) A unique ID is assigned to the candidate event. This UID is reported to the entity submitting the candidate event. The UID takes the form GXXXX, where XXXX is a number with a minimum of four digits. Extra digits will be used as needed. The UID is intended for internal use only. 2) The submitter, the search group, the search type that generated the event are recorded. 3) A web area is created to hold information about the candidate event and any follow-ups that are performed. These directories are accessible via web browsers and by logging into any of the submit machines at UWM, in particular hydra.phys.uwm.edu:/archive/gracedb/data/GXXXX. The general directories have the same permissions as /tmp in a Unix file system, so any LVC users can add content under that directory. The use of directories based on ligo.org usernames is encouraged to keep things organized. 4) A wiki page is created to allow easy entry of information about the candidate event. 5) An alert is published to the corresponding node in the LVAlert system; subscribers to that node receive the alert and initiate follow-ups. Alerts are also sent to the gracedb@ligo.org mailing list. The system continues to evolve to support the joint gravitational-wave and electromagnetic observing campaigns planned for fall 2010.

As with the DQ and Segment Database, the database lacks robust monitoring and fail-over solutions. The production database is backed up, but no hot spare exists. An automated fail-over solution must be developed, along with replication of event information to redundant off-site systems. In the intermediate term, production-level support will be provided for the LSC and Virgo collaborations through the end of the S6 run and during the era that VSR3 and GEO-HF will be operating.

LARS will be a collection of tools and services that provides archival storage for LIGO. A prototype has been delivered. The user tools are simple programs that are intended to allow LIGO scientists to catalog and share search results. Users may add descriptions and locations of search results to a simple database. This database may be queried by others to discover result locations. Individual results may be narrowed within a search by specifying a description in the result's LALCache. When found, query results may be listed or data may be presented to the user in a local directory, if sshfs is available. Work has started to leverage cli.globus.org services to provide transparent and efficient transport and distribution of data products by users. When implemented, we anticipate LARS will become an important tool for scientists.

7. **Multi-Site Scheduling and Brokering** The ability to plan, schedule, and monitor large workflows simultaneously across multiple LDG sites is becoming increasingly necessary in order to load balance across the computational resources distributed throughout the LDG and to support ever larger workflows which cannot easily or always be serviced within time constraints at a single LDG site. A number of intermediate-term development activities are focused on supporting LIGO data analysis workflows across multiple LDG sites as well as other "grid" sites external to LDG.

   One such activity focuses on leveraging the "Grid Universe" available with the Condor High Throughput Computing system and in particular "Condor-C", the Condor Grid type. Currently Condor manages most LDG computational resources (Linux clusters) at a site level. That is, each Linux cluster resource is its own Condor pool and jobs submitted to be run and managed at any single site only run within that same Condor pool. When properly configured, however, the jobs submitted at one site and into one Condor pool may migrate and be run and managed by a remote Condor pool, with the results and output being staged back to the original submission site as if the jobs had ran at the submitting site. An earlier attempt by Condor to support this type of migration of jobs was the Condor "flocking" mechanism. This newer approach known as Condor-C promises to scale better. LDG staff are evaluating Condor-C throughput and scaling behavior and providing feedback to the Condor team, as well as working to understand how best to abstract the details of Condor-C job submission and management away so that LDG users do not have to manage the details themselves.

   While Condor-C provides the "plumbing" to allow jobs to flow between clusters, LSC-Virgo workflows must be written to take advantage of the transport mechanisms Condor-C provides. One approach to solving this problem is leverages the Pegasus workflow mapping engine. GLUE has the ability tooutpu LSC-Virgo workflows in the abstract directed acyclic graphs (DAX) format (this is now the standard format used by CBC workflows). The Pegasus workflow planner can then be used to render these "abstract" workflows to "concrete" Condor DAGs. The actual management of the workflow is handled by Condor DAGMan. At present, Pegasus translates CBC workflows to Condor DAGs consisting of standard and vanilla Condor universe jobs targeted for a single LDG cluster. Pegasus provdes services such as bundling short running jobs into larger jobs and better management of input and output data products. LDG scientists are currently investigating Pegasus' ability to render workflows as Condor-C jobs which would allow execution accross multiple LDG sites connected with Condor-C.

Pegasus can plan also plan workflows for execution across sites that do not run Condor pools as well as into sites that do run Condor pools. LDG staff are evaluating Pegasus and working to understand how to tune Pegasus to schedule LIGO workflows across non-LSC sites (such as the OSG) most efficiently.

Finally, the use of pilot servers to provide a simple interface for the users that want to submit jobs on the LDG, but have them run on other resources including those available to the Virgo collaboration. An existing test system will be duplicatd and extended to provide efficient resource sharing across the LDG.

8. **Test and Build** To ensure the successful analysis of LIGO-Virgo data, it is increasing important to automate the validation of LIGO software and infrastructure. With continuous advancements in scientific analyses and computing technology, LIGO's software and computing infrastructure is growing in size and complexity. This trend is driving the need for more automated validation.

   As a result, automated build and test systems such as the NSF-sponsored Metronome framework can be of enormous benefit to LIGO. Such automated testing is also critical to the validation of changes to LDG system architecture, operating systems, and runtime environment. However, regression testing a distributed software stack is a computationally demanding task—an apparently harmless update of one component can cause subtle failures elsewhere in the system. And in the event of a critical security patch to one or more components, regression validation of the entire system absolutely must happen very quickly.

   Enabling an automated testing solution tailored to the needs of LIGO's distributed computing environment, will help ensure that changes to LIGO code or to LDG system architecture, operating systems, and runtime environment do not cause unexpected and undesirable changes to scientific results. Additional testing resources would also support testing the reproducibility of past results in the face of such changes. Automated test and build is essential to enable higher-quality software and prevent "bitrot" as LIGO scales past the capabilities of largely manual software and infrastructure validation processes.

9. **gstLAL** gstLAL in a software project to wrap the GW data analysis machinery of LALSuite in GStreamer "elements". GStreamer is a free C library that provides the infrastructure required to build complex realtime and non-realtime digital signal processing pipelines. GStreamer is primarily intended to be used in multimedia applications for the desktop, but it is of high quality with many features and easily satisfies our own needs for data analysis pipelines.

   Using gstLAL, a prototype application is being developed to search for the gravitational-waves from collisions of very low-mass primordial black holes. The PBH templates used by the search are up to 30 minutes long, and so the completion of this search will serve as a technology demonstrator for Advanced LIGO, proving that we have the software infrastructure required to handle the long templates in the flagship binary neutron star search. At present the PBH trigger generator program runs, but many problems remain to be solved before a full PBH search can be completed, for example how to perform background estimation with such long templates and how to practically construct banks of such long templates. These problems are outside the scope of gstLAL itself but the further development of gstLAL will be driven by their solution.

7.5 Preparing for the advanced detector era

The requirements for hardware, software and services neede for gravitational-wave astronomy in the advanced detector era has been ongoing for about a year now. A number of discussions and presentations have allowed the CompComm and DASWG to build a task list and determine the approximate FTE count to meet the needs. This section lays out guiding principles for some of the larger projects that will need to be completed in order to leverage aLIGO for the maximal scientific productivity. The details of the plan will be fleshed out during August and September 2010 in preparation for submission of a proposal to support the continued operations of the LIGO Data Grid including software support, enhancement, and release.

7.5.1 Software

1. **I/O Libraries** Because of the volume of data involved and the complexity of the algorithms we use to process it, searches for gravitational waves can quickly transform from problems of astrophysics and astronomy to problems of data management. Experience has taught us that the ease and speed with which data analysis challenges are solved is often closely related to the quality of the software libraries used to read and write data products, and so the selection of an I/O library is an important step in the development of a search for gravitational waves. Libraries with well-designed interfaces and robust bug-free internals allow us to spend more time doing science and less time solving I/O problems.

   Today, our searches rely on a combination of I/O libraries developed in-house and libraries maintained by third parties. Libraries developed in-house provide the benefit of being under our control — bugs that affect us are fixed when we need them to be, and we choose when to change software interfaces and file formats — but suffer by requiring people within the collaborations to do the design and maintenance work, people who often do not have a great deal of experience engineering and maintaining complex software projects and whose time is principally allocated to other tasks. Libraries developed externally, on the other hand, are often designed and maintained by people with greater software engineering experience, but sometimes see interface changes occur at times that are not convenient for us.

   We have seen a trend within the collaboration to transition from in-house libraries to libraries developed externally. For example, much of our XML I/O now relies on professionally-maintained XML parsers like expat instead of the metaio library developed in-house. In the future, this trend should continue. Whenever a new I/O challenge presents itself every effort should be made to research existing solutions, and use them when possible. In particular, we forsee a growing need for the network transport of many different types of data including astronomical alerts, audio-frequency time-series data in both realtime and non-realtime, database queries and other kinds of remote procedure calls. An enormous variety of technologies has already been developed for solving problems of these types and more. It is important to use those existing solutions whenever possible to allow the expertise and time of their designers and maintainers to streamline our own work, and to help drive the development of those projects so that gravitational wave astronomy can contribute to technological progress in other areas as well.

2. **Low-latency tools** It is not yet clear whether or not a network of ground-based gravitational-wave

	Task	Support	Programming	Architect	FTE total
Applications	DQ pipelines	TBD	TBD	TBD	TBD
	Low-latency analysis	TBD	TBD	TBD	TBD
	Offline analysis	TBD	TBD	TBD	TBD
	Simulations	TBD	TBD	TBD	TBD
	Other applications	TBD	TBD	TBD	TBD
	Open Data Workshops	TBD	TBD	TBD	TBD
	Task	Support	Programming	Architect	FTE total
Data Handling and Analysis Software	Architect	0	0	0.5	0.5
	Software R&D	0	0.5	0.5	1
	Support	0.6	0	0	0.6
	I/O Libraries	0.2	0.8	0	1
	Low-latency tools	0.5	0.5	0.5	1.5
	MatApps	0.3	0.7	0	1
	Service Clients	0.2	0.4	0	0.6
	LDAS	0.3	0.7	0	1
	LAL Suite (LAL, Glue, Pylal)	0.6	1.4	0	2
	DMT	0.4	1.4	0.2	2
	NDS	0.2	0.2	0	0.4
	LIGO DV	0.3	0.7	0	1
	Open Data Software Support	TBD	TBD	TBD	TBD
	Open Data Documentation	TBD	TBD	TBD	TBD
	Task	Support	Programming	Architect	FTE total
Data Handling and Analysis Services	Architect	0	0	0.5	0.5
	Middleware R&D	0	0.5	0.5	1
	Support	0.6	0	0	0.6
	Build & Test	0.4	0.3	0.3	1
	Workflow Service	0.4	0.3	0.3	1
	OSG/EGEE Integration	0.6	0.2	0.2	1
	Monitoring	0.4	0.2	0.2	0.8
	LDR	0.5	1.3	0.2	2
	DQ Database	0.5	1.2	0.3	2
	GRaCEDb	0.5	1.2	0.3	2
	Open Data Web Services	TBD	TBD	TBD	TBD
	Auth/Roster	0.5	1.2	0.3	2
	h(t) production	0.2	0.4	0.1	0.7
	RDS generation	0.2	0.4	0.1	0.7
	Open Data Support Services	TBD	TBD	TBD	TBD
	Open Data Cleaning	TBD	TBD	TBD	TBD
	Task	Support	Programming	Architect	FTE total
Data Center Operations	UWM-WebS	0.6	0.2	0.2	1
	UWM-Tier2	1	0	0	1
	SYR-Tier2	1	0	0	1
	LLO	1.5	0	0	1.5
	LHO	1.5	0	0	1.5
	MIT	1	0	0	1
	CIT	3	0	0.5	3.5
	Open Data Centers	TBD	TBD	TBD	TBD
	Totals	18	14.7	5.7	38.4

