LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	LIGO–T1100322 VIR-0353A-11

The LSC-Virgo white paper on gravitational wave data analysis

Science goals, status and plans, priorities

(2011–2012 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 2011/10/13

Contents

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 2011–2012. 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

Compared to previous years, a greater fraction of 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 September/October 2010 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 Ice Cube is allowing 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.

In October 2010, S6/VSR3 ended. The LIGO detectors were shut down to allow advanced LIGO installation while Virgo underwent a commissioning period to better understand the behavior of the monolithic suspensions installed between VSR2 and VSR3 in the 2 arms cavities. In the meantime the GEO detector was able for the first time to run with squeezed light, improving the sensitivity at high frequency. In 2011, given the sensitivity improvments at high frequency in GEO and at low frequency in Virgo, the two collaborations decided to take science data in coincidence during the summer having demonstrated that low frequency pulsar searches with Virgo data could be improved and that transient GW search triggered by an astrophysical source signal at high frequency could be interesting. S6e/VSR4 run took place between June and September 2011.

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

A thorough characterization of the LIGO, Virgo and GEO600 detectors and their data is critical for being able to make confident detections. This includes the identification and reduction of (stationary, non-stationary, transient and periodic) noise coupling into the detector outputs, accurate and stable calibration and timing, and careful monitoring of the interferometer state.

In the years 2011-2012 the LIGO detectors will be in the midst of an upgrade to Advanced LIGO, and the GEO600 and Virgo detectors will be undertaking coincident running and then further upgrades towards GEO-HF and Advanced Virgo, respectively. Meanwhile searches will continue on the data sets from the initial detectors. This will be a particularly important era for detector characterization. Noise transients, upconversion, features and lines had important negative impacts on searches for gravitational waves with the initial detectors. A collaboration-wide effort in detector characterization is required to ensure that these disturbances are mitigated in the advanced detectors, allowing detections. Detector characterization efforts should be focused on work that directly i) improves the sensitivity, calibration, and/or performance of the detectors and/or ii) improves the sensitivity and/or false alarm rate for LSC-Virgo gravitational-wave searches. LIGO, Virgo and GEO will be focused on different styles of characterization depending on the status of each instrument, as described in the coming sections.

Subsection 2.1 describes the LSC-Virgo-wide priorities for the characterization of the detectors and data. The detector characterization efforts specific to the LIGO, Virgo, and GEO600 are described separately in subsections 2.2, 2.3, and 2.4, respectively.

2.1 LSC-Virgo-wide detector characterization priorities

The LSC and Virgo detector characterization groups have the following overall priorities for 2011-2012.

- Characterize Advanced LIGO, Advanced Virgo and GEO-HF subsystems as they are built to maximize our knowledge of these and improve their glitch and noise performance as much as possible during the installation and commissioning phases.
- Characterize and provide run and search support for the GEO600 and Virgo detectors particularly for the S6e and VSR4 runs.
- Develop and begin implementing the infrastructure needed for enabling low-latency and multi-messenger searches with high data quality (calibration, flags, vetoes, timing, etc).
- Provide data quality support to search groups for all searches in progress on the collected S6/VSR2,3,4 data sets.
- Press for a breakthrough in glitch reduction techniques to greatly clean the background in future searches for burst and short-duration CBC signals. Use S6/VSR2,3,4 data to target much better efficiency (>90% after current category 3 data quality) for low additional deadtime ($\approx$10%).
- Implement and prototype improved detector characterization tools, techniques, and figures of merit at running detectors (GEO600 and Virgo) to test how they can be used to interact with commissioning.
- Document the detector characterization work that had an impact on the S6/VSR2,3,4 detectors and searches. This includes contributing to analysis papers, writing an overview detector characterization paper for LIGO S6, VSR2,3,4, and GEO S6, and completing papers describing impactful methods.
- Transfer characterization knowledge and experience into automated tools and machine learning approaches.

- Audit, calibrate, sys-ID, and document the physical environmental monitoring (PEM) systems.

2.2 LIGO Detector Characterization

2.2.1 Introduction

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is an NSF-funded project with the mission of directly detecting gravitational waves from astrophysical sources. LIGO has completed six periods of scientific running, S1-S6, over the past decade, amassing several years of coincident observation with its detectors at design sensitivity, and often in coincidence with its international partners GEO600 and Virgo. LIGO is in the midst of a major upgrade, called Advanced LIGO, that will increase the strain sensitivity of each interferometer by more than a factor of 10 and the volume of the universe observable by gravitational waves by a factor of 1000. Advanced LIGO is expected to make the first gravitational-wave detections and begin an era of gravitational-wave astronomy.

The LSC detector characterization group [1] directly supports LIGO's mission. A thorough characterization of the detectors is required to confidently detect gravitational waves. Gravitational-wave searches require accurately calibrated data with precise timing. The collaborations ability to make detections and the level at which upper limits for gravitational-wave emission are set depend critically on detector performance characteristics, such as the overall level of the noise-limited detector spectrum, the probability distribution of transients in the detector output, the degree to which the noise components are stationary, and lines and features that are present in the data. Detector characterization is also an important aid to the commissioning process. Characterization efforts identify issues and provide clues to commissioners, who use these to improve the instruments.

Detector characterization is carried out by a variety of scientists focused on different problems. Commissioners operate at the cutting edge of detector characterization, modifying the detector systems to increase their performance in terms of noise, lines, and robustness. Their investigations may focus on interferometer-based detector characterization, such as investigation of noise sources, lines and features, or environmental disturbances. Members of the analysis groups also make important contributions to detector characterization. They often direct their efforts toward impediments to astrophysical searches, such as coherent or accidentally coincident glitches that pollute inspiral and burst searches, features that could blind searches for periodic sources, or wandering line features that could mimic a pulsar.

During intense commissioning, it is difficult to evaluate the long-term performance of the instruments. Science and engineering runs serve as testing grounds for interferometer stability. As experience has accumulated, tools to evaluate the search backgrounds and instrument stability have improved and the latency of diagnostic feedback has decreased dramatically. Rapid feedback useful for noise and transient mitigation and commissioning is becoming routine.

However, even after years of commissioning and detector characterization the data recorded for scientific analysis contains unforeseen artifacts that decrease the sensitivity of or even blind some searches if left unchecked. For that reason, the detector characterization group has a strong effort to identify and remove non-astrophysical artifacts from the recorded data. For transient searches this is done using data quality flags and vetoes. For periodic and stochastic searches, times and/or specific frequency ranges are identified and removed from the analyses. These efforts have led to improved upper limits in the searches performed to date.

As new artifacts are found, new characterization methods are developed. If the artifacts persist, the group works to automate the relevant methods for more rapid detection of problems. For initial LIGO, the online monitoring systems included the Data Monitoring Tool (DMT)[2] with a number of targeted monitors, the controls system software (EPICS)[3], and search-oriented monitors such as the trigger generators Omega [4] and KleineWelle [5] for the burst search, and daily iHope [6] for CBC, as well as a variety of customized tools written in e.g., python, C++, and Matlab. It also included a human element, namely the attentiveness and active data exploration by interferometer operators and scientific monitors (scimons). For the advanced detector era, as much of this work as possible should be automated.

The LSC Detector Characterization community has a broad membership, including full-time commissioners, full-time data analysts, and many in between.The DetChar working group has concentrated most of its effort in the initial detector era on providing online characterization tools and monitors and on providing characterization (most directly data quality flags and vetoes, and diagnostic information) of interferometer data in science runs for astrophysical analysis. Every search working group has active DetChar members, and information flows in both directions, to mutual benefit. *The Det Char group has few members exclusively or even mostly dedicated to the group. This brings a beneficial diversity of ideas, but reduces efficiency with respect to full-time members. A goal of our group is to recruit more members that commit a substantial amount of their time to the group. The new LSC-Virgo MOU includes a publication policy that supports this goal.*

In the following subsections, we describe the requirements for the advanced detector era (2.2.2), the group priorities for LIGO detchar (2.2.3), status and plans for data run support (2.2.4) and software infrastructure (2.2.5), as well as the activities and priorities of the different working groups, noise transients (2.2.6), spectral features (2.2.7), calibration (2.2.8) and timing (2.2.9).

2.2.2 Preparing for the Advanced Detector Era

The following science requirements for detector characterization in the advanced detector era are driven by the science requirements of the search groups described in further sections of this whitepaper.

1. Accurate calibration and timing is required by all searches.

2. Searches for neutron star binaries will be relatively robust against short noise transients in advanced LIGO, due to minutes-long waveforms to me matched filtered. To maximize detections, these searches require that egregious data quality problems are identified and removed in the detectors.

3. Searches for shorter duration, less well modeled waveforms such as burst sources, and higher mass CBC systems share the above requirements, plus require a significant reduction (possibly $> 90\%$ of single-detector noise transients with respect to the rate observed in initial LIGO.

4. Low-latency multimessenger searches require high data quality in near real time. The best way to achieve this is through mitigations in the detector, but automating data quality products such as flags and vetoes is also required.

5. For continuous waves and stochastic searches the detector characterization group should identify and help remove any lines or spectral features found coherent in the instruments and/or occurring at frequencies targeted by the searches.

6. For detection confidence, all searches require a well-understood and documented physical environmental monitoring (PEM) system. We should be able to say with confidence that potential coincident signals did not arise from terrestrial, atmospheric, etc, effects.

In 2011-2012 we will work toward all of these areas, but will focus on investigating 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.

2.2.3 Priorities for LIGO Detector Characterization

Detector characterization for the LIGO detectors over the past five years has been squarely focused on thorough characterization of the initial and enhanced detectors, their calibration, timing and noise performance. The data runs in particular have driven work on identifying sources of glitches, features and lines and working with commissioners to ameliorate them, or cutting them out of the data with quality flags and vetoes. This work has now reached a turning point. The LIGO detectors are being upgraded to Advanced LIGO and will not be fully operational for at least several years. New priorities are needed. The following are the priorities for LIGO detector characterization this year:

1. Characterize the Advanced LIGO subsystems as they are commissioned, with the goals identifying and removing problems early, and more producing more scientists familiar with the instruments.

2. Make use of data sets from S6 to develop better strategies for identifying and removing features and glitches from the data.

3. Collaborate with DASWG and the LIGO Lab to develop and begin implementing the infrastructure needed for enabling low latency searches with high data quality (calibration, flags, vetoes, etc).

4. Provide data quality support to search groups for all remaining S6 analyses.

5. Document the detector characterization work that had an impact on the S6 detector and searches. This includes contributing to S6/VSR2,3 analysis papers, writing an overview detector characterization paper for LIGO S6, documenting the PEM system, and completing papers describing impactful methods.

2.2.4 Data Run Support

For most of the engineering runs and all of the science runs, LSC policy has been to staff the control rooms around the clock with one 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.

In 2009 the LSC approved a policy of a 7-day minimum stay for scientific monitors. This resulted in improved data monitoring in S6. We will critically evaluate the performance of the scimon system and develop a plan to be used in the Advanced LIGO runs. Relatedly, having a detector characterization expert at the Hanford Observatory for extended periods of time over the past years has proved critical to the understanding of many artifacts, especially those related to the coupling of the physical environment. Because longer stays have proven beneficial, we will investigate programs, such as long-term fellowships, to encourage even longer stays of scimons and other LSC scientists at the Observatories.

Another important aspect of data run support is injection of simulated astrophysical signals into the detector hardware to validate data analysis pipelines with high confidence. LIGO Laboratory and LSC scientists have provided the human resources to set up the injection infrastructure and carry out the injections.

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

Although there will be no science runs in 2011-2012, there will be engineering runs, short time intervals (one to several days) dedicated to study the data quality of the different Advanced LIGO subsystems. We expect the first of these in January 2012. We will use this period to test data monitoring systems and interactions with the searches, and to gain an early understanding of the artifacts in each of the subsystems. This will require a close communication between commissioning teams and detector characterization groups, a relationship we would like to foster.

2.2.5 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[3], DTT[7] 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[7], is mostly used as a background-process environment for continuous monitoring. The DMT system provided many critical functions in S6, including the online production of calibrated strain files that are 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.

Data quality monitoring involves reviewing the results produced by monitors such as those run in DMT, veto selection programs such as hveto [8] and UPC [9], coherence and line monitoring, results of data analysis such as search background, and other scripts running in the LIGO Data Grid clusters at the Observatories. This review is done continuously during science runs by scientific monitors at a basic level, and periodically at a deeper level by members of the Glitch and Spectral features subgroups. These investigations result in the diagnosing, identification 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 run. The information is used to form time intervals which are flagged with data quality flags, which are incorporated into a database (which also has the information on which times are in nominal science mode). The database can be queried by the astrophysical search programs, as well as by members of the group diagnosing problems in the data. The database infrastructure was first used in S5, and was significantly revamped for S6 by the LSC Software Working group (7). This resulted in a reliable system where flags are introduced online by well tested DMT monitoring programs, and offline by the DetChar group and by scimons.

In the initial LIGO data access to the data from outside the observatory control rooms was not optimal. Frames were the primary means of data access, while streaming/server data access from the Network Data Server (NDS) was limited to few people, since heavy volume requests could cause problems with the online data archiving. 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. This situation was greatly improved in 2010 with the LIGO Laboratory's development of a secure network data server (NDS2) that serves live and archived data and is independent from the live archiving processes (and thus safe for satisfying the greater demand). This has enabled widespread remote data access to detector characterization software such as the data retrieving and manipluation tools ligoDV [10], mDV [11], and pynds. Detector characterization requires not only the gravitational wave channels, but also many other data channels (and their archived trends). Because

Advanced LIGO is more complex than initial LIGO (multiple suspensions, active seismic isolation, signal recycling readout with output mode cleaners, etc), it will have more channels, and demand for served data will be greater.

Despite the critical importance of Detector Characterization, there is currently only a single member with full time dedication to detector characterization software (dedicated to NDS2 and DMT), and only a few other members partially dedicated to develop and maintain software tools. To thoroughly characterize the more complex Advanced LIGO detectors and enable the first detections, the detector characterization group requires a more robust software infrastructure and human resources to support it.

Software priorities This section describes priorities for detector characterization software work over the next few years. In general we want to build on the successful software infrastructure from S6 and expand it to meet the demands of searches in the Advanced detector era. These activities will be coordinated with the Software Working Group, as described in Section 7.

- Develop the infrastructure for the continuous online running of an improved (in sensitivity, and SNR and frequency accuracy) online trigger generator for the detector output and roughly 200 fast channels per detector from 1Hz to 6kHz.
- Update the existing online segment generation for science, lock boundary, calibration, overflows, egregious environment, etc to work with the Advanced LIGO detectors and software infrastructure.
- Produce software to monitor the first subsystems of Advanced LIGO that will form the foundation for data quality flags in the first runs.
- Continue development of the LIGO segment database to increase input and output speed, robustness and to improved user interface tools.
- Develop a new trigger database appropriate for the storage of short-duration veto information. This should be able to store parameters such as central time, duration, central frequency and SNR.
- Develop an improved channel information system containing channel names, sample frequencies, and descriptions.
- Migrate data quality and veto flags that have proven useful in S6 and are likely to be useful in Advanced LIGO to on-line production.
- Automate and improve upon current data quality flag and veto segment validation tools. For advanced LIGO these should be capable of running daily (and on longer timescales) for all data quality and vetoes and report individual and cumulative efficiency, deadtime, used percentage and safety with respect to hardware signal injections.
- Improve the current dead channel monitor with a lower false alarm, integrated reporting, and more direct ties to the segment database.
- Maintain appropriate reduced data sets for Advanced LIGO to be used for detector characterization and for data analysis.
- Write a requirements document for the DMT to help clarify its role in Advanced LIGO. Continue upkeep of existing DMT infrastructure and identify and upkeep the core DMT monitors that will also be useful for Advanced LIGO.
- Develop a LIGO daily report system similar to the GEO summary pages.

2.2.6 Noise Transients

The largest detector characterization subgroup, the Glitch Group[12], carries out studies of interferometer noise transients, or "glitches". Composed of experimentalists and analysts, the working group has broad expertise, and its work is closely coupled to the burst and CBC searches.

The goals of the Glitch Working Group are i) 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. ii) To provide information to experimentalists and builders of future detectors to achieve interferometer noise that stationary and Gaussian.

To help accomplish these goals, the Glitch group will 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 monitored the detectors in "glitch shifts" and reported weekly on issues and investigations of them[12]. This rapid-feedback analysis was based on transients found in the gravitational wave channel and in auxiliary channels (using 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 this year, we envision short engineering runs dedicated to study the data quality of the different Advanced LIGO subsystems as they are installed. We expect this activity will start in January 2012 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 are:

- organize (i.e., guarantee the people to staff) glitch shifts during the advanced LIGO engineering runs, with prompt delivery of conclusions to Laboratory staff and to LSC scientists.
- Assist the integration of subsystem status segment generation into the interferometer front end computers.
- Identification of glitches in subsystem channels that would/will affect the interferometer in future science runs. Special focus should be given to glitches rare enough that they traditionally would have only be found in extended science runs.
- Automate production of graphical visualization of the data products needed for evaluating data quality and identifying transients.
- Automate on *and* off-line identification of time intervals to be flagged as having uncertain or poor data quality;
- Implement Advanced LIGO event trigger generator (Omega) on a (large) subset of auxiliary channels.
- 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.

In addition a plan should be devised on deriving the requirements and methods for the diagnosing of glitches in Advanced LIGO.

2.2.7 Spectral Features

Another working group of the LSC Detector Characterization group is charged with investigating spectral features of the gravitational wave spectral noise density, which is especially important for the searches of gravitational waves from rotating stars and stochastic background. Many of the spectral features are due to environmental disturbances, including 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 2011-2012 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.

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

2.2.8 Calibration

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[13] available to the Collaboration, as well as recorded in the electronic logs, software repositories, and LIGO documents[14].

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)[15, 16]. 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 §2.2.9.

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 necessary 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 delivery of the vetted calibration. In addition, for students this provide valuable instrumental training. It would be highly desirable to sustain this broad participation.

The work of the calibration team is currently focused on finalizing the review of S6 calibration, and preparations for the advanced detector era. Specifically for the advanced detector era, several independent new techniques are being developed to produce h(t) data with sub-second latencies (during S6 the latency was 1 minute), on-line procedures to monitor the quality of the data produced on the fly, and the development of pre-processed h(t) products (e.g. whitened, cleaned, and coherent data streams).

2.2.9 Timing

Traceable and closely monitored timing performance of the GW detectors is mission critical for reliable interferometer operation, astrophysical data analysis, and discoveries. For example, (a) timing jitter of digitization of the GW signal could directly contribute to the noise level degrading the astrophysical reach of the LIGO interferometers, (b) coincident and coherent observation using the network of GW detectors is only possible if the absolute timing of the data streams agree within a high degree of accuracy, (c) a network of interferometric GW detectors can only recover both the polarization and sky direction information for a detected event if the absolute timing of their data-streams are known, and (d) multimessenger astronomy with external observatories also require traceable and accurate absolute timing.

The Timing Stability Working Group (TSWG) includes scientists from both the LSC and the LIGO Laboratory. The group shall be responsible for (a.) the availability and diagnostics of timing information and signals provided for various subsystems (e.g., LSC, OMC, etc.), (b.) measuring and documenting the timing performance of mission critical digital subsystems such as LSC and OMC DAQs, (c.) in close collaboration with the Calibration team (also see 2.2.8), the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator[20], characterization of analog modules, etc.), and (d.) the documented review and certification of the physical/software implementation and verification of the availability of precise documentation of timing related parts of mission critical subsystems. While it is quite likely that issues with the timing performance of subsystems are discovered by the timing team, it is the responsibility of the subsystems to address the problem; the timing team is responsible only for the certification that the issue was indeed eliminated.

The construction, testing and diagnostics tasks have already provided fertile ground for undergraduate and graduate student research involvement and diversity in the program is strongly encouraged for the future.

The next challenge in timing diagnostic is long term. Several projects will be executed in preparation of the advanced detector era, such as:

- Further develop and test injection techniques to determine accurate timing through direct test mass excitations
- Augment and expand the capabilities of data monitoring tools related to timing and phase calibration
- Enhance the availability of timing diagnostics capabilities provided for various subsystems
- Measure and document the timing performance of mission critical digital subsystems
- Measure and document the end to end timing and phase calibration measurements of the detectors (e.g., through the photon calibrator, characterization of analog modules, etc.)
- Review and certify the physical/software implementation and verify of the availability of precise documentation of timing related parts of mission critical subsystems

2.3 GEO Detector Characterization

2.3.1 Introduction

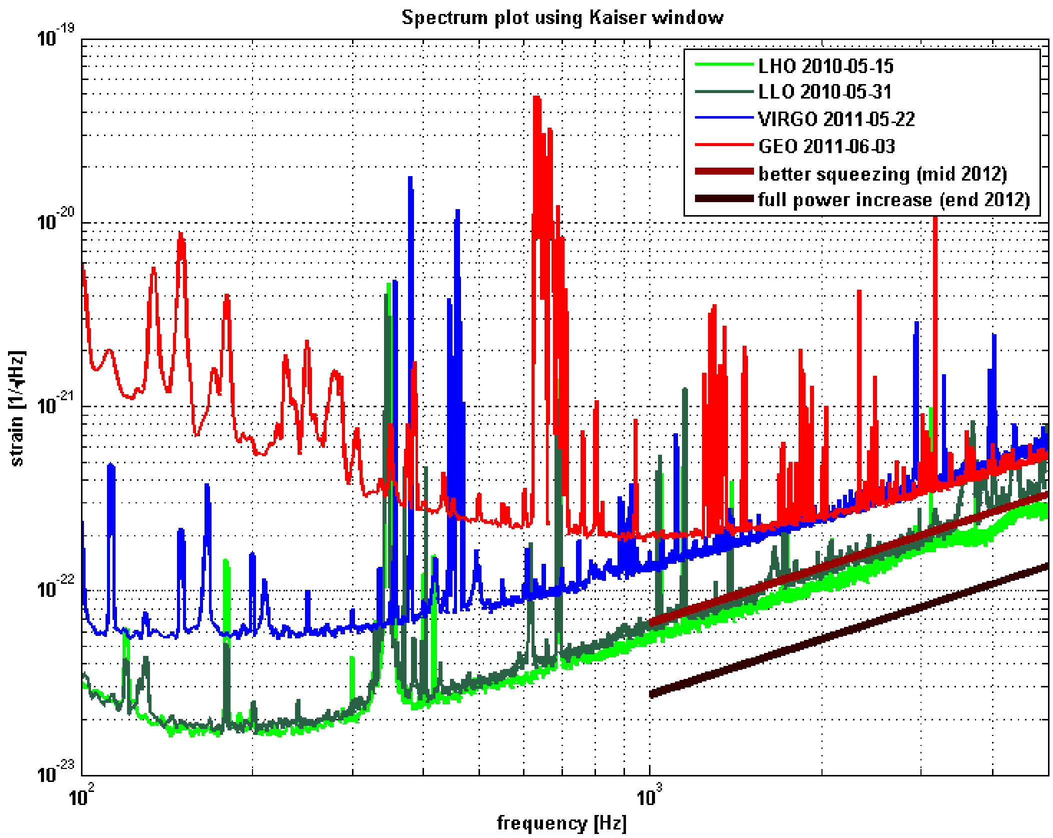

Figure 1: Recent interferometer sensitivities and GEO 600 projections at high frequencies. The projections are only plotted only above 1 kHz because the GEO-HF upgrades have mainly targeted this region.

Currently GE0 600 is operating in a joint science run with Virgo (S6E/VSR4). This run makes use of Virgo upgrades that affect the low frequency part of the h-spectrum as well as increases in the GEO 600 sensitivity at high frequencies which are a part of the the GEO-HF upgrade program [23]. The LIGO detectors have already been decommissioned for installation of the AdvLIGO instruments. See Figure 1 for the high-frequency sensitivities of GEO 600 and Virgo at the beginning of the run in comparison to the LIGO detectors right before they were decommissioned.

In September 2011 after completion of S6E/VSR4, Virgo will start a series of relatively long commissioning breaks that will ultimately lead to the completion of Advanced Virgo. During this period in which the larger gravitational wave detectors are undergoing major upgrades, GEO 600 will be mostly carrying out only minor commissioning projects during the day interspersed with data taking periods during the nights and weekends. In addition to this there are two foreseen longer commissioning periods that will prevent operation for at least an entire month each as well as a few possible 1–2 week long disruptions. We hope to gradually achieve the high frequency sensitivity projections shown in Figure 1. Notice that above 1 kHz we are already in a scientifically interesting regime, but after completing the GEO-HF upgrades we will surpass all previous gravitational wave detectors in that frequency range. As we come nearer to our sensitivity goals, there will thus be longer breaks between commissioning times where science data will be recorded.

To summarize, in the coming years, GEO 600 will operate as a fully functioning, scientifically interesting interferometer with a relatively high duty cycle interspersed with light commissioning. This commissioning will focus not only on improving the stationary noise floors but also on characterizing the detector for transients that can spoil astrophysical searches. GEO 600 can then be viewed as the sole facility in which one can find this combination of commissioning so tightly interwoven with a running gravitational wave detector.

In light of this situation, in order to best contribute to the community at large, GEO 600 detector characterization efforts over the next few years should have as its main emphasis the *implementation of new characterization techniques for the study of how they interact with the commissioning process and the interferometer behavior as a whole*. These techniques will mostly leverage off work that is already carried out by members of the LIGO detector characterization group, however we hope to also develop some completely new techniques as well.

2.3.2 Transient Studies

Due to GEO 600's focus on the high frequency regime of a few kHz our main target for astrophysical searches will be un-modeled burst searches. Over the next few years when Virgo is undergoing large commissioning operations and is not recording scientifically relevant data, GEO 600 will be the only operational laser interferometric gravitational wave detector. We will then have to carry out searches in coincidence with triggers from other types of astrophysical observatories like gamma ray telescopes and neutrino detectors. To increase the strength of the searches, since we cannot rely on multiple interferometer coherences to clean the data, it becomes important to eliminate as many instrumental artifacts in the data as possible and then have a good understanding of those remaining. For this task, it is important to have a strong relationship between detector characterization and commissioning. This need in combination with GEO 600's unique position within the community outlined in Section 2.3.1 argues for transient studies to be the highlight of GEO 600 detector characterization.

The most important way in which we want to strengthen the relationship between characterizers and commissioners is to implement already existing and currently being developed glitch characterization tools. We will then study how these tools aid the commissioning process of eliminating the glitches or determining their sources. Here is a list of tools that we will be implementing and testing in the future:

- Hierarchical Veto (HVeto) is a veto segment generation algorithm based on a Poisson statistical significance of coincidences between auxiliary channels and the gravitational wave channel. Veto segments are constructed around all triggers in the auxiliary channel with the most significance. The next round uses those triggers remaining after these veto segments are applied. This continues until a stopping significance threshold is reached [24].

- Used Percentage Veto (UPV) is a vetoing tool which finds auxiliary channels that couple into the gravitational wave channel when a glitch of large enough amplitude occurs. It then sets veto thresholds on that channel based on how often coincidences are found above that threshold [25].

- Bilinear Veto (B-veto) is a veto method that assumes a bilinear coupling of two auxiliary (normally one slow and one fast) into the gravitational wave channel [26].

- MultiVariate Statistical Classification (MVSC) is an automatic technique that takes in a large number of inputs, for example details of the gravitational wave and various auxiliary channels, and organizes the information in useful ways. One possible output is a continuous weighting of whether or not the detector was behaving well. MVSC is also capable of describing which inputs were most useful in the decision-making [27].

- WaveMon is a tool to find time-frequency correlations between the gravitational wave channel and auxiliary channels. Found correlations can be used to identify a coupling between the gravitational wave and other channels and generate a list of environmental and instrumental vetoes [28].
- Critical Coupling Likelihood (CCL) is a technique which is designed to identify a potential coupling between an auxiliary channel and the gravitational wave channel by using a likelihood analysis [29].

In working with these tools, we will identify which elements serve as useful pieces of information for the commissioning process. This identification will help to develop the tools further from this perspective and thus pave a pathway to characterization in the advanced detector era.

Another way in which we hope to aid the commissioning process is by implementing new Figures Of Merit (FOMs) that display a few simple measures of detector performance. One FOM that became important for initial LIGO was called SenseMonitor [30]. This measured the average range that a defined neutron star binary coalescence (strongest signals in the 200 Hz regime) would be just visible based on the stationary noise floor of the interferometer. Although a very useful tool for initial LIGO, SenseMonitor does not take into account transient noise very well and thus directs commissioning more in the direction of lowering the stationary noise floor. Further more since GEO 600 will be focusing on the high frequency regime, SenseMonitor is not well tuned for our needs. In addition to finding a source that produces gravitational waves in frequencies relevant to GEO 600, we also want to take a holistic approach to characterization where we view the detector as including a possible network of detectors with associated analysis pipelines. A first attempt at a FOM for GRB and supernovae signals which use the output of standard glitch identification pipelines already exists [31]. We will implement this and also work on providing a network FOM for times that Virgo is taking scientifically interesting data. We then plan to study the ways in which these FOMs direct the commissioning efforts.