Table 1: A list of tasks and FTE requirements for LIGO Data Grid operations and software/service design, development, release and support. The support activity includes administration, help desks, packaging, testing, release. The architect activity refers to high-level architecture development. Notice that the applications layer is considered separately from core operations and support activities. All open-data activities remain TBD until the plan is formulated and accepted.

antennas can be used to successfully provide alerts of transient events to non-GW observatories, there is a significant probability that useful alerts will continue to flow the other way for many years into the advanced detector era. However, one of the challenges facing the search for GWs from binary neutron star collisions in the advanced detector era is the length of the template waveforms required by the search and the number of them. Advanced LIGO BNS templates might be up to 30 minutes in length and be more than an order of magnitude more numerous than the 45 s long templates used by initial LIGO. The increase in the BNS search's computational complexity indicates the need for a new approach to the problem of matched-filtering, in particular the desire is to develop techniques that allow data to be processed in small chunks *less* than the length of a single template. We have been addressing this need by developing a new software project named gstlal. See `http://www.lsc-group.uwm.edu/daswg/projects/gstlal.html`. Although the development of this technology is motivated by the need to reduce the memory requirements of the analysis pipeline, a side-effect of the effort has been the creation of a suite of data analysis software tools that allow the creation of pipelines in which the time delay between data going in and answer coming out is short.

The data analysis machinery used by gstlal continues to reside within the lalsuite of libraries (see below). gstlal wraps the lalsuite machinery in GStreamer "elements". GStreamer is a free C library providing the infrastructure required to assemble digital signal processing pipelines, and although it is primarily used to implement multimedia recording and playback on desktop computers, the GStreamer library is of very high quality and easily satisfies all of our own needs for such a library. By using it, not only do we leverage the design experience of GStreamer's developers, but the bug fixes and feature enhancements we have provided back to the project can now be found in Nokia cell phones where GStreamer provides the multimedia playback software, making the gstlal project one of the few places where GW data analysis can be said to have provided industrially-relevant spin-off technology.

A prototype application has been constructed using the tools provided by gstlal to search LIGO and Virgo data for GWs from compact object collisions. Because gstlal-based applications also have access to all the machinery of GStreamer, they are easily interfaced to network protocols, sound cards and multimedia file formats, and so in the future gstlal might be useful for outreach activities. For example, one could imagine writing software to demonstrate what GWs sound like, allow users to add simulated GWs to simulated detector data streams to hear how different detector configurations make it easier or harder to find different GWs, and so on.

3. **MatApps** With Advanced LIGO comes the prospect of the first direct detection of gravitational waves and the beginning of the field of gravitational wave astronomy. As a consequence, real-time data analysis will have increased importance as will rapid prototyping of code and visualization of results. While MATLAB is not the only choice users have to achieve these goals, MatApps intends to support this effort by building its infrastructure through coordination and communication with the MATLAB-using community. Coordination needs to be developed between MATLAB-rich repositories that exist outside of MatApps (e.g. LigoDV) to promote ease of code development and to reduce duplication of efforts. Communication is the foundation of user support in MatApps. While we will continue to address individual user questions and concerns, we want to develop the MatApps community to be a clearinghouse of best practices to achieve computational speed and ease of use. We also intend to communicate MATLAB knowledge through documentation. MATLAB is a powerful tool for use in the grid computing environment and we intend to promote its use in this way by keeping complete documentation in a centralized location and offering training to those who wish to gain experience. MathWorks, the author of MATLAB, often updates MATLAB several times a year and we intend to streamline our vetting of new versions and updating documentation about any issues or other consid-

erations so that users may take advantage of the latest features. These new initiatives, combined with our ongoing efforts, will help scaffold the increased demand for data analysis results that Advanced LIGO will introduce.

4. **LAL Suite (LAL, Glue, Pylal)** The LAL Suite of tools has grown beyond its initial scope to include I/O libraries, time and frequency series analysis tools, and domain-specific functionality that enables scientists to access and analyze gravitational-wave data. The development model adopted during the first generation LIGO science runs was deliberately agile. It allowed the developers, largely the same group as the user base, to be remarkably fleet-footed. The LAL Suite user base continues to expand. Indeed, the software has been used by scientists involved in the LISA mock data challenge demonstrating the utility of the software beyond LIGO. It is timely, as advanced instrumentation is installed in the LIGO facilities, to rework and retool LAL Suite to meet the decade-long scientific campaign that lies ahead by providing LSC scientists as well as the wider community of gravitational-wave astronmers with a toolbox for performing gravitational wave data analysis and simulation.

   As LAL Suite developed organically without an imposed final design, the code is not as clean, or general, as it could be. Therefore one of the first steps in improving the sustainability of LAL Suite for the future is to ensure that it has a clean, and easy to understand API (Application Programming Interface). Another effect of the organic development of LAL Suite is that there are numerous functions that are no longer used and that there are many functions that perform similar tasks. The code base will be simplified by unifying these similar functions, thereby decreasing the amount of code redundancy.

   While having a clean code base will greatly improve the maintainability and sustainability of the code, another critical aspect is adequate documentaiton of the software. The LAL code has now been restructured, but unfortunately the documentation has not; therefore the documentation sectioning does follow the current structure of LAL Suite. The documentation will be unified and restructured to improve the clarity and usefulness.

   LAL Suite was originally written using the C89 standard, as at the time the C99 standard had been approved but there where no shipping compilers that supported the standard to an acceptable level. This is no longer the case. C99 provides many improvements and features to the C language which will help in the maintenance of LAL Suite. The adoption of the C99 standard has already started in several minor, but key, areas: the first of which is the use of the C99 fixed-width integer datatypes. The C89 standard did not define the size of the integer datatypes, and therefore they are platform and compiler dependent. As LAL Suite is used on multiple platforms with different compilers, a way was needed to ensure that the integer datatypes were consistent across the different platforms. This led to custom code that determined the size of the various integer datatypes and made the appropriate typedefs. The C99 standard provides fixed width integer datatypes that are of the same size regardless of platform and compiler. Using these greatly simplifies the code base which leads to increased maintainability.

   There are many functions in LAL that are very similar and only differ in the datatype on which they operate. This leads to a lot of similar code that needs to be maintained consistently so that errors are not introduced by updating one function and not another. Ways in which this duplicated code can be reduced will be investigated.

   Another key feature that is provided by the C99 standard is support for complex numbers and complex arithmetic. Currently LAL defines is own complex datatype as a structure with two floating point fields. While this accomplishes the task of representing complex numbers, it complicates matters as helper functions need to be written to perform simple arithmetic. This greatly complicates the code base, and a transition to the built in complex type will alleviate a lot of problems. The C99 complex type is however not entirely a drop in replacement for the current LAL complex structure therefore, so an in depth study will be done in order to determine the optimal way to transition to the native C99 complex datatypes.

The ability to simulate the gravitational wave strain that would be produced from various types of astrophysical sources, e.g., coalescing compact binaries, continuous waves from distorted pulsars, random gravitational-wave noise from the early universe, is an important feature of the LAL libraries. However, the simulation software is currently integrated into individual searches, and is not exposed in a general, well documented, and easy-to-use API. This situation is unsatisfactory: one of the major functions that LAL Suite should perform is to provide the community with vetted software for gravitational waveform simulation. Therefore, one of the significant goals is to extract the routines that perform gravitational wave simulation from the individual search packages and combine them into a LALSimulation library. The routines will be re-implemented where necessary so that they have a common and useful interface. They will also be carefully documented community vetted so that their correctness is assured. This library will be the primary contribution of LAL Suite to the community of scientists outside of the LSC.

While it is important to have a clean and well-documented code base, it is also important that this code base is tested on a regular basis to ensure that the code works as expected, and that no code modifications lead to unexpected changes in behaviour. One way towards achieving this is to implement unit tests which aim to isolate each part of the code and shows that each of these "units" behaves as expected. Ideally every function inside the LAL libraries should have a test associated with it, therefore individual functions can be regularly tested to ensure correct behavior. Unit tests best work when the library functions perform one simple task that can be easily tested; many of the core library functions are now being written to perform such single tasks, and are therefore amenable to effective unit testing. The unit tests will be developed for these routines within the existing testing environment. Testing of individual functions is a step in the right direction but to ensure that the code works as expected complete workflows need to be tested in addition to the unit tests. Therefore an investigation into a build and test systems, such as Metronome, will be made to determine how complete LAL Suite workflows can be tested on a regular basis.

Increasingly, programs are being writing in scripting languages, such as python, as these provide a quick and easy method to accomplish tasks. We are finding that we are frequently needing to access many of the LAL Suite functions within such scripting languages. To date, required functions have been manually wrapped by hand as needed, an approach which clearly will not scale and a task that will need to be done for each scripting language that needs access to these functions. SWIG (Simplified Wrapper and Interface Generator), is a tool that can be used to automate the generation of bindings and one of the main advantages is that once things are setup bindings for any supported language can be automatically created. It will therefore be investigated how SWIG can be used to automate generation of language bindings.

5. **DMT** With the upgrade of the Ligo Control and Data System (CDS) for advanced Ligo the reference platform for the DMT will formally move from Solaris to Linux. In fact, because of the enormous offline computing pawer available from the Linux clusters, much of the recent DMT development has been tested on both Linux and Solaris insuring relatively simple porting to the new platform. Futher development and rigorous testing will still be necessary for the online components, especially those involved in distributing online data to all the processes. Although at present, the plan is to leave the frame broadcasting mechanism much the same as for initial ligo, the opportunity to receive data more quickly by way of an direct connection to the CDS data network should be evaluated.

   Additional DMT software development will also be needed to monitor and characterize the new and more complex aLigo interferometry.

The use of the same operating system online and offline, provides the opportunity to unify the packaging and distribution of the online and offline DMT software. Already, much work has been done to unify the packaging of all software from the DMT/GDS package used by CDS DMT-online and DMT offline.