Lastly, we also would like to develop some control room tools for directly helping the hunting and removal of transient behavior in the interferometer. Glitch hunting is often aided when commissioners are able to hear the problem real time. One can then probe the interferometer in many ways quickly while listening to some output which contains the desired glitches. For some very strong glitches, it is easy to hear the problem in the raw photo diode output of the interferometer. For quieter glitches something more sophisticated must be done. One path that we are pursuing is to develop a very compact line following code library. This library can be used to track the amplitude and frequency of various lines in data and can be applied to, among many other things, the real time removal of the violin modes from the audio stream output of an interferometer. A tool such as this would already be a huge improvement as the violin modes are the dominant sound in the interferometer's output and often distract the listener from possibly more important sounds. For even quieter glitches, we may investigate the possibility of creating a very small latency, 1 s at the most, "glitch stream". This would be a whitened and cleaned timestream that contains only the transient behavior of the interferometer. Ideally we would like to use the same procedures that the analysis pipelines would use to find events in single detector data to create this timestream so that commissioners would be able to listen to the output of the analyzes at almost real time.

2.3.3 Stationary Studies

Transients in our detector output are not the only type of noise that hinders our ability to see gravitational waves. We must also push forward by lowering the stationary noise floor of the detector. Part of this push is purely in the realm of commissioning. This is in the high frequency regime where the GEO-HF upgrade program is focused. Here we have planned installations that will have a significant effect on the stationary noise of GEO 600. In the middle frequency range from 100–600 Hz however, GEO 600 has been limited by a yet unexplained noise source that has an approximate $1/f$ behavior. We are very interested in finding the source of this noise and then hopefully eliminating it. This frequency range is not as scientifically interesting as the high frequencies because the LIGO and Virgo detectors have already probed this area at much higher sensitivities for long periods of time. However, GEO 600 in some ways plays the role in the community of a kilometer scale prototype that produces scientifically interesting data. Many of the techniques employed in GEO 600 today will end up in AdvLIGO and it is therefore important for us to completely understand our noise floor. In addition to this it is important to note that as we improve our high frequency performance this unexplained noise, if it really has a $1/f$ shape, will start dominating even our sensitivity above 1 kHz.

The most powerful technique for studying stationary noise floors is to project technical noises into the gravitational wave output using a measured transfer function from an actuator which that noise sits on [32]. We add to these projections, calculated fundamental noise levels and sum all the contributions to give a total projected noise floor. We will continue using this technique to access the current noise budget of GEO 600 where we see the discrepancy between 100–600 Hz. However, in order to search for this discrepancy we must expand this technique either by improving the way in which the projections are calculated, or by including more potential noise sources of both technical and fundamental nature, or both. One way in which we may improve on the noise projection technique in general is to investigate the effects of bilinear couplings or other types of interaction between two different noise coupling points in the interferometer. We plan to keep increasing the scope and accuracy of the noise projection technique until the projected noise matches with the actual noise and we will have understood the source of the unexplained noise in GEO 600.

2.3.4 Stability Studies

Another aspect of detector characterization that we will be looking into indirectly affects our scientific and technical goals. Studying the stability of GEO 600 is also very important because a drop in our duty cycle due to lock losses and noisy conditions results in less data output and sometimes prevents us from carrying out our commissioning tasks. We hope to study stability through collecting detailed statistics on many lock losses which pinpoint a possible cause for each loss of lock. This is not an easy task as it is sometimes very difficult to pinpoint the cause of a single lock loss, however as we uncover certain types of lock losses we can start cataloging when that type of lock loss occurs again in the future. These stability statistics will help to guide commissioning aimed at improving our duty cycle and also feed back into our transient studies. Lock losses sometimes result from very extreme versions of fluctuations that are happening all the time. Knowing where these fluctuations are will hint at which parts of the interferometer we should pay attention to when searching for causes of transient behavior.

2.3.5 Calibration

Calibration of an interferometric gravitational wave detector consists of two very different measurements. A relative calibration allows us to determine the shape of the $h(t)$ curve from the electrical outputs of the detector. Once we have this, we must pin down the absolute amplitude of the $h(t)$ timestream requiring a separate measurement [33]. Since GEO 600 will be mostly operating alone in the coming years the 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 will continue to provide periodic checks of our relative calibration as well as absolute amplitude measurements. Over the last year we have developed a good method for verifying the relative calibration. We may continue on a very low level trying to understand some systematic errors in the relative calibration that show up with this method. For our absolute amplitude calibration, we will carry out little to no development work and simply monitor the level periodically as stated.

2.3.6 Resources

Over the last year we have been building up our small, core GEO 600 detector characterization team. This core consists of two tiers out of a three tiered group which is involved in GEO 600 detector characterization as a whole. The core exists to service the interaction between a greater community of people who can provide tools to organize the interferometer data in useful ways (the third tier) and the commissioners who will eventually attempt to get rid of instrumental artifacts or analysts who will carry out searches on our data. The two tiers in the core group in principle have very different focuses. One tier is made up of mostly on-site members whose major function is to ensure that implemented tools are being studied and used. This is done by first having a working knowledge of how each tool works and then using them to direct probes on the detector through commissioning. The second tier of the core GEO 600 detector characterization group consists mostly of off-site members that will focus on providing data inputs to various tools provided by the third tier described above. This group should also have a working knowledge of the tools and carry out investigations that can help to guide the commissioning process, only in a way that is more data-oriented due to a lack of access to the instrument. Both of these core tiers have been growing over the last year and we hope to continue this growth.

Without third tier of the GEO 600 detector characterization group, however, the outlined plan above cannot succeed. We would like to call for people to help implement their tools on GEO 600 data! Some of the items listed in this document already have members of the team pushing them forward. Others only represent hopes that someone will step in to implement them. In addition, what is listed here is only a subset of the vast number of possible ways in which we can organize GEO 600 data in ways that may be useful for either the commissioning process or search analyzes. Please come and contact us with your ideas so that we can make GEO 600 a more well-behaved instrument, and through this process gain valuable experience for commissioning and data characterization in the advanced detector era!

2.4 Virgo Detector Characterization

2.4.1 Introduction

The search for gravitational signals requires a careful characterization of the detector and of the noise sources. The detector's response functions have to be well estimated through a suitable calibration, timing must be accurate and all noise sources, both internal ones and those due to the external environment, must be understood and reduced or controlled. In practice, the characterization of the Virgo interferometer is carried out at various levels, by different groups which are interconnected:

A first analysis is performed within the commissioning activities. it is mainly devoted to the understanding of noises, their coupling in the GW channel (dark fringe) and to the interferometer's calibration. Noise analysis is performed both to evaluate the noise budget and to eliminate eventual noise sources. Those studies are performed either when the detector is taking "Science" data (no operation is carried out) or when the interferometer is excited with suitable noise sources to evaluate, typically, transfer functions from environmental noise sources to the dark fringe channel. The activities and plans of these detector characterization commissioning activities are described in Section 2.4.2 for the calibration and Section 2.4.3 for the environmental noise hunting. Finally, the monitoring tools developed over the years to help commissioning and noise hunting groups are described in Section 2.4.5.

In addition to these efforts, two other groups, more connected to the LSC/Virgo data analysis groups are also participating to the Virgo detector characterization:

- The Virgo data quality group (VDQ) activities are focused on the investigation of any sources of non stationarity and non Gaussianity that particularly generate loud transient events that can be confused with genuine GW events from a compact binary system or a burst-like source (core collapse supernova, high mass binary black hole merger, neutron star oscillation, cosmic super-string, ...) The VDQ group is especially in charge of producing data quality flags and vetoes which are used by transient GW search groups. To do this, many tools have been developed and used over years (see Sections 2.4.4 and 2.4.7). During science run, the VDQ group also monitors and investigates the origin of the loudest glitches and any abnormal excess of background events. The goal is to provide fast feedback and hints to the commissioning team and noise hunting group to eliminate the instrumental or environmental source of glitches as soon as possible.

- The Virgo noise group activities are focused on the determination and the understanding of any narrow spectral features that appear in the Virgo sensitivity curve ("lines"). Stationary or non-stationary noise lines can mask the GW signal emitted by a continuous wave source like a neutron star. The noise group has especially developed recently tools to monitor online all lines in the Virgo sensitivity curve as explained in Section 2.4.6.

Before any science run's start, noise hunting campaigns are particularly important to eliminate seismic and acoustic sources of noise, that can generate spurious signals mostly by modulating the position of elements that scatter light inside the interferometer cavities. Commissioning shifts are also dedicated to the careful calibration of the instrument. The Virgo calibration is presently using the same actuators used to control the interferometer. The possibility of using an external actuator to produce the excitation signal, the photon calibrator, is still under evaluation (see Section 2.4.2.)

During 2010-2011, the Virgo detector characterization has been devoted to both the analysis of the VSR3 data and to the appropriate upgrades of the tools in view of VSR4 and Advanced Virgo (AdV). In the following subsections the main activities are summarized, with reference to papers or internal notes for the details. In some cases, a brief description of the foreseen developments for AdV are described. More information may be found in the Virgo Data Analysis Report [34].

2.4.2 Calibration and h-reconstruction

The "standard" Virgo calibration has been automated and extended to have some redundant measurements during VSR1, in 2007. It includes measurement of the absolute time of the Virgo data, measurement of the transfer function of the dark fringe photo-diode readout electronics, measurement of the mirror and marionette actuation transfer functions and monitoring of the finesse of the arm cavities. All these measured parameters are then used in the frequency-domain calibration algorithm, resulting in the Virgo sensitivity curve, and in the time-domain calibration algorithm, producing the h(t) strain digital time series.

Between VSR1 and VSR2, parts of the detector have been modified. In particular, the Virgo timing system has additional monitoring tools to have absolute timing calibration of the Virgo data with respect to the GPS clock, and to monitor the timing stability [35]. The mirror actuation coil drivers have been changed as well as the readout of the photo-diode of the dark port of the interferometer: all calibration measurements had thus to be reprocessed before VSR2 [36, 37]. The calibration methods are described in [38]. The installation of the monolithic suspensions during spring 2010, between VSR2 and VSR3, also induced changes in the actuation response and measurements were performed again before VSR3, in June 2010 [39, 40]. The last calibration campaign to prepare VSR4 has shown that the actuation responses have been stable at the level of the percent during a one-year period, from June 1010 to May 2011. The calibration methods that have been developed and automatized for Virgo should be still valid for AdV calibration. However, it has still to be checked whether the method used for the mirror actuation will still be sensitive enough with the new optical configuration.

The improvements of the h(t) reconstruction that were achieved for VSR1 data have been used for the online reconstruction during VSR2 and VSR3. Online reconstruction (using preliminary calibration parameterizations) was valid from 10 Hz to 10 kHz, with errors estimated to the order of $10 - 15\%$ in amplitude and $150\,\mathrm{mrad}/15\,\mu\mathrm{s}$ in phase and timing. Using final calibration, the reprocessed VSR2 and VSR3 h(t) has errors estimated to 6% in amplitude and $50\,\mathrm{mrad}/8\,\mu\mathrm{s}$ in phase and timing [41, 42]. The same reconstruction method is expected to work for the first steps of AdV, without signal recycling cavity. Studies will be needed to update the h(t)-reconstruction method for the interferometer with the signal recycling.

In order to perform coherent hardware injections to simulate gravitational wave signals in both LIGO and Virgo detectors, a system has been setup in 2009, based on GPS times and using the calibration of the timing system. It was successfully used during VSR2 and VSR3 for hardware injections (normal and blind) of gravitational-wave like signals. It should be possible to keep the same architecture for injections in AdV.

An independent calibration method ("photon calibrator"), developed to measure the mirror actuation gain, using the radiation pressure of a secondary laser which pushes the input mirrors, is being used since 2007. It is now used to determine the sign of h(t), defined as the sign of Lx- Ly where Lx and Ly are the lengths of the north and west arm cavities. in 2010, the calibration of this system has been improved [43]. It has then been used to estimate the errors of h(t) below 400 Hz with a completely independent method compared to the standard calibration. No difference were found within the systematic errors of this system, of the order of 8% in amplitude and 30 microsec in timing. Finite-element modeling are on-going in order to better understand the measurements at higher frequencies and to plan future developments of this setup.

More information and documents about the calibration and h-reconstruction are available in:
https://workarea.ego-gw.it/ego2/virgo/data-analysis/calibration-reconstruction

2.4.3 Environmental noise

The Environmental Noise (ENV) team of Virgo is in charge of the identification of noise sources of environmental origin and the reduction of their degrading effects on the detector sensitivity. Sources include natural events, detector infrastructure machineries and human activities in the site surrounding area. The ENV team works in tight collaboration with the commissioning team, the Virgo Data Quality group and the EGO laboratory groups (optics, electronics, vacuum and infrastructure). Hereafter we summarize major achievements for the VSR3 run (July to October 2010) and the near term plans for the preparation of the VSR4 run.

Major achievements during VSR3 Virgo VSR3 was a joint science run with LIGO detectors. Virgo detector sensitivity at low frequency was worse than during VSR2 because of vibration noise coupling through diffused light [44]. In particular, the output window of the dark fringe beams is an important source of light diffusion. This was due to the asymmetry in the radius of curvature of the mirrors of the two new monolithic suspensions, and the consequent bad contrast defect and large residual light power carried by these beams (3 Watts in the B1s auxiliary beam).

A second critical diffused light spot (ring shaped) was located on the wall of the West End vacuum tank: a small fraction of the arm beam was passing between the mirror and the new reference mass. Besides degrading the detector sensitivity up to about 150Hz, diffused light caused an increase of the glitch rate when natural micro-seismic activity (i.e. rough sea conditions) becomes intense. Light mitigation actions carried out during the run consisted of:

- reducing the vibration noise of the output window (slowing down the DAQ HVAC air fan, and switching to one intrinsically less noisy HVAC water circulation pump)
- installing one in-vacuum-ring-shaped absorbing glass baffle behind the WE mirror and connecting it to ground through a rigid support structure. The problem of diffused light due to the output window was solved at the end of the run by dumping a large fraction of the B1s beam power using an in-vacuum suspended and high-power-resistant dumper.

One major VSR3 goal was to decrease the glitch rate and to reduce non-Gaussian event distribution tails for transient GW event searches. The large glitch rate that appeared at the start of VSR3 was correlated to one source of high frequency EM noise which coupled to the dark fringe signal demodulation chain and then down-converted into the Virgo detection bandwidth. Culprits were probe data servers whose clock rate harmonics fall within 10kHz from the Virgo 6MHz modulation frequency. The temporary solution has been to increase the DAQ room temperature by about 2^oC, which proves sufficient to drive the temperature sensitive clock-rate more than 10kHz away from the 6MHz. Consequent to this action the glitch rate reduced from 10Hz to about 1Hz. We realized then the importance of improving the ITF robustness against Radio-Frequency environmental noise. A more careful cabling and shielding of the whole modulation-demodulation chain is needed to reduce RF noise coupling and has been set as a priority requirement for AdV.

The Continuous Wave (CW) data analysis group notified the ENV team about a noise line at the frequency of the expected Vela pulsar signal (22.38Hz). The GW strain sensitivity at the Vela frequency was worsened by 16% during VSR2. The noise source was an electromechanical component of the two TCS laser chillers whose vibration noise propagates probably to the TCS optical tables through the cooled water pipes. We implemented one mitigation solution which had a minimal impact on the run and minimized the risk for such a crucial and delicate system. The noise has been moved in frequency sufficiently far away from the Vela frequency by supplying a slightly lowered mains (216Vac instead of the usual 228Vac) to the chillers with the use of a stepping-down autotransformer.

VSR4 preparation Virgo low frequency sensitivity has sensibly improved with respect to VSR3, also thanks to the installation of two CHROCCs (Central Heater Radius Of Curvature Corrector) which allow to correct the radius of curvature mismatch of both end mirrors, and also allow to reduce the mismatch of power losses in the two arms. The VSR4 science run, in coincidence with the GEO detector, is interesting on CW searches. The most promising CW sources are in the low frequency (10-100Hz) band, in particular at 22.38Hz (Vela pulsar). One major goal for the ENV team is to keep the known pulsars Doppler bands clean from environmental noises. Noise monitoring tools, such as the NoEMi line miner, and the multiple coherence tools (see 2.4.6) help in this noise hunting.

In preparation of the run, we improved the detector environmental monitoring by replacing noisy accelerometers. We added also current sensors to better monitor switching electrical loads (i.e. infrastructure machinery) which are a possible source of glitches. As a side work, environmental sensors have been archived in the hardware inventory software tool, which provides up-to-date maps and easy access to the documentation of the sensors.

Before the start of VSR4, we extensively verified the detector coupling to environmental noises via seismic, magnetic and (for the first time) Radio Frequency noise injections. Extreme care has been put in checking possible couplings of diffused light by exciting critical components such as vacuum tank walls, vacuum links and optical benches. The noise measured transfer functions have been used to update the on-line noise budget.

Protection and monitoring of Virgo site "environmental noise climate" Soil vibration noise produced by the four 2 MWatts wind turbines farm recently installed near the EGO site has been thoroughly studied in collaboration with expert geophysicists and seismologists from the Instituto Nazionale di Geofisica e Vulcanologia (INGV). This study pointed out the presence of a seismic wave field at 1.7Hz that is detectable above background noise at all Virgo Experimental Buildings (located between 6 km and 11 km from the turbines). Details on the measurements and propagation model can be found in [45].

This result has stressed the importance of preserving our site "noise climate" by keeping under control the rise of new potentially noisy activities in our surroundings. The ENV group in collaboration with the EGO laboratory directorate elaborated a new version of a document (a first version was written in 2005) that characterizes the present noise at the site and defines noise quality requirements for new noise emissions at the site. We have taken into account seismic, acoustic and electro-magnetic noise (magnetic fields radiated by HV power lines, radio communications and disturbances induced in the electric power voltage). Following this document, EGO has been invited to the discussion table to evaluate some proposals: a solar energy power plant that inputs its generated power on the same 15kV HV line feeding the Virgo North building and a recreational area including a motor-speedway 3.5km away from the Virgo North building.

At the same time, we aim to improve the monitoring of the environmental noise at the site which originates in the surrounding territory. To this end, we plan to install one new "external environmental monitoring station" which will be located in a particularly quiet location at some distance of experimental buildings. This station shall be equipped with a tri-axial low frequency seismic sensor, one acoustic sensor and one magnetic sensor. Possibly, it will host also the present weather probes and lightning sensor. This data has several uses:

- to guarantee continuity of environmental noise monitoring during the AdV construction period (when the Virgo environmental monitoring will be shut-off)
- to study noise of external sources
- to provide useful triggers of environmental noise transient events to AdV data quality studies (ie: aircraft events).

More detailed information can be found in the specific web page dedicated to environment monitoring [46].

2.4.4 Virgo Data Quality and vetoes

The Virgo Data Quality group (VDQ) [47][48] is in charge of providing to the GW search groups the Data Quality (DQ) segment vetoes and event-by-event vetoes for each science run. The DQ segments are generated online from the data quality information provided by different online sources: the online h(t) reconstruction algorithm, the online detector monitoring, the online data quality flags processes and the Kleine Welle (KW) triggers. For VSR3 run, alike VSR2, the DQ flag segments were generated online and stored in the Virgo DataBase (VDB) and in the LIGO database (segdb) in order to provide vetoes to real time GW searches.

In addition to the online data quality information production, the VDQ group is investigating, offline, the different sources of noise that play a role in the analyzes (mainly burst and inspiral searches). It concerns primarily the source of glitches (non Gaussian transient events that are typically shorter than one second), but, as well, any source of non-stationarity that can generate an increase of the glitch rate during the scientific run. Various possible sources are the seismic activity (human activity, sea, wind, earthquake), interferometer misalignments, electronic problems, electromagnetic and acoustic disturbances, etc ...

More exhaustively, the VDQ group is in charge of the following tasks:

- Identification of the glitches found in the GW channel: this is done mainly by using algorithms searching for transient or online monitoring tools: e.g. Omega-online and OutlierMoni. The Omega-scan tool is also heavily used to find hints of coincident glitches in auxiliary channels. For VSR3 and VSR2 runs, glitch shifts have been set up and a list of loudest triggers found by online GW analyses and not DQ flagged as well as a list of glitch families were tentatively built. Such glitch shifts give also the opportunity to strengthen the interactions with the scientists in shift (Scimons) and the commissioning team.
- Investigation of the origin of the glitches and the noise coupling: this activity is done in collaboration with the Noise group (see Sections 2.4.6 and 2.4.3) and the commissioning team. We especially investigate the possible correlation between the GW channel and all the auxiliary channels. One

useful tool for that task were the triggers generated by the Kleine Welle (KW) algorithm for all the auxiliary channels recorded in Virgo.

- Generation of the DQ segments: using the same basic libraries as the Data Acquisition system, a set of DQ monitors generated online DQ flags which are collected by a specific process (SegOnline) which produces the DQ segments for the Virgo and LIGO segments data base (VDB and segdb). Additional flags are produced offline or reprocessed whenever needed using the same tools as online. Offline and online production include several cross-checks that are necessary to validate the DQ segment lists for the search groups.

- DQ performances: on a weekly basis, we study the effect of the DQ vetoes on the triggers produced by the online burst and inspiral GW searches. Statistical information such as efficiency, use percentage and dead-time of each DQ list is computed each week.

- DQ segments categorization: the VDQ group, in collaboration with the search groups assign category to each DQ that tells how the DQ vetoes should be applied by the GW search pipelines.

- KW-based vetoes: Using the KW triggers produced for all of the auxiliary channels, event by event vetoes are generated daily. Those vetoes are intended to get rid of the glitches that have not been suppressed by the DQ flags. Two algorithms originally developed in LIGO are applied: UPV (Use Percentage Veto) and hVeto (hierarchical Veto). They are often very helpful to develop new DQ flags as they reveal channels with transient in coincidence with h(t) glitches. Studies to understand the noise coupling help to assign a category to each event by event vetoes.

- Validate the safety property of all DQ and event by event vetoes using loud hardware injections.

A very useful for glitches investigations and glitchiness monitoring is the Omega burst search algorithm running online (Omega-online). Omega-online reads data from disk with a latency of few minutes. A web interface provides figures of merit using the omega triggers stored in ASCII files. Those various results and time-frequency plots are available from https://wwwcascina.virgo.infn.it/MonitoringWeb/Bursts.
For each interesting trigger, an other process, Omega-scan, runs over several auxiliary channels and provides interesting information about possible origin of the trigger. Omega-scan results of all loud triggers are available in a centralized web page for further human investigation. Omega-online and Omega-scans are tools developed by LIGO and that we adapted to the Virgo needs.

The VDQ group uses also the Virgo logbook where are reported any incident or any noise disturbance during a run. This allows to define, offline, some DQ segment that flags bad quality data periods and/or to understand some mysterious noise increase. In addition, the VDQ group has set up a Scimon interface to allows Scimons in shift to create DQ segments for any event that could affect the data quality.

All DQ segments are stored in the Virgo Data Base (see section 2.4.7). This tool allows to archive the DQ segments lists, keeping a trace of any change. It also contains all basic segments, such as the Science mode and hardware injection periods. VDB provides also useful facilities to perform requests and logical operations on lists. All DQ segments can be accessed from VDB by anyone.

In addition, during VSR3, "in-time" vetoes produced daily or hourly (like PQ veto) have been written automatically in the LIGO data base, thanks to a GRID service process.

The event-by-event vetoes based on KW triggers have also to be generated with a low latency. Actually, given the low statistics, it can only be done on a daily basis. However, these stringent requirements include that all the studies to validate the correctness of the vetoes (use of "safe channels", reasonable dead-time, etc ...) should be done and published within a few hours after the data taking.

After VSR2, the results obtained by the VDQ group were described in [49]. In time coordination with LSC, a publication is in preparation, describing the Virgo detector characterization results obtained over the scientific runs VSR1, VSR2 and VSR3. It will contain all the results obtained by the VDQ group [50].

For Advanced Virgo era, the VDQ group does not plan to do large upgrades on the computing and software that generates online DQ flags or KW-based vetoes. The main improvements will concern glitch investigations and new veto developments. For instance, noise non linear coupling investigations (using BCV tools for instance) are foreseen. We also plan to improve the identification of glitch families to determine which are the main characteristics of glitches which could be combined to give a weight to GW triggers instead of vetoing time period as we do now. The veto safety assessment has also to be improved. To compute the safety probability, we assume now that hardware injections and veto segments are randomly distributed in time. This is not the case for hardware injections (done each 5 seconds in a series of 10) and not always the case for the veto segments. More work will be done on VDB API's and utilities and we will try to provide more tools to help investigations on glitches. The organization of the glitch shifts and the interactions with Scimons and commissioning team will also be some of the main items to be improved for Advanced Virgo.

2.4.5 Monitoring Tools

Several monitoring tools have been developed along the VSR1, VSR2 and VSR3 runs. Below are described only the main ones. For the VSR4 run and in preparation of Advanced Virgo, we implemented a full archiving of this monitoring information and the tools to access to this archive. Also, a first level of information (Figures Of Merits) has been set up and several command line tools have been set up to help the Scimons to get more rapidly and easily the needed information.

Detector Monitoring: since the beginning of the Virgo interferometer's commissioning, several online detector monitoring algorithms have been implemented. The aim of such algorithms is to provide information (mainly quality flags) tagging the time epochs where data are not usable for analysis [51]. Technically, the detector monitoring implementation uses the same software tools as the data acquisition system. All the detector monitoring algorithms use the same input and output software interface [52] and the same configuration syntax [53]. Moreover, the quality tests done within each detector monitoring algorithm can depend from the interferometer's state provided by the automation system.

Are monitored the timing system, the online processing tasks and all the various interferometer's subsystems. In addition a dedicated monitoring process checks the status of each software process involved in the data acquisition or interferometer's controls.

The quality flags generated by those algorithms are collected and used to create a summary quality flag representing the general detector's quality. In addition, a "Science Mode" flag is generated upon operator's request, when a basic set of quality checks are fulfilled [51]. Finally, each algorithm generates an xml file containing the flags values and those files are used to build online web pages showing red and green flags to inform operators in control room about the interferometer's behavior and operators actions to be done:
https://pub3.ego-gw.it/itf/qcmoni/V5/index.php

Data Display: since the beginning of the Virgo commissioning, a tool has been developed [54] to provide an easy way to display data online or offline. It allows to follow in the control room the changes made on the interferometer or to do some minimal data processing (FFT, histograms, filtering, ...) This software allows to combine signals and to do several types of plots: time plots, spectra, coherence, transfer functions, 1D or 2D distributions, time-frequency plots, etc... It runs under Unix-like operating systems like SL4, SL5, Suse or cygwin and is widely used within the Virgo collaboration.

Web Monitoring: A tool based on web display, VEGA scripts [55] and bash scripts have been set up some years ago to generate periodically updated web pages. It has been recently updated to provide more features, a standardization of the scripts and an automatic archiving with html links managed to read it. This monitoring provides, in about 20 web pages, a status of each detector's subsystems, including data acquisition, online processing and online GW searches [56]. This includes a set of plots showing the noise budget of the detector and a set of spectrograms computed over days scale [57] in order to monitor the spectral behaviour of the dark fringe signal as well as several interesting auxiliary channels

Noise Budget This tool, using model or measured transfer functions, produces information and plots, periodically updated, which summarize various noises contributions to the sensitivity curve:
https://wwwcascina.virgo.infn.it/MonitoringWeb/NoiseBudget

outlierMoni: the online whitening It computes a moving averaged spectrum of the dark fringe and uses it to whiten the data. Then, it searches any output signal above 6. Its results are in the trend data and are used to show the distribution of the whitened dark fringe signal (its gaussianity) and to show lists, time-frequency plots and statistics about the events detected:
https://wwwcascina.virgo.infn.it/MonitoringWeb/outlierMoni

2.4.6 Noise monitoring tools

The goal is to have an easy to use instrument to collect and display information for scientist in the control room and to record information. These set of monitors could be used also by commissioners during daily activities to check the status of data. All the monitors' plots are gathered in a web page
https://wwwcascina.virgo.infn.it/MonitoringWeb/Noise

We produce web pages which are automatically updated with the results of running algorithms. Some of these monitors are linked to online chain and produce results in real time. Some other tools, which require to accumulate data, run on the last data written on disk and produce results with some latency of at least several minutes.

Each of these algorithms generates plots and web pages, periodically refreshed, useful for scientist who is in shift during science run.

The noise monitors which run during VSR4 archive their results in a mysql database; During VSR3, we have setup a system to archive plots and results produced by noise monitoring and we have organized access in a calendar day by day or week by week. This is useful in particular for the lines search where it is important to follow the history of lines.

For the VSR4 run, we setup the web interface to the Noise DataBase [58] and developed the NMAPI [59] which let interface to database directly via the web. The noise monitoring web pages are: https://pub3.virgo.infn.it/Monitoring

- Transient events

  - WDF: The Wavelet Detection Filter looks for transient signals in a wavelet based map. The input data pass trough a whitening filter. A wavelet transform is applied to the whitened data and a set of wavelet coefficient are selected following a thresholding criteria with respect to the noise background. The highest values for wavelet coefficients are supposed to be linked to the transient signal. The receiver is built summing all the squared coefficients and dividing this value by the noise RMS. This in principle is proportional to the SNR of our signal. After the selection of the events above a fixed threshold the events are clusterized if they are closest than a given time window. The filter gives indication also on frequency content of the identified trigger. The frequency content and trigger time of the event is associated to the maximum value of wavelet coefficient. The online report is given at:

https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=10

- Glitchiness: The glitches rate is estimated on the results of the different pipelines which run on lines. A page reporting the glitches rate in different frequency band is available at
https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=405

- Non stationarity

  - NonStatMoni: It computes RMS over various frequency bands of the dark fringe signal. Such RMS value versus time is a new signal on which spectra can be computed to observe low frequency evolution of the dark fringe frequency bands. Its results are in the trend data and are used to show the evolution of various bands RMS:
  https://wwwcascina.virgo.infn.it/MonitoringWeb/Noise/NonStatMoni

  - BRMSMoni: It computes for a set of channels, the RMS over a set of frequency bands. In addition, it can compute the amplitude of a spectral line at a given frequency. Its results are in the trend data and can be used to determine the level of noise in a control loop (looking for instance at the 1111Hz line injected in the laser frequency control loop) or to follow the level and type of seismic properties on the site:
  https://wwwcascina.virgo.infn.it/MonitoringWeb/Environment

- Lines behavior

  - NoEMi The NoEMi (Noise Event Miner) framework is used for the in-time analysis of the frequency lines in the Dark Fringe, h(t) and a set of environmental monitors spectra.

  NoEMi exploits some of the algorithms implemented for the CW search to extract, on a regular time basis, the peaks (Frequency Events or EVF) from the spectra of the Virgo gravitational wave channel and of a subset of the environmental sensors. From the persistency distributions of the peak maps NoEMi identifies the noise lines and looks for coincidences among the channels. The lines are compared with those extracted in the previous iterations, in order to track the non-stationary lines, i.e. whose frequency changes with time. NoEMi produces every day a set of summary plots of the peak maps and of the list of identified lines, and publishes them on the Virgo monitoring pages [60]. In particular, the tool raises an alarm if noise lines are found close to the doppler band of the most "interesting" pulsars.

  The peak maps and line parameters are stored into a MySQL database and can be accessed through a web user interface [61], which allows data filtering, plotting, adding user-defined metadata, etc... in order to fully characterize the archived lines.

  There are currently two instances of NoEMi running on the VSR4 data: a "low resolution, fast update" one, with a 10 mHz frequency resolution and an update of the lines database every 2 hours, and a "high resolution, slow update" one with a 1 mHz frequency resolution and an update once a day.

- High-SNR events follow up

  - WDF-Followup The last hour five loudest events are automatically analyzed to check their features in time and time-frequency domains. This is done on the Dark Fringe signal and on a set of auxiliary data.
  https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=34

- Omega-scan Omega-Scan scan in Time Frequency domain the most energetic events cross-checking with a given set of auxiliary data.
https://wwwcascina.virgo.infn.it/DataAnalysis/Burst/wonline/V1/

- Coherence

- Coherence with environmental noise auxiliary channels This tool computes the coherence between dark fringe channel and all other ITF channels. The computation is performed on segments of 5 minutes of data, decimated to 1 kHz (using a 8th order Butterworth filter) and produces a web page with two tables. The first table is a list of all channels analyzed: each entry in the table is a link to a plot of the corresponding coherence. The second table shows for each frequency bin a list of the 10 channels which gave the largest coherence. https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=40

We are working on coherence tool to insert also the results of the coherence in a database and to implement the production of coherence plots just querying the database through a web interface. This tool will be linked to Noemi to identify the source of noise for spectral peaks.
https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=406

2.4.7 The Virgo Data Base

The Virgo Database (VDB) system is based on a mysql v5 server and on a set of user and administration tools. The aim of this project is to store and manage different kinds of information that are important for data analysis:

- bookeeping: Virgo frame files geographical position (SITE, PATHNAME, FILENAME, ...)
- metadata information: data about frame data (science mode, data quality and ITF status)
- segments information: ITF specific segments (e.g. science mode) and user defined segments
- events: inspiral, burst and others
- triggers and veto: ITF specific or user defined

VDB is one of the main tool used by the Virgo Data Quality group (VDQ). It is the official location of Virgo data quality segments lists. One peculiarity of the VDB system is, on one hand the possibility to upload online DQ segments as well as Scimon or offline DQ segments, on the other hand the fact that it provides interesting tools to upload, manage and download segments lists:

- The VDB Web UI - This is the Web interface used to query and combine together data archived into the DB. In particular a dedicated section about data quality and science mode segments list has been developed. In this section it is possible to use several features like showing each single DQ list with its properties or combine together several segments lists using user defined logical expression or access to the online documentation about each segments list.
- The VDBtk_segments.exe - That is the command line tools for the DQ segments lists. This extends the functionalities available in the Web UI, in particular it is used to upload DQ segments list into VDB or to combine them with logical operators.
- The API's of the VDB library: Using only the two functions *VDB_set_online* and *VDB_segments_download*, you can download online DQ segments within your data analysis code if you linked it to the VDB library.

Currently, VDB hosts more than few hundred stable Virgo DQ lists from VSR1, VSR2, VSR3 and most of the L1, H1 DQ lists from S5 or S6 runs. Since VSR2 run, we have improved the VDB system in order to share versioning with the LSC Databases. The tools **VDB Web UI** and **VDB_segments.exe** have also been improved to manage the offline and online DQ segments lists.

In addition, a Scimon web interface has been developed, linked with the Virgo logbook. It allows Scimons to do manual entries which tag potential time periods where the quality of the data may be impacted. Each Scimon DQ entry is managed by VDB and transformed into a DQ segment.

We plan to maintain VDB development for AdV, with the help of the EGO computing department, and to make it more largely used within the LSC-Virgo collaboration. Providing simplified access tools and API's for online DQ segments will be one of our priority task.

2.4.8 Virgo detector characterization next steps

In the coming months, a publication which summarizes the main detector characterization results of the Virgo science runs VSR1, VSR2 and VSR3 will be written. It will focalize on the results obtained in noise investigations (glitches and non-stationary lines) and vetoes production and their impact in GW searches.

The VSR4 run, from June to September 2011, will be a good occasion to improve the monitoring tools and to help commissioning team to reduce the detector glitchiness or to improve the understanding of the noise budget. Preparation of the Advanced detector is already a concern for detector characterization and the use of new tools investigating non linear couplings for instance (like BCV) or new veto categorization will be one or our priorities. But no major reorganization is foreseen for advanced Virgo.

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 [62]. 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 [62].

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 [64, 65, 66] 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 to enable the flow of observed or candidate events 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 [67].
- Establish the relationship between gamma ray bursts and compact binary coalescences involving neutron stars [67, 127].
- 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 [68] and alternative theories of gravity, such as scalar-tensor theories which can result in modified phasing of the gravitational waves from binary inspiral [69].
- Bound the mass of the graviton by direct observation of the gravitational waves from binary inspiral [69].
- 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 [70, 71].

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 $\sim 1M_{\odot}$ 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., [72, 73]).

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 [72] 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 20-30 $M_{\odot}$.

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, [74]). The ringdown waves will be in the detector's sensitive band for black holes of masses above $\sim 100M_{\odot}$, up to several thousand solar masses for advanced detectors [75]. 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 [74]. 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 [76, 77, 78, 79]. Analytical models, including the effective one body (EOB) [80, 81, 82, 83] framework, and hybrid or phenomenological approaches [84], 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) are now available [125, 126]. 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 [85, 86, 87]. 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 [88, 89, 90]. This is an area of ongoing development and research [91, 92].

Eccentric orbits will also modify the gravitational waveform [93, 94]. 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 [113]. 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 [115]. 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. (Some candidate objects which may form BH-BH binaries have been recently identified; however, predictions of their future evolution still depend on population-synthesis models [116].) 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 [95] 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 [96, 66], 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 addition to the two pipelines mentioned above, there is ongoing research and development into an ultra-low latency and advanced detector analysis infrastructure for CBC detection within the gstlal project [130]. gstlal provides gravitational wave analysis elements for the open source GStreamer [129] signal processing framework. It allows modular stream-based pipelines to be constructed for various gravitational-wave tasks. Currently a pipeline to analyze long inspiral signals such as what is anticipated in advanced LIGO is under development. The CBC group will use this pipeline to search for sub-solar mass compact objects with S5 data as a test ground for Advanced LIGO searches.

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 [97]. 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% [98]. The template banks for the mass-space are well in hand [99, 100] for searches using Post-Newtonian approximations; work is ongoing to include spins in the most efficient manner [101].

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 [102]. 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 [103]. Other signal dependent discriminators are being used in current analyses [104, 105], and many others are being explored. Fully coherent methods that combine information from multiple detectors, extracting or making use of source sky position [128], 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 are 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[95] 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 [129], with CBC-specific plug-ins based on code in the LIGO Algorithm Library (LAL [131]). 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 [96, 66]. The current pipeline, implemented in the publicly available LIGO Algorithm Library suite [131], 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 [103], 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 [66] and has been used in the S5 analyses [106, 107]. 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 [75] (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.

The two-stage pipeline described above is effective at reducing the computational load associated with computing the χ^2 and other signal-based vetoes. However, it does add complexity to the analysis. Furthermore, when when trying to estimate very low false alarm rates it is beneficial to have access to all triggers, with signal-based vetoes calculated. For the blind injection challenge, this necessitated a re-analysis of the data around the injection with a single stage pipeline which calculates signal consistency tests up front. Although this adds to the overall computational burden, it seems likely that this approach will be necessary in the advanced detector era. Consequently, new methods to speed up the analysis, such as the use of GPUs and more efficient coverage of the signal space, become important.

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 [67, 127]) 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 [84] or higher [82]. More recently, waveform families have been developed which also incorporate aligned spins [125, 126]. 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 [144] has used the waveforms from [81, 82] as search templates, and waveforms from [82, 84], and [125] 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 [109], 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 [127], 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 instrumental 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, and 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 [113].

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 [114]. Therefore, we chose 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 [64, 65], 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 [67]). We can also make statistical statements about the presence of detectable gravitational waves from the entire population of GRBs examined by the searches [127].

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 [117], S2 [118, 119], S3 and S4 [66] data
- A search for primordial black holes using post-Newtonian templates in S2 [120], S3 and S4 [66] data
- A search for stellar mass black holes (with component mass greater than $3M_{\odot}$) performed using phenomenological templates [121] in S2 [122], S3 and S4 [66] data
- A search for spinning binary black hole systems in the S3 data [91]
- A search for the ringdown of black holes formed after coalescence of binary systems, in S4 data [75].

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 CBC sensitivity of the S5/VSR1 data is documented in [146]. The search of these data for CBC signals, using the methods described in the previous sections, is essentially complete. As discussed below, the analysis of GRB051103 is still under final review, the search for gravitational waves associated to GRBs identified by the interplanetary network (IPN) and the ringdown search are still in progress.

The CBC all-sky 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 [123] 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. 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 [80, 82, 83], 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 [84, 124, 125] have been used to tune, test, and evaluate the efficiency of the search pipeline. Despite the great advances in modeling inspiral, merger and ringdown waveforms, the uncertainties in the waveform amplitudes can be as large as 50% or more at high masses ($\sim 75M_{\odot}$ 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.

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; 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 for the high mass search.

At even higher masses, a search for black hole ringdowns is being performed, making use of both S5 and S6 data. In addition, a search for gravitational waves associated to short GRBs which occurred during S5-VSR1 has been completed and published. A search for 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 complete and under final review for publication.

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

- Gravitational waves associated to Gamma Ray Burst 070201 whose allowable sky location range overlapped the Andromeda galaxy [67].
- Search for gravitational-wave inspiral signals associated with short Gamma-Ray Bursts during LIGO's fifth and Virgo's first science run [127].
- Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data [106].
- Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run [107].
- Search for Gravitational Waves from Compact Binary Coalescence in LIGO and Virgo Data from S5 and VSR1 [108].
- Search for gravitational waves from binary black hole inspiral, merger and ringdown [144].
- Implications For The Origin Of GRB 051103 From LIGO Observations [145].

3.7 CBC Searches with the LIGO S6 and Virgo VSR2/3 data

The LIGO S6 run collected data between July 7, 2009 and October 20, 2010, ending in order to prepare the sites for Advanced LIGO. Virgo Science Run 2 (VSR2) was from July 7, 2009 through January 10, 2010, and VSR3 was from August 7 through October 20, 2010. The CBC sensitivity of the S6/VSR2/3 data is documented in [147].

Our original analysis goals for S6 were:

- A broad suite of analysis tools to identify detector artifacts, glitches, and other mis-behavior.
- Tighter coupling / interaction between CBC group and DetChar, operators, commissioners.
- Low mass pipeline software that is more robust and easier to run than in S5.
- Latency on the order of two weeks with our low mass pipeline.
- The use of an effective detection statistic to suppress the effect of glitches (the so-called "combined chisq-re-weighted SNR", ρ_c).
- Rapid estimation of the false alarm probability (FAP) for any low mass pipeline detection candidate, down to FAP $\simeq$ 1/100.
- Rapid event candidate follow-up for all coincident events with FAP $\simeq$ 1/100.
- A very low-latency pipeline to pass promising (low false alarm rate) events on for electromagnetic followup.

All these goals were achieved. Most were achieved throughout S6; the very low-latency pipeline was in place for the last month of S6. In addition, as described below, the blind injection challenge pushed our detection readiness far beyond out original goals.

The S6 and VSR2/3 runs have brought significantly improved broadband noise, but also new classes of detector glitches (for example, those associated with the LIGO OMC) and increased sensitivity to low-frequency seismic disturbances such as ocean storms. These directly impacted our ability to make reliable detections near the noise threshold, and also the speed with which our search pipelines were able to keep up with the data taking. On the other hand, much of S6 brought the best data we have ever acquired, with SenseMon BNS ranges approaching 20 Mpc, low glitch rates, and the most powerful and sensitive gravitational wave detection network ever assembled. We are employing the most rapid, well-understood and sensitive detection pipelines that we've ever had. We have increasingly well developed detection confidence follow-up pipelines and increasingly accurate and useful parameter estimation tools. We expect to be setting the most stringent event rate confidence bands ever obtained, over a broad range of binary parameters.

The main CBC group activities in S6/VSR2/3 primarily followed 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.

3.7.1 Detector Characterization

There was 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/VSR2/3. In particular:

- CBC group members developed comprehensive weekly data quality summary pages to supplement the work of the glitch group. Information obtained from these pages was promptly fed back to the DetChar and glitch groups, and to operators and commissioners when applicable.
- Known glitch-producing detector anomalies were rapidly and automatically identified and inserted into the database. New anomalies were discovered, investigated and vetoes developed through careful analysis. Their safety (against false dismissal, tested using hardware injections) were carefully evaluated. This included: (i) identification of auxiliary channels containing noise events which are correlated with inspiral triggers in the gravitational-wave channels; (ii) investigation of the couplings between these channels and the gravitational-wave channel; (iii) understanding and identifying particular artifacts by the signature across all appropriate channels; (iv) development of methods to clean the data of such artifacts, e.g., by vetoing noisy times; and (v) automation of these methods within the pipeline and the on-line analysis.
- Correlations between glitches in the auxiliary channels and gravitational wave channel were studied using UPV and hveto. These tools have been used to identify instrumental artifacts which affect the CBC searches, as well as producing veto times to be excluded from the analysis.
- A new class of "bilinear" vetoes was 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 were produced daily (*daily ihope*). This information played 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, produced 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 were used for rapid EM follow-up within the LUMIN system (as described below).

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 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.

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 3.12.

3.7.2 Low Mass CBC Search

A search using LIGO S6 and Virgo VSR2/3 data has been performed for binaries with masses between 2 $M_{\odot}$ and 25 $M_{\odot}$ and component masses greater than 1 $M_{\odot}$ which covers the binary neutron star, neutron star-black hole and lower mass binary black hole space.

The CBC group ran the *ihope*-based low mass and pipeline, largely unchanged from the ones employed for S5, on a weekly or bi-weekly basis throughout S6. We continued 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 was on the order of 1 month throughout S6.

Coherent (H1L1V1) low mass and high mass CBC hardware injections (up to total mass less than 35 $M_{\odot}$) were performed regularly according to a predetermined schedule, in close coordination with the hardware injection teams at all sites. These injections were identified with the *MBTA*, *daily ihope*, *low mass weekly ihope* and *high mass weekly ihope* pipelines, and deviations from expected performance were promptly investigated.

The group endeavoured to finalise the search within 3 months of the data being taken, and prepared "quarterly reports" [137] on the results. In some cases, results have been available in less than 3 months. The up to $\sim$ 6 month delay reflected 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.

Aside from the blind injection (see below), no detection candidates were found. The sensitivity of the search is consistent with the expectations given the S6/VSR2 noise curves [147]. Results from the analysis of the entire S6, VSR2/3 data have been combined with those from S3-S5, incorporating previous results as Bayesian priors, to produce upper limit statements for different mass regions. The search and paper [148] are in the final stages of review for publication.

3.7.3 The S6 GW100916 ("Big Dog") Blind Injection event

A hardware injection was performed during the run without the knowledge of the data analysis teams as apart of a "blind injection challenge." This challenge was intended to test the data analysis procedures and processes for evaluating candidate events. The injection was performed by actuating the mirrors on the LIGO and Virgo detectors to mimic a gravitational-wave signal.

The blind injection occurred on September 16, 2010 and was identified by multiple searches. It was initially observed by the cWB low-latency search for unmodeled transients. This search identified the injection as a gravitational-wave candidate with high significance in both LIGO detectors, chirping upward in frequency from 40 to 400 Hz. Virgo was observing with lower sensitivity and did not show a significant signal. Coherent analyses showed that the lack of a significant event in Virgo was consistent with the signal seen by the two LIGO detectors. The observed chirp waveform was consistent with gravitational waves radiated in the final moments of a coalescing compact binary system. The detectors were in normal operation at the time of detection and no evidence that the signal was of instrumental or environmental origin was found. In the LIGO detectors, the signal was louder than all time-shifted H1L1 coincident events in the same mass bin throughout S6. However, with only 100 time shifts, we could only bound the false alarm rate (FAR) to $< 1/23$ years, even when folding in all data from the entire analysis.

This spurred an enormous effort to develop confidence in the event as, potentially, our first detection of a gravitational wave. GW100916 forced the CBC group to:

- Estimate the False Alarm Probability (FAP) of the event to better than 1/100 - down to 10^{-6}.
- Understand the consequences of the presence of one (or more) loud events on the background estimation.
- Understand the power and limitations of both Bayesian and non-Bayesian parameter estimation (of masses, sky location, distance, spin,), and promptly write a paper [149] summarizing the results.
- Promptly analyze our follow up optical observations.
- Understand the significance of the event in Virgo and the consistency of the network response.
- Establish a rate estimate based on one event.
- Learn how to write a first detection paper [150] that is acceptable to everyone in the collaboration, and understandable to the larger community (astrophysicists, scientists, general public).
- Exercise the joint LIGO-Virgo Detection Committee [151] and obtain valuable input into the issues they grappled with.
- Develop open data release products and mechanisms.
- Develop and assemble lots of great EPO resources.
- Keep a secret for 6 months.

As a result of this exercise, we are much better prepared for the first detections in the advanced detector era. However, many of the above were achieved specifically for GW100916 and require additional work in order to be consolidated as standard pieces of the CBC analysis. The associated action items for the CBC group in preparation for the advanced detector era are discussed in more detail in a later section.

After the analysis of the event was finished it was revealed to be a blind injection and removed from the data. With the injection removed, there were no gravitational waves observed above the noise background. Although the exercise did slow down progress in completing the low mass search at some level, it was time well spent.

3.7.4 High mass Inspiral–Merger–Ringdown CBC:

A search using LIGO S6 and Virgo VSR2/3 data has been performed 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. Much work has gone into the development of data quality vetoes appropriate for the high mass search, which has templates that can be as short as 100 ms or less. Both UPV and bilinear coherence vetoes (BCV) are being employed.

A new "multivariate" detection statistic (*MVSC, pronounced "music"*) has been developed for the CBC *ihope* highmass search, 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. The use of this statistic has uncovered subtle correlations in the zero-lag coincident trigger data that does not seem to be present in the time-slide (background) trigger data, leading to a significant difference between the two data sets. This is under investigation.

In the meantime, the more conventional "combined chisq-reweighted SNR" ρ_c is used as an intermediate detection statistic, and the comparison of the value of ρ_c for a zero-lag coincident trigger with the distribution for time-slide (background) trigger data allows us to construct an inverse false alarm rate (IFAR), used as a final detection statistic. This is exactly the procedure used also for the low mass search in S5 and S6. The division of triggers into several bins depending on template duration, and the development of an improved detection statistic that incorporates template duration, are S6 innovations with respect to the S5 search.

The S6 high mass search, as in S5, uses EOB-NR templates, and both EOB-NR and PhenomB waveforms to test the pipeline and measure the efficiency. Both waveform families have been updated in the past year to give more accurate matching with the latest numerical relativity CBC simulations. The EOB-NR waveform family does not incorporate the effects of spin, either aligned or non-aligned with the orbital angular momentum. As discussed below, new waveform families that incorporate the effects of spin are under intense development by the gravitational wave community; there will be crucial for the detection of astrophysically realistic signals. When available these will be used to produce simulations which are subsequently used to evaluate the efficiency of the search.

Blind analysis boxes from all data periods during S6 and VSR2/3 have been opened, and no detection candidates have been found. The blind injection discussed above was seen, but was not sufficiently significant to be a viable detection candidate (its masses put it at the edge of the parameter range for the high mass search). The sensitivity of the search is consistent with the expectations given the S6/VSR2 noise curves [147]. Results from the analysis of the entire S6, VSR2/3 data using the conventional detection statistic are now in hand, and a full rerun is in progress to make use of the latest, consistent set of data quality vetoes. As in the low mass search, the results will be combined with those from S3-S5, incorporating previous results as Bayesian priors, to produce upper limit statements for different mass regions. The review of the analysis procedures and results is in progress.

3.7.5 Externally Triggered Searches:

The preferred progenitor scenario for short-hard gamma ray bursts (GRBs), is the merger of a neutron star or black hole-neutron star binary system. This motivates a search for gravitational waves using LIGO S6 and Virgo VSR2/3 data for CBC signals associated with GRB events; this search is in progress. The search targets gravitational wave signals within a six second window of the GRB and originating from the same place in the sky. In the S6-VSR2/3 search, GRB triggers were obtained from both the Swift and Fermi satellites. Although the interplanetary network (IPN) is also able to detect and localize GRBs, manual followup is required to analyze the IPN data to compute a sky localization. A list of IPN GRBs detected during S5-VSR1 has now been produced and a CBC search for gravitational waves associated to the short IPN GRBs is underway. A search for GW signals from additional GRB detections from the IPN network for the S6/VSR2/3 data is in planning.

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 25 $M_\odot$. To take advantage of the known sky position, this search uses a fully coherent search pipeline [63]. The Fermi satellite provides a sky localization region several degrees wide. Consequently, the targetted analysis has been extended to cover the appropriate patches of the sky for these GRBs. Forty minutes of data around the time of the GRB is analyzed with six seconds classified as the on source data, while the remaining time is used as off source data for background estimation. For those bursts where an interesting candidate is observed, a time shift analysis is performed to allow for a more accurate (by about 2 orders of magnitude) estimate of the false alarm probability. If no gravitational wave signal is observed, a series of simulations are performed in order to draw exclusion distances to the source, assuming a merger of two compact objects. Two astrophysically interesting populations are considered; a population of two neutron stars in the mass range between 1 and 3 $M_\odot$, and a population consisting of a neutron star with a black hole whose mass is around $\sim 10 M_\odot$.

Although the prime outcome from mergers are short-hard gamma ray bursts, a detection only search is conducted for all long-soft GRBs as well. Such a search is scientifically interesting, as the classification is not univocal, and sometimes even controversial among GRB experts. For these GRBs we will provide only a (non-)detection statement, and not an exclusion distance.

3.7.6 Low-latency search and EM follow-up:

The CBC group performed a low-latency search during the triple-coincident data collection of S6/VSR3. The low-latency pipeline was based on MBTA (see section 3.4.2 and was run at Cascina. The search covered sources with component mass between 1 and 34 $M_\odot$ and total mass below 35 $M_\odot$, although subsequent processing only selected events from a template that included a neutron star (mass between 1 and 3 $M_\odot$). Single-detector triggers were registered for SNR > 5.5, clustered, and subject to a chisq test with 2 degrees of freedom. Multi-detector coincident triggers were found in time and in template chirp mass, allowing for up to 40 ms of time delay between Virgo and the LIGO detectors. The pipeline selected triple-coincident CBC triggers passing standard data quality vetoes, and passed them on to the sky localization and skymap generation stage. Meaningful false alarm rates were assigned to events without the use of time slides. The expected rate of coincident triggers arising from background is estimated as the product of the individual detector rates above a fixed threshold, multiplied by a factor that accounts for the coincidence windows in time and template mass. The trigger time, properties, and sky map were passed via GraCEDb [134] and LVAlert [135] to the LUMIN [132] and GEM [133] EM-followup programs developed by the Burst group.

The entire process, including the gathering of data from the LIGO and Virgo detectors at Cascina, proceeded with a latency of about three minutes, including $h(t)$ generation and sky localization.

The entire pipeline was tested and validated using hardware and software injections. The timing accuracy, efficiency versus source distance, and sky localization accuracy were all measured.

During the last month of S6/VSR3 running, 23 triple-coincident and sky-localized triggers were passed on to GraceDB and LUMIN. Thirteen of these had a false alarm rate of less than 0.25 events/day. Of these, only 3 were from a template that included a neutron star (mass between 1 and 3 $M_\odot$). Of these 3, one was during the initial testing period before CBC triggers were sent out as LUMIN alerts, and one was sky-located too close to the sun to be observable by EM telescopes. A single triple-coincidence CBC trigger was passed on to GraceDB and LUMIN, and was subsequently imaged by several telescopes (Quest, ROTSE, SkyMapper, TAROT and Zadko). The image analysis is in progress.

A LIGO-Virgo paper documenting this effort [152] is currently under review for publication.
Goals for the future include:

- Improvements in low-latency sky localization accuracy.
- More accurate false alarm rate estimation.
- The use of stream-based pipelines to achieve lower latencies and handle longer templates.
- Developing better relationships with astronomical partners to ensure the image analysis is performed quickly by experts.
- Maintaining and updating GraCEDb [134] and LVAlert [135] to provide a clearinghouse for all relevant GW trigger information, including that from electromagnetic observatories, neutrino detectors, etc.

3.7.7 Parameter estimation

Accurate parameter estimation on detection candidates is crucial for the astrophysical interpretation of search results. It can also be used as part of follow-up studies, providing another mechanism for understanding signal consistency.

During S6, the CBC group has significantly advanced the readiness of Bayesian parameter estimation tools. Multiple tools using different sampling techniques had been developed and successfully compared with each other. These tools were applied to a number of triggers from the low-mass and high-mass pipelines, including the Big Dog event. Although the first parameter-estimation results were available within a few days after the trigger, several months were needed to carry out a full investigation and a fast-track review, resulting in a summary letter on Big Dog parameter estimation [149]. This exercise was extremely useful in highlighting areas where more work is required. In particular, it demonstrated that systematic errors from using the wrong waveform family can be very significant, sometimes dominating measurement uncertainties, stressing the need for the development of accurate model waveforms.

Further development is rapidly progressing. Several parameter-estimation codes have recently been merged into a single LAL code-base, allowing for better development prospects. Common processing tools have been developed and are being improved. Codes are being significantly sped up. New waveform families are being incorporated into the parameter-estimation toolkit, which is also growing through the implementation of new techniques for evidence calculation. A full parameter-estimation code review is about to commence. A full-author list paper describing the parameter-estimation tools developed by the CBC will be written to include examples of the performance on S6 hardware and/or software injections.

3.8 Ongoing CBC Searches and Studies

3.8.1 Ringdown search:

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 a known and simple waveform (a damped sinusoid), whose parameters (frequency and damping time) can be related to the final black hole mass and spin. 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. In addition to binary coalescence, other mechanisms for strongly perturbing an astrophysical black hole can also result in an observable ringdown.

This search has been completed on the S4 data [75], and is now being run on the S5/VSR1 (H1,H2,L1) and S6/VSR2-3 (H1,L1,V1) data using an improved and re-tuned *ihope* pipeline. The new pipeline incorporates some tools originally developed for other CBC analyses, including coherent signal-to-noise, null stream information, and automatic tuning using the multi-dimensional classifier MVSC (see section 3.7.4). With these new techniques the S5 and S6 ringdown pipeline will perform significantly better than the previous S4 pipeline. A stripped down version of the ringdown search (without coherent statistics or machine classifiers) was used in the IMR and NINJA2 projects discussed below, and the results of the IMR study showed that the ringdown search outperforms other techniques at the high end of the mass range.

3.8.2 Joint CBC-Burst IMR detection 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 is in progress and still needs to be written up and published.