6. **NDS** The version 2 Network Data Server (NDS2) allows Ligo-Virgo collaboration members to access current and archived Ligo data remotely. The network protocol uses Kerberos to allow nearly transparent authentication by the Ligo-Virgo scientists while preventing access by unauthorized persons. Offline and online NDS2 servers are currently running at Caltech and LHO, respectviely, with the offline server making all Ligo raw data aquired since the start of S5 available. The NDS2 client code has been interfaced to matlab, octave, python, C and C++. An example client application is the ligoDV viewer, described in the following section.

   This year the server and client have advanced considerably. The focus of recent development has been to:

   - Improve server reliability: preventing hangups when requested data are temporarily not available or the server is swamped with requests.
   - Improve error reporting and fail-over: Produce more meaningful error status returns and allow sucessful return of partial data if some channel is unavailable.
   - Improve portability: Client interfaces have been added for several interpretive languages (matlab, octave and python) and building anpackagind has been devloped and tested on many platforms (centos, debian, solaris, Mac).

   We expect that use of data from the NDS2 server will increase significantly in the future. NDS2 provides an exceptionally fast and convenient means to fetch data for real-time analysis. It may also provide a distribution mechanism for the proposed Open Data Initiative.

   Future improvements will include ports to additional platforms (e.g. Windows) and improved dynamic data finding by the server.

7. **LIGO DV** The ligoDV (LIGO Data Viewer) project (`https://www.lsc-group.phys.uwm.edu/daswg/projects/ligodv.html`) is aimed at increasing the accessibility of LIGO data and standard data processing algorithms to off- and on-site scientists within the LSC. The primary software tool in the project, ligoDV, is a Matlab-based graphical user interface that allows LSC members to connect to LIGO data servers, specifically the network data servers NDS and NDS2, and retrieve data. Furthermore it provides a platform for applying mathematical manipulations such as filters, Fourier transform, coherence, transfer functions and others to this data and finally exporting and/or plotting the results. The package is essentially operating system independent, since it consists of a collection of m-files that require only a graphics-capable Matlab installation. Owing to the portability of the NDS client, ligoDV is also location independent allowing users to query data and do studies while at meetings or anywhere with an internet connection. The ligoDV user-base has grown over the past few years and it is now used widely within the LSC. This in turn has aided detector characterization, commissioning and data analysis studies by lowering the hurdles required to access LIGO data.

   Over the past year ligoDV was upgraded with an interface that allows users to access data from all detectors and times served by NDS2 - a significantly larger set than the approximately 1 month of raw data available from the Livingston and Hanford NDS1 servers. Several smaller fixes were also implemented. On the management side a Gnats bug tracking system was set up at UWM, and a project website was set up on the DASWG homepage. There is however need for further work on ligoDV. Until recently the NDS2 client tools were under very active development. LigoDV does not yet take full advantage of these tools. In addition there have been a number of requests for enhancements to ligoDV. The following are future development goals for ligoDV:

- Improve the robustness of the ligoDV interface to the NDS2 client. This involves, e.g., developing a solution to handle the large (several hundred thousand) channel lists returned from the server for low sample rates.
- Follow NDS2 client developments by updating the ligoDV interface, installation instructions and examples.
- Investigate packaging and release options. A promising option is the Mac DMG format (which would also be an attractive option for NDS client installation).
- Implement user-suggested improvements. Some examples include, a revised channel list interface with the option to save/load channel lists, adding an omega-scan plotting feature, streamlining exported data structures, adding an oscilloscope option for online data, improving the filter interface to allow automated reading and usage of LIGO control system and calibration filters, and a variety of smaller bug fixes and suggestions.

The continued development of ligoDV will lead to a much improved tool. This will benefit the detector characterization and analysis work remaining for Enhanced LIGO. It will also be a central component of the Advanced LIGO detector characterization program that will be actively monitoring data during the establishment of the first Advanced LIGO subsystems.

7.5.2 Services

1. **Network Data Simulator** It will be important to continually test and assess the data analysis infrastructure for the advanced detector era as it is developed. A new project will be established to simulate the data streams from the netowrk of gravitational-wave detectors and deliver it to analysts by the same means they can expect during future observing runs. This will allow users to develop their analysis tools with knowledge of the key infrastructure and an operational testbed against which to test. A key component of this project will be to run regular and long-lived mock data challenges of increasing complexity which will allow the collaborations to efficiently benchmark analysis codes against each other. This project will be initiated in the coming year. Details should be available by mid 2011.

2. **Monitoring** The main advantage of having a customized solution to provide LDG metadata is that it can be integrated, redesigned and reconfigured at will, with almost any other tools designed to gather information about complex infrastructures. The LSCGIS prototype can be integrated with less flexible tools such as Nagios, ReSS (Re_source Selection Service, used by OSG), BDII, Relational DataBases, etc., which can help to improve the information service. Implemented as a RESTFul Web Service, LSCGIS is flexible and scalable enough that it can even use web tecnologies such as Google Maps API, PHP dynamic server scripting, be displayed in Internet enabled cellphones, etc.

   Under the umbrella of this project several studies are being carried out to choose among the best Grid monitoring technologies and to integrate them in a customized monitoring environment for LDG. The **OSG ReSS** is particularly interesting since it can also be integrated with Condor-G, which could be useful once LDG transitions from a Data Grid towards a Computational Grid, with the aid of other Grid technologies.

Also, studies about integration of Identity Management technologies (Kerberos, MyProxy, LIGO Auth Project, etc.) with Monitoring services are being considered. We are convinced that no a single solution will be enough to cover all the needs of a complex VO such as LSC/VIRGO and that the integration of several customized proposals will be the best approach to keep the computational infrastructure as flexible and scalable as possible. Besides, intellegint workflows will need of all the best available information gathering and customized solutions in order to retrieve useful and relevant LDG metadata and use it as input for the analysis pipelines.

3. **LIGO Data Replicator (LDR)** Initial and enhanced LIGO have clearly demonstrated the need for bulk replication of interferometer data sets to computing centers around the world during the advanced detector era. The growing development of "real time" or stream based analysis in addition to file based analysis does not diminish the need for robust replication of curated interferometer data sets to computing sites for efficient consumption by data analysts and their codes.

   The experience gained during the LIGO S6 science run with data replication provided as a service ("LDR as a service") to the Cardiff University and Birmingham University groups demonstrated that the software as a service (SAS) model for bulk data replication in the advanced interferometer era is not only viable but preferred. Individual computing sites are simply not staffed at a level that allows deep knowledge at each local site of all the necessary details for robust and continuous replication of data. Instead, the LDR as a service model demonstrated the efficiency of requiring at a computing site only a standards-based interface to local site storage (usually GridFTP) and then locating the rest of the necessary logic and tooling at a centrally managed collaboration facility, where local experts can monitor and tune the replication network over time.

   Of course moving all of the LDR functionality except for the interface to local site storage to one physical facility carries the risk that the entire replication network could fail if that one central facility should be cut off from the internet or go down for whatever reason. In the advanced detector era, hence, LDR as a service must evolve so that it leverages the necessary tools and infrastructure to provide a high availability service with failover capabilities supported by a small but expertly staffed geographically distributed set of data centers.

   The S6 and earlier science runs have also demonstrated that data discovery, as provided most recently by the server tool LDRDataFindServer and the client tool ligo_data_find, should be to some extent decoupled from bulk data replication. While the server tool needs to be able to consume state information from the replication service, it need not be tightly coupled to the replication service and should be capable of consuming information about the location of data files from many sources in a "pluggable" fashion, and then delivering that information via a variety of standardized interfaces and APIs to various data analysis tools and other services.

4. **GRaCEDb** The current Gracedb/Lumin system is excellent prototype service accepting a number of injected event streams (MBTAOnline, Omega, Ringdown etc) and a small number of event subscribers (ROTSE, LOFAR, QUEST, TAROT, SWIFT etc).

   As we move to the era of Data Challenges and eventually Advanced LIGO, the set of pipelines reporting events will change with time, with new pipelines, modifications to existing pipelines, and changes in personnel leading to "black-box" code. The LIGO event system should have a well-defined protocol for working with the pipelines that create event triggers, so that streams can be added and throttled in a well-defined way. On the other side, the set of subscribers to LIGO events will grow, and the upgraded system should streamline the process of adding subscribers and handling the individual requirements.

   The current event model is an "imperative", where LIGO software decides what each follow-up telescope should do. As the number of subscribers grows, this individual attention will become an undue burden on LIGO staff. Furthermore, we expect stiff competition for robotic follow-up facilities, as new, prolific event streams come on line (LOFAR, LSST, etc). It will become more difficult to get the best facilities if LIGO expects to take immediate, imperative control of the telescope. The new model will need to shift to *informational* rather than *imperative*, meaning the subscriber gets what LIGO has observed, and decides what to do. Thus the telescope scheduling code (much of Lumin) will be run and modified by the event subscriber rather than the event author.

Past events are also important, for studies of correlation with other event streams. Already astronomical events are being collected into repositories (PTF, Galex, CRTS, SWIFT, etc), and interoperability of these archives will be needed for these important scientific studies. The International Virtual Observatory Alliance has already defined a standard event representation (VOEvent), and many authors and subscribers are exchanging these standard packets. The LIGO event distribution system would be well-positioned for the future by adopting VOEvent.

Currently, Lumin delivers messages by a protocol customized for each subscriber, and as the number of subscribers grows, this will be more and more difficult. Therefore the event distribution from LIGO should adopt a standard transport protocol. Some requirements for this may include buffering and reliable delivery, strong security (preferebly linked to the LIGO Auth system), integrity signatures, indication of presence and readiness, broadcast, multiple implementations, and widespread adoption.

In addition to delivering LIGO observations to follow-up facilities (Gracedb and LoocUp), LIGO is acting as a follow-up facility for external triggers from other observatories. These two sides of the same coin could be unified by handling *all* relevant astronomical event streams in the same way, whether they are from LIGO or not.

5. **Identity Management** Moving into the advanced detector era, the size of the gravitational wave community (which includes the LIGO Laboratory, the LIGO Scientific Collaboration (LSC), Virgo and other collaborators) will continue to grow. Having centralized identity management for members of the community will be essential for a satisfactory user experience, for effective resource administration and for computer security. LIGO Laboratory and the LSC have initiated a joint effort called the Auth (Authentication and Authorizations) Project to develop a unified identity management infrastructure to serve the needs of the community. The project is currently funded on grid operations funding from the Physics at the Information Frontier award **insert PIF award number**. It focuses on four areas - core infrastructure to collect and store user information and create credentials, web services, grid services and general computing and shell access.