3.8.3 NINJA 2 and NR-AR:

The Numerical Injection Analysis (NINJA) project [136] 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. Although the NINJA 2 project is not part of the LSC, there is now an MOU between the project and the LSC, to facilitate the use of LSC resources such as LIGO and Virgo data, and software.

Currently, the NINJA 2 project has amassed a long list of waveforms supplied by numerical relativity groups, joined to post-Newtonian inspiral waveforms at early times ("hybridization"). There are also numerous data analysis teams, many drawn from the LSC CBC group, poised to search through simulated data to detect and measure the parameters of these waveforms. The project aims to collect and publish the results in the coming year.

The related NR-AR project [154] fosters collaboration between numerical and analytical relativists to develop and improve analytical and phenomenological parameterized waveform families which capture ultra-relativistic corrections and the effects of aligned and non-aligned (precessing) component spins, with increasing accuracy. These waveform families can serve to interpolate (and perhaps even extrapolate) between numerical relativity waveforms which sparsely sample the 7-dimensional parameter space of compact binary coalescence, and to produce physically realistic waveforms in far less computational time than can be achieved with numerical simulations. Such improved waveform families are required and crucial tools for (i) templates for search pipelines; (ii) search pipeline tests and validation; (iii) measurement of the sensitivity of detectors and search pipelines to signals as a function of luminosity distance and binary parameters; (iv) estimation of those parameters from observed signals. Because of the difficulties associated with accurately reproducing waveforms over the parge parameter space, this is envisioned as a long-term project. Although several LSC members participate, the NR-AR project is not formally affiliated with the LSC. However, the plan of the NR-AR collaboration is to provide the LSC with the parameterized waveforms that will be produced.

3.8.4 Sub-solar mass search

Compact binaries with components below 1 $M_\odot$ were searched for in S2, S3 and S4. An S5/S6/VSR1/2/3 search offers a new region of possible compact binary parameter space to make a detection with existing data, but it also tests the scale of binary neutron star searches in advanced LIGO/Virgo. By searching for component masses of 0.2-1.8 $M_\odot$ the sub-solar mass search has templates lasting $\sim 10^3$ s starting at 40Hz and requires about 10^5 templates to complete the search. This is a similar computational scale to the anticipated Advanced LIGO/Virgo searches with component masses 1-35 $M_\odot$ starting at 10 Hz. Searching this parameter space now is an invaluable opportunity to tackle some of the technical problems associated with large template numbers and long templates that we will have to face in Advanced LIGO/Virgo.

3.9 CBC Group Priorities and Goals for FY2012 - before the Advanced Detector Era

The *highest priorities* for analyses to be performed on the existing S5-6 and VSR1-3 data for the CBC group during FY2012 are listed belo. In section 3.12, we list the development tasks required to prepare for the advanced detector era (ADE). This section lists those goals aimed at science before the Advanced Detector Era. The “CBC” code is intended to aid in referring to these tasks in group MOUs to the LSC.

3.9.1 S5 and S6/VSR2/3 Analyses

- **GRB 051103 Search - CBC-GRB051103** Complete of the search of (pre-)S5 data at the time of GRB 051103. Complete the review and publication of full author list results paper.
- **IPN GRB - CBC-IPN-GRB** Perform the search of S5 data at the time of short GRBs identified by the IPN satellites. Complete the review and publication of full author list results paper, jointly with the burst group.
- **S5/S6/VSRn ringdown search - CBC-S5ringdown**. Complete the S5/VSR1/S6/VSR2/3 ringdown search analysis. Write, review and publish full author list results paper.
- **IMR comparision - CBC-IMR** Completion of the comparison of IMR methods used on the S5 data, and publication of the results in full author list paper.
- **NINJA 2 - CBC-NINJA2** Perform NINJA 2 analysis on S5/VSR1 data, review results, publish findings in full author list paper.

- **S6 Data Characterization - CBC-Detchar**. Write and publish one or more papers documenting the identification and development of data quality characterization and vetoes used by the CBC group during the S6/VSR2/VSR3 run.

- **S6/VSR3 Low Latency Analysis and EM follow-up - CBC-S6lowlatency**. Complete, review and publish full author list paper documenting the S6/VSR2/3 CBC low-latency analysis and EM follow-up work.

- **S6/VSR2/3 high mass CBC search - CBC-S6highmass**. Complete the S6/VSR2/3 high mass CBC search analysis and review, and publish full author list results paper.

- **S6/VSR2/3 GRB Externally Triggered Search - CBC-S6GRB**. Complete the S6/VSR2/3 CBC GRB/external trigger search, identify and evaluate detection candidates, evaluate detection sensitivity and compute rate limits, prepare publishable results and full author list paper(s) jointly with the Burst group, and conduct a thorough review.

- **S6/VSR3 Parameter estimation - CBC-S6paramEstimation**. Write, review and publish one or more papers documenting the Bayesian and non-Bayesian analyses developed to estimate parameters for CBC detection candidates, including the GW100916 blind injection. At least one of these will be a full author list paper.

- **Followup of Candidates - CBC-S6followup**. Write, review and publish a full author list paper documenting the examination and detection confidence procedures for CBC detection candidates performed in the S6/VSR2/3 low mass, high mass, and GRB analyses, including the GW100916 blind injection.

- **S6/VSR2/3 sub-solar mass search - CBC-S5S6SSM.** Write, review and publish a full author list paper with the results of this search and a discussion of the technical aspects explored for Advanced detector analysis.

- **Searches using new techniques - CBC-S6newSearches**. Apply newly developed searches, covering regions of CBC parameter space not yet covered by our existing searches, to S5/S6/VSRn data. If it can be established that a new search technique greatly enhances our sensitivity (with respect to our existing searches) for sources in a significant region of parameter space, write and publish full author list observational papers. This may include searches that employ templates incorporating the effects of aligned or general spin; higher harmonics; intermediate mass ratio inspirals (IMRI); eccentricity; etc.

3.9.2 Searches "Before the 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 "Before the Advanced detector Era" (BADE).

- **GRB-triggered searches - CBC-BADE-GRBs**: 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-BADE-other**: 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. At the moment, we do not anticipate that the effort required to mount such searches is justified.

3.10 First science papers in the Advanced Detector Era

Initially, we anticipate following science goals (analysis and paper plans) that to was done for the S5-6 and VSR1-3 runs. Specifically, we plan to perform low and high mass CBC searches, a targetted search using external, astronomical triggers and low latency search for both data quality and telescope pointing. Each of these analyses will result in at least one full author list observational paper. We have a long list of issues that need to be addressed to improve the analysis pipelines and infrastructue and make them ready for the advanced detector era. For the most part, the required enhancements are known and in many cases a strategy has been developed and now needs implementing.

3.10.1 Low mass CBC search:

The low mass CBC search has been the flagship search of the group since the first science run. It is an all-sky, all-time search for CBCs, in the mass range for which the inspiral phase dominates: from $m_1 = m_2 = 1$ $M_\odot$ up to a total mass that depends on the noise spectrum, but is likely to be in the range of $M_{tot} = 10 - 20$ $M_\odot$. In addition to being well modelled by the post-Newtonian approximation, waveforms in this mass range will be sufficiently long in the advanced detector band that we can expect that signal based vetoes such as the chisq test will be as effective as they have been in iLIGO/eLIGO for suppressing non-Gaussian glitches. This is our most mature search. We can use an improved version of the ihope pipeline, or move to newer filtering methods.

This mass range includes binary neutron star (BNS), neutron star-black hole (NSBH) and the least massive binary black hole (BBH) systems. The biggest challenge in performing this search will be handling the long templates (several minutes duration) necessary for the advanced detector sensitivities and covering the large space of templates. This problem is only exacerbated by the necessity of modelling the component spins when computing the waveforms. For the NSBH systems, this may also require taking into account binary precession during the inspiral. It may prove to be advantageous to have a dedicated analysis pipeline for BNS systems; but if so, we need to weigh the benefits of an additional search against the "trials factor" associated with multiple searches in different regions of parameter space.

We plan to run this search offline, in, eg, 2-week intervals, as we did in S6-VSR2/3. We will need to run the time shifted analysis and software injections concurrently, in order to promptly evaluate our efficiency and compute upper limits or rates in the presence of signal events. We need to be prepared for the first detections! Consequently, much of the development work in the next two years will be devoted to implementing and automating the lessons learned and new methods derived during the S6 blind injection.

The first paper can be prepared when we either accumulate significantly more (likely around an order of magnitude) search volume times time (V$\times$ T) than our previous analyses or when we have our first detection(s).

3.10.2 High mass CBC search:

This is an all-sky, all-time search for CBCs, in the mass range for which the inspiral, merger and ringdown phases are all important: total mass above $\sim$10 $M_\odot$, as high as 100 $M_\odot$ for the merger, and as high as 600 $M_\odot$ for the ringdown phases. High mass systems spend only a short time in the detectors' sensitive band so improved suppression of detector glitches is required; rapid DetChar, and the use of a multivariate statistical classifier is called for. It is not clear whether the search for systems with $M_{tot} > 10$ M$_{\odot}$ should be pursued with one, two, or more pipelines / template banks; it depends on the noise characteristics of the detectors. A CBC task force will address this issue.

Accurate knowledge of the binary coalescence waveforms in this mass range need to be developed, validated, implemented and carefully vetted. This will be an ongoing project performed in collaboration between those running numerical simulations, those analytically and phenomenologically modelling the waveforms, and those developing and running the searches. The development, implementation and testing of full binary coalescence waveforms is a key goal of both the NINJA 2 and NRAR projects. Component spin will be important in this mass range, and a method of searching the full parameter space including spin will need to be developed. Currently, there are efforts to include aligned spins in a 3-D template bank for this search, if it is technically sensible. Further study is needed to determine those regions of parameter space where an aligned spin search is necessary and sufficient. Development of searches that perform better for non-aligned, precessing spins are ongoing.

3.10.3 Low latency CBC search:

A low latency search of the data will provide both rapid feedback on the performance and glitchiness of the detectors as well as providing rapid localization of CBC signals for possible followup by electromagnetic observers. This search will focus on templates for low-mass systems incorporating at least one neutron star as most models of EM emission from binary systems require the presence of matter in the binary. We will require a latency and sky localization performance at least as good as in S6/VSR3, although we will aim for improvements in both latency and sky localization. These can be achieved through the use of streaming-data, time domain search techniques, and more sophisticated localization techniques that go beyond timing alone to reconstruct the signal location.

3.10.4 Externally triggered CBC search:

We will perform a search for BNS, NS-BH binary inspiral signals in time- and spatial coincidence with an external trigger (primarily, short-hard GRBs). Additionally, we will extend the "external triggers" to other sources beyond GRBs that might signal mergers, like neutrinos, radio transients, and optical transients. The implementation of this search in the advanced detector era will face many of the same challenges as the all sky, all time low mass CBC search. Specifically, we will have to implement a search which is capable of covering the full parameter space, including the longest template waveforms. The current analysis makes use of a coherent combination of the data from all detectors, and this should be maintained in the advanced detector era.

If will be of critical importance to perform this search with low latency and provide a fast response to the astronomy community so they can perform any necessary additional observations. This will require the development of methods for distributing triggers, including GW triggers, to and from astronomers and GW physicists (i.e. VOEvent or similar), with as low latency as possible.

3.10.5 Milestones

The milestones and timeline are common for the four main searches described above:

- Determine how to search the parameter space for early/design advanced detector runs, taking into account detector noise spectra, astrophysical information on neutron star and black hole masses and spins. Use first software engineering runs to help answer these questions. Hope to complete by March 2012.

- Fully review all waveform families used for templates, testing, efficiency and parameter estimation, by Spring 2012. New waveform families arising from NINJA 2 and NR-AR work, incorporated and reviewed when available.
- Implement most of the needed improvements to the pipeline in the following 6 months to have a functional pipeline in place by late 2012.
- Run on mock data during the upcoming software engineering runs in 2013.
- Continued improvements towards perfection throughout 2013-2014.
- Fully reviewed pipeline by early-2014.

3.11 Science and key analysis issues in the 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 in detail in the remainder of this section, 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, incorporating spin effects into the waveform templates, performing coherent network analysis, and developing additional features in order to improve detection sensitivity and parameter estimation.

We will have a worldwide network of advanced detectors which will enable us to confidently detect of gravitational waves, locate the sources on the sky, decompose the polarizations, and enable tests of the predictions of general relativity and beyond-GR theories. 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, sky localization, parameter estimation, and fundamental tests of GR.

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 available.

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. 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.

As the Earth-based gravitational-wave detectors move into a phase where they are making routine astronomical observations, the nature of our analysis efforts will change and we must be prepared to meet a broad range of ambitious science goals. 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; prompt source localization and parameter estimation, and close connections with associated electromagnetic observations.

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". The threshold of sensitivity that should trigger the public release of the strain data has been defined [153], and plans are progressing to define exactly what the open data products will be and how they will be delivered and managed.

In the following subections, we list the identified tasks, innovations and top-level goals that we plan to pursue, to prepare for the coming Advanced Detector Era of intensive search, the first detections, then subsequent exploration of the physics of gravitational waves and astrophysics of compact binary sources with routine detections.

3.11.1 Analysis development required prior to first detections

1. **Understanding the parameter space of binary coalescences and searching over it**

   The range of parameters for waveforms from CBC sources is huge, ranging from total masses of $2-1000\,\mathrm{M}_{\odot}$, and including spin, higher harmonics, matter effects in mergers containing neutron stars, extreme mass ratios, and maybe even eccentricity. It is computationally impossible to optimally cover the parameter space with template banks, and many carefully-considered compromises must be made, requiring study. This will only be possible by combining our best astrophysical understanding of these systems, the most up to date models of the waveforms and an systematic study of the sensitivity of the detectors to the different features of the waveforms. The issues depend upon the system in question and also vary across the mass space:

   - **BNS**: How can we best handle waveforms that last up to ten minutes in our data? How important is spin for binary neutron stars in the ADE? Can we see evidence of tidal effects in the inspiral phase, and tidal distruption at high frequencies (400 to 1500 Hz)? What latency (h(t), DQ, coincident triggers, background estimation) do we need and can we get? Can we really detect a significant signal before merger? We currently do externally triggered searches with short GRBs. Evaluate the possibilities for searches to follow up other electromagnetic transients that are likely to be associated with binary mergers and implement such searches when they become feasible. The timescale for this activity will depend in part on when proposed all-sky surveys, such as LSST, become operational.

- **NS-BH**: Here, the spin of the BH has significant effects on the waveform, and spin-orbit precession becomes important. We need to understand how to cover this space. An aligned spin (a 3D template bank) search is being developed and should be tested. How much of the space of non-aligned (precessing) sources can such a search capture? We expect there are regions of parameter space that we cover poorly, especially for high mass ratios. Does a search which accurately models the precession of the binary increase our sensitivity? Can searches employing Physical Template Family (PTF) waveforms [90, 101], appropriate for one highly spinning component, and incorporating precessional effects, with 4D template banks be useful in recoving this part of the parameter space? Can we develop improved parameterized waveform families for systems with high mass ratios, non-aligned spins, tidal effects, and higher harmonics?
- **BBH**: Here, the spins of both bodies can be important. For high mass BBH, do we need a better model of the ringdown (including other harmonics/modes)? The NINJA2 work will be valuable in developing such models. Do we have good models that incorporate beyond-GR effects, to test GR? Can we develop improved parameterized waveform families (from the NR-AR collaboration) for systems with high mass ratios, non-aligned spins, and higher harmonics? For high mass ratios (100:1 or more), we may need to implement a search for intermediate mass ratio inspirals (IMRI) using waveforms from perturbation theory.

2. **Developing and incorporating improved waveform models**:

   CBC pipelines use binary coalescence waveforms for tuning and testing our search pipelines; as templates for matched filtering at the heart of the search pipeline; evaluating the efficiency and sensitive volume of our searches as a function of luminosity distance and binary parameters; and as templates for parameter estimation after detection. Much work has gone into developing waveforms of increasing accuracy, covering all phases of coalescence, 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 our analysis , and used in the high mass searches: 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. New CBC waveform families and template banks will be incorporated into our software base and our advanced detector era searches as soon as they become available.

3. **Development of search pipelines appropriate for the Advanced Detector Era**:

   As in the past, we will be searching for signals from low mass CBCs, high mass CBCs, BNS and NSBH signals associated with GRBs, and ringdown signals. We will require search pipelines that are robust, automated, full-featured, fast, easily configurable, self-documenting, and fully reviewed. A natural approach is to revisit our *ihope* pipeline, improving and rewriting parts of it in order to better achieve the above goals. There are a number of specific challenges which need to be addressed:

   (a) **Handling minutes-long waveforms**: CBC waveforms will spend as much as ten minutes above 15 Hz (for 1+1 $M_\odot$). The current CBC *ihope* pipeline analyzes data in 40 minute, with a small overlap, and thus will not reliably capture the full duration of the signal. Various approaches can be taken to overcome this problem: For detection, we may only be interested in capturing signal SNR above, say, 40 Hz, before triggering more complex procedures to capture the full duration at a later follow-up stage. It may be possible, with minor changes, to better optimize the existing analysis it for these long duration signals. Alternatively, it may prove necessary to to restructure the existing pipeline, perhaps along the lines of the Virgo MBTA pipeline, in order to handle such long signals coherently.

(b) **Hierarchical Search Methods**: The size of the parameter space increases rapidly as the low frequency sensitivity of the detectors improves. Additionally, the incorporation of new features in the waveforms, particularly spin, only serves to further expand the search space. In order to efficiently cover the whole space, it may be necessary to employ hierarchical search methods. 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.

(c) **Coherent analysis**: The current all sky CBC searches filter data from each detector in the network independently, find single detector triggers, then look for multi-detector trigger coincidence in time and in template parameters. In principle there is a sensitivity gain from performing a fully coherent analysis that explicitly requires a consistent signal, comprised of only two gravitational wave polarizations, between all detectors in the network. The power of the coherent analysis increases as more detectors are added to the network. However, performing a fully coherent analysis is computationally costly as each point on the sky must be searched independently. Additionally, the full analysis must be repeated for each time-shifted background. To date, only the triggered GRB searches have been performed coherently. We need to understand how to implement a coherent analysis stage in the all sky, all time pipeline. It may well be that there is a need for an initial coincidence stage and only loud triggers are analyzed in a fully coherent manner.

(d) **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. Consequently, the analysis is time consuming and the use of a batch mode analysis introduces an inherent latency in the pipeline. 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.

4. **Development of methods to distinguish signal from detector noise and estimate the false alarm rate of events**:

While they achieved remarkable sensitivities, the first generation of gravitational wave detectors did not produce data that was stationary or Gaussian. The output data contained numerous non-stationary times and glitches caused by both environmental and instrumental effects. There is no reason to believe that the situation will be different in the advanced detector era. Thus, a significant effort is required to allow us to efficiently separate signals from the noise background and to characterize the false alarm probability of candidate events. This requires work on a number of fronts, including understanding the detectors, using smarter ways of distinguishing signal and noise, estimating FAR and dealing with FAR estimates with lots of signals.

(a) **Identifying and dealing with glitches**: As discussed throughout this white paper, detector glitches can fake CBC and burst triggers, and even fake coincident triggers if the rates are sufficiently high. This can severely limit our search sensitivity and detection ability. The advanced detectors may have very different, and time-varying, behavior in this respect. We need to be prepared with a full suite of tools to help identify detector glitches, and systematically search for their instrumental causes. Rapid turn-around and feedback to the detector commissioners will result in eliminating many glitches in the science data. For those that remain, we must prepare effective data quality vetoes. Since CBC waveforms may be many minutes long, there may always be glitches on top of signals. These need to be identified and "notched out" in order to be confident in the candidate detection, and in order to estimate waveform parameters from the data. The glitches may also affect the estimation of the underlying noise spectrum, and the calibrated detector response to signals. These effects must be understood and corrected for.

(b) **Improved Classification of Signal and Background**: 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, …). The pure "signal-based" measures of loudness ignore other information such as local measures of the data quality and physical environment. 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.

(c) **Estimating false alarm rates**

Our confidence in the first detections will only be sufficiently high if the false alarm probability (FAP) for them is very low, down to 10^{-6} (or "5σ"). This in turn is based on estimating the "local" false alarm rate down to very low values. As we make more CBC detections, we will gain more confidence and may set the bar lower for confident detection (eg, FAP on the order of 0.1%). We were able to reach a FAP on the order of 10^{-5} for the Big Dog after much special effort and computing, including the use of a single-stage version to perform millions of time-shifts. We need to put in place automated, rapid, unbiased, reliable and fully tuned and tested methods for estimating the FAR and FAP for the most significant coincident triggers in the advanced detector network. 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 the necessary changes to the analysi to make use of these alternative methods;
(iv) investigate possible sources of bias in the background estimate from environmental or other correlations.

(d) **Estimating background in the presence of signals**: In the Big Dog background estimation, we were ultimately limited by the presence of the Big Dog triggers in the data sample used to estimate the background from accidental coincident triggers. We need to deal with this "little dog" problem in an acceptable way. This might get even more tricky in the happy event of a "signal rich environment". What do we mean by the usual statement "having many similar detections will give us more confidence"? It may be that there's no right answer to this question, with two approaches (assuming that all loud time-slide triggers coincident with a zero-lag candidate are signal and leaving them out of the background estimate; or assuming that they are all background and leaving them in) bracketing the truth.

(e) **Event candidate follow-up procedures**:
Events output by the CBC pipelines that are significantly inconsistent with the background estimates are subjected to a series of checks and tests to look for signs that the events were caused by instrumental glitches. These include 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. We developed and employed an elaborate set of checks and procedures to follow up on the loudest detection candidates in the initial detector science runs, as described in section 3.4.10. However, these procedures are time-consuming and require considerable effort. Most of these detection confidence checks are at best semi-quantitative, and in many cases, are wholly qualitative and subjective, requiring experience and good judgment in their application. However, 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. We will further develop these procedures with the goal of making them better automated, faster, less labor intensive, and less subjective.

5. **A simulation campaign**

The use of simulated signals added to the detector data is a vital part of developing and validating the analysis pipelines. These simulations are performed both in hardware, by actuating the mirrors of the detectors, and in software, by adding the signal to the recorded data.

(a) **Hardware injections**: Coherent injections of signals into the detectors (typically via actuation on the end test masses) are powerful tools for understanding many aspects of the detection process, including: low-statistic estimates of the search sensitivity; the calibrated detector response; the coherence of the data streams from the detector network; the safety of data quality vetoes with respect to vetoing signals; the comparability of the (infrequent) hardware injections and the (much more frequent) software injections that are used to estimate the search sensitivity; the accuracy of the expected SNR and parameter estimation; and as a full end-to-end test of the analysis pipeline. We will continue to perform a broad range of hardware injections in the advanced detector era, carefully optimized for maximum benefit with minimum disruption of science data taking. This will require significantly more manpower and expertise than we have been devoting in previous science runs.

(b) **Evaluating search sensitivity with software injections**: We need to understand our detection SNR thresholds versus false alarm probability as a function of detector glitchiness, etc, for the range of CBC source parameters. 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, particularly those incorporating spin. 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.

3.11.2 The first detections

1. **Parameter estimation**: The continued development of the multi-project analysis efforts and joint analysis projects is of central 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. All of our tools need development in order to be more efficient, more accurate, better understood and more useful in the event of the first detections. There are several areas of active development:

   - Use Bayesian parameter estimation techniques to determine posteriors on the parameters of a detected signal. Use MCMC and nested sampling 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 other harmonic structure. It is known that other harmonics can significantly enhance our ability to detect gravitational waves from binary systems as the total mass increases. 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. On the other hand, failure to include spin in the models can result in significant systematic biases in the measurement of other parameters, such as masses.
   - Make a quantitative comparison of tradeoffs between the full parameter estimation codes and rapid follow-up codes regarding the accuracy to which the position of the source can be localized on the sky.

2. **Effect of calibration uncertainties**: Calibration uncertainties translate directly into undertainties in detector sensitivity and therefore they need to be propagated into the astrophysical rate estimates that are the end products of our searches, and into the errors in waveform parameter estimation. At present, we only account for the overall normalization error in h(t) generation due to calibration uncertainties. However, frequency-dependent amplitude and phase errors will also affect the matched filtering that is at the heart of our search pipelines and parameter estimation. Tools are needed to perform this error propagation robustly, in the era of potentially high SNR signals and potentially highly accurate parameter extraction. This will be especially important for precision tests of GR and for doing cosmology with GW standard sirens.

3. **Going beyond the "loudest event statistic"**: Currently, we estimate upper limits on CBC rates in the local universe using the loudest event statistic [64, 65], in which the loudest event in the search (in terms of some detection statistic such as combined SNR or IFAR) defines the threshold at which software injections are counted as "detections" for purposes of computing the sensitivity of the search.

This is only sensible in the presence of at most one detection. The method therefore needs to be generalized or replaced by an alternative one (such as the simpler approach of defining a fixed detection threshold).

4. **Development of low latency online CBC search pipeline**: As described in section 3.4.2, the CBC group ran a low latency online CBC search pipeline in the last two months of S6 / VSR3, based on MBTA. The pipeline passed triple-coincident low mass CBC detection candidates on to GraceDB and thence to Looc-Up for EM followup. We intend to prepare a robust low latency online CBC search pipeline for the advanced detector network, with real-time false alarm rate estimation, sky localization, and basic parameter estimation (component masses). We may be able to move away from triple-coincidence and only require a threshold on significant network-coherent SNR.

5. **Astronomical alerts**: Implement prompt methods for the flow of GW triggers from low-latency search pipelines at gravitational-wave observatories to EM observatories. 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 and reliably. Also implement methods to associate GW event triggers with events observed with neutrino detectors; this does not require low latency, since neutrino detectors, like GW detectors, keep all their data and do not need to be pointed. These methods are being implemented jointly with the Burst group; see section 4.4.5.

6. **Electromagnetic followup of GW triggers**: GW triggers generated by a low-latency CBC pipeline will be used to point telescopes across the EM bands: radio, IR, optical, UV, X-ray and gamma ray. The products of those EM observations must be analyzed to look for signals that can be associated with the GW, as evidence that they were produced by the same underlying engine. The resulting associations will serve to greatly strengthen confidence in the GW signal, and provide powerful and unique information on the astrophysics of the source engine. In order to properly interpret the results, the sensitivity of the EM observations (in terms of magnitude or EM energy release) and the false alarm probabilities for the associations much be carefully evaluated. CBC group members will work closely with the EM astronomers to aid in the discovery, analysis and interpretation of the GW-EM associated triggers.

7. **Determining the nature of short GRB progenitors**: The current working hypothesis for the nature of short, hard GRB progenitors is that they are binary mergers involving at least one neutron star. Such systems will produce gravitational waves, prompt gamma ray bursts (within seconds of the GW merger signal), as well as prompt and/or extended emissions in the x-ray, UV, optical and radio. Coincident observations in some or all of these bands can confirm or disprove this working hypothesis, and greatly constrain progenitor models.

8. **Release of data associated with confirmed detections**: We will define the data products to be released associated with confirmed detections (including time, sky location, false alarm rate and false alarm probability, parameter estimation, snippets of h(t), noise PSDs, the status of the detector network, etc), along with contextual information and FAQs to help consumers interpret the data. We will develop tools to promptly publicly release and publish this information.

3.11.3 Analysis development for routine detections

1. **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 astrophysical and cosmological models and predictions and address important issues

such as the compact binary population of the Universe, details of stellar and binary evolution, dynamical interactions in globular clusters and other dense environments, evolution of the star formation rate, etc.

2. **Determining the EOS of neutron stars**: As binary systems containing at least one neutron star approach merger, tidal effects will cause the neutron star(s) to deform and eventually be disrupted. The tidal distortion changes both the system's binding energy as well as the radiative flux in such a way that the gravitational waveform frequency increases with time faster for more deformable neutron stars. The tidal deformability of the neutron star (essentially, the size) can thus be measured, and this parameter in turn depends on the neutron star equation of state (EOS). The EOS describes the nuclear physics of matter at these uniquely extreme densities. Studies have shown that observable effects may be present starting at frequencies as low as a few hundred Hz, for high SNR detections. Such observations can constrain the neutron star mass and radius, and thereby the EOS of the nuclear matter, providing much insight into the nature of these extraordinary objects.

3. **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).

4. **Standard siren**: A binary inspiral is an astronomer's ideal standard siren: 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, and to measure the Hubble constant in the local universe. However, for binary black hole sources at cosmological distances the mass is "blue-shifted"; requiring observations of the host galaxies to break the mass-redshift degeneracy. This is best done on an individual detection basis, but can also be done statistically, using many detections with poor association with host galaxies. In addition, binaries containing at least one neutron star will emit gravitational waveforms which encode tidal effects at characteristic frequencies that can be determined using relatively nearby binary mergers. For the more distant mergers, the tidal effects on the observed waveform can thus be used to determine the redshift from the gravitational wave alone, without reference to the host galaxy.

5. **Testing the post-Newtonian expansion**: Comparison of the observed inspiral signal in binary black hole mergers to the parametrized post-Newtonian signal model allows one to 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.

6. **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.

3.11.4 Computational challenges

Some of the big computational issues that we must address are:

- The traditional approach of matched filtering noisy data with banks of templates comes at significant computational cost, growing rapidly as the parameter space increases. New techniques are being developed to speed this up, including the use of methods to reduce the redundancy in the banks via SVDs or reduced basis catalogs).

- The coherent analysis cannot be done on all the data; instead, we only run it on the loudest detection candidates. And it is even more challenging to run it with a large number of time slides. We need to develop ideas for overcoming the computational obstacles.
- 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. Our computing clusters will have GPUs, and work to employ them efficiently in our search codes is in progress.
- Parameter estimation is extremely compute-intensive, as we have seen from the analysis of the Big Dog using fully spinning 15-dimensional parameter spaces; it may be much slower for the longer waveforms in the ADE. There are many potential ways of speeding this up, including the use of improved sampling, such as more efficient jump proposal distributions; restricted frequency ranges or parameter spaces; hierarchical methods; MPI; GPUs, etc.
- There are also numerous computational challenges associated with achieving extremely low latency, which are being addressed in DASWG and the CBC group.

3.12 CBC Group Priorities and Goals for FY2012 - in preparation for the Advanced Detector Era

Section 3.11 above lists numerous key analysis issues to be addressed in order to fully achieve our science goals in the Advanced detector era. Based on those, we here briefly enumerate the specific R&D goals that the CBC group will pursue in the coming year.

- **Advanced detector analysis pipeline development - CBC-ADE-pipeline**. Development and improvement of our workhorse search pipeline software.
- **Coherent Analysis - CBC-ADE-coherent**: Development and incorporation of coherent multi-detector network information to more effectively distinguish signal from background.
- **GPU-enhanced CBC pipelines - CBC-ADE-GPUs**: Developing the necessary software modifications, and designing next-generation computing clusters to implement GPU-enhanced CBC search pipelines.
- **Next Generation Low Latency Pipelines - CBC-ADE-lowLatency**: 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.
- **Detector Characterization and data quality techniques for the advanced detectors - CBC-ADE-DetChar**. 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. Includes the implementation of multivariate statistical techniques to combine data quality information into one measure.
- **Improved signal based vetoes - CBC-ADE-signalBasedVetoes**: Development and implementation of improved "signal-based" quantities like the χ^2 and r^2.

- **Improved Classification of Signal and Background - CBC-ADE-detectionStatistic**: Improved detection statistic, multivariate classification, improved IFAR estimation, etc.
- **Improved Background and FAR estimation - CBC-ADE-background**: Improved estimation of the false alarm rate (FAR) for the most significant coincident event candidates, using large numbers of time slides, single-detector trigger rates, and other methods. Automation of the procedures for the loudest events.
- **Improved detection confidence procedures - CBC-ADE-detectionConfidence**: Development of the follow-up pipeline and checklist. Automation, faster, less labor intensive, and less subjective.
- **Detection sensitivity in the advanced detector era - CBC-ADE-sensitivity**. Understanding our detection SNR thresholds versus false alarm probability as a function of detector glitchiness, etc, for the range of CBC source parameters. Large-scale simulation efforts.
- **Hardware Injections - CBC-ADE-HWinjections**: Development and analysis of hardware injections as tests of detector response, veto safety, pipeline sensitivity, network coherence, etc.
- **Studies of the advanced and third generation detector network - CBC-ADE-network**. This includes understanding the impact of detectors in India, Japan, Australia, and elsewhere on the detection sensitivity, detection confidence, sky localization, etc.
- **Development of new waveform families - CBC-ADE-waveforms**: Development of parameterized waveform families suitable for CBC searches, extending our coverage of parameter space. Implementing these waveforms into LAL, and deploying them in searches (as templates or as siftware injections).
- **Searches incorporating spin - CBC-ADE-spinsearch**: Identify regions of parameter space where spin waveforms differ significantly from the non-spinning templates. Development of spin-aligned template banks and searches. Evaluation of the efficacy of a PTF-based search. Implementation of parameterized IMR waveform families including aligned, and non-aligned (precessing) spins.
- **Overlap with Burst searches for high-mass systems - CBC-IMR-NINJA2**: 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-ADE-paramEstimation**: 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.
- **Effect of calibration uncertainties - CBC-ADE-calibration**. Develop tools to properly propagate the estimated calibration errors to our search SNRs and parameter estimates.
- **Extension of externally triggered search - CBC-ADE-ExtTrig**: 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).
- **Electromagnetic followup of GW triggers - CBC-ADE-EMfollowup**: Develop, test and document strategies and procedures for associating EM transients (in all EM bands; and also neutrinos) with GW

candidate CBC events (either from the low-latency online pipeline or the offline analysis), including estimation of sensitivity and false alarm rates. Interpreting and extracting astrophysical information from the associations.

- **Advanced detector era astrophysics and GR - CBC-ADE-astrophysics**: Developing the tools that the LVC will use to turn LVC observations into astrophysical conclusions: models of short GRB progenitors; measuring parameters of the EOS of neutron stars; astrophysical source population parameters; tests of the post-Newtonian expansion; tests of GR vs alternative theories of gravity; etc. Emphasis is not on the theory or "toy" studies, but on tools that we will use to extract the physics or astrophysics from our actual GW observations.

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*, or short duration transients. As discussed in Section 4.1, 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. For this reason, the Burst Group must utilize robust detection methods. Initially, the group focused on signals much shorter than 1 second in duration, with little or no assumption about their morphology [155]. 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 4.2 summarizes recent GW burst search results, while Section 4.3 describes the methods used in our searches.

In section 4.4 we describe the Burst Group scientific goals and plans for GW burst searches with the LIGO, GEO and Virgo instruments over the next $\sim$2 years. During this time we will complete the analysis of data from S6/VSR2+3 and prepare for burst searches in the advanced detector era, while some specialized searches will also use S5/VSR1 and Astrowatch data. While this document includes all ongoing or planned burst analyses and investigations, the Burst Group remains open to new, scientifically motivated activities that will support its goals; the analysis techniques and tools that we will continue to refine over the next few years will form the foundation of burst data analysis in the Advanced LIGO/Virgo era.

Special attention is dedicated in Section 4.5 to the planned usage of new data that will be collected during the transition era before the advanced detectors come online. Within the $\sim$2 year horizon, GEO-HF and Virgo data will provide a testing ground for evolving techniques, as well as new, good results in certain cases, such as searches for high-frequency GW bursts associated with external triggers.

Finally, in Section 4.6 we discuss the Burst Group's top priorities and the papers we commit to publishing early in the advanced detector era, with a timeline of our preparations.

The Burst Group's ultimate goal is the unequivocal detection and astrophysical interpretation of GW burst signals. In the absence of a detection, our goal is to set physically significant upper limits on gravitational-wave emission and constrain astrophysical or cosmological mechanisms responsible for that, using complementary approaches. The broadest net is cast by the so-called **untriggered burst search**, which "listens" to the whole sky for statistically significant excursions of signal power within the LIGO-GEO-Virgo sensitive frequency band (see Section 4.4.1). In some cases, astrophysically motivated assumptions on the sources, such as the GW waveform or the source population, point to **customized untriggered searches**, with improved sensitivity and ability to interpret the results (see Section 4.4.4). We also pursue GW bursts associated with energetic transient astrophysical events that are 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 a rich potential for interpretation within an astrophysical context. These are the so-called **externally triggered burst searches** (see Section 4.4.2). During the S6/VSR2+3 run we also began **seeking electromagnetic (EM) counterparts** to GW burst candidates, through collaborations with observer groups and Target of Opportunity (ToO) observations across the electromagnetic spectrum (see Section 4.4.5).

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 the Burst and CBC 3 Groups. Longer duration transients ($\sim$minutes or longer) may also be pursued using methods developed by the Stochastic 6 and Continuous Waves 5 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 Gravitational-wave Burst Science

A broad range of astrophysical systems and phenomena that potentially emit gravitational-wave bursts have been targets for burst searches with the initial LIGO, GEO and Virgo detectors. The primary goal of these searches was to make a first detection, but along the way first astrophysically relevant observational limits on the gravitational wave emission from the targeted sources were made.

The list of potential burst sources 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.

Searches for gravitational-wave bursts in the advanced detector era will be a combination of untriggered, all-sky searches and externally triggered localized searches. Untriggered, all-sky searches (See 4.4.1) have the greatest potential of finding electromagnetically dark sources (see Section 4.1.6) and may discover unexpected sources (see Section 4.1.8). They also provide triggers for follow-up studies of candidate events with EM observations (see Section 4.4.5). Externally triggered searches will have electromagnetic 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 continue to work with our external partners and to extend collaborations from trigger exchange to, for example, full data sharing and joint analysis, parameter estimation and physics extraction wherever it facilitates better scientific output.

While we can expect to learn more astrophysics about known sources (even non-detections, translated into improved upper limits, will have important consequences for astrophysics) 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) that is completely different from model predictions; and (ii) a high-significance event that is detected with unexpected characteristics and with no EM or neutrino counterpart. We must develop strategies of how to handle both scenarios.

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 possibilities and potential. This discussion is also relevant to the analysis of the joint GEO-Virgo S6E/VSR4 run (June-Sept 2011; see Section 4.5) and of subsequent GEO runs in the "dark ages" before the advanced LIGO and Virgo detectors come online.

4.1.1 Soft Gamma Repeaters (SGRs) and Anomalous X-ray Pulsars (AXPs)

SGRs and AXPs 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 [156]. SGRs/AXPs are most likely strongly magnetized neutron stars (magnetars) that experience recurring dynamical events in their crust or magnetosphere, which lead to the observed outbursts, though the details of the outburst mechanism remain to be understood.

If magnetar outbursts/giant flares are due to magnetic-stress induced crust quakes (e.g., [157, 158]) and if crust and core are efficiently coupled in magnetars, SGRs/AXPs will emit gravitational waves via the excitation of the magnetar's non-radial oscillation modes [159, 160, 161, 162, 163, 164, 165, 166]. The pulsating X-ray tail of SGR1806−20 also revealed the presence of quasi-periodic oscillations (QPOs) (e.g. see [167, 168, 169]) that were plausibly attributed to seismic modes of the neutron star, thus suggesting associated associated gravitational-wave emission [170, 171].

Together with glitching radio pulsars, bursting magnetars are the closest known potential sources of gravitational waves. Less than two dozen Galactic SGRs and AXPs have been identified, at distances from $\sim 1.5\,\mathrm{kpc}$ to $\sim 50\,\mathrm{kpc}$. SGRs 1806-20, 1900+14, and 1627-41 are located in our Galaxy near the galactic plane, between 6 and $15\,\mathrm{kpc}$ distance. Another, SGR 0526-66, is located in the Large Magellanic Cloud, about $50\,\mathrm{kpc}$ away. Two others, SGR 0501+4516 and SGR 0418+5729, are located opposite the galactic center and may be an order of magnitude closer to us ($\sim 1\,\mathrm{kpc}$) than the other known Galactic SGRs .

A network of advanced detectors should set energy upper limits for gravitational-wave emission by magnetar outbursts that are a factor of 100 better than the S5y2/A5 search [172]. This allows us to probe the energetics of these sources several orders of magnitude below the typical electromagnetic energy output of giant SGR flares, constraining an unexplored and interesting region [163, 173, 164, 165, 159]. If provided with reliable gravitational-wave emission models that go beyond analytic estimates (i.e., ringdowns), advanced detectors may be able to constrain the outburst mechanism by putting astrophysically interesting upper limits on emitted energies.

4.1.2 Gamma-Ray Bursts

Gamma-ray bursts (GRB) 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 [174, 175]. Long ($\gtrsim 2\,\mathrm{s}$) GRBs are associated with star-forming galaxies of redshifts of $z \lesssim 9$ and core-collapse supernovae [176, 177, 178, 179]. Short GRBs ($\lesssim 2\,\mathrm{s}$) 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 [180, 181], while up to 15% may be due to extragalactic SGRs [182, 183] (see Section 4.1.1).

Short GRBs: NSBH, NSNS Coalescence and Postmerger Evolution

Burst searches generally have the capability of picking up also compact binary (inspiral-merger-ringdown) events (e.g., [184, 185]). Recent analytic and computational work suggests that constraints on the nuclear equation of state (EOS) are possible by matched filtering of advanced detector data from the intermediate to late inspiral of NSNS (and NSBH) binaries [186]. 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 [184] this may be possible with burst methods alone to $\sim 100\,\mathrm{Mpc}$ with advanced detectors.

Interesting science potential is not restricted to the inspiral phase. 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 place constraints on MHD spindown and the nuclear EOS. The postmerger signal (which 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 $1-4\,\mathrm{kHz}$. With an expected energy emission of up to $\sim 0.001-0.01\,M_\odot c^2$, these signals will most likely be detectable only for nearby ($D \lesssim \mathrm{few}\,\times\,10\,\mathrm{Mpc}$) events. It will therefore be worthwhile to perform a targeted search on the postmerger evolution for the most nearby events. This search should be informed by the next generation of predictions from self-consistent computational models that include all the necessary physics.

The majority of nearby NSNS/NSBH inspiral/merger events are likely to be gamma-weak/silent due to the expected beamed prompt emission, but more isotropically emitted precursors or afterglows (e.g., [187, 188, 189, 190, 191, 192, 193]) are expected to be emitted in bands from radio to X-ray. Discovering the EM counterpart of a NSNS/NSBH merger will be a major breakthrough. Joint EM-GW observations of short GRBs in the advanced detector have great science potential and will provide answers to many open astrophysics questions connected to short GRBs and binary mergers (e.g., [194, 181, 180]).

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. There is overwhelming evidence from EM observations that long GRBs are related to massive star death and core-collapse supernovae [195] (see also Section 4.1.3), 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 and/or accretion energy via magnetohydrodynamic processes (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 [196]. Hence, long GRBs should be approached with similar modeling input as supernova searches. During the GRB stage, gravitational-wave emission may come from accretion disk instabilities (clumping, fragmentation) [197, 198, 199] or nonaxisymmetric magnetar deformation [200, 197]. 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) [184]. 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., [201, 195]). 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 gravitational-wave emission. An engine-driven supernova within a few tens of Mpc would be an interesting science target for advanced detectors.

Current and future searches for long GRBs rely on external triggering by the prompt gamma emission observed by satellites. As in the case of short GRBs, joint EM-GW observations may help answer pressing questions regarding the central engine and progenitors of long GRBs.

Additional long GRB science opportunities are in joint observations/searches for gravitational waves and high-energy neutrinos (HENs) from GRBs. HENs are expected to be generated by interactions of accelerated protons in relativistic shocks in GRBs [202, 203, 204, 205, 206, 207, 208, 209, 210, 211, 212] and may be observable by HEN detectors. Joint GW+HEN searches could discover EM weak/silent or so-called *choked* GRBs (e.g., [213, 206]) and add another important component to multi-messenger GRB science. An important component in probing the structure of joint GW+HEN sources is the connection between the time-of-arrival of GW and HEN signals [214]. Besides long GRBs, other phenomena are also expected to produce HENs and GWs simultaneously. Other sources include core-collapse supernovae [208], short-hard GRBs [202], soft gamma-repeaters [215, 216, 217] and microquasars [218, 219].

4.1.3 Core-Collapse Supernovae and Accretion Induced Collapse

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 [220]. It has been argued [221] 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 [196]. Based on the S5y2/VSR1 reach and all-sky search sine-Gaussian upper limits [222], 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 current simulations (see, e.g., [196, 223] for reviews). More optimistic predictions based on analytic models suggest gravitational-wave energies of up to $\sim 0.1\, M_{\odot}c^2$ [197]. 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 fully exploit the improved reach/sensitivity that advanced detectors will offer.

The rate of nearby core-collapse supernovae is known only to a factor of $\sim$two. Galactic events are expected once or twice per century and the rate for the entire local group, including Andromeda, is roughly two to four per century [224, 225]. The rate increases dramatically at a $3 - 5\,\mathrm{Mpc}$ where star forming galaxies produce $\sim$0.5 core-collapse supernova per year [212, 226]. Within $10\,\mathrm{Mpc}$ the core-collapse supernova rate is $\sim 1/\,\mathrm{year}$. While the expected event rate is moderate, the physics that may be learned from a detection of gravitational waves from a core-collapse supernova is very reach and 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 [227]. 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 [222]. Sensitivity improvements (by factors of order unity) can be expected from a targeted 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.

Triggers from low-energy neutrinos (emitted copiously in core collapse) will provide much better timing than EM triggers (e.g., [228]) and current initiatives to collaborate with neutrino observatories will need to be intensified in the advanced detectors era. Triggering by neutrinos will be trivial for any Milky Way or Magellanic cloud event. More distant events in galaxies of the local group (e.g., in Andromeda) will lead to only very few events in current and near-future neutrino detectors [229, 230] and a joint GW-neutrino search for sub-threshold triggers may increase the detection efficiency by up to a factor of two [231].

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 an “unnova” (an event that leads to a weak explosion or no explosion at all [232] 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. Unnovae or obscured supernova may make up $\sim$50% of all core collapse events [233].

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., [234]). 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 [234, 197].

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.

4.1.4 Phase-transition Induced (Mini-)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 waves by ringdown oscillations following a "minicollapse" of the neutron star and its stabilization at a higher-density equilibrium [235, 236]. 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 [196]. In the former case, typical gravitational-wave 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) [222, 237] 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.

4.1.5 Pulsar Glitches

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

There exist two main candidates for the mechanism behind pulsar glitches. One suggestion 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. [238]) that once the differential rotation reaches a critical value, an instability sets in, causing a dramatic unpinning of the superfluid vortices. This transfers angular momentum from the superfluid component to the crust and causes 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 [239] 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 starquake-driven glitches, it also seems reasonable to expect that the neutron star f-modes may be excited.

Given 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 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.

The excited post-glitch pulsar may also emit gravitational waves at r-mode frequencies (of order tens of Hz for most glitching pulsars). The time decay constant for a quasi-sinusoidal gravitational-wave ring-down at r-mode frequencies is likely dominated by damping from the viscous boundary layer between the outer crust of the neutron star and its fluid core. According to the model proposed by Levin and Ushomirsky [240], the expected time constants are of order 10^4 seconds or shorter for neutron stars in the frequency band of terrestrial gravitational-wave detectors.

4.1.6 Mergers of Black Hole Binary Systems

The coalescence of binary compact objects consisting of neutron stars (NS) and/or black holes (BH) is the most promising source for detection by LIGO and Virgo. Burst analysis of the NSNS and NSBH sources has been already discussed in section 4.1.2 in the context of GRB searches. Burst searches are important in particular for coalescence events of binaries for which no accurate post-Newtonian (PN) or numerical relativity (NR) waveforms are readily available due to theoretical or computational limitations. Since no templates are needed by burst searches, they are also insensitive to possible differences between the PN/NR models and nature. Therefore, burst searches are an important complement to the conventional matched-filtering searches conducted by LIGO and Virgo [241], which use the knowledge of complete PN/NR waveforms. Astrophysics wise, burst search are particularly important for detecting the coalescence of intermediate-mass black hole (IMBH) binaries and binaries with high eccentricity.

Coalescence of IMBH binary systems (total mass $\gtrsim 50\,M_{\odot}$) produce gravitational waves where the inspiral portion of the waveform is outside of the LIGO-Virgo sensitive band and the detectable signal is dominated by the merger phase [185]. IMBHs have been posited to form a transition in the black hole mass hierarchy. As such, IMBHs cover several decades in the black hole mass spectrum between stellar mass black holes of a few tens of $M_{\odot}$, formed from stellar collapse, and supermassive black holes of $10^5 M_{\odot}$ or more present in the center of galaxies. Some models of IMBH formation include runaway stellar collision scenarios [242] in globular clusters, failing core-collapse supernovae in very massive primordial stars and the progressive accumulation of mass into a large $\gtrsim 50\,M_{\odot}$ seed black hole via coalescence of a population of lower-mass black holes. The potential existence of IMBHs has been investigated by EM astronomy and a few candidates exist [243]. It has been suggested that IMBHs are the engines powering ultraluminous X-ray (ULX) sources [244] such as M82 X-1 [245] or NGC 1313 X-2 [246]. Most commentators agree that the primary hosts of these objects would be globular clusters [247]. These objects are thought to grow from accretion of smaller compact objects [248], hence IMBHs could be the sites of frequent intermediate-mass-ratio inspiral events, emitting gravitational waves in the frequency range accessible to ground-based detectors. The observation of an IMBH binary could have important consequences for theories of the formation and growth of supermassive black holes and the dynamics and evolution of globular clusters [249].

Current approaches to the detection of gravitational-waves from binary black holes expect the binary to have assumed a circular orbit by the time it enters the frequency band of ground-based detectors. However, black hole binary systems may form through other scenarios, such as dynamical interactions in galactic nuclei containing a supermassive black hole (SMBH), or in globular clusters [250]. If stellar-mass or intermediate-mass black holes form a dense population around SMBHs, the probability of close encounters of two black holes is high [251]. Such encounters may lead to the formation of binary black hole systems. The initial eccentricity of such binary system is likely to be close to unity and remain large all the way to the merger. The merger may happen within hours, and such short lived systems are expected to have a unique gravitational-wave signature: a series of short bursts. There are no accurate eccentric binary black hole waveforms available and so burst searches may be the only way to detect and study binary sources with high eccentricity. Such sources are exciting targets for the advanced detectors and their detection would be an event of great astrophysical importance.

4.1.7 Cosmic (Super-)String Cusps and Kinks

Cosmic strings are one-dimensional topological defects that may have formed during one of the early symmetry-breaking phase transitions in the universe [252]. Superstrings are the supposed basic constituents of matter in fundamental string theory or M-theory, which is at present the leading contender for a unified theory of all fundamental interactions including gravity. Both cosmic strings and superstrings are still purely hypothetical objects. There is no direct observational evidence for their existence even though they may produce a variety of astrophysical signatures. One of the most promising ways of detecting the presence of cosmic strings and superstrings is via their gravitational-wave emission. This is the primary mechanism of energy loss from strings, so strings are believed to contribute significantly to the cosmological gravitational-wave background.

An important role in the emission of gravitational-wave is played by the presence of cusp features on strings resulting from an oscillating loop. At such points the string reaches the speed of light for a brief moment and a burst of gravitational-waves is emitted, with strong beaming in the direction of the string's motion. Damour and Vilenkin showed this yields a very characteristic gravitational-wave signature which depends on the reconnection probability p, the size of the loop given by the parameter ε, and the string tension $G\mu$ [253] [254] [255]. This last parameter is crucial to characterize a string network and its evolution. The best current constraint is given by the CMB observations of WMAP: $G\mu < 2.1 \times 10^{-7}$ (68% C.L.) [256].

The Damour and Vilenkin work motivates a dedicated search for a gravitational-wave burst signal in LIGO and Virgo data. The waveform is well-modeled by a power law in frequency ($\sim f^{-4/3}$) and a frequency cutoff given by the angle between the cusp feature and the line of sight. This offers the possibility of performing a templated search with a better efficiency than the typical all-sky burst search. A first analysiss has been performed on S4 data but the short live-time and the weak sensitivity of the detectors did not allow to constrain the cosmic string parameter space significantly [257]. The S5/6-VSR1/2/3 analysis is in progress and, if no detection is made, it is expected to set upper limits on $G\mu$ about 10 times more stringent than the current CMB constraints. Another factor of 10 (or more) may be expected from searches in the advanced detector era, especially since signals from string cusps are strongest at low frequencies where the sensitivity improvement is expected to be the greatest.

Another gravitational-wave signal is expected from cosmic strings. It would originate from a string kink produced after the reconnection of two string segments. As for cusps, the kink signal obeys a power law in frequency ($\sim f^{-5/3}$) but with a smaller strain amplitude, which disfavors detection. However, recent studies (e.g., [258]) show this could be compensated by a higher production rate. A kink search would be a straightforward extension of the current LIGO-Virgo analysis.

4.1.8 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 possibility 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.

4.2 Recent Observational Results

In the past year, the LSC-Virgo Burst Group has published results of two externally triggered searches; in addition, a number of analysis are mature, under review, and expected to be published in the coming months. This section provides a brief overview of published and forthcoming results from the Burst Group.

Vela Pulsar Glitch [259]. In August 2006, during the first half of S5, the Hartebeesthoek radio observatory (HartRAO) observed a glitch in the radio emission of PSR B0833-45, the Vela pulsar. Both of the LIGO-Hanford detectors were operating continuously over the estimated on-source period of the glitch. As the mechanism responsible for pulsar glitches may also excite neutron star f-modes, where the main source of damping is gravitational-wave emission (see Section 4.1.5), we performed a Bayesian time-frequency analysis to search for an accompanying ring-down signal with frequency $\sim$kHz and duration ~ 100 ms. Collaboration with S. Buchner of HartRAO resulted in a refined estimate of the uncertainty in the glitch epoch allowing a more sensitive and computationally feasible analysis. The search found no evidence for a gravitational-wave ring-down signal; we set Bayesian 90% confidence upper limits of 6.3×10^{-21} to 1.4×10^{-20} on the peak intrinsic strain amplitude of gravitational-wave ring-downs, depending on which spherical harmonic mode is excited. The corresponding range of energy upper limits is 5.0×10^{44} to 1.3×10^{45} erg. This was the first search for gravitational-wave emission associated with a pulsar glitch and, by folding in the known orientation and expected emission pattern of spherical harmonic modes rather than assuming isotropic emission, allowed us to probe considerably lower than previous f-mode upper limits.

Six Magnetars [172]. The group has continued its efforts to search for gravitational-wave bursts from soft gamma repeaters (SGRs) and anomalous X-ray pulsars (AXPs), objects thought to be neutron stars powered by extreme magnetic fields, or magnetars (see Section 4.1.1). The analysis of LIGO, Virgo and GEO 600 data using 1279 electromagnetic triggers from six magnetars occurring between November 2006 and June 2009 (i.e., S5 year 2 and A5) was published in 2011. Four of the sources in this search were targeted for gravitational wave signals for the first time. These include SGR 0501+4516 and SGR 0418+5729, both of which are likely ~ 1 kpc from the Earth, an order of magnitude closer than objects targeted in previous searches. In addition, one of the targets gave an electromagnetic burst with estimated isotropic energy $> 10^{44}$ erg, comparable to giant flares. No evidence for gravitational-wave bursts was found in the data around the triggers. Upper limits on gravitational wave burst energy were set using simulated short-duration burst signals. The lowest upper limits for band- and time-limited white noise bursts in the detectors sensitive band and for f-mode ring-downs (at 1090 Hz) were $3.0\times10^{44} d_1^2$ erg and $1.4\times10^{47} d_1^2$ erg, respectively, where $d_1 = \frac{d_{0501}}{1\mathrm{kpc}}$ and d_{0501} is the distance to SGR 0501+4516. These limits are an order of magnitude lower than in previous magnetar searches and approach the range of electromagnetic energies seen in SGR giant flares.

GRB 051103. On 3 November 2005, one day before the start of S5, a short-duration gamma-ray burst was detected and localized to the outskirts of M81, estimated to lie at 3.61 Mpc from the Earth. The 2 km Hanford detector (H2) and the 4 km Livingston detector (L1) were up and recording data, with BNS horizon distances of ~ 4 Mpc and ~ 8 Mpc respectively. If the progenitor was indeed a compact binary coalescence *and* M81 was its host galaxy, the inspiral gravitational wave signal from GRB 051103 should be detectable in LIGO data. The event has been analyzed by the matched-filtering analysis pipeline used to search for inspiral signals from short GRBs in S5/VSR1 data, and also by Burst searches using reviewed S5/VSR1 versions of the coherent burst analysis packages used for GRB burst searches and SGR flare searches. The analysis is mature and internal review is well underway; both analysis and review are expected to be complete by the end of the year.

S6/VSR2+3 Gamma-Ray Bursts. During the S6/VSR2+3 run GCN notices reported approximately 150 GRBs detected by satellite-based gamma-ray experiments at times when two or more GW detectors were operating. A search for GW bursts using the same coherent analysis method used for S5/VSR1 [260, 261] has been performed. No evidence for GW burst signals was found. Simulations of short-duration ($<$ 1 s) waveforms set 90% confidence upper limits on the amplitude of gravitational waves associated with each GRB. These were then translated into lower bounds on the distance, assuming a fixed energy emission in gravitational waves. The median distance exclusion for short-duration signals in the detectors' sensitive band (around 150 Hz) is 17 Mpc $(E^{\mathrm{rot}}_{\mathrm{GW}}/0.01\mathrm{M}_{\odot}\mathrm{c}^2)^{1/2}$, where E^{rot} is the energy emitted in circularly polarised gravitational-waves from face-on rotating systems, as expected in astrophysical models for GRB progenitor systems. We also place distance exclusions on binary progenitors, with median distances of $\sim$ 7 Mpc for NSNS systems and $\sim$ 15 Mpc for NSBH systems. This analysis introduced several new features to GRB searches, including the ability to target poorly localized GRBs (from e.g., the Fermi satellite), a more powerful, weighted binomial population detection method and a novel population-based GRB-rate exclusion statement. The review is at an advanced stage and a joint publication with the compact binary coalescence working group is in preparation.