   The core infrastructure includes a custom MySQL database with PHP user interface to collect and store user information, two Kerberos realms (LIGO.ORG and LIGO-GUEST.ORG) to provide single sign-on (SSO) authentication to community members and collaborators, an LDAP directory service to provide authorization information and a second to provide user directory information, and an Internet2 product called Grouper which allows creates easy and flexible organization of LDAP entries into groups with an ability to delegate group management. Group membership forms the basis for all subsequent authorization decisions (eg members of the LSC Computing Committee group can view and edit the minutes of the LSC Computing Committee meeting).

   Web services leverage Internet2's Shibboleth software for authentication and authorizations. This integrates with Kerberos to provide an SSO web solution with fine-grained authorization capabilities to LIGO and LSC web services. We currently provide our own Shibboleth Identity Provider (IdP) service because too few community members participate in InCommon (or other Shibboleth federations) to make using external IdPs feasible, but our plan is to start leveraging external IdPs as they become more ubiquitous. The use of Shibboleth is starting to spread throughout LSC web resources and should be generic in the advanced detector era.

Grid services will leverage MyProxy, a product developed at NCSA, to provide transparent grid access to the community leveraging Kerberos authentication. MyProxy will issue short-lived X.509 certificates and proxy certificates underwritten by the LIGO.ORG Certificate Authorities (CAs). Distribution of Grid Map Files with the MyProxy generated user credentials will be handled by an in-house product. At present, the LIGO CAs are in operation, and detailed planning documents for the rest of the infrastructure have been written. In the advanced detector era, the grid services will be fully deployed and operational. Based on extensive discussions with the grid CA community, we expect the LIGO.ORG CAs to be accredited by TAGPMA by that time as well.

General computing and shell services will leverage kerberos via SSH or PAM for login. Account provisioning will be serviced by LDAP and NSSwitch, with LDAP transparent proxy overlays to augment the LSC managed information with site-specific information. This model, which is intended for major community compute centers but not individual community institution workstations and laptops, is currently deployed for LHO and LLO general computing.

As well as the four major areas, there are a number of other IdM related services that the Auth Project provides support and development for, including mailing lists, request tracking systems, version control systems, special needs environments such as LIGO control room CDS systems, and others. A more comprehensive description of the plans and activities is available at the project wiki (`https://www.lsc-group.phys.uwm.edu/twiki/bin/view/AuthProject/WebHome`).

References

[1] An introduction and overview of DMT monitors can be found at
http://blue.ligo-wa.caltech.edu/scirun/S5/scimoncamp06/talks/dmtintroduction.html
The DMT infrastructure is described in detail at
http://www.ligo.caltech.edu/

[2] EPICS home page: http://www.aps.anl.gov/epics/index.php.

[3] D. Sigg, R. Bork, and J. Zweizig, "Detector Characterization and Global Diagnostics System of the Laser Interferometer Graviational-wave Observatory (LIGO)", Proceedings of the Ninth Marcel Grossman Conference, July 2-8, 2000, Rome, Italy (2001).

[4] ROOT home page: http://root.cern.ch.

[5] LIGOtools home page: http://www.ldas-sw.ligo.caltech.edu/ligotools/.

[6] Home page for QScan program:
http://www.ligo.caltech.edu/ shourov/q/qscan/scimon.html

[7] Home page of Data Set Reduction Working Group:
http://darkwing.uoregon.edu/ ileonor/ligo/s5/rds/s5rds.html

[8] LIGO calibration web page:
http://blue.ligo-wa.caltech.edu/engrun/Calib_Home/.

[9] 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).

[10] Home page for time domain calibration:
http://www.lsc-group.phys.uwm.edu/∼siemens/ht.html

[11] X. Siemens *et al.*, "Making h(t) for LIGO", Class. Quantum Grav. 21 (2004) S1723.

[12] Home page for time domain calibration:
http://www.lsc-group.phys.uwm.edu/∼siemens/ht.html
Description of procedure: X. Siemens *et al.*, "Making h(t) for LIGO", *Class. Quant. Grav.* **21** S1723 (2004).

[13] Reports on timing stability:
S. Marka and D. Sigg, "Summary of LSC Timing Performance during the First Scientific Run (S1)", LIGO-T030070 (2003);
S. Marka and D. Sigg, "LSC Timing Performance during the Second Scientific Run (S2)", LIGO-T030080 (2003);
S. Marka and D. Sigg, "Report on the LSC Timing Performance during the Science Mode Segments of the Third Scientific Run (S3)", LIGO-T040012 (2004).

[14] S. Marka and D. Sigg, "Atomic Clock Proposal", LIGO-T030098 (2003).

[15] Y. Aso et al., "Accurate measurement of the time delay in the response of the LIGO gravitational wave detectors", to be submitted to CQG, LIGO-P080072-00 (2008).

[16] Advanced LIGO timing wiki page:
http://ilog.ligo-wa.caltech.edu:7285/advligo/Timing

[17] Home page for Glitch Working Group:
http://lancelot.mit.edu/ cadonati/S5/online/S5reports/.

[18] Glitch Working Group's S5 electronic notebook:
http://www.lsc-group.phys.uwm.edu/glitch/investigations/s5index.html

[19] Glitch Working Group's S6 electronic notebook:
https://www.lsc-group.phys.uwm.edu/twiki/bin/view/DetChar/GlitchStudies

[20] Home page for Interchannel Correlations Working Group:
http://virgo.physics.carleton.edu/Hans/coherence/peaks/.

[21] S Chatterji, L Blackburn, G Martin and E Katsavounidis, "Multiresolution techniques for the detection of gravitational-wave bursts", Class. Quantum Grav. 21 No 20 (21 October 2004) S1809-S1818 http://ligo.mit.edu/ lindy/kleineWelle/doc/kwdoc.pdf LIGO-T060221-00.

[22] L.S. Finn, "A Veto Selection Criteria for Gravitational Wave Burst Searches", LIGO-T030181 (2003).

[23] J. Camp, J. Cannizzo and K. Numata, "Application of Hilbert-Huang Transform to the Search for Gravitational Waves", *Phys Rev D (rapid comm.* **75**, 061101 (2006)
N. Huang et al., "The Empirical Mode Decomposition and the Hilbert Spectrum for Nonlinear and Non-stationary Timeseries Analysis", *Proc. R. Soc. A* **454** 903 (1998).

[24] R. Schofield, "Acoustic Mitigation Status for S3", LIGO-G030640 (2003).

[25] Home page for Upconversion Working Group:
http://ilog.ligo-wa.caltech.edu:7285/upconv/

[26] Home page for S5 Spectral Line Working Group:
https://gravity.psu.edu/ psurg/detchar/LineFinding/S5_Spectral_Lines.html.

[27] Home page for Data Quality Working Group:
http://www.ligo.caltech.edu/ jzweizig/S5_Data_Quality/index.html
Information page on Data Quality database:
http://www.lsc-group.phys.uwm.edu/daswg/projects/glue/segment_db/ Information page on Data Quality flags:
http://gallatin.physics.lsa.umich.edu/ keithr/S5DQ/

[28] Home page for Hardware Injections Working Group:
http://blue.ligo-wa.caltech.edu/scirun/S5/HardwareInjection/

[29] Used Percentage Veto for LIGO and Virgo Binary Inspiral Searches, T. Isogai for the LIGO Scientic Collaboration and the Virgo Collaboration, P1000050 in LIGO Document Center https://dcc.ligo.org/

[30] K S Thorne. Gravitational radiation. In S. W. Hawking and W. Israel, editors, *Three hundred years of gravitation*, chapter 9, pages 330–458. Cambridge University Press, Cambridge, 1987.

[31] Patrick R. Brady and Stephen Fairhurst. Interpreting the results of searches for gravitational waves from coalescing binaries. *Class. Quant. Grav.*, 25(10):1050002, 2008.

[32] Rahul Biswas, Patrick R. Brady, Jolien D. E. Creighton, and Stephen Fairhurst. The loudest event statistic: General formulation, properties and applications. 2007.

[33] B. Abbott et al. (The LIGO Scientific Collaboration). Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Phys. Rev. D*, 77:062002, 2008.

[34] B. Abbott et al. (The LIGO Scientific Collaboration) and K Hurley. Implications for the origin of grb 070201 from ligo observations. *Astroph. Journal*, 681:1419–1430, 2008.

[35] K. G. Arun, Bala R Iyer, B. S. Sathyaprakash, and Pranesh A Sundararajan. Parameter estimation of inspiralling compact binaries using 3.5 post-Newtonian gravitational wave phasing: The nonspinning case. *Physical Review D (Particles, Fields, Gravitation, and Cosmology)*, 71(8):084008, 2005.

[36] Clifford M. Will. The confrontation between general relativity and experiment. *Living Rev. Rel.*, 9:3, 2005.

[37] F. D. Ryan. Gravitational waves from the inspiral of a compact object into a massive, axisymmetric body with arbitrary multipole moments. *Phys. Rev.*, D52:5707–5718, 1995.

[38] D. A. Brown et al. Prospects for Detection of Gravitational Waves from Intermediate-Mass-Ratio Inspirals *Phys. Rev. Lett.*, 99:201102, 2007.

[39] Luc Blanchet. Gravitational radiation from post-Newtonian sources and inspiralling compact binaries. *Living Rev. Rel.*, 9:3, 2006.

[40] E. Berti et al. Inspiral, merger, and ringdown of unequal mass black hole binaries: A multipolar analysis *Phys. Rev. D*, 76:064034, 2007.

[41] E. Berti, V. Cardoso, and C. M. Will. Gravitational-wave spectroscopy of massive black holes with the space interferometer LISA. *Phys. Rev. D*, 73(6):064030–+, March 2006.

[42] B. Abbott et al. (The LIGO Scientific Collaboration). Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data. *Phys. Rev. D* 80, 062001, 2009.

[43] Frans Pretorius. Evolution of binary black hole spacetimes. *Phys. Rev. Lett.*, 95:121101, 2005.

[44] Frans Pretorius. Binary black hole coalescence. In M. Colpi, P. Casella, V. Gorini, U. Moschella, and A. Possenti, editors, *Physics of Relativistic Objects in Compact Binaries: from Birth to Coalescence*. Springer, Heidelberg, Germany, 2009.

[45] Sascha Husa. Numerical modeling of black holes as sources of gravitational waves in a nutshell. *Eur. Phys. J. ST*, 152:183–207, 2007.

[46] Mark Hannam. Status of black-hole-binary simulations for gravitational-wave detection. 2009.

[47] Alessandra Buonanno and Thibault Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, 1999.

[48] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev. D*, 62:064015, 2000.

[49] 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.

[50] Thibault Damour and Alessandro Nagar. Comparing Effective-One-Body gravitational waveforms to accurate numerical data. *Phys. Rev. D*, 77:024043, 2008.

[51] Parameswaran Ajith et al. Phenomenological template family for black-hole coalescence waveforms. *Class. Quant. Grav.*, 24:S689–S699, 2007.