S5/VSR1 ANTARES GW+HEN. Cataclysmic cosmic events such as gamma-ray bursts can be plausible sources of both gravitational waves and high energy neutrinos. A new search for coincident gravitational wave bursts and high energy neutrinos is currently under review, using $\sim$200 neutrino triggers collected by the ANTARES neutrino telescope from February to September 2007. The GW analysis follows that used for the S6/VSR2+3 GRB burst search [260], with an on-source time window based on astrophysical expectations for the relative delay between neutrino and gravitational wave emission [214] in GRBs. This analysis is mature and currently under review.

S6/VSR2+3 All-Sky Search. The all-sky search for unmodeled gravitational wave bursts in S6/VSR2+3 was completed using the coherent analysis pipeline employed for S4 and S5/VSR1 [262, 263, 264, 222]. The search covered the 64 – 5000 Hz band and had a livetime of 202 days across all detector networks. The analysis is mature and its review is in progress.

Electromagnetic Follow-up. During S6/VSR2+3, low-latency analysis pipelines for both burst [262, 265] and inspiral signals [266] were used to identify gravitational-wave candidates and reconstruct possible sky locations. The locations were then sent to a network of optical, radio and X-ray observatories during two observation periods (17 December 2009 to 8 January 2010, and 2 September to 20 October 2010) with a latency of about 30 minutes. Simulations show that the median sky area for reconstructed signals with signal-to-noise ratio at the detection threshold to be between 10 and 100 square degrees. A methodological paper on the low-latency analysis and sky localization is in the late stages of collaboration review.

S5/VSR1 Intermediate Mass Black Hole (IMBH) Mergers. As discussed in Section 4.1.6, the gravitational-wave signal from the merger and ringdown of IMBH coalescence lies in the most sensitive frequency band of LIGO-Virgo. A first search for IMBH mergers has been conducted on S5/VSR1 data, focusing on IMBH binaries with a total mass of 100 to 450 solar masses and mass ratios between 1:1 and 1:4. The standard unmodelled search pipeline [262] is augmented with an elliptical polarization constraint to improve rejection of background glitches. No potential gravitational wave signal candidates were observed and an upper limit on the rate of IMBH mergers has been set at 0.6 $\mathrm{Mpc}^{-3}\ \mathrm{Myr}^{-1}$, averaged across all search masses, with a 90% confidence level.

4.3 Methods Used in Burst Searches

This section summarizes the approaches used to search for gravitational-wave bursts. We explain how we search for signals without knowing their form, how we select good-quality data, how we evaluate the background rate of false triggers from detector noise fluctuations, and how we estimate the sensitivity of our searches to plausible gravitational-wave signals.

4.3.1 Signal Extraction

Despite rapid progress in relativistic astrophysics, there is still significant uncertainty in predicted waveforms for most gravitational wave burst sources (see Section 4.1). Therefore, burst searches implement a variety of methods to find transients in the data that are inconsistent with the baseline noise level, and rely heavily on coincidence between multiple detectors to discriminate gravitational waves from noise fluctuations. 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 gravitational-wave burst events.

In a few special cases when an accurate signal model is available, such as for cosmic string cusps or neutron star or black hole ringdowns, a search can be done using matched filtering with a bank of templates. Otherwise, un-modeled bursts can be identified in the detector output data as excess-power localized events 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 [267] (Symlets, Meyer wavelets) and continuous wavelet transforms [268] (Q-transform). These transformations have been actively used in the burst search algorithms. At the same time, the Burst Group is open to other promising approaches such as the Hibert-Huang Transform (HHT) [269], an adaptive time-frequency decomposition which can be used for more detailed study of the TF content of gravitational wave events, and the fast Wilson-Daubechies transform. A handful other burst search methods which do not start with a TF representation have been implemented in past searches. These include a change-point analysis of bandpassed data [270] and a cross-correlation analysis using pairs of detectors [271].

Access to detector *networks* is especially important for the detection of gravitational-wave bursts. The identification of a consistent signal in multiple instruments will increase confidence in an event, especially for a burst signal which otherwise may not be distinguished from 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 classified as *incoherent* or *coherent*. Incoherent algorithms [270, 268, 272, 273] identify excess-power events in individual detectors; a time coincidence is then required between events in different detectors. Incoherent algorithms are particularly useful for detector characterization and they are actively used for studies of enviromental and instrumental artifacts (see Section 4.3.2). Coherent algorithms are based either on cross-correlating pairs of detectors [274, 275, 276] or on a more general coherent network analysis approach [262, 260].

Coherent network analyses address the problem of detection and reconstruction of gravitational waves with networks of gravitational-wave detectors. In these methods, a statistic is built as a coherent sum over the detector responses, which, in general, yields better sensitivity (at the same false alarm rate) than individual detector statistics. The development and implementation of more sophisticated coherent burst algorithms, in preparation for the advanced detector era, remains a high priority of the Burst Group. Several coherent approaches have already been adopted employed in past analyses, including a constrained likelihood method for untriggered searches [277, 262, 278] and a likelihood method for triggered searches [260]. In preparation for the analysis of advanced detector data, these methods need to be upgraded and tested both on initial detector data and engineering runs. The group is currently exploring a Bayesian formulation of coherent network analysis [279] (already used in the analysis of initial detector data [265]) and on maximum entropy methods [280]. Also dedicated coherent algorithms which may use partial information about burst sources and incomplete source models are being investigated by the Burst Group (see Section 4.4.4).

Coherent algorithms enable not only the detection of gravitational waves, but also astrophysical analyses of gravitational-wave signals and measurements of source properties, including reconstruction of the gravitational-wave waveforms and source coordinates. These are necessary tools for the nascent field of gravitational-wave astronomy. Prompt detection of gravitational-wave signals and estimation of source coordinates enables coincident observations with other astronomical instruments, which can significantly increase the confidence of detection [281]. Such measurements may not only aid the first detection of gravitational waves but also they will give us fundamentally new information about the sources and their distribution. The Burst Group has made significant progress in the development of source localization methods, which were extensively tested during the Position Reconstruction Challenge and used during the S6-VSR2/3 run. The group will continue to work on the source localization problem and apply the coordinate reconstruction methods to the advanced detector networks (see also Sections 4.4.3, 4.4.4).

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 desideratum 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 (see also Sections 4.4.3, 4.4.4).

.

4.3.2 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.

For the S6/VSR2+3 run, we followed the strategy of data quality and event-by-event vetoes originally developed in S5/VSR1 [282, 283]. Many people contributed to this work, though it is still generally manpower-limited. In a coordinated effort of Burst and CBC Group members, the Glitch Group and the Virgo Data Quality (VDQ) Group 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 and the VDQ 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 (a few percents) 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 4.4.7.

- Category 4 flags specifically tag the times where hardware injections were performed. In most searches they are to be used as Category 1 vetoes (before running the analysis pipeline) since they might affect the PSD computation.
- Category 5 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 [284, 265]: they are not tuned on multi-detector outputs, but once the tuning is complete, their effectiveness is tested on coherent network analysis 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 single detector triggers in auxiliary channels and event candidates, with a software package which considers many possible veto conditions in a hierarchical classification process. 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 [285].

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

The Virgo-GEO science data collected in the summer 2011 will be used by the Burst group to search for gravitational-wave signals associated with external triggers. Only the high frequency region will be considered (> 1 kHz) where the two detectors have a similar sensitivity. The DQ work to achieve will be similar to what has been done for the previous runs except that the attention has been shifted toward the high frequency region. The standard tools of LIGO/Virgo detector characterization are run at higher frequency to better monitor the band of interest. New DQ flags will be produced accordingly.

Despite the intense detector characterization effort, the burst search sensitivity in initial detector data is still dominated by glitches. As we prepare for the advanced detector era, new, more sophisticated techniques need to be developed and implemented for an effective mitigation of non-gaussian noise transients. This is a high priority both for the Burst Group and for the Detector Characterization team.

4.3.3 Accidental Background Estimation

The key to a burst search is the discrimination between real gravitational-wave signals and *false alarms* or *background*, noise fluctuations which satisfy the analysis selection criteria. The False Alarm Rate (FAR) or, alternatively, the False Alarm Probability for the observation time, depends on the detector noise properties and on all selection criteria in the pipeline, including coincidence or coherence tests among the detectors. The FAR is typically estimated with an *off-source* resampling of the observation data, equivalent to switching off any gravitational-wave signals.

In a network of detectors with independent noise, the off-source resampling is performed by *time-shifting* the data of detectors relative to each other by more than the maximum light travel time between the sites.

This destroys the coincidence conditions for any real gravitational-wave signal 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. This procedure is repeated many times with 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. If there is a known trigger time for the possible GW, the background estimation procedure can take advantage of it 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. For an objective, non controversial assessment of the FAR, we adopt a *blind* statistical procedure, where we use the off-source data (without examining the on-source data) to tune the procedures to compute the test statistics and their thresholds. Once the on-source has been disclosed, a second set of independent off-source resampling can be drawn to re-evaluate the false alarm rate. This helps avoiding possible biases from over-fitting to the first off-source sample.

This method cannot discriminate between gravitational-wave signals, other foreground signals or correlated noise sources at different detectors. If the accidental events are not compliant with a Poisson point process, it is not obvious how to evaluate the uncertainty of the empirical off-source distribution, which propagates into the uncertainty of the false alarm rate. Another problem with the time-shift method occurs if strong signal events in the detectors produce accidental coincident events in off-source resamples: these would result in a positive bias on the False Alarm estimates. This risk is mitigated in coherent network analyses: their stringent consistency checks between detector responses makes them more robust against this kind of bias than incoherent analyses. Up to this point we have seen no evidence that signal injections in the data can affect the result of coherent burst searches.

4.3.4 Simulations

Software signal injections in the data are used in burst pipelines to tune the analysis, to assess its *detection efficiency* once the selection criteria are fixed, and to interpret the analysis results against different signal models. Software simulations have been performed, in the analysis of initial detector data, using the *Mock Data Challenge* (MDC) technique, where special analysis runs are performed after adding simulated signals to the detector data at pseudo-random times. The simulated signals are chosen to span the expected range of signal properties (frequency, duration, etc.), but they do not exhaustively test all plausible signals, since 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. In addition to the more traditional random position distribution with uniform density in the sky, results are now interpreted also using models of galactic mass distribution and of extragalactic mass distributions. The selected waveforms has been routinely extended as well to improve the astrophysical interpretation, e.g. including elliptically polarized signal models and random distribution of the inclination angle of the source with respect to the line of sight.

Systematic effects of calibration uncertainties on the detection efficiency are measured by performing 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 used in the analysis of initial detector data is well-established, exploratory work is being done for a more flexible, on-the-fly simulation mechanism in the advanced detector data. Aslo, much work is still needed to expand the set of waveforms to include modeled or astrophysically motivated signals (see Section 4.4.4).

4.3.5 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 the safety of vetoes, i.e. testing the limits on the cross-coupling of loud GW signals into auxiliary data channels that might be used to define vetoes.

4.4 Science Goals

The Burst Group's fundamental goals are to detect gravitational-wave transient signals and use them to test the general theory of relativity, learn new things about astrophysical objects, and, in the advanced detector era, enable statistical analyses of source populations and emission mechanisms. The Burst Group's special role, of searching the whole space of possible transient GW signals, compels us to utilize every possible handle to distinguish real signals, even those whose form is unknown, from detector noise fluctuations. This section introduces seven broad goals which together form our vision of how to search for GW burst signals. These goals motivate the searches we implemented on data from the initial generation of detectors, and our planning for the advanced detector era.

4.4.1 Search as Broadly as Possible for Gravitational Wave Bursts

There is strong astrophysical motivation to search for burst-like gravitational-wave signals using ground-based laser interferometers [286]. The emphasis has historically been on astrophysical systems for which the resulting burst waveforms are either poorly modeled or 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 [287], as well as important new information on the features of signals accompanying core-collapse [196]. Burst sources with well-modeled waveforms include emission from neutron star ringdowns following a pulsar glitch, black hole ringdowns or cosmic string cusps (see Sections 4.1 and 4.4.4).

Typical signals predicted by these models last from less than a millisecond to hundreds of milliseconds, with signal power in the frequency range from 50 Hz to few kHz. Various models of gravitational-wave emission from core-collapse supernovae [196] and gamma-ray burst progenitors [288] may result in signals lasting up to several seconds. Although the best sensitivity of ground-based interferometers is achieved around $\sim$150 Hz, first generation 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. 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*[1]. 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 (see Section 4.4.7). 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 of S5/S6-VSR1/2/3 has been split into "low" ($\sim$50–2000 Hz) and "high" ($\sim$2000–6000 Hz) frequency bursts; such search over the full frequency spectrum is allowed by the instrument and motivated by source phenomenology (see Section 4.1). The analysis of the S6/VSR2+3 data set has been split into four intervals, dictated by changes in the individual detectors' and the global network's operational characteristics, and a single paper describing the search over the entire frequency regime and over the entire S6/VSR2+3 data set is in preparation. In preparation for the advanced detector era, the group will revise the set of simulations used for tuning and astrophysical interpretation of the results of this analysis, and the role of the all-sky analysis in presence or absence of detections.

In addition to this high priority search, the group is expanding analysis coverage to signals that last up to several seconds or minutes, with tailored modification to existing all-sky algorithms or new cross-correlation techniques. Some of these efforts are pursued jointly with the Stochastic Group 6.

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? Also, is the gravitational-wave burst sky uniform, and how can the ability of the detector network to reconstruct event sky positions be used to perform sky map comparisons between foreground and background? The implementation of a directional analysis and a study of directional anisotropy can in principle improve sensitivity and provide hints for possible astrophysical models behind candidate signals. These techniques have been used routinely by cosmic ray experiments in the search for astrophysical sources of high energy photons and neutrinos [289, 290, 291, 292].

4.4.2 Look for GW Transients Triggered by Energetic Astrophysical Events (ExtTrig)

Many of the plausible gravitational-wave sources introduced in Section 4.1 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 a burst search compared to the untriggered all-sky search described in Section 4.4.1. The association with a known astrophysical event may be critical in establishing our confidence in a candidate GW burst detection. Perhaps most importantly, joint studies can enable scientific insight that cannot be accessed through gravitational waves or other messengers alone.

As we enter the advanced detector era, it remains a high priority for the Burst Group to build bridges to

[1] It should be noted that many of the model-dependent searches (see Section 4.4.4) are effectively "all-times, all-sky" ones with the difference, of course, that they optimize their sensitivity to specific waveform models.

the astrophysics and astronomical communities. Past successes indicate that these connections are critical for extracting the most astrophysics from advanced detector data, as discussed also in Section 4.4.4. In this section, we list the externally triggered searches the group is focusing its attention on, and will continue to pursue in preparation for the advanced detector era. Some of these goals are common with the CBC working group (see Section 3) and the stochastic working group (see Section 6).

Triggers from Gamma-ray bursts (GRBs)

The Burst Group has a long history of searching for GWBs associated with GRBs, both long and short [293, 294, 295, 296, 275, 261]. 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 or weakly modeled gravitational wave transients. The GRB burst analysis is complementary to the compact binary coalescence search, providing the ability to detect or constrain any gravitational wave emission that does not precisely match the inspiral waveform, or is advanced by a few minutes of the GRB trigger. Even non-detections can be significant; for example, the analysis of GRB 070201 supported the hypothesis that the event progenitor was an extragalactic SGR giant flare [294].

GRB triggers are collected primarily from the GRB Coordinate Network (GCN, [139]), and from the Third Interplanetary Network (IPN, [297]). While GCN notices are released by email in real time, the localization of IPN GRBs requires manual effort by IPN scientists and therefore these triggers are obtained via collaboration with the IPN with a time lag of months to years. As a result, GCN and IPN GRBs are analysed separately. The Burst Group will analyse all available GCN/IPN GRBs from S5/VSR1, Astrowatch, and S6/VSR2-3 periods for which we have coincident data; current efforts focus on the ~ 50 GCN Astrowatch GRBs and ~ 300 IPN S5/VSR1 GRBs. Ongoing collaboration between LIGO/Virgo and the IPN should streamline IPN trigger generation in the advanced detector era.

Additional GRB triggers are generated by analysis of data from the *Swift* satellite. This returns candidate GRBs which are below the nominal *Swift* trigger threshold. Some of these GRBs may be be relatively near the Earth, and appear as low luminosity due to beaming or from being members of a low-luminosity subclass of GRBs [298, 299]. However, many of these triggers are likely not GRBs.

Recent improvements to the GRB analysis include the ability to handle the large or irregularly shaped error boxes associated with Fermi, sub-threshold, and IPN GRBs, as well as incorporation of prior information on the expected GW polarization due to GRB beaming.

An alternative approach is to search collectively over a population of GRB triggers, to infer properties of the population of astrophysical sources as a whole rather than any one member [300]. An important feature of the method is a monotonic improvement in the astrophysical constraints obtained as the trigger sample increases [301, 300]. This strategy has been successfully used to improve upper limits on GW emission from GRBs [275]. Future 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, and possibly other information obtained from the non-GW observations.

Triggers from Magnetars

ExtTrig searches look for GW associated with bursts from the magnetar candidates, soft gamma repeaters (SGRs) and anomalous X-ray pulsars (AXPs) currently following three distinct search strategies.

The first strategy looks for GW emission associated with the prompt gamma ray burst from individual magnetar bursts and giant flares. An emphasis is placed on GW emission from the damping of f-modes in the magnetar which may be excited in the burst, but the entire detector band is searched up to 3 kHz. This strategy has been used to analyze over 1400 electromagnetic SGR and AXP triggers, including the 2004 giant flare [276, 172].

A second strategy looks for the collective signature of weak GWBs from repeated flares from a single

SGR source, stacking the GW data to dig deeper into the noise at the expense of additional but plausible model dependence. This strategy was used to analyze the 2006 SGR 1900+14 storm [302].

The third strategy looks for long quasiperiodic GWs associated with quasiperiodic oscillations (QPOs) observed in the electromagnetic tails of the giant flares, with special emphasis on frequencies and durations matching those seen in the electromagnetic emission [303]. This strategy was used to analyze the 2004 giant flare [304], and is being applied to an ensemble of regular flares as well, using observed QPO frequencies from more energetic bursts in a statistical search.

High-Energy Neutrinos as Astrophysical Triggers

High-energy neutrinos (HEN) could be emitted in some scenarios such as *choked* GRBs (see Section 4.1.2). The IceCube [305] and ANTARES [306] HEN 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 HEN events. These detectors can also reconstruct the energy of individual neutrino events [307], which can enhance background rejection as the spectrum of astrophysical HENs differs from that of atmospheric events [307]. Joint GW+HEN searches benefit from reduced false alarm rates and would allow the study of astrophysical objects via both messengers synergistically. For example the relative times of arrival of GW and HEN signals can be important in reconstructing some of the properties of the astrophysical source [214].

The feasibility and expected performance of a hypothetical IceCube-LIGO-Virgo network was demonstrated [308, 309], 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.

Supernovae as Astrophysical Triggers

Core collapse supernovae are interesting candidates as gravitational-wave sources and can be studied in conjunction with both neutrino and optical messengers (see Section 4.1).

Most optical triggers carry the burden of a large uncertainty on the derived event time (order of several hours or more), making the GW data analysis task challenging. Well-known sky locations are a significant aid to the analysis. An optical supernova search in initial detectors data is underway, and might be able to constrain the most extreme core collapse models. The enhanced reach of advanced detectors will be able to constrain a more significant swath of the model space.

Supernova triggers from detectors sensitive to low energy (tens of MeV) neutrinos can be used in GW searches as well. For example, a core-collapse supernova near the galactic center is expected to produce a large flux of ~ 8000 detected neutrinos in the Super-Kamiokande detector [310] with a pointing accuracy of 4°. Unlike photons, which can take up to $\sim$day to break out, neutrino bursts and gravitational waves mark the moment of core collapse, and are expected to be coincident in time to < 1 s. 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. With cooperation from Super-K and other detections, we hope to participate in joint GW-neutrino supernova searches. However, the agreements are still under negotiation at the time of writing.

Pulsar Glitches as External Triggers

A search for a GWB associated with the August 2006 Vela pulsar glitch has been performed [259], using a Bayesian model selection algorithm [311]. This search strategy could be applied to other pulsar glitches coincident with initial detector data, and also to pulsar glitches occurring in the advanced detector era.

The excited post-glitch pulsar can also emit gravitational waves at r-mode frequencies (of order tens of Hz for most glitching pulsars).; preliminary studies on the feasibility of detecting these signals have been carried out, and a feasibility study is being performed on anomalous X-ray millisecond pulsars (AMXPs) as potential sources of the same r-mode emissions. A search strategy is now being formulated with members of the Continuous Waves working group to search for these long-lasting transient signals.

Triggers from the low mass X-ray binary (LMXB) Sco X-1
The Rossi X-ray Timing Explorer (RXTE) mission has been extended to the end of 2010 and observations therefore overlapped 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. 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. We propose to perform radio-triggered follow up searches in our data, starting with existing radio transients detected during S5/VSR1 and S6/VSR2-3. 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, allowing us to better predict when the gravitational wave should have arrived at our detectors and also to estimate the absolute gravitational-wave luminosity of a detected event.

Sub-second radio pulses without an identified cause are present in a number of pulsar searches, including the 350 MHz drift scan and northern celestial cap surveys using the Robert C. Byrd Green Bank Telescope [312] as well as 1.4 GHz surveys conducted with the Arecibo Observatory by the ALFA Pulsar Consortium [313]. Some plausible origins of these radio transients, e.g. compact binary coalescence scenarios or cosmic string cusps [314], could lead to an identifiable gravitational wave signal accompanying the radio burst. We will use a coherent network analysis [260] to search for gravitational waves near the location and (dispersion measure adjusted) arrival times of these radio pulses. The search will resemble burst GRB analyses in methodology, but with parameters such as on-source time window, frequency coverage, veto method, and set of software injections adjusted as appropriate.

4.4.3 Burst Parameter Estimation

The detection of an un-modeled burst of gravitational wave radiation will present the astrophysics community with a wonderful mystery to solve: What produced the signal? Answering this question will require an assessment of the degree of confidence that one process or another was at work, and learning as much as possible about the underlying dynamics of the system from the gravitational wave data.

The process of characterizing gravitational wave signals for modeled systems such as binary black hole mergers is fairly straightforward and well understood (see for instance [315]). Physical parameters such as the mass and spin of the component black holes affect the waveforms in a known way, and by comparing the parameterized waveforms to the data it is possible to develop posterior probability distributions for the model parameters. The CBC group has a Bayesian inference working group that is actively pursuing techniques such as Markov Chain Monte Carlo and Nested Sampling to perform parameter estimation and model selection studies for modeled signals.

The problem of burst parameter estimation for un-modeled signals is more difficult, and has received less attention. At first sight the question may seem ill-posed since the techniques used to search for un-modeled signals generally do not use physically parameterized models for the signals. For example, the output of burst search may be a list of best-fit wavelet amplitudes for the reconstructed signal and an estimate of the sky location. However, from this output it is possible to derive physical quantities such as the duration, rise and decay times, peak frequency, frequency band containing 90% of the power, etc.. The choice of physical quantities is flexible, and can be tailored to address particular astrophysical questions. In addition to characterizing the time-frequency and energy content of the signal, the technique of coherent network analysis allows the sky location and polarization content to be estimated.

The goal of burst parameter estimation is to go beyond finding best fit point estimates for the quantities that characterize a signal, and to produce full posterior probability distributions for the parameters. It is only when the spread in the parameter estimates are available that it becomes possible to meaningfully compare the predictions of different astrophysical models.

One approach to burst parameter estimation currently under development is a Bayesian method [316], where wavelets are used to model both the network-coherent gravitational wave signals and instrument glitches. A transdimensional Markov Chain Monte Carlo algorithm is used to find the most parsimonious model (both the amplitude and number of wavelets used to model the excess power are allowed to vary), and full posterior distributions characterizing the signals and glitches are produced. The algorithm naturally incorporates Bayesian model selection, and provide odds ratios for the detection and non-detection hypotheses. Algorithms like this may prove useful for automated detector characterization, and will likely be useful in assigning a degree of confidence in a putative detection. An additional benefit of using Bayesian inference in the analysis is that it is easy to incorporate priors, including strong priors on the signal morphology that can be used to produce targeted searches.

There are also efforts under way to incorporate information from the myriad of gravitational waveform catalogues from numerical relativity simulations of supernovae. The main method currently employed is to decompose all waveforms in each catalogue into an orthonormal set of basis vectors or Principal Components. These Principal Components are then used in a Markov Chain Monte Carlo (MCMC) to reconstruct a detected or injected gravitational wave signal. This approach has been shown to correctly reconstruct supernova signals in simulations [227]. It also allows one to infer the physical parameters of the supernova progenitor based on the best fit waveform from the applied supernova waveform catalogue.

In addition to supernova signal reconstruction and parameter estimation, a Bayesian model selection pipeline is being developed to distinguish between supernova waveforms from different catalogues. This pipeline uses the Principal Components described above and calculates the evidence that the signal belongs to a catalogue using Nested Sampling. Since each waveform catalogue is a result of simulations that evolve different supernova models, the ability to associate a detected gravitational wave signal with one catalogue will allow us to infer the astrophysical processes behind the observed supernova. Initial tests of a prototype model selection pipeline on simualted Gaussian noise have shown that it is indeed able to distinguish between waveforms from different catalogues.

The MCMC analysis has been extended to run on data acquired by multiple detectors. This will allow for better signal reconstruction and the ability to make probabilistic statements about the location of the signal source on the sky. A similar enhancement is being developed for the model selection pipeline.

4.4.4 Use Knowledge of Astrophysically Plausible Signals

Burst searches are un-modeled by definition, to be open to the broad spectrum of sources discussed in Section 4.1. However, in the era of gravitational wave astrophysics we are committed to make the best use of current knowledge in relativistic astrophysics, to increase our discovery potential and our ability to extract scientific content from GW data. Recent progress in astrophysics and numerical relativity provide information on the intensity and morphology of expected gravitational wave signals from some sources, such as the coalescence of compact objects over a broad parameter space, core collapse supernovae and neutron star oscillations.

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 [317], 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 $\sim$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 [318]. 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 [196]. The Burst Group aims to exploit such knowledge for an astrophysical interpretation of its results.

Burst searches do not typically use calculated waveforms as templates for matched filtering (with some notable exceptions, such as cosmic string cusps and neutron star ringdowns). 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. An ongoing collaborative effort with the CBC group is studying the use of full coalescence waveforms, including the inspiral, merger, and ringdown phases. The same set of phenomenological and EOBNR waveforms [124, 82, 319, 320] is being analyzed by burst search methods, inspiral and ring-down matched filtering to compare their detectability 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 for systems with 25-350 $M_\odot$ total mass and (anti)aligned 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 [321, 251]. 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, as discussed in section 4.4.3. Progress towards waveform reconstruction has already been made via coherent techniques, but more progress is needed to compare a candidate to waveform parameters. Bayesian [322] and MCMC [323] techniques are currently being explored by the CBC Group as well as members of the Burst Group. We are open to exploration of new techniques which may prove useful for reconstructing waveforms, and the Burst Group expects to increase its efforts in this direction in the coming years.

Finally, 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 COHERENTWAVEBURST to search for elliptically polarized signals such as black hole mergers [185]. Or the configuration file for OMEGA could be modified to target high frequency, long waveforms from certain models for core collapse supernovae [324, 325, 326, 327, 328]. Matched filtering can be used for some classes of GW burst signals, such as cosmic string cusps [257] 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.

Aside from the use of astrophysically plausible GW waveforms, the Burst Group pursues other kinds of customized analyses which exploit the source sky locations. We 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 4.4.2 for a full discussion. Dedicated searches have been designed 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 [292], 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*.

4.4.5 Seek Confirmation of GW Candidates with Follow-up Observations

Section 4.4.2 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 [329]; 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 [330, 331]. Thus there may be short-lived "orphan afterglows" from nearby energetic events which are going undetected [181]. 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 [332, 333, 192] and supernovas [197]. More details on these sources and their expected observational signatures can be found in [334, 329, 335].

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 [336, 337, 338].

The scope of the Burst effort has evolved over time from optical observations [334] to include other EM bands and to be a joint effort with the CBC Group, who wish to follow up promising candidates from (well-modeled) inspiral searches. During the S6/VSR2+3 run, we successfully implemented and operated a program of rapid electromagnetic follow-up observations using several ground-based optical telescopes—most with fields of view of 3 square degrees or more—the LOFAR radio telescope, and the XRT (X-ray) and UVOT (UV/optical) instruments on the *Swift* satellite. Although the total live time of the project was only about 7 weeks and only 9 triggers were followed up during that time, it was successful as a demonstration of the methods and basic capabilities, and a great learning experience to build on for the advanced detector era.

As of June 2011, analysis of the images obtained in the S6/VSR2+3 run is still in progress. Two basic approaches are being used: image subtraction (a target image minus a reference image should reveal a variable object), and object identification followed by comparing object lists extracted from target vs. reference images. Several collaboration members have gained experience with image processing tools with some amount of help from our professional astronomer partners. We are currently working on a paper presenting the *Swift* follow-up results and will write a paper about the ground-based results when the image processing has been completed.

In addition to the rapid-response EM followup mentioned above, there is also a wealth of all-sky high-energy photon survey data which can be searched offline for EM counterparts to GW events. Joint search pipelines are currently in development to search specifically for GW-triggered EM counterparts in Fermi GBM (20 keV$-$40 MeV) and RXTE ASM (1$-$10 keV) archival data. This is done in coordination with scientists from both missions. The search targets prompt gamma-ray emission from GBM which may be below threshold for public reporting to the GCN, as well as x-ray afterglow signals in the ASM.

It is already time to start planning for EM follow-up observations in the advanced detector era, when the astrophysical motivation will be even stronger since event candidates will be more likely to be real GW signals. We need to re-think our relation with astronomers—will we continue MOU-based arrangements with selected astronomers, or will we make triggers publicly available from the beginning of the run?—and what set of instruments (optical including provisions for spectroscopic follow-up, radio, X-ray) are needed. We'll need to re-engineer the many technical steps, from data transport and data discovery to the online search algorithms, usage of data quality information, selection of significant triggers, communication with telescopes (hopefully standardized), monitoring and/or coordination of telescope observing, and faster/better image analysis. Although the details can be worked out over the next 3 or 4 years, the general plan should be in place by the middle of 2012, with sufficient flexibility to adjust to any changes to the LSC-Virgo policies on open data.

4.4.6 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 [339, 340, 341, 342]. 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 [343]. The LSC-Virgo network 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 [344]. Coordinated electromagnetic and gravitational wave searches could place 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 [344], and the most general gravitational wave can have up to six polarization states [339, 340]. For instance, Brans-Dicke theory and other scalar-tensor theories predict an additional scalar transverse polarization state. 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 [340], 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.

4.4.7 Be Prepared to Detect a GW Signal with Confidence

Confident detection of gravitational wave signals is a challenging problem in burst searches. Burst searches aim detection of unmodeled GW signals and therefore they are more affected by a non-stationary detector noise than modeled (template) searches. Burst searches use various tests and techniques to identify whether a candidate event should be considered a genuine gravitational wave detection, including:

- consistency tests among coincident signals from different interferometers, usually performed in a context of the coherent network analysis,
- extensive detector characterization studies aiming to link event candidates to known instrumental and environmental transients,
- extensive background studies to establish the significance of observed candidate events.

All these three directions of research were in the focus of the burst group analysis of data collected by the initial GW interferometers. Powerful consistency tests have been developed and used in most of our search methods. Section 4.3.2 outlines the data quality studies and vetoes that have been used to handle various sources of glitches. Every burst search needs to carefully understand its background. There have been significant improvements and advances in the burst algorithms during the analysis of the initial data, and yet, they are not sufficient for a confident detection of expected gravitational-wave signals. During the data runs of initial detectors there were two blind burst injections, which have been promptly discovered and reconstructed by the burst algorithms. However, due to excessive non-stationary background noise the burst search could not identify these events as gravitational-wave signals with high confidence.

In preparation for the analysis of data from advanced detectors, and to solve this extremely hard problem of non-stationary background noise, the Burst Group identifies several directions of research aiming a confident detection of low rate burst gravitational-wave signals.

Coherent network analysis: this approach has been already used in burst searches and proven very effective; it require further development to fully utilize capabilities of future advanced detector networks.

Inclusion of models into burst searches: one of the main priorities in the burst analysis (see section 4.4.4), the inclusion of models (which may not be accurate) helps to divide assorted un-modeled events into wide weakly-modeled classes with significantly reduced background. For example, such classification of burst triggers into different polarization states has been already used in the analysis of initial detector data. Other models targeting particular classes of burst signals can be also developed and used in the analysis. Such weakly-modeled algorithms are quite computationally intensive and also a development of robust source models will require a close collaboration with a wider astrophysical community.

Improved statistical analysis is required to combine different runs and searches in order to establish a significance of observed candidate events. For example, during the S6/VSR2+3 burst analysis there were 4 different time epochs and four different network configurations with significantly different rates of background events, for a total of 16 analysis configuration. All these configurations need to be combined in a single measurement with a clear statement of statistical significance of the candidate events. Such statistical approaches based on likelihood, false alarm density, and other methods are under development in the Burst and CBC groups. They need to be implemented and tested, and perhaps other, more advanced statistical algorithms, need to be developed.

Advanced background studies need to be performed for better understanding of low rate non-stationary tails in the background rate distributions. Burst searches are capable to accumulate hundreds of years of effective live time of the background sample (thousands of time lags). However this is not sufficient for an accurate estimate of the false alarm rates. New approaches for accumulation of million lags need to be developed.

Detector characterization studies is one of the most important activities in the Burst Group. Burst Group members actively participate in detector characterization studies (as part of the LSC Detector Characterization group) and the burst algorithms are used for identification of data quality flags and vetoes. During the analysis of initial detector data, the burst group also developed a "detection checklist" - a follow-up procedure for candidate events. We will continue to work with members of other LSC working groups to improve the identification of spurious events in the detectors. New DC approaches need to be developed to make a breakthrough in this area and to dramatically improve the efficiency of the data quality selection cuts and vetoes.

Best use of the initial detectors data set: We do not need to wait for advanced detector data to improve and test advanced analysis algorithms, new statistical procedures and novel data quality methods which are required for a confident detection of burst GW signals. All this work can be done on the existing data set, before a data from advanced detectors is collected. Detection challenges and tests of confident detection approaches can be performed as a part of the software engineering run plan.

Efficient use of trigger event properties and source distribution: Beyond perfecting our GW search methods, it is also crucial to understand and incorporate information from other messengers in the analysis. The anticipated source distribution based on the available galaxy catalogs is a good example where we can significantly enhance our sensitivity by weighting results with the probability distribution of potential sources. For external triggers, trigger properties beyond source direction (e.g. source distance for GRBs or neutrino energy for HEN searches) can be an important addition to the information used in the analyses.

4.5 Using Data from the Transition Era

With the LIGO detectors shut down for the installation of advanced LIGO and with the imminent shutdown of Virgo, the GEO detector will play a special role over the next three years as the only large-scale interferometer in regular operation. In the short term, coincident GEO-Virgo data from the 2011 summer run will be searched for high-frequency GW bursts associated with external triggers. In the longer term (post

September 2011) GEO-HF will operate alone in astrowatch mode, to be ready in case of an exceptional astrophysical event such as a galactic supernova.

For the **GEO-Virgo S6e/VSR4 run** (June-Sept 2011), the two detectors have similar sensitivity at frequencies near and above 1 kHz. This is the frequency range of interest for signals from perturbed compact objects, such as neutron-star f modes (1-3 kHz). The latter may be excited in SGR flares, making large or nearby flares potentially interesting targets for an externally triggered search on S6e/VSR4 data [170, 302, 172]. Similarly, a galactic core-collapse supernova may give detectable signals and would be of tremendous interest. A gamma-ray burst (GRB) directionally consistent with a nearby galaxy, such as GRB 070201 [294], would be another interesting target. The Bursts group therefore intends to conduct a search for purposes of publication if such a trigger is identified in the S6e/VSR4 time.

GEO in the "dark ages" (2011-2014) will serve as the only large interferometer in regular operation. It will be particularly challenging to have confidence in a GW candidate found by a single detector, so work with the GEO-HF data will need to focus on the development of new and aggressive background mitigation methods. Such methods should be developed with an eye to their application to the advanced LIGO and Virgo detectors when they come online.

Estimated effort neded: FTE requirements are included in the relevant sections on externally triggered searches (Section 4.4.2) and detector characterisation (Section 4.3.2).

4.6 Preparing for the Advanced Detector Era

The Burst Group aims to extract a broad range of science results from early data of the advanced gravitational wave detector network, by building on our online and offline analysis experience and the data analysis infrastructure developed for the initial detectors. The analysis for S6/VSR2+3 first science targets is approaching completion, with several papers in preparation, and more papers on other searches are expected over the next two years. In the meantime we expect the group will devote significant attention to preparations and infrastructure developments for the advanced detector era.

First Science Targets

We plan to be able to perform rigorous and sensitive core analyses on "day one" of the advanced detector era, including being ready to detect and fully characterize a GW signal when it arrives. Specifically, the Burst Group is committed to deliver the following papers:

1. An all-sky statement on generic gravitational wave bursts: upper limits if there is no detection, a rare-event detection significance if we have one candidate, or population studies if we have several detections.

2. Prompt analysis for electromagnetic followup of early detections.

3. Prompt reports on interesting astrophysical triggers, such as GRBs, SGRs, and supernovae.

In preparation, the group will pursue necessary ingredients such as data access, data quality, simulations, interpretation and statistics, many lags, as outlined in Section 4.3 and summarized in Section 4.4.7.

Optimize the Burst Science Output

Beyond simply detecting gravitational-wave bursts, we wish to study astrophysical sources such as black hole mergers, core collapse supernovae, perturbed neutron stars, eccentric binaries, gamma-ray burst engines, cosmic string cusps, joint emitters of GWs and neutrinos, as well as test alternate theories of gravity 4.1. Early planning in the next few years will allow us to think deeply about these science goals and how to prioritize and address them. Searches will be embraced by the group once they have proven their potential (with astrophysical motivation and a viable implementation plan). Simulations and astrophysical input will help decide whether a specific source requires a targeted analysis or a different tuning and interpretation of the standard all-sky and externally triggered searches. Over the next 3–5 years, the group will continue to 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. To accomplish that, the analysis teams will need to begin their work rather early and establish milestones for the development and testing of analysis software and analysis techniques, as well as documenting and pre-reviewing them. The planned engineering runs provide a natural context for pacing the work and for testing the analyses under realistic conditions, eventually including real AD engineering-run data. Here are the tasks to start working on now:

- Identify the first science targets for the AD era – the analyses we want to be sure to complete as soon as possible. They won't be all the searches we will ultimately do, but these are the ones which will be under scrutiny to make sure they are on track to succeed in good time.
- Form teams for each analysis, who will consider the science goals and determine what specifically is needed to meet them. Specifically, the team should consider the likely signals; what astrophysical interpretation(s) should be pursued for a non-detection or for a detection; the suitability of existing analysis tools and techniques; what specific improvements need to be made; and what new tools or techniques (if any) need to be developed.
- The team should lay out a timeline with milestones for being ready for the first advanced detector data, using the engineering runs schedule to plan on testing, review accomplishments and report to the whole group.

External Collaborations

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. These interactions remain of high priority for the group; the specifics of the format will be affected by the collaboration's resolution of open data policy.

Data Quality

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 closely coupled with the detector characterization effort, with at least one, or more, liaison 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.

Software

Future needs for hardware computing infrastructure will 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. The development of new software will be coordinated with DASWG; the group encourages making best possible use of existing and validated code. We remain open to better ways to do things but aim to avoid unnecessary duplication of effort, maximize efficiency, recognize practical development of new tools.

Organization

The scientific goals of the burst group will be pursued in working groups, each with a defined goal, with updates in the weekly group teleconference. In addition to the weekly working group updates, each analysis project is encouraged to provide an in-depth report to the burst group on a quarterly base, more frequently if the need arises.

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.[2] These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [349, 350, 352], magnetic deformations [351, 355], unstable r-mode oscillations [353, 349, 354], and free precession [359], all of which operate differently in accreting and non-accreting stars. We present a review of these emission mechanisms in [348]. Indirect upper limits on gravitational wave emission inferred from photon astronomy are more optimistic for non-accreting stars, but according to present theories accreting neutron stars are more likely to be emitting at or near the indirect limits.

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

The breadth of investigation is fundamental to our search method. Given the large uncertainties in neutron star demographics (only $\sim$2000 of 10^8-10^9 neutron stars in the galaxy have been detected), evolution, and structure, we cannot confidently predict which type of source will provide our first continuous-wave discovery. Prudence demands an eyes-wide-open approach and enough flexibility to exploit unexpected waveform models. 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 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 (see section 5.6.1). 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

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

need search only a small parameter space and are not computationally limited (see section 5.1.1 below). 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 spindown 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 [348]

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

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 "spindown limit" by a factor of 7 for the Crab pulsar (59.45 Hz) and by $\sim$40% for the Vela pulsar (22.38 Hz). Other pulsars for which the spindown limit may be reached in S6/VSR2/VSR4 include PSRs J0205+6449 (30.45 Hz), J1833-1034 (32.33 Hz), J1813-1749 (44.74 Hz), J1913+1011 (50.59 Hz), J1952+3252 (55.69 Hz), J0737−3039A (88.11 Hz) and J0537−6910 (123.95 Hz) [356].

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 spindown 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.

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

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

5.1.1 Time domain Bayesian method

The time-domain Bayesian method has been applied successfully to data from the first five LSC science runs [345, 346, 365, 363, 364] and to the Virgo VSR2 run. It is also currently being used for searches in S6 data and will be used for searches of VSR4 data. A detailed discussion of the method can be found in [362], with the implementation of the inclusion of binary system parameters in [366].

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 spindown limit. We are also compiling timing data on new γ-ray pulsars discovered with the Fermi satellite.

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 led 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 [365, 366]. 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 [363] (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 spindown 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 spindown limit by an extra factor of 1.3. This result has allowed us to constrain the amount of the pulsar spindown power budget released by gravitational radiation to be less than about 2%.

Results from a search for the Vela pulsar in VSR2 data have been obtained for both uniform and restricted priors on its orientation. These results beat the spin-down limit for Vela and have been accepted for publication [408], 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/VSR4 data sets for all accessible known pulsars with available precise timing.

In the longer term, the time-domain method will be used in a full S6/VSR2/VSR4 search for known pulsars, searching not only at the nominal $2f_{\rm rot}$ frequency, but also at $f_{\rm rot}$ (using the model described in [407]), and potentially in the vicinity of $(4/3)f_{\rm rot}$ for *r*-modes, as part of a broader effort to widen the parameter space in targeted searches. An upgrade to the current MCMC parameter estimation code is being implemented that makes use of the nested sampling algorithm [406]. In addition to providing marginal posterior probability distributions for signal parameters this will provide a signal vs. noise odds ratio that can be used as a detection statistic. The algorithm will also be more robust when applied to wider band searches. A methods publication describing and characterising the nested sampling algorithm, the odds ratio and its use for deriving limits simultaneously from $f_{\rm rot}$ and $2f_{\rm rot}$ searches is planned for the coming year. In addition, more use will be made of Fscans (discussed below) and other diagnostics to optimize searches for the Crab and Vela pulsars.

5.1.2 Narrowband 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 uncertainty 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 complement 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, whereas 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 [363]. 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 spindown 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 methods paper comparing the nested approach above with the templated F-statistic approach will be submitted in the coming year. A paper describing the results of these wider searches in S5/S6/VSR2/VSR4 data is expected to be submitted for publication in the following year. In the longer term, the narrowband search method will be explored as a follow-up step to semi-coherent all-sky searches.

5.1.3 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[371]. The method is based on the computation of the analytical signal [367]. 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, spindown 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.

Results from a search for the Vela pulsar in VSR2 data have been obtained both with and without constraints on orientation. These results have been 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/VSR4 data sets for all pulsars with reachable spindown limits, including searches of narrow bands surrounding 2*f, with publication expected in the following year.

5.1.4 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 [394]. 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 [395]. 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 was beaten in VSR2. For this analysis we are using updated ephemeris data provided by radio-astronomers (Hobart & HartRAO) 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 [377]. Final results of the VSR2 have been submitted for publication, along with those of the other two targeted pipelines described above. In the coming year searches will be carried out in S6/VSR2/VSR4 data for all known pulsars with accessible spindown limits.

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 spindown parameters. Since we must search over many frequencies and spindown 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 spindown 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 24 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 [396] and shortly will start searching for other interesting locations on the sky, such as star-forming regions and globular clusters. 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 5.2.7).

The first search for a source with no timing (Cas A) used the $\mathcal{F}$-statistic code with a single integration of $\mathcal{O}(10)$ d. Our estimate of computational cost and sensitivity [397] 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/S6 data sets (and future sets of comparable length) rather than $\mathcal{O}(10)$ d.

A similar but deeper search was carried out for Calvera, which is an X-ray source originally detected by ROSAT and confirmed with Swift and Chandra measurements. Until fall 2010 it had no detected optical or radio counterpart and was 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 was carried out for Calvera over the 90-360 Hz band [399], assuming an age of O(10^6) years which severely restricts the spindown range to be searched. Because the sky location was known, the search benefitted dramatically in reduced computational cost from the resampling step during preprocessing. Preliminary upper limits were obtained in summer 2010, but the subsequent observations of pulsations in X-rays [398] outside the sensitive LIGO band deterred pursuing publication. Nonetheless, the search served as a useful benchmark for directed, deep searches for old objects and as a proving ground for the barycentric resampling method.

5.2.1 Coherent directed searches

The coherent search for Cas A will be extended to top targets from each category of directed search (see Sec. 5.2.7), 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.

5.2.2 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 near 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 $2\,\mathrm{deg} \times 0.8\,\mathrm{deg}$ 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 direction of the galactic center involves searching over a small sky-region but over large frequency and spindown ranges.

The entire S5 data has been scanned using the `Einstein@Home` infrastructure. This search is meant to be more sensitive over a different region of parameter space – it searches a larger range of spindowns and a narrower frequency band – and uses a larger data set.

A single sky position is used for this search, a single template pointed directly to the galactic center; the frequency band is 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.

Production running has finished, and follow-up of outliers is under way. The search will be completed and final results obtained for publication submission in the coming year.

5.2.3 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 [411], can be found in [385]. 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 $T_{\rm sft}$ covering a total observation time $T_{\rm obs}$ followed by: i) a choice of the minimum coherence time $T_{\rm coh-min}$ 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 $T_{\rm coh-min} \approx 1\,{\rm hr}$. 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 $T_{\rm coh-min} \approx 1\,{\rm hr}$, 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.

5.2.4 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) [369]. The removal (or the duplication) of the samples takes place typically each few seconds. The list of samples to be removed or duplicated (named *mask*) is thus not huge and can be easily computed by simple geometrical consideration. As detailed in [369] 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 spindown 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 spindown 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 spindown 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 spindown).

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 spindown. 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, spindown values, sampling frequencies and source frequencies. In all cases, the expected phase-locking and peak reconstruction accuracy has been found. Similar tests have been performed injecting signals in the VSR1 data. All the results are described in Torre's graduation thesis [386], or (more in summary) in [387]. The resampling of the data requires less than 10^{-5} of the investigated time (for a single direction and spindown on a few Hz band), that is negligible with respect to the time required to read the HF stream (of the order of 10^{-4}). A method to read the data and apply the resampling technique directly to the down-sampled data (making negligible the computing time for reading data) is in progress. The amplitude modulation will be taken into account using a matched filtering in the frequency domain, in a way similar to the one developed by the Rome group. 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 methods paper was recently published in Phys. Rev. D [401]. Correction for amplitude modulation is under study, and an exploratory search for sources near the galactic center has begun. A search for Cas A will begin in the coming year, with a results publication expected in the following year.

5.2.5 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 ~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 2011. A methods paper has been accepted for publication [402]. In the coming year, the pipeline will be used to study sensitivity to pulsar post-glitch transients, and a benchmark search of S5/S6 data will begin.

5.2.6 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. 5.3.7). 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.

5.2.7 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. 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 Fermi point sources and HESS TeV gamma-ray sources that are followed up with Chandra and XMM-Newton X-ray observations to yield pulsar wind nebulae and sometimes the neutron stars themselves.

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

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 [348]. Note, however, that a more recent simulation analysis finds significantly lower expectations that depend on the assumed source frequency and ellipticity [383].

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

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

Five all-sky pipelines 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-interferometer Hough transform method starting from 30-minute SFTs, 3) a hierarchical algorithm using `Einstein@Home`, based on phase-preserving demodulation over many $\sim$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 [384] using the PowerFlux infrastructure; and 2) a cross-correlation method [385] 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 [357], Hough [378, 347], and PowerFlux [379]. 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 [380] 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 [381]. 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.

In parallel, a"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 [384]. 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. Preliminary upper limits for the 50-800 Hz band based on a search of the full 2-year S5 data have been produced, and outliers followed up with the new loose coherence step. These results have been presented at spring 2011 conferences and will be submitted for publication in summer 2011, including the outcome of outlier follow-ups.

In the coming year preliminary S6/VSR2/VSR4 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 [382], 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 [380]) 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 [388] 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.

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 [391]. A similar search was run on S5 data (S5R1), and the paper presenting these results published in PRD [392].

S5 R3/R5/R6 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. 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. S5R5 covered the frequency range 50-1000 Hz, with the follow-up S5R6 run covering the 1000-1190 Hz band for the same data segments. These run have finished and the work on postprocessing of their outliers has been completed. A paper describing the S5R3/S5R5/S5R6 searches will 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 [393]. An engineering run (S5GCE) to test this improvement was completed in spring 2010, and a new production run (S5GC1) over the full S5 data has recently completed. Post-processing has begun, with final results and submission for publication expected in the coming 1.5 -2 years.

S6 Bucket search: A new production run over the most sensitive S6 data (in the "bucket") was recently begun, where the relatively low frequencies (50-450 Hz) permit searching with a longer coherence time (90 segments of 60 hours each) and up to spindowns corresponding to an age of 600 years. Results from this search are expected in the coming year with submission for publication in the following year.

Data selection and run set-up procedures: The best exploitation of the limited computing resources for a given parameter space is an optimization problem that depends on the noise level of the data and on their gaps relative to all the usable detectors. Up to now eac `Einstein@Home` run has required an ad hoc iterative study of the set-up parameters. The team is trying to quantify through figures of merit the relevant choice parameters and to set up automatic procedures to define run parameters. Part of the data selection and preparation procedures also involves characterizing the effects of the Rome time-domain cleaning on LIGO data and making it possible for this cleaning to be incorporated in our automatic procedures. This is part of the effort.

Speed-up of the F statistic search: Incorporate the resampling-based code developed by Chris Messenger in the `Einstein@Home` hierarchical pipeline.

Sliding coherence scheme We're studying a novel hierarchical search technique for all-sky surveys for continuous gravitational-wave sources. The present work proposes to break the data into subsegments shorter than the desired maximum coherence time span (size of the coherence window). Then matched-filter outputs from the different subsegments are efficiently combined by sliding the coherence window in time: Subsegments whose timestamps are closer than coherence window size are combined coherently, otherwise incoherently. Compared to the standard scheme at the same coherence time baseline, data sets longer by about 50Ð100

Signal-based veto procedures: Based on the final results of ongoing post processing efforts, signal-based vetoes will be incorporated in the hierarchical `Einstein@Home` pipeline allowing to discard spurious disturbances from the "top list" of returned candidates hence ultimately improving the sensitivity of the search.

Local maxima, clustering of correlated candidates, detection tests over portions of the parameter space Based on the final results of ongoing post processing efforts, identification of correlated candidates will be studied and possibly incorporated in the hierarchical `Einstein@Home` pipeline allowing to save follow-up cycles. For the same purpose, i.e. identification of interesting regions in parameter space, we will also investigate the application of standard hypothesis testing methods, for example the Bayes ratio.

Longer coherent integration time baselines Currently we are not able to perform searches with a coherent time baseline longer than about 60 hours. The reasons are twofold: 1 – the hierarchical search software cannot handle spindown orders greater than 1, and 2 – we do not know how to set up template grids with integration times in this range. Both problems are being studied, and we aim for an enhancement of the hierarchical search code in the next 12 months.

Sensitivity estimates and comparisons We maintain and develop an analytical framework to estimate the sensitivity of the complex search procedures that the `Einstein@Home` performs as well as other searches, most notably the Powerflux search.

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

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 numerical studies carried out on a small band of S5 data. 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 Instrumental line-veto statistics for wide-parameter searches

One of the standard methods of CW data analysis is the multi-detector $\mathcal{F}$-statistic. In a typical search, the $\mathcal{F}$-statistic is computed over a range in frequency, spin-down and sky position, and the candidates with highest $\mathcal{F}$ values are kept for further analysis. However, this detection statistic is susceptible to a class of noise artifacts, strong monochromatic lines in a single detector. Conventionally, these artifacts are removed manually in post-processing. By assuming an extended noise model - standard Gaussian noise plus single-detector lines - we can use a Bayesian odds ratio to derive a generalized detection statistic, the line veto (LV-) statistic. In the absence of lines, it behaves similarly to the multi-detector $\mathcal{F}$-statistic, but it is much more robust against line artifacts, reducing the need for manual removal in post-processing.

We are currently in the process of testing and tuning this statistic on simulated data of Gaussian noise and lines, and on candidates returned from the Einstein@Home S5R3 run. If successful, we are planning to implement this statistic directly on the Einstein@Home clients in the next wide-parameter search in order to improve the sensitivity of the search. Namely, by avoiding saturation of the top-list of candidates returned by triggers from line artifacts, the detection threshold is effectively lowered.

5.3.6 Hierarchical Hough search

A hierarchical pipeline with an incoherent adaptive Hough step [370, 371, 372] will be applied to Virgo and LIGO 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[374]). A new Hough procedure, based on the transformation between the time-frequency plane and the source frequency/spindown plane, has been developed[376] 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 had been slowed by parallel work on the Vela search, now completed, but in the past year the search pipeline has been augmented by a refined Hough step. In the coming year a coincidence analysis of candidates will be implemented and the pipeline applied to VSR2 data. In the longer term the search will be applied to the full VSR2/VSR4 and S6 data.

5.3.7 $\mathcal{F}$-statistic all-sky search

Another analysis method currently used for all-sky searches is based on the $\mathcal{F}$-statistic [394]. It consists of a coherent stage over 2-day data segments, each covering a 1-Hz bandwidth plus a follow-up analysis of candidates in a 4-day data segment. We shall assume the minimal spindown 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 can 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 section 5.1.4. 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 Einstein@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 are analyzing 25000 2-day 1Hz band sequences.

VSR1 roduction running is nearly complete, and a paper describing the all-sky search 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 VSR4 data, with improvements based on the global correlations code used in `Einstein@Home`.

5.3.8 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. It is not yet clear whether or not this approach can be applied to a full-sky search.

5.3.9 Cross-correlation search

The cross-correlation method has been described 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 ongoing or previous accretion onto the neutron star. In this scenario, first proposed by [360], 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 [348]. The resulting indirect limit can be put in terms of X-ray flux F_x and spin frequency $f_{\rm rot}$ as

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

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 published a single coherent analysis on the accreting neutron star in the LMXB Sco X-1 using S2 data [348]. 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 [411].

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 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 [389, 390]. This is equivalent to convolving the detection statistic frequency series with a “comb” of unit amplitude spikes separated in frequency by the inverse of the orbital period. The incoherent summing makes this a non-optimal strategy, but one that can have greater sensitivity to GW signals than a matched–filter approach because its observation length is not computationally limited. When using this approach, the parameter space resolution (and hence the number of search templates) is significantly reduced. It should also be noted that the sensitivity of this search to GWs scales with $T^{-1/2}$, as with a coherent search (and unlike other incoherent searches), however, the sensitivity also scales as $M^{-1/4}$ (M is the number of sidebands) and hence high frequency, large orbit sources will be harder to detect with this method.

Of the LMXBs it is those of unknown spin frequency to which this search is most suited. This includes the Z and atoll sources (rather than the accreting millisecond X-ray pulsars) which have known sky position, and for some, a reasonably well known orbital period. The remaining orbital parameters, semi-major axis, time of passage through the ascending node, eccentricity etc. are generally quite poorly known. This scenario suits this search, as the sensitivity is relatively insensitive to all orbital parameters except for the orbital period.

The search code and associated pipeline are complete, and preliminary results have been obtained from S5 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 [357].

A paper describing the exploratory S5 search for Sco X-1 is expected to be submitted for publication this year.

5.4.2 Cross-correlation searches for known binary systems

The cross-correlation search described in section 5.2.3 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 [348], and a more sensitive search than the incoherent cross-correlation search done in [411].

We will extend the cross-correlation code written for isolated neutron stars, described in section 5.2.3, 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.

5.4.3 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 has been developed. An extension of the coherent search method with extra parameters to describe the orbital motion of the neutron star is not computationally feasible (for coherence times in the order of 1 h, the extra number of parameter values needed to cover all likely Keplerian orbits exceed a factor of 10^9). A hierarchical search method like the stack-slide or Hough transform methods as discussed in Ref. [357] 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.

5.4.4 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. 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 a narrowband S6 search for unknown binaries have been produced, and a broadband production run over the S6 run will begin in summer 2011, with publication expected in the following year. A methods paper has been submitted for publication [403].

In parallel, a study will be carried out in the coming year of optimizing the TwoSpect method for directed searches for known LMXB's and unknown binaries in globular clusters, for which elimination of the IHS step should be feasible and should improve sensitivity.

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 [404] and the Gregory-Loredo method [405]. 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.

Because the method is computationally intensive, it will likely be applied only in directed searches at interesting points on the sky, such as the galactic center and globular cluster cores, where high stellar density might lead to periodic transients.

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 Support infrastructure

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

5.6.1 Software injections

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

The parameters of the initial 4000 injections were settled upon by the CW group in spring 2011. Infrastructure for creating parallel data streams (frames and derived SFT's) has been developed and will be used in summer 2011 to create standard S6/VSR2 data sets for testing. Corresponding VSR4 data sets will be generated in fall 2011. These parallel data streams are identical to the original calibrated strain data sets, but with the addition of the 4000 signals. These signals will include isolated stars, stars in binary systems, and a small number of glitching stars, with the binary stars in separate streams from the isolated stars.