[52] M. D. Duez et al. Evolving black hole-neutron star binaries in general relativity using pseudospectral and finite difference methods *Phys. Rev. D*, 78:104015, 2008.

[53] J. S. Read et al. Measuring the neutron star equation of state with gravitational wave observations *Phys. Rev. D*, 79:124033, 2009.

[54] M. Shibata and K. Taniguchi Merger of black hole and neutron star in general relativity: Tidal disruption, torus mass, and gravitational waves Phys. Rev. D, 77:084015, 2008.

[55] T. A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev. D*, 52:605, 1995.

[56] Alessandra Buonanno, Yanbei Chen, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Phys. Rev. D*, 67:104025, 2003. Erratum-ibid. 74 (2006) 029904(E).

[57] 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. *Phys. Rev. D*, 69:104017, 2004. Erratum-ibid. 74 (2006) 029905(E).

[58] B. Abbott et al. (The LIGO Scientific Collaboration). Search of S3 LIGO data for gravitational wave signals from spinning black hole and neutron star binary inspirals. *Phys. Rev. D*, 78:042002, 2008.

[59] D Fazi PhD thesis. Development of a Physical-Template Search for Gravitational Waves from Spinning Compact-Object Binaries with LIGO. LIGO Document P0900057-v2 (2009).

[60] T. Cokelaer and D. Pathak. Searching for gravitational-wave signals emitted by eccentric compact binaries using a non-eccentric template bank: implications for ground-based detectors. *Class. Quant. Grav.*, 26:045013, 2009.

[61] Karl Martel and Eric Poisson. Gravitational waves from eccentric compact binaries: Reduction in signal-to-noise ratio due to nonoptimal signal processing. *Phys. Rev.*, D60:124008, 1999.

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

[63] B. A. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton. Findchirp: an algorithm for detection of gravitational waves from inspiraling compact binaries. 2005.

[64] C. W. Helmstrom. *Statistical Theory of Signal Detection, 2nd edition*. Pergamon Press, London, 1968.

[65] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, 1999.

[66] 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. Quant. Grav.*, 23:5477–5504, 2006.

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

[68] 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. *Phys. Rev. D*, 72:084027, 2005.

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

[70] 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.

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

[72] B. Abbott et al. (The LIGO Scientific Collaboration). Tuning matched filter searches for compact binary coalescence. Technical Report LIGO-T070109-01, 2007.

[73] B. Abbott et al. (The LIGO Scientific Collaboration). Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. arXiv:0901.0302, *Phys. Rev. D* 79:122001, 2009.

[74] B. Abbott et al. (The LIGO Scientific Collaboration). Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run. Phys. Rev. D 80 (2009) 047101.

[75] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Search for Gravitational Waves from Compact Binary Coalescence in LIGO and Virgo Data from S5 and VSR1. LIGO Document P0900305. Submitted to PRD, 2010.

[76] Ehud Nakar. Short-hard gamma-ray bursts. *Phys. Rep.*, 442:166–236, 2007.

[77] V. Kalogera, C. Kim, D. R. Lorimer, M. Burgay, N. D'Amico, A. Possenti, R. N. Manchester, A. G. Lyne, B. C. Joshi, M. A. McLaughlin, M. Kramer, J. M. Sarkissian, and F. Camilo. The cosmic coalescence rates for double neutron star binaries. 601:L179–L182, 2004. Erratum-ibid. 614 (2004) L137.

[78] R. O'Shaughnessy, C. Kim, V. Kalogera, and K. Belczynski. Constraining Population Synthesis Models via Empirical Binary Compact Object Merger and Supernova Rates. *apj*, 672:479–488, January 2008.

[79] R. O'Shaughnessy, C. Kim, T. Fragos, V. Kalogera, and K. Belczynski. Constraining Population Synthesis Models via the Binary Neutron Star Population. *apj*, 633:1076–1084, November 2005.

[80] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. LIGO Document P0900125, arXiv 1003.2480, to appear in CQG (2010).

[81] E. S. Phinney. The rate of neutron star binary mergers in the universe: Minimal predictions for gravity wave detector. *Astrophysical Journal*, 380:L17, 1991.

[82] Ravi Kumar Kopparapu, Chad Hanna, Vicky Kalogera, Richard O'Shaughnessy, Gabriela Gonzalez, Patrick R. Brady, and Stephen Fairhurst. Host Galaxies Catalog Used in LIGO Searches for Compact Binary Coalescence Events. *Astrophys. J.*, 675(2):1459–1467, 2008.

[83] B. Abbott et al. (The LIGO Scientific Collaboration). Analysis of LIGO data for gravitational waves from binary neutron stars. *Phys. Rev. D*, 69:122001, 2004.

[84] B. Abbott et al. (The LIGO Scientific Collaboration). Search for gravitational waves from galactic and extra-galactic binary neutron stars. *Phys. Rev. D*, 72:082001, 2005.

[85] B. Abbott et al. (The LIGO Scientific Collaboration). Joint LIGO and tama300 search for gravitational waves from inspiralling neutron star binaries. *Phys. Rev.*, D73:102002, 2006.

[86] B. Abbott et al. (The LIGO Scientific Collaboration). Search for gravitational waves from primordial black hole binary coalescences in the galactic halo. *Phys. Rev. D*, 72:082002, 2005.

[87] Alessandra Buonanno, Yanbei Chen, and Michele Vallisneri. Detection template families for gravitational waves from the final stages of binary-black-hole inspirals: Nonspinning case. *Phys. Rev. D*, 67:024016, 2003. Erratum-ibid. 74 (2006) 029903(E).

[88] B. Abbott et al. (The LIGO Scientific Collaboration). Search for gravitational waves from binary black hole inspirals in LIGO data. *Phys. Rev. D*, 73:062001, 2006.

[89] Serge Droz, Daniel J. Knapp, Eric Poisson, and Benjamin J. Owen. Gravitational waves from inspiraling compact binaries: Validity of the stationary-phase approximation to the fourier transform. *Phys. Rev. D*, 59:124016, 1999.

[90] P. Ajith et al. A template bank for gravitational waveforms from coalescing binary black holes: non-spinning binaries. *Phys. Rev. D*, 77:104017, 2008.

[91] P. Ajith et al. "Complete" gravitational waveforms for black-hole binaries with non-precessing spins. LIGO-P0900085, arXiv:0909.2867 (2009).

[92] L. Santamaria et al. Matching post-Newtonian and numerical relativity waveforms: systematic errors and a new phenomenological model for non-precessing black hole binaries. LIGO Document P1000048, submitted to PRD (2010).

[93] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO's fifth and Virgo's first science run. arXiv:1001.0165; *Astrophys. J.* 715, 1453, 2010.

[94] Citation for coherent methods in CBC.

[95] http://www.gstreamer.net/

[96] The LIGO Algorithm Library. https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html

[97] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Search for gravitational waves from black hole binary inspiral, merger and ringdown. LIGO-P1000025, in preparation for submission to PRD (2010).

[98] Lumin processor reference. P1000062-v3, https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=11492 . LIGO-G0900951-v4, https://dcc.ligo.org/cgi-bin/private/DocDB/ShowDocument?docid=6495 .

[99] Swift processor reference. http://swiftoo1.mit.edu/∼online/public/

[100] GraceDB: https://www.lsc-group.phys.uwm.edu/daswg/projects/gracedb.html .

[101] B. Aylott et al. Testing gravitational-wave searches with numerical relativity waveforms: Results from the first Numerical INJection Analysis (NINJA) project. *Class.Quant.Grav.* 26:165008, 2009. http://arxiv.org/abs/0901.4399.

[102] https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/S6a_Quarterly_Report and https://www.lsc-group.phys.uwm.edu/ligovirgo/cbcnote/S6b_Quarterly_Report (password required).

[103] https://ldas-jobs.ligo.caltech.edu/∼xpipeline/S6/grb/online/triggers/S6grbs_list.html (password required).

[104] S. D. Barthelmy, The GCN Web site, http://gcn.gsfc.nasa.gov/gcn/

[105] Minutes of the JRPC meeting on Monday 17-May-10, LIGO Document L1000274-v1.

[106] H. Lueck et al. The upgrade of GEO600. In the proceedings of Amaldi8. arXiv:1004.0339. LIGO Document P0900122-v1.

[107] P. Brady and I. Mandel. Estimating the impact of first advanced detector science runs. T1000243-v1 (2010).

[108] Einstein Telescope Design Study: Vision Document, P. Amaro-Seoane, et al, https://workarea.et-gw.eu/et/WG4-Astrophysics/visdoc

[109] B. Abbott et al. First upper limits from ligo on gravitational-wave bursts. *Phys. Rev. D*, 69:102001, 2004.

[110] B. Abbott 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.

[111] B. Abbott et al. Search for gravitational-wave bursts in the first year of the fifth ligo science run. *Phys. Rev. D*, 80(10):102001, 2009.

[112] B. Abbott et al. Search for high frequency gravitational-wave bursts in ligo data from the fifth science run. *Phys. Rev. D*, 80(10):102002, 2009.

[113] B. Abbott et al. All-sky search for gravitational-wave bursts in the first joint ligo-geo-virgo run. *Phys. Rev. D*, 81(10):102001, 2010.

[114] C. D. Ott, A. Burrows, L. Dessart, and E. Livne. . *Phys. Rev. Lett.*, 96:201102, 2006.

[115] B. Abbott et al. First ligo search for gravitational wave bursts from cosmic (super)strings. *Phys. Rev. D*, 80(6):062002, 2009.

[116] S. Klimenko and Guenakh Mitselmakher. A wavelet method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 21:S1819–S1830, 2004.

[117] 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.

[118] J. Camp et al. Application of hilbert-huang transform to the search for gravitational waves. *Phys. Rev. D*, 75:061101(R), 2007.

[119] Shantanu Desai, Sam Finn, John McNabb, Amber Stuver, Tiffany Summerscales and Keith Thorne. Block Normal. *Class. Quantum Grav.*, 21:S1705–S1710, 2004.

[120] L. Cadonati and Sz. Marka. CorrPower: a cross-correlation-based algorithm for triggered and untriggered gravitational-wave burst searches. *Class. Quant. Grav.*, 22:S1159–S1167, 2005.

[121] S. Klimenko et al. Performance of the WaveBurst algorithm on LIGO data. *Class. Quant. Grav.*, 21:S1685–S1694, 2004.

[122] L. Blackburn et al. Glitch investigations with kleineWelle. *LIGO-G050158-00-Z*, 2005.

[123] L. Cadonati. Coherent waveform consistency test for LIGO burst candidates. *Class. Quant. Grav.*, 21:S1695–S1704, 2004.

[124] B. Abbott et al. Search for gravitational-wave bursts from soft gamma repeaters. Phys. Rev. Lett., 101(21):211102, 2008.

[125] 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.

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

[127] R. A. Mercer and S. Klimenko. Visualising gravitational-wave event candidates using the Coherent Event Display. *Class. Quantum Grav.*, 23:184025, 2008.

[128] 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.

[129] Antony C. Searle, Patrick J. Sutton, and Massimo Tinto. Bayesian detection of unmodeled bursts of gravitational waves. *Class. Quantum Grav. (to appear)*, 2009.

[130] T. Z. Summerscales, Adam Burrows, Lee Samuel Finn, and Christian D. Ott. Maximum entropy for gravitational wave data analysis: Inferring the physical parameters of core-collapse supernovae. ApJ, 678(2):1142–1157, MAY 10 2008.

[131] Omega Pipeline documentation and wiki, https://geco.phys.columbia.edu/omega/.

[132] C. Culter and K. Thorne. An overview of gravitational-wave sources. In N. T. Bishop and D. M. Sunil, editors, *General Relativity and Gravitation*, page 72, 2002.

[133] J. G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower. Comparisons of binary black hole merger waveforms. *Classical and Quantum Gravity*, 24:25–+, June 2007.

[134] C. D Ott. TOPICAL REVIEW: The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, 26(6):063001–+, March 2009.

[135] M. H. P. M. van Putten. Gravitational Wave Frequencies and Energies in Hypernovae. ApJ, 583:374–378, January 2003.

[136] The MACRO Collaboration. Neutrino astronomy with the macro detector. ApJ, 546:1038, 2001.

[137] The IceCube Collaboration. Search for point sources of high energy neutrinos with final data from amanda-ii. Phys. Rev. D, 79:062001, 2009.

[138] The Super-Kamiokande Collaboration. High energy astronomy using upward-going muons in super-kamiokande-i. ApJ, 652:198, 2006.

[139] The Pierre Auger Collaboration. Correlation of the highest energy cosmic rays with nearby extra-galactic objects. *Science*, 318:938, 2007.

[140] Abadie, J. and others. Predictions for the Rates of Compact Binary Coalescences Observable by Ground-based Gravitational-wave Detectors. *ArXiv e-prints*, March 2010. arXiv1003.2480.

[141] S. Ando, J. F. Beacom, and H. Yüksel. Detection of Neutrinos from Supernovae in Nearby Galaxies. *Physical Review Letters*, 95(17):171101–+, October 2005.

[142] P. Ajith et al. A template bank for gravitational waveforms from coalescing binary black holes: I. non-spinning binaries. *Phys. Rev.*, D77:104017, 2008.

[143] 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.

[144] Yi Pan et al. A data-analysis driven comparison of analytic and numerical coalescing binary waveforms: Nonspinning case. *Phys. Rev. D*, 77:024014, 2008.

[145] 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.

[146] 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.

[147] 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. ApJ, 648:411–429, September 2006.

[148] John Veitch and Alberto Vecchio. Assigning confidence to inspiral gravitational wave candidates with Bayesian model selection. *Class. Quant. Grav.*, 25:184010, 2008.

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

[150] N. Huang et al. The empirical mode decomposition and the hilbert spectrum for non-linear and non-stationary timeseries analysis. *Proc R Soc Lond A*, 454:903, 1998.

[151] C. Pankow et al. A burst search for gravitational waves from binary black holes. 2009.

[152] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, B. Zink, I. Hawke, and E. Schnetter. Rotating collapse of stellar iron cores in general relativity. *Classical and Quantum Gravity*, 24:139–+, June 2007.

[153] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, I. Hawke, B. Zink, and E. Schnetter. 3D Collapse of Rotating Stellar Iron Cores in General Relativity Including Deleptonization and a Nuclear Equation of State. *Physical Review Letters*, 98(26):261101–+, June 2007.

[154] C. D. Ott, A. Burrows, L. Dessart, and E. Livne. A New Mechanism for Gravitational-Wave Emission in Core-Collapse Supernovae. *Physical Review Letters*, 96(20):201102–+, May 2006.

[155] K. Kotake, W. Iwakami, N. Ohnishi, and S. Yamada. Stochastic Nature of Gravitational Waves from Supernova Explosions with Standing Accretion Shock Instability. ApJ Lett., 697:L133–L136, June 2009.

[156] A. Marek, H.-T. Janka, and E. Müller. Equation-of-state dependent features in shock-oscillation modulated neutrino and gravitational-wave signals from supernovae. A&A, 496:475–494, March 2009.

[157] The LIGO Scientific Collaboration and The Virgo Collaboration. 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.

[158] LIGO Scientific Collaboration and K. Hurley. Implications for the Origin of GRB 070201 from LIGO Observations. ApJ, 681:1419–1430, July 2008.

[159] Abbott, B. and others. Search for Gravitational Wave Bursts from Soft Gamma Repeaters. Phys. Rev. Lett., 101:211102, 2008.

[160] Abbott, B. and others. Search for gravitational wave radiation associated with the pulsating tail of the SGR 1806-20 hyperflare of 27 December 2004 using LIGO. *Phys. Rev. D*, 76(6):062003, September 2007.

[161] C. Kouveliotou, C. A. Meegan, G. J. Fishman, N. P. Bhat, M. S. Briggs, T. M. Koshut, W. S. Paciesas, and G. N. Pendleton. Identification of two classes of gamma-ray bursts. *Astrophysical Journal Letters*, 413:L101–L104, August 1993.

[162] N. Gehrels et al. A new γ-ray burst classification scheme from GRB060614. *Nature*, 444:1044–1046, December 2006.

[163] S. Campana et al. The shock break-out of grb 060218/sn 2006aj. *Nature*, 442:1008–1010, 2006.

[164] D. Malesani et al. SN 2003lw and GRB 031203: A Bright Supernova for a Faint Gamma-Ray Burst. *Astrophysical Journal Letters*, 609:L5–L8, July 2004.

[165] Jens Hjorth et al. A very energetic supernova associated with the gamma-ray burst of 29 march 2003. *Nature*, 423:847–850, 2003.

[166] T. J. Galama et al. Discovery of the peculiar supernova 1998bw in the error box of grb980425. *Nature*, 395:670, 1998.

[167] J. S. Bloom et al. A Putative Early-Type Host Galaxy for GRB 060502B: Implications for the Progenitors of Short-Duration Hard-Spectrum Bursts. *Astrophysical Journal*, 654:878–884, January 2007.

[168] E. Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442:166–236, April 2007.

[169] E. Nakar, A. Gal-Yam, T. Piran, and D. B. Fox. The Distances of Short-Hard Gamma-Ray Bursts and the Soft Gamma-Ray Repeater Connection. *Astrophysical Journal*, 640:849–853, April 2006.

[170] R. Chapman, R. S. Priddey, and N. R. Tanvir. Two populations are better than one: Short gamma-ray bursts from SGR giant flares and NS-NS mergers. *[arXiv:0709.4640]*, 709, September 2007.

[171] Anthony L. Piro and Eric Pfahl. Fragmentation of collapsar disks and the production of gravitational waves. *The Astrophysical Journal*, 658(2):1173–1176, 2007.

[172] M. H. van Putten, A. Levinson, H. K. Lee, T. Regimbau, M. Punturo, and G. M. Harry. Gravitational radiation from gamma-ray burst-supernovae as observational opportunities for LIGO and VIRGO. Phys. Rev. D, 69(4):044007–+, February 2004.

[173] R. Chapman, N. R. Tanvir, R. S. Priddey, and A. J. Levan. How common are long gamma-ray bursts in the local Universe? *MNRAS*, 382:L21–L25, November 2007.

[174] GCN. GCN: The Gamma-ray bursts Coordination Network. *http://gcn.gsfc.nasa.gov/*, 2007.

[175] LSC and Virgo Collaboration. https://www.lsc-group.phys.uwm.edu/twiki/bin/view/Bursts/S5VSR1GRB051103Ove rview.

[176] P. M. Woods, C. Kouveliotou, M. H. Finger, E. Göğüş, C. A. Wilson, S. K. Patel, K. Hurley, and J. H. Swank. The Prelude to and Aftermath of the Giant Flare of 2004 December 27: Persistent and Pulsed X-Ray Properties of SGR 1806-20 from 1993 to 2005. ApJ, 654:470–486, January 2007.

[177] C. Thompson and R. C. Duncan. The soft gamma repeaters as very strongly magnetized neutron stars - I. Radiative mechanism for outbursts. *Mon. Not. R. Astron. Soc.*, 275:255–300, July 1995.

[178] S. J. Schwartz, S. Zane, R. J. Wilson, F. P. Pijpers, D. R. Moore, D. O. Kataria, T. S. Horbury, A. N. Fazakerley, and P. J. Cargill. The Gamma-Ray Giant Flare from SGR 1806-20: Evidence of Crustal Cracking via Initial Timescales. ApJ Lett., 627:L129–L132, July 2005.

[179] N. Andersson and K. D. Kokkotas. Towards gravitational wave asteroseismology. MNRAS, 299:1059–1068, October 1998.

[180] J. A. de Freitas Pacheco. Do soft gamma repeaters emit gravitational waves? *Astronomy and Astrophysics*, 336:397–401, August 1998.

[181] K. Ioka. Magnetic deformation of magnetars for the giant flares of the soft gamma-ray repeaters. *MNRAS*, 327:639–662, October 2001.

[182] R. X. Xu. Solid quark stars? *The Astrophysical Journal Letters*, 596(1):L59–L62, 2003.

[183] B. J. Owen. Maximum Elastic Deformations of Compact Stars with Exotic Equations of State. *Physical Review Letters*, 95(21):211101–+, November 2005.

[184] J. E. Horvath. Energetics of the Superflare from SGR1806–20 and a Possible Associated Gravitational Wave Burst. *Modern Physics Lett. A*, 20:2799–2804, 2005.

[185] T. E. Strohmayer and A. L. Watts. The 2004 Hyperflare from SGR 1806-20: Further Evidence for Global Torsional Vibrations. ApJ, 653:593–601, December 2006.

[186] A. L. Watts and T. E. Strohmayer. Detection with RHESSI of High-Frequency X-Ray Oscillations in the Tailof the 2004 Hyperflare from SGR 1806-20. ApJ Lett., 637:L117–L120, February 2006.

[187] T. E. Strohmayer and A. L. Watts. Discovery of Fast X-Ray Oscillations during the 1998 Giant Flare from SGR 1900+14. ApJ Lett., 632:L111–L114, October 2005.

[188] R. Khan. Searching for gravitational wave fingerprints of SGR QPOs, poster at 2008 APS April .

[189] E. Waxman and J. Bahcall. High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physical Review Letters*, 78:2292–2295, March 1997.

[190] J. P. Rachen and P. Mészáros. Photohadronic neutrinos from transients in astrophysical sources. *Phys. Rev. D*, 58(12):123005–+, December 1998.

[191] J. Alvarez-Muñiz, F. Halzen, and D. W. Hooper. High energy neutrinos from gamma ray bursts: Event rates in neutrino telescopes. *Phys. Rev. D*, 62(9):093015–+, November 2000.

[192] E. Waxman and J. N. Bahcall. Neutrino Afterglow from Gamma-Ray Bursts: 10^{18} eV. *Astrophys. Journal*, 541:707–711, October 2000.

[193] P. Mészáros and E. Waxman. TeV Neutrinos from Successful and Choked Gamma-Ray Bursts. *Physical Review Letters*, 87(17):171102–+, October 2001.

[194] D. Guetta and J. Granot. Neutrinos of Energy $\tilde{1}0^{16}$ eV from Gamma-Ray Bursts in Pulsar Wind Bubbles. *Physical Review Letters*, 90(20):201103–+, May 2003.

[195] S. Razzaque, P. Mészáros, and E. Waxman. High Energy Neutrinos from Gamma-Ray Bursts with Precursor Supernovae. *Physical Review Letters*, 90(24):241103–+, June 2003.

[196] 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.

[197] C. D. Dermer and A. Atoyan. High-Energy Neutrinos from Gamma Ray Bursts. *Physical Review Letters*, 91(7):071102–+, August 2003.

[198] K. Murase and S. Nagataki. High Energy Neutrino Flashes from Far-Ultraviolet and X-Ray Flares in Gamma-Ray Bursts. *Physical Review Letters*, 97(5):051101–+, August 2006.

[199] K. Murase, K. Ioka, S. Nagataki, and T. Nakamura. High-Energy Neutrinos and Cosmic Rays from Low-Luminosity Gamma-Ray Bursts? *Astrophys. Journal*, 651:L5–L8, November 2006.

[200] M. Vietri. Ultrahigh Energy Neutrinos from Gamma Ray Bursts. *Physical Review Letters*, 80:3690–3693, April 1998.

[201] E. Waxman. High Energy Cosmic-Rays and Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physica Scripta Volume T*, 85:117–+, 2000.

[202] W. H. Lee and E. Ramirez-Ruiz. The progenitors of short gamma-ray bursts. *New Journal of Physics*, 9:17–+, January 2007.

[203] S. R. Kulkarni et al. 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.

[204] A. M. Soderberg et al. The sub-energetic γ-ray burst GRB 031203 as a cosmic analogue to the nearby GRB 980425. *Nature*, 430:648–650, August 2004.

[205] 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.

[206] E. Pian et al. An optical supernova associated with the X-ray flash XRF 060218. *Nature*, 442:1011–1013, August 2006.

[207] A. M. Soderberg et al. Relativistic ejecta from X-ray flash XRF 060218 and the rate of cosmic explosions. *Nature*, 442:1014–1017, August 2006.

[208] E. Liang, B. Zhang, F. Virgili, and Z. G. Dai. Low-Luminosity Gamma-Ray Bursts as a Unique Population: Luminosity Function, Local Rate, and Beaming Factor. *Astrophys. Journal*, 662:1111–1118, June 2007.

[209] N. Gupta and B. Zhang. Neutrino spectra from low and high luminosity populations of gamma ray bursts. *Astroparticle Physics*, 27:386–391, June 2007.

[210] X.-Y. Wang, S. Razzaque, P. Mészáros, and Z.-G. Dai. High-energy cosmic rays and neutrinos from semirelativistic hypernovae. *Phys. Rev. D*, 76(8):083009–+, October 2007.

[211] S. Ando and J. F. Beacom. Revealing the Supernova Gamma-Ray Burst Connection with TeV Neutrinos. *Physical Review Letters*, 95(6):061103–+, August 2005.

[212] A. M. Soderberg, D. A. Frail, and M. H. Wieringa. Constraints on Off-Axis Gamma-Ray Burst Jets in Type Ibc Supernovae from Late-Time Radio Observations. *Astrophys. Journal*, 607:L13–L16, May 2004.

[213] J. Granot and E. Ramirez-Ruiz. The Case for a Misaligned Relativistic Jet from SN 2001em. *Astrophys. Journal*, 609:L9–L12, July 2004.

[214] P. A. Mazzali et al. An Asymmetric Energetic Type Ic Supernova Viewed Off-Axis, and a Link to Gamma Ray Bursts. *Science*, 308:1284–1287, May 2005.

[215] Shunsaku Horiuchi and Shin'ichiro Ando. High-energy neutrinos from reverse shocks in choked and successful relativistic jets. *Phys. Rev. D*, 77(6):063007, Mar 2008.

[216] K. Ioka, S. Razzaque, S. Kobayashi, and P. Mészáros. TeV-PeV Neutrinos from Giant Flares of Magnetars and the Case of SGR 1806-20. *Astrophys. Journal*, 633:1013–1017, November 2005.

[217] A. Achterberg et al. Limits on the High-Energy Gamma and Neutrino Fluxes from the SGR 1806-20 Giant Flare of 27 December 2004 with the AMANDA-II Detector. *Physical Review Letters*, 97(22):221101–+, December 2006.

[218] IceCube. http://icecube.wisc.edu.

[219] ANTARES. http://antares.in2p3.fr/.

[220] J. Braun, J. Dumm, F. de Palma, C. Finley, A. Karle, and T. Montaruli. Methods for point source analysis in high energy neutrino telescopes. *Astroparticle Physics*, 29:299–305, May 2008.

[221] 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.

[222] Aso, Y. and others. Analysis Method to Search for Coincidence Events between the LIGO-Virgo Gravitational-wave Detector Network and the IceCube Neutrino Detector, 2008 APS April Meeting, St. Louis, Missouri, 2008.

[223] K. Scholberg. Supernova neutrino detection. *ArXiv Astrophysics e-prints*, January 2007.

[224] L. Cadonati et al. 2010.

[225] James Clark, Ik Siong Heng, Matthew Pitkin, and Graham Woan. Evidence-based search method for gravitational waves from neutron star ring-downs. *Physical Review D (Particles, Fields, Gravitation, and Cosmology)*, 76(4):043003, 2007.

[226] L. S. Finn, S. D. Mohanty, and J. D. Romano. Detecting an association between gamma ray and gravitational wave bursts. *Phys. Rev. D.*, 60(12):121101–+, December 1999.

[227] Soumya D Mohanty. Improving the sensitivity of searches for an association between gamma-ray bursts and gravitational waves. *Classical and Quantum Gravity*, 22(18):S1349–S1358, 2005.

[228] B. Abbott et al. Search for gravitational waves associated with 39 gamma-ray bursts using data from the second, third, and fourth ligo runs. Phys. Rev. D, 77(6):062004, 2008.

[229] S. D. Mohanty. Population study of gamma ray bursts: detection sensitivity and upper limits. *Classical and Quantum Gravity*, 23:723–+, October 2006.

[230] M. van der Klis, J. H. Swank, W. Zhang, K. Jahoda, E. H. Morgan, W. H. G. Lewin, B. Vaughan, and J. van Paradijs. Discovery of Submillisecond Quasi-periodic Oscillations in the X-Ray Flux of Scorpius X-1. ApJ Lett., 469:L1+, September 1996.

[231] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt. Nuclear-powered millisecond pulsars and the maximum spin frequency of neutron stars. Nature, 424:42–44, July 2003.

[232] R. V. Wagoner. Gravitational radiation from accreting neutron stars. ApJ, 278:345–348, March 1984.

[233] N. Andersson. A New Class of Unstable Modes of Rotating Relativistic Stars. ApJ, 502:708–+, August 1998.

[234] L. Bildsten. Gravitational Radiation and Rotation of Accreting Neutron Stars. ApJ Lett., 501:L89+, July 1998.

[235] J.G. Jernigan. *AIPC*, 586:805, 2001.

[236] A. Drago, G. Pagliara, and Z. Berezhiani. Gravitational wave bursts induced by r-mode spin-down of hybrid stars. A&A, 445:1053–1060, January 2006.

[237] E. Coccia, F. Dubath, and M. Maggiore. Possible sources of gravitational wave bursts detectable today. Phys. Rev. D, 70(8):084010–+, October 2004.

[238] V. M. Lipunov and I. E. Panchenko. Pulsars revived by gravitational waves. A&A, 312:937–940, August 1996.

[239] B. M. S. Hansen and M. Lyutikov. Radio and X-ray signatures of merging neutron stars. MNRAS, 322:695–701, April 2001.

[240] J. Moortgat and J. Kuijpers. Indirect Visibility of Gravitational Waves in Magnetohydrodynamic Plasmas. *ArXiv General Relativity and Quantum Cosmology e-prints*, March 2005.

[241] M. D. Duez, Y. T. Liu, S. L. Shapiro, and B. C. Stephens. Excitation of magnetohydrodynamic modes with gravitational waves: A testbed for numerical codes. Phys. Rev. D, 72(2):024029–+, July 2005.

[242] V. V. Usov and J. I. Katz. Low frequency radio pulses from gamma-ray bursts? A&A, 364:655–659, December 2000.

[243] C. W. Stubbs. *Class. Quantum Grav.*, 25:184033, 2008.

[244] D. B. Fox et al. *Nature*, 437:845–50, 2005.

[245] N. Gehrels, J. K. Cannizzo, and J. P. Norris. *New Journal of Physics*, 9, 2007.

[246] L. Li and B. Pacyznski. *Astrophys. J.*, 507:L59–62, 1998.

[247] S. Kulkarni. Modeling supernova-like explosions associated with gamma-ray bursts with short durations. 2005.

[248] C. L. Fryer, D. E. Holz, and S. A. Hughes. *Astrophys. J.*, 565:430, 2002.

[249] J. Kanner, T. L. Huard, S. Marka, D. C. Murphy, J. Piscionere, M. Reed, and P. Shawhan. *Class. Quantum Grav.*, 25:184034, 2008.

[250] J. Sylvestre. *Astrophys. J.*, 591:1152–6, 2003.

[251] J. S. Bloom et al. Astro2010 decadal survey whitepaper: Coordinated science in the gravitational and electromagnetic skies. 2009.

[252] S. Desai, E. O. Kahya, and R. P. Woodard. *Phys. Rev. D*, 77:124041, 2008.

[253] N. Dalal, D. E. Holz, S. Hughes, and B. Jain. *Phys. Rev. D*, 74:063006, 2006.

[254] Douglas M. Eardley, David L. Lee, Alan P. Lightman, Robert V. Wagoner, and Clifford M. Will. Gravitational-wave observations as a tool for testing relativistic gravity. *Phys. Rev. Lett.*, 30(18):884–886, Apr 1973.

[255] Douglas M. Eardley, David L. Lee, and Alan P. Lightman. Gravitational-wave observations as a tool for testing relativistic gravity. *Phys. Rev. D*, 8(10):3308–3321, Nov 1973.

[256] Teviet Creighton, Fredrick A. Jenet, and Richard H. Price. Pulsar timing and spacetime curvature. 2008.

[257] Atsushi Nishizawa, Atsushi Taruya, Kazuhiro Hayama, Seiji Kawamura, and Masa-aki Sakagami. Probing non-tensorial polarizations of stochastic gravitational-wave backgrounds with ground-based laser interferometers. 2009.

[258] Stephon Alexander, Lee Samuel Finn, and Nicolas Yunes. A gravitational-wave probe of effective quantum gravity. *Phys. Rev.*, D78:066005, 2008.

[259] Clifford M. Will. The confrontation between general relativity and experiment: A 1998 update. 1998.

[260] McGill SGR/AXP Online Catalog.

[261] J. Abadie et al., 2010. in preparation.

[262] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller. Theory of core-collapse supernovae. Phys. Rep., 442:38, 2007.

[263] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Class. Quant. Grav.*, 26(20):204015, October 2009.

[264] C. Röver, M.-A. Bizouard, N. Christensen, H. Dimmelmeier, I. S. Heng, and R. Meyer. Bayesian reconstruction of gravitational wave burst signals from simulations of rotating stellar core collapse and bounce. Phys. Rev. D, 80(10):102004, November 2009.

[265] E. B. Abdikamalov, C. D. Ott, L. Rezzolla, L. Dessart, H. Dimmelmeier, A. Marek, and H.-T. Janka. Axisymmetric general relativistic simulations of the accretion-induced collapse of white dwarfs. Phys. Rev. D, 81(4):044012, February 2010.

[266] S. E. Woosley and J. S. Bloom. The Supernova Gamma-Ray Burst Connection. *Ann. Rev. Astron. Astrophys.*, 44:507, September 2006.

[267] A. Corsi and P. Mészáros. Gamma-ray Burst Afterglow Plateaus and Gravitational Waves: Multi-messenger Signature of a Millisecond Magnetar? ApJ, 702:1171, September 2009.

[268] A. M. Soderberg, S. Chakraborti, G. Pignata, R. A. Chevalier, P. Chandra, A. Ray, M. H. Wieringa, A. Copete, V. Chaplin, V. Connaughton, S. D. Barthelmy, M. F. Bietenholz, N. Chugai, M. D. Stritzinger, M. Hamuy, C. Fransson, O. Fox, E. M. Levesque, J. E. Grindlay, P. Challis, R. J. Foley, R. P. Kirshner, P. A. Milne, and M. A. P. Torres. A relativistic type Ibc supernova without a detected γ-ray burst. *Nature*, 463:513, January 2010.

[269] E. B. Abdikamalov, H. Dimmelmeier, L. Rezzolla, and J. C. Miller. Relativistic simulations of the phase-transition-induced collapse of neutron stars. MNRAS, 392:52, January 2009.

[270] H. Dimmelmeier, M. Bejger, P. Haensel, and J. L. Zdunik. Dynamic migration of rotating neutron stars due to a phase transition instability. MNRAS, 396:2269, July 2009.

[271] N. Andersson, G. L. Comer, and R. Prix. The superfluid two-stream instability. 354:101–110, October 2004.

[272] A. Melatos, C. Peralta, and J. S. B. Wyithe. Avalanche Dynamics of Radio Pulsar Glitches. ApJ, 672:1103–1118, January 2008.

[273] T. Hinderer, B. D. Lackey, R. N. Lang, and J. S. Read. Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral. *ArXiv e-prints*, November 2009.

[274] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004).

[275] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. Lett. **94**, 181103 (2005).

[276] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 102004 (2005).

[277] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 082001 (2007)

[278] L. Bildsten, Astrophys. J. **501**, L89 (1998).

[279] G. Ushomirsky, C. Cutler and L. Bildsten, Mon. Not. Roy. Astron. Soc. **319**, 902 (2000).

[280] C. Cutler, Phys. Rev. D **66**, 084025 (2002).

[281] B. J. Owen, Phys. Rev. Lett. **95**, 211101 (2005).

[282] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio and N. Andersson, Phys. Rev. D **58**, 084020 (1998).

[283] N. Andersson, K. D. Kokkotas and N. Stergioulas, Astrophys. J. **516**, 307 (1999).

[284] A. Melatos and D. J. B. Payne, Astrophys. J. **623**, 1044 (2005).

[285] B. J. Owen, Class. Quant. Grav. **23**, S1 (2006).

[286] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).

[287] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).

[288] D. I. Jones and N. Andersson, Mon. Not. Roy. Astron. Soc. **331**, 203 (2002).

[289] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).

[290] K. S. Thorne, in *Three hundred years of gravitation,* eds. S. W. Hawking and W. Israel (Cambridge University Press: 1987).

[291] R.J Dupuis, Ph.D. Thesis, University of Glasgow (2004); `LIGO-P050014-00-R`.

[292] B Abbott *et al.* (The LIGO Scientific Collaboration), Ap J. Lett. **683** 45 (2008).

[293] B Abbott *et al.* (The LIGO Scientific Collaboration), Ap. J. **713** 671 (2010).

[294] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D, **76**, 042001 (2007).

[295] M. Pitkin and G. Woan, Phys. Rev. D, **76**, 042006 (2007).

[296] P. Astone *et al.*, Phys. Rev. D **65**, 022001 (2001).

[297] S. Frasca and C. Palomba, Class. Quant. Grav. **21**, S1645 (2004).

[298] D. Passuello and S. Braccini, Virgo Note (VIR-046A-07).

[299] S. Frasca, P. Astone and C. Palomba, Class. Quant. Grav. **22**, S1013 (2005).

[300] P. Astone, S. Frasca and C. Palomba, Class. Quant. Grav. **22**, S1197 (2005).

[301] C. Palomba, P. Astone and S. Frasca, Class. Quant. Grav. **22**, S1255 (2005).

[302] F. Acernese *et al.*, Class. Quant. Grav. **24**, S491 (2007).

[303] S. D'Antonio, S. Frasca, C. Palomba, submitted to CQG (2009).

[304] S. Frasca and C. La Posta, Il Nuovo Cimento, **14 C**, (1991).

[305] F. Antonucci, P. Astone, S. D'Antonio, S. Frasca, C. Palomba, CQG **25**, 184015 (2008).

[306] P. Jaranowski and Andrzej Królak, submitted to CQQ (2010), arXiv:1004.0324.

[307] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca and C. Palomba, Phys. Rev. D **70**, 082001 (2004).

[308] V. Dergachev, "Description of PowerFlux Algorithms and Implementation", *LIGO Technical Document* LIGO-T050186 (2005), available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[309] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008).

[310] B. Abbott *et al.* (The LIGO Scientific Collaboration), "All-sky LIGO Search for Periodic Gravitational Waves in the Early S5 Data", Phys. Rev. Lett. **102**, 111102 (2009).

[311] J. Zhang, C. Jarman, and K. Riles, "Adding Coherent Summing to PowerFlux", *LIGO Technical Document* LIGO-T080061 (2008), available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[312] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008).

[313] V. Dergachev, "On blind searches for noise dominated signals: a loosely coherent approach", arXiv:1003.2178, March 2010.

[314] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Phys. Rev. D **77**, 082001 (2008).

[315] O.Torre - Pisa University, Graduation Thesis: "Resampling method for semi-targeted continuous gravitational wave search in Virgo interferometer" https://pub3.ego-gw.it/itf/tds/index.php?callContent=2&callCode=7279&author =torre&startPage=, Virgo-LIGO note (VIR-0013A-10)

[316] S.Braccini, G.Cella, I.Ferrante, D.Passuello, O.Torre: "A resampling technique to correct for Doppler effect in gravitational wave search",Virgo-LIGO note VIR-0090A-10

[317] L. Sancho de la Jordana, A.M. Sintes, Class. Quant. Grav. **25** 184014 (2008).

[318] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007)

[319] Ransom, S. M., Cordes, J. M., & Eikenberry, S. S. 2003, Astrophys. J. **589**, 911

[320] C. Messenger and G. Woan, Class. Quant. Grav., **24**, 469 (2007).

[321] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009).

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

[323] H. Pletch and B. Allen, Phys. Rev. Lett. **103** 181102 (2009).

[324] P. Jaranowski, A. Krolak, and B. F. Schutz, "Data analysis of gravitational-wave signals from pulsars. I. The signal and its detection", Physical Review D **58**, 063001-24, (1998).

[325] C.-Y. Ng and R. W. Romani, "Fitting Pulsar Wind Tori. II. Error Analysis and Applications", arXiv.:0710.4168v1

[326] J. Abadie *et al.* (The LIGO Scientific Collaboration), "First search for gravitational waves from the youngest known neutron star," arXiv:1006.2535.

[327] K. Wette *et al.* [LIGO Collaboration], Class. Quant. Grav. **25** 235011 (2008).

[328] P. Patel *et al.* "Implementation of barycentric resampling for continous wave searches in gravitational wave data", LIGO P0900301, January 2010.

[329] R. E. Rutledge, D. B. Fox and A. H. Shevchuk, arXiv:0705.1011 [astro-ph].

[330] J.H. Taylor and G.R. Huguenin, Nature **221**, 816 (1969).

[331] P.C. Gregory and T.J. Loredo, ApJ. bf 398, 146 (1992).

[332] B. Abbott *et al.* (The LIGO Scientific Collaboration), Ap. J. **659**, 918 (2007).

[333] B. Abbott *et al.* (The LIGO Scientific Collaboration and the ALLEGRO Collaboration), Phys. Rev. D **76**, 022001 (2007).

[334] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007)

[335] Flanagan É É 1993 *Phys. Rev.* **D48** 2389; astro-ph/9305029

[336] Allen B and Romano J D 1999 *Phys. Rev.* **D59** 102001; gr-qc/9710117.

[337] G. Cella *et al.*, Class. Quant. Grav. **24**, S639 (2007)

[338] B. Abbott et al., "Analysis of first LIGO science data for stochastic gravitational waves", Phys. Rev. D **69**, 122004 (2004).

[339] Ballmer S W 2006 *Class. Quant. Grav.* **23** S179; gr-qc/0510096

[340] N.V. Fotopoulos, "Identifying Correlated Environmental Noise in Co-Located Interferometers with Application to Stochastic Gravitational Wave Analysis", in preparation for GWDAW 10 proceedings.

[341] B. Abbott et al, "An upper limit on the stochastic gravitational-wave background of cosmological origin", Nature **460**, 990 (2009).

[342] E. Thrane et al. "Probing the anisotropies of a stochastic gravitational-wave background using a network of ground-based laser interferometers", Phys. Rev. D 80, 122002 (2009).

[343] http://wwwcascina.virgo.infn.it/DataAnalysis/Noise/nap_index.html