5.6.2 CW detector characterization

The CW Group has a strong interest in monitoring (and mitigating) instrumental spectral lines (see the detector characterization chapter) with low latency. In preparation for S6 the nascent "F-Scan" infrastructure developed during S5 for generating high-resolution spectrograms to detect wandering lines visually was automated and expanded to auxiliary channels and to Virgo channels. These spectrograms and data files proved invaluable in quickly spotting and tracking down instrumental lines in S6 data. The F-Scan output has also been mined for spectral coincidences between interferometers and between strain and auxiliary channels. Further improvements included dedicated F-scans for bands relevant to special pulsars, such as the Crab and Vela. More such bands will be defined and monitored for the VSR4 run.

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

On the Virgo side, a new infrastructure called NOEMI has been developed for identifying/mining transient events and stationary spectral lines. It was run offline on VSR2 data, but was run in real-time on VSR3 data. A further improved version will run in real-time on VSR4 data. A lines database has been developed, and a web interface to the database is under development with additional support for coherence calculations.

Because of sidereal modulation of CW source amplitudes (as seen by the interferometers), it is good to be aware of instrumental or environmental effects that could lead to apparent sidereal artifacts. Ongoing investigations have indeed revealed apparent artifacts in both strain and auxiliary environmental channel. Further analysis will be carried out of the full S6/S6/VSR2/VSR4 data sets to quantify the effects of these artifacts on CW searches.

In addition, the Gregory Loredo method will be used in the coming year to investigate instrumental artifacts that lead to periodic transients, and spectral correlation methods (different from spectral coherence) will be used for studying environmental contamination of the strain channels.

5.7 Scientific goals and plans for the advanced-detector era

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

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

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

1. Targeted search for $>$100 known pulsars (radio, X-Ray, γ-ray).

2. Narrowband search for the same known pulsars, in which a slight frequency mismatch between electromagnetic and gravitational wave mismatch is permitted.

3. Directed search for interesting objects & sky locations assumed to be isolated, such as Cassiopeia A (if pulsations not yet detected), SN1987A, and isolated stars near the galactic center and in nearby globular clusters.

4. Directed search for known binary sources without pulsations, such as Sco X-1, and assumed binary sources near the galactic center and in nearby globular clusters.

5. All-sky search for unknown isolated neutron stars ($\sim$10-2000 Hz)

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

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

2. Search for transients from known stars exhibiting electromagnetic glitches

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

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

In the following, we address for each flagship search and for each "desirable" search the following issues:

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

Flagship Searches:

1. *Targeted search for >100 known pulsars.*

    - We plan to have at least two independent targeted-search pipelines ready to search for O(100) pulsars with known ephemerides from radio, X-ray or gamma-ray observations.
      Three such targeted pipelines (Glasgow, Rome, Warsaw) have produced upper limits on Vela from the VSR2 data, including one that has produced upper limits on 116 pulsars from S5 data. All three pipelines have been reviewed with results submitted for publication.

- In the event of a discovery, each pipeline already provides parameter estimation for the star's inclination and polarization angles, along with phase offset with respect to electromagnetic observations. For some stars (*e.g.*, Crab & Vela), we can compare GW orientation angles directly to those observed electromagnetically. The quite different techniques used in the parameter reconstructions will allow cross-checks of systematic uncertainties. Any inferred ellipticities would warrant comment on implied neutron star equations of state, citing the literature.
- We expect to produce journal-ready search papers based on S6 / VSR2-4 data at the end of 2011 or early in 2012 using the current (or slightly improved) pipelines. Those reviewed pipelines should be ready for advanced detector data, and snapshots will be taken of them. Software injections will also be used to evaluate and compare the performance of the pipelines in the coming year.

2. *Narrowband search for the same known pulsars, allowing EM/GW mismatch.*

- We plan to have at least one stable and proven narrowband pipeline to search for signals only slightly mismatched in frequency from that implied by electromagnetic pulsations. One such pipeline was used in a published search near the Crab frequency in early S5 data. That pipeline has been adapted to use Doppler resampling code, to gain in computational speed, code which was largely reviewed in the context of the recent Calvera search (aborted when out-of-band pulsations observed by x-ray astronomers).
- In the event of a discovery, it should be straightforward to "zoom in" to determine source parameters well enough for detection with targeted search pipelines, in which case SNR would likely increase, and parameter estimation prove to be powerful.
- We expect to produce a journal-ready publication using the current (or slightly improved) pipeline on the time scale of the end of 2012. That pipeline should be ready for the advanced detector era with a snapshot taken. Software injections will be used to evaluate its performance.

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

- We plan to have at least one stable and proven directed-search pipeline for isolated, well modeled neutron stars (known direction, but unknown frequency) based on long coherence times. Recent or ongoing searches for Cas A, Calvera, and the galactic center, all based on the $\mathcal{F}$-statistic, provide prototype approaches. For stars satisfying the nominal source model, Doppler resampling will be valuable in reducing the computational cost of these searches. While more than one resampling method has been used within the group, none has been completely reviewed yet. In addition, we plan to have a pipeline robust against source modeling errors, one based on cross correlation methods, now under review for an S5 SN1987A search.
- In the event of a discovery, zooming in is, again, likely to prove effective and permit powerful parameter estimation.
- We expect both types of directed, isolated-star searches (Fstat and cross correlation) to be fully reviewed with pubished results on the time scale of 2012 to 2013 and available for use with advanced detector data. Again, software injections will be used to evaluate their performance.

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

- We plan to have at least three independent, stable and proven pipelines able to search for Sco X-1 and other LMXBs. The current Sideband pipeline and the Radiometer pipeline (used in stochastic GW searches) provide defaults, but we expect other pipelines (Cross-correlation and TwoSpect) to be available for Sco X-1 application, too. Given Sco X-1's relatively high likelihood of producing detectable radiation and its likely phase wandering during accretion, it is prudent to attack it with a variety of algorithms.
- In the event of a discovery, zooming in could prove challenging for accreting binaries, but at the same time, if a source is strong enough to be seen over short coherence times, then it may be possible to track time-varying accretion effects directly, which could yield extremely interesting insights into both the accretion process and neutron star structure. Existing tools in targeted search pipelines should be adequate for tracking, if the signal is strong enough to be detected in the first place.
- The Stochastic group's Radiometer pipeline has been fully reviewed, and the Sideband review is well under way. By the end of 2011, both should have yielded published results and be available for use with advanced detector data. The cross-correlation and TwoSpect algorithms (optimized for a directed search) are much less mature, with reviews likely to begin in 2012. Nonetheless, it is likely that both pipelines will be available for use with advanced detector data. Software injections will be important in validating and comparing all four pipelines.

5. *All-sky search for unknown isolated neutron stars.*

- We plan to have 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 provides a default pipeline. In addition, we plan to have a stable and proven version of Einstein@Home able to carry out production running and finalize post-processed results within two years of the start of data taking, allowing a search over a large parameter space volume. Another mature all-sky search is the weighted Hough pipeline used in S4 and S5 searches. In addition, there are several other pipelines nearing maturity (Rome frequency Hough, Warsaw time-domain F-statistic, Einstein@Home with global correlations) and a new pipeline under development based on stand-alone loose-coherence (developed for a sky-patch search, may be applicable to an all-sky search). The reviews of the PowerFlux, weighted Hough and Einstein@Home searches used for the full S5 data set should all be complete by fall 2011. The other pipeline reviews are expected to begin in 2011 or 2012.
- In the event of discovery, systematic zooming in on the source should increase SNR dramatically, perhaps by as much a factor of 30 if the source phase is stable over a year or more. We should be able to point electromagnetic astronomers to the source with $\sim$arcsecond precision and with frequency precision better than $\sim$ μHz. How to arrange that pointing merits careful thought. Our inclination now is to do as well as we think can do using only our own data and then issue an open Astronomer's Telegram (ATel) with our best information so that the entire astronomical community has the opportunity to contribute observations that could be included in the discovery paper. Depending on what the GW and EM observations (including lack thereof) tell us about the nearness of the source, it may be appropriate to include inferences about neutron star poplulation densities in the discovery paper.
- Reviewed snapshots of all of the currently mature all-sky pipeline should be available during 2011-2012 for use in the advanced detector era. Refinements are likely to continue for at least PowerFlux and Einstein@Home, however, and the other pipelines discussed above will continue development for use with S6 / VSR2-4 data. As the start date of advanced detector data taking

approaches, it will be important to ensure that updated snapshots are taken for those pipelines that warrant use with the advanced data, snapshots taken with sufficient time to complete review of improvements before data taking begins. Deciding which pipelines merit continued use in advanced detector data will depend critically upon the results of standardized software injection studies.

Additional Planned Searches:

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

   - We plan to have at least two stable and proven all-sky searches for unknown neutron stars in binary systems. The Polynomial and TwoSpect pipelines now approaching maturity on S6/VSR2 searches provide defaults. Review of the Polynomial pipeline has begun, and review of TwoSpect will begin in 2011.
   - In the event of discovery, zooming in should be possible. Because significant sensitivity must be sacrificed (even more than for the all-sky isolated searches) to gain computational tractability, potential SNR gains are very large after initial detection. At the same time, zooming in on a binary may prove difficult and will need serious investigation between now and the advanced detector era. As for the discovery of an new isolated star, we favor strong interaction with electromagnetic astronomers via an ATel once we have thoroughly explored the GW data. Again, statements in the discovery paper on neutron star population densities may be appropriate, depending on GW and EM observations.
   - Both the Polynomial and TwoSpect searches are expected to produce publication-ready articles and reviewed snapshots in mid to late 2012, ready for use with advanced detector data. As is the case for all-sky isolated-star searches, however, further improvements are likely, and new (reviewed) snapshots will be needed as the start of data taking approaches. Software injections will be used to evaluate and compare the two pipelines.

2. *Search for transients from known stars.*

   - We plan to have at least one pipeline to search for long-lived transients associated with observed electromagnetic glitches in neutron stars, including magnetars. (It should be noted that both the Burst and Stochastic search groups have developed or are developing pipelines for post-glitch searches, using more generic approaches.) A transient CW pipeline has been developed recently for use in S5-S6 / VSR2-4 searches, but has not yet been reviewed.
   - In the event of discovery, it may not be possible to zoom in to increase SNR, as the source may go quiet again for years. In that respect the transient search is similar to a burst search. Signal interpretation will depend on reconstructed signal strength, frequency and spindown. One could observe f modes, r modes or torsional modes, for example. Depending on how quickly a signal is detected, it may be appropriate to issue an ATel before publishing the discovery paper.
   - We expect publishable results from the new pipeline in 2012-2013, in time for a snapshot for use with advanced detector data. The pipeline will be validated and its performance assessed via software injections.

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

New astrophysical targets include:

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

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

But we can already identify some future algorithmic needs:

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

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

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

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 detectors' frequency band, 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.

The group started looking for long-lasting transients, i.e. signals that are present for a duration of minutes, hours, or longer. The scientific motivation for searches for long transients was recently laid out in [428]. Possible sources include: instabilities in and excitation of accretion disks, instabilities in newborn neutron stars, glitches/flares associated with isolated neutron stars and dynamically formed black holes binaries. Possible external triggers include: long GRBs, short GRBs, supernovae, pulsar glitches, soft gamma repeaters (SGRs) and anomalous X-ray pulsars (AXPs). For more information, including references to specific models, see [428]. The duration of those signals is long enough to make the stochastic method described below an efficient way to search for them.

Finally the possibility of implementing a detection pipeline for non-Gaussian stochastic background are under study. The main motivations for a non-Gaussian search are the following:

- There are many theoretical predictions for non-Gaussian stochastic backgrounds, either cosmologicals (cosmic strings) or astrophysicals (BH ringdowns, supernovas, bar modes)
- For astrophysical backgrounds, the non-Gaussian contribution created by the closest sources is expected to be several times larger than the Gaussian contribution.
- It is quite unlikely that the anisotropic contribution of known astrophysical sources falls into the Gaussian regime (typically the Universe becomes isotropic after 150-300 Mpc) and it may be interesting to combine anisotropic and non-Gaussian searches.

6.2 Stochastic Search Method

The stochastic search method has evolved from 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 three 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 Isotropic 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 the 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) [412], which can be written as

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

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 $P^{\rm TT}_{abcd}(\hat{n})$ is a projector onto traceless symmetric tensors transverse to the unit vector $\hat{n}$ (see [413], p. 10).

We deploy a cross-correlation method to search for the stochastic GW background, following [414]. 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 s_A and s_B, and T is the measurement time. Optimization of the signal-to-noise ratio leads to the following form of the optimal filter [414]:

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

$S_{GW}(f)$ 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 $\langle Y_{AB} \rangle = \Omega_\alpha T$. The signal-to-noise ratio for a pair of interferometers for an ideal measurement of length T in stationary noise can be written as

$$\text{SNR} = \frac{3H_0^2}{10\pi^2} \left(2T \int_0^\infty df \gamma_{AB}^2(f) \frac{\Omega_{GW}^2}{f^6 P_A(f) P_B(f)} \right)^{1/2} \tag{14}$$

where H_0 is the present value of the Hubble expansion rate. The largest contribution to this integral comes from the frequency region where $P_{A,B}$ is minimum, which is between 50Hz and 150Hz.

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 [416].

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 [417], one can write:

$$Q(t, f, \hat{\Omega}) = N(t, \hat{\Omega}) \, \frac{\int d\hat{\Omega}' \, \gamma(t, f, \hat{\Omega}') \, A(\hat{\Omega}, \hat{\Omega}') \, H(f)}{P_1(f) \, P_2(f)}, \tag{15}$$

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:

$$\mathcal{P}(\hat{\Omega}) = \mathcal{P}_\alpha \mathbf{e}_\alpha(\hat{\Omega}) \, , \tag{16}$$

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

$$C(f,t) = \frac{2}{\tau} \, \tilde{s}_1^*(f,t)\tilde{s}_2(f,t) \, , \tag{17}$$

one can show that its expectation value is given by

$$\langle C(f,t) \rangle = H(f)\gamma_\alpha(f,t)\mathcal{P}_\alpha \, , \tag{18}$$

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

$$\begin{aligned} N_{ft,t't'} &= \langle C_{ft} C^*_{f't'} \rangle - \langle C_{ft} \rangle \langle C^*_{f't'} \rangle & (19) \\ &\approx \delta_{tt'} \delta(f-f') P_1(f,t) P_2(f,t), & (20) \end{aligned}$$

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

$$\begin{aligned} p(C_{ft}|\mathcal{P}_\alpha) \;\propto\; \exp\Big[&-\frac{1}{2}((C^*_{ft} - \langle C^*_{ft} \rangle) N^{-1}_{ft,f't'} \\ &(C_{f't'} - \langle C_{f't'} \rangle) \Big] \end{aligned} \tag{21}$$

where $\langle C_{ft} \rangle$ is given by 18 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

$$\hat{\mathcal{P}}_\alpha = (\Gamma^{-1})_{\alpha\beta} \, X_\beta \tag{22}$$

where

$$X_\beta = \sum_t \tau \int_{-\infty}^{\infty} df \, \gamma^*_\beta(f,t) \, \frac{H(f)}{P_1(f,t)P_2(f,t)} \, C(f,t) \, , \tag{23}$$

$$\Gamma_{\alpha\beta} = \sum_t \tau \int_{-\infty}^{\infty} df \, \gamma^*_\alpha(f,t) \, \frac{H^2(f)}{P_1(f,t)P_2(f,t)} \, \gamma_\beta(f,t) \, . \tag{24}$$

The matrix inversion in 22 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 [414], 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 [415] 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. The contribution in the low frequency range is instead small, owing to the overlap reduction factor in Eq. (14).

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 Isotropic Search

The isotropic 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 isotropic 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

$$\Gamma_{XY}(f) = \frac{|P_{XY}(f)|^2}{P_{XX}(f)P_{YY}(f)} \tag{25}$$

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 [419], it is possible to estimate the instrumental correlations between interferometers 1 and 2 by

$$\Gamma_{instr,12} \approx \ \max_i(\Gamma_{1Z_i} \times \Gamma_{2Z_i}) \tag{26}$$

where Z_i are the numerous environmental channels, including microphones, seismometers, accelerometers, power-line monitors etc. As discussed in [419], 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 isotropic 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 $\sim$ 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 developed a pipeline called STAMP (Stochastic Transient Analysis Multi-detector Pipeline) for the study of long gravitational-wave transients lasting from $\mathcal{O}(\mathrm{s})$ to weeks.

These searches are a natural extension of the cross-correlation techniques described above. In particular, the group started generating the Stochastic Intermediate Data (SID) frames, 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.

STAMP utilizes frequency-time (ft)-maps (spectrograms) of cross-power created from two or more spatially separated detectors (e.g., H1 and L1). The stochastic narrowband radiometer (see [411, 417, 418]) sums ft-map pixels at a fixed frequency for the duration of a science run in order to produce an estimator for GW power. STAMP extends the principle behind the narrowband radiometer in order to consider signals with shorter durations and/or more complicated spectral content. The search for long transients, therefore, is a natural extension of stochastic tools developed for searches for persistent point sources. (Where possible, we rely on previously reviewed stochastic code.)

STAMP calls upon a variety of pattern recognition algorithms in order to identify structure in ft-maps. Current clustering algorithms in development include the box search, the Radon search, the Hough search, the Locust search, and `burstCluster`. (Some of these clustering algorithms are also used by the Burst Group.) Typically, choosing between different pattern recognition algorithms involves assessing the trade off between characteristics such as sensitivity to a given model, flexibility in case the model is not exactly right, and ease of implementation. Much of the foundational work so far has relied on the Radon algorithm, which is very simple in its implementation and suitable for narrowband signals. STAMP searches can be externally triggered (e.g., using long GRB triggers). It is also possible to perform an untriggered "all-sky" analysis, though, they are expected to be computationally expensive.

STAMP is also used for detector characterization purposes. Instead of cross-correlating two gravitational-wave channels, we cross-correlate one gravitational-wave channel with a physical environmental monitoring channel to look for noise that is coherent in the gravitational-wave channel. This "STAMP-PEM" technique has been used to investigate airplane/helicopter noise, chiller-pump "Crab" noise, and noise from lightning and thunder.

6.2.7 Statistics of cross-correlation data

From our earliest searches, which used relatively long data segments, we have essentially assumed the continuum limit in determining the effects and impact of coarse graining and zero padding during our data manipulation. Now that STAMP has several searches planned which are interested in shorter and shorter data segments, it has become necessary to revisit the continuum assumption and to explicitly allow for some finite correction in describing the distributions of CSD, PSD and even SNR. This work, begun in [428], is now well established and will continue during the coming year, with a positive impact expected on all searches that aim to use very short data segments to allow for better time resolution.

6.3 Results and Plans

6.3.1 Status of S5 Searches

Isotropic Search: The stochastic group finished 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 [409], 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 [420], 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 and Spherical Harmonics: The stochastic group has analysed the S5 data set with both the radiometer and the Spherical Harmonics decomposition analysis described in sections 6.2.2 and 6.2.3. A corresponding paper is in the last stages of the LSC internal review. The radiometer analysis produced maps of the GW sky 30 times (in strain power) more sensitive than those produced using the S4 data [411], as well as targeted narrow-band limits on GW radiation coming from Sco X-1 and the Galactic Center. It also confirmed the previouly published isotropic result. The second directional analysis produced the first spherical-harmonic decomposition map of the gravitational-wave sky, similarly to what is done in the field of Cosmic Microwave Background. This method targeted complex source distributions on the sky, and has been summarized in a method paper published last year in Physical Review D [421].

Isotropic Search using Co-located 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 developped 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. A detailed study of this residual contamination required a re-calculation of all PEM correlation estimates for the S5 run, and is still in progress. 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.

LIGO-VIRGO Searches The stochastic group also conducted 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 can improve the overall network sensitivity even at lower frequencies. The analysis is the internal review process.

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.

Stochastic Intermediate Data: Since several searches rely on similar quantities (such as strain cross and power spectral densities of different interferometers), the group has produced 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 developed the infrastructure necessary for searching transient GW signals on the time scales of minutes, hours or longer. This Stochastic Transient Analysis Multi-detector Pipeline (STAMP) uses the SID frames, i.e. 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 on S5 data. The group is also drawing from experiences of other groups (such as the burst group) performing searches for short duration transients.

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 SID, likely in the collapsed form (to one sidereal day).

The strain sensitivity of the LIGO 4-km interferometers (H1 and L1) during S6 have improved compared to S5. However, the additional observation time was limited due to a low duty cycle, particularly early in the run. In addition, the sensitivity of the VIRGO interferometer during S6 did not reach a level similar to the LIGO interferometers at low frequencies (around 100 Hz). However, 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. Thus we will analyse and publish the S6/VSR2 data using the network SHD algorithm as an extention to the S5 run, but we do not expect to significantly improve the S5 result.

STAMP: Using the framework of STAMP, several searches for both triggered (by GRB's or neutrino coincidences) and untriggered GW signals are currently being commissioned. In particular triggered searches for long GRBs, neutron star r-mode instabilities and pulsar glitches, as well as untriggered searches e.g. for narrow-band transient signals are currently at various stages for development. The group plans to run the mature STAMP searches on the combined S5/S6/VSR1-3 data set.

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.

6.4 Stochastic Group's Plans for Advanced Detector Era

Advanced detectors are expected to have about 10x better strain sensitivity than the initial interferometric detectors, as well as to extend the sensitive band from 40 Hz down to 10 Hz. Furthermore, the number of detectors operating in a world-wide network is expected to increase, eventually including sites LIGO-Hanford, LIGO-Livingston, Virgo, GEO-HF (at high frequencies), LCGT (Japan), and potentially even LIGO-South in Australia. These significant strain sensitivity improvements and the increase in the number of available detector-pairs (baselines) that could be cross-correlated will enable real breakthroughs in the searches for the stochastic GW background and for long-lasting transients. We summarize below the scientific targets and the corresponding searches that the Stochastic Group plans to perform in the advanced detector era.

6.4.1 Searches for Isotropic Stochastic Gravitational-Wave Background

Isotropic stochastic GW background will be one of the prime targets of the Stochastic Group in the advanced detector era. The flagship search targeting this background will be the multi-baseline cross-correlation search, using non-collocated GW detectors. It is likely that this search will be the limiting case of the anisotropic search discussed below. The improved detector strain sensitivities and the increased number of baselines will improve the sensitivity of the detector network to the level of $\Omega_{GW} \sim 10^{-9}$ or better. The code/pipeline for these searches already exists and has been applied to S4/S5 LIGO data [409, 410, 420]. Minor adjustments will be needed to apply it to advanced detector data (such as extending the sensitive band down to 10 Hz). Further tuning of the search may be pursued in the form of adjusting the spectral index to target particularly promising astrophysical and cosmological sources.

Another isotropic stochastic search could be performed using the colocated H1-H2 detector pair (in the case when H2 does not move to LIGO South in Australia). As discussed above, due to the shared location/environment, this detector pair is much more susceptible to instrumental and environmental correlations, which could potentially compromise the stochastic GW background measurement. Substantial efforts have been made in the past to perform this search with S5 data. While the S5 search is not completed yet, two conclusions can already be made. First, at high frequencies (above 400Hz) the instrumental/environmental correlations are small and can be effectively identified and removed from the analysis. Consequently, isotropic stochastic search at these frequencies should be possible with aLIGO H1 and H2, providing sensitivity of order $\Omega_{GW} \sim 10^{-5}$ at 1 kHz (a factor of 100x improvement over the S5 result which is not yet published). This analysis will be of particular interest for astrophysical models of stochastic background, which are expected to peak at few hundred Hz or 1 kHz. Second, at frequencies below 300Hz, the environmental contamination may be substantial, and it might not be possible to completely identify and remove it. While the S5 analysis will give important information about this when it is completed, it is clear that the analysis method could be improved to make better use of the PEM channels - currently only the PEM channel contributing the most to cross-correlation is used in the analysis, which is clearly not the optimal use of available information. Depending on the available manpower in the coming years, the Stochastic Group is interested in pursuing such improvements in the method, possibly using the expected software-engineering runs to pace the progress.

It is important to note that at these expected sensitivities, the advanced detector network may be sensitive to a number of astrophysical models (such as magnetars [423], double neutron stars [425], or rotating neutron stars [424]) and of cosmological models (cosmic strings [426], alternative cosmology models such as pre-big-bang models [427]). While some work has been done on understanding the possible implications of the isotropic stochastic searches for these models, a more systematic study is required to make a full assessment of the scientific potential of these searches. For example, it is currently not known whether the isotropic stochastic searches can be used to constrain the magnetic field in magnetars, or the density of compact binary objects etc. A systematic study of such questions may also reveal possible ways of tuning the search to optimize its sensitivity for answering them. Furthermore, such a study should investigate possible ways of distinguishing between different models based on their frequency content and on the predicted anisotropy levels and angular scales.

6.4.2 Searches for Anisotropic Stochastic Gravitational-Wave Background

Anisotropic stochastic searches have already been performed on S4 data [411] and S5 data (currently undergoing internal review). This includes both the radiometer searches that target point-sources on the sky, and the spherical-harmonic decomposition search which is capable of searching for complex spatial distributions in the GW sky. Both of these methods also produce the isotropic result as a limiting case, so could be combined with the isotropic search discussed above. The improved strain sensitivities of detectors, as well as the larger number of detector baselines, will lead to at least 100-fold improvement in the sensitivity to anisotropies in the stochastic GW background as compared to the S5 result (which is yet to be published). Additional improvements are expected by extending the frequency band of the search down to 10 Hz.

The increased number of detector baselines will also significantly improve the angular resolution of the search. A dedicated study should be made to quantify such improvements for different detector network configurations: with or without aLIGO South in Australia, with GEO-HF at high frequencies etc.

The radiometer search can be used to target specific point-like sources, as was done in the case of Sco-X1 and the galactic center with S4/S5 data. Another detailed study is needed to determine which sources would be suitable targets for such searches (especially since the new band 10-40 Hz will be opened), as well as to understand the relative advantages and disadvantages as compared to Continuous-Waves searches. However, it is also possible that the simplicity of the method will drive these targeted searches despite of possible lower sensitivity than the CW searches.

Finally, we note that the current anisotropic searches average over the two GW polarizations. The algorithm could be extended to estimate the two polarization components. Depending on the available manpower, a detailed study could be made to add this functionality to the search code/pipeline and to quantify its performance using a series of signal simulations/injections.

6.4.3 Searches for Long-Duration Transients

As discussed above, STAMP ("Stochastic Transient Analysis Multi-detector Pipeline") is a pipeline devoted to the study of long $\gtrsim 1\,\mathrm{s}$ gravitational-wave (GW) transients [428]. The pipeline is designed to be flexible, so that it could be used to search for a variety of different astrophysical signals. Consequently, multiple searches (and multiple papers) are being planned for the advanced detector era. We select several analyses currently considered as priorities:

- Targeted search for long GRBs. In the event of non-detection, constrain the models of van Putten [430] and Piro and Pfahl [429].
- Targeted search for newborn neutron star emission (e.g., following supernovae triggers). Constrain models of r-modes [431], PNS convection [432] and secular instabilities [434, 433]. A dedicated study is needed to establish which of these models are indeed accessible to advanced detectors.
- Targeted search for transients from isolated neutron stars. For pulsar glitches, we constrain the fraction of the change in rotational energy that is emitted in GWs. For SGR/AXP flares, we constrain the ratio of GW energy compared to the observed electromagnetic energy.
- Un-triggered narrow-band searches. Due to the likely highly anisotropic emission of electromagnetic radiation from the above targets—e.g., long GRB jets may be confined to a $\theta_{\text{jet}} < 4^\circ$ beam [435], which means there is only a $\sim 0.25\%$ chance that it points directly at Earth—it is worth looking for the same objects with an untriggered search. We also propose untriggered searches for dynamically formed black-hole binaries, which are not expected to be associated with any electromagnetic trigger.
- Depending on the available manpower and interest of the involved researchers, other STAMP-based studies may also be possible. For example, in tensor-scalar theories, a potential damping mechanism for the bulk vibrational modes in neutron stars would be through the emission of scalar gravitational waves [437, 438]. In the case of extended (low density) neutron stars, this could occur over an interval of more than 80 seconds (see [437]), resulting in a long gravitational-wave transient, making it ideally suitable for a STAMP search.

The STAMP code suite is reaching a mature stage. Most of the infrastructure is already complete, and numerous checks have been performed. The first analysis (targeted search for GRBs) is under way, and we expect preliminary results using S5/6 data by June 2011. We expect the key STAMP infrastructure to be reviewed internally in the context of S5/6 - VSR2-3 analyses during 2011-12, well before advanced detectors begin taking high-quality science data. Minimal algorithm development will be necessary after this review is completed. However, for each of the specific analyses mentioned above, the following "roadmap" must be completed [436]:

- Astrophysical justification.
- Background studies, understanding false alarm rates, possible statistical bias etc.
- Signal (injection) studies.
- Presentation of dry run with time-shifted data.
- Open box on non-shifted data.

We also note here that STAMP has already found detector characterization applications (identification of data segments corrupted by airplanes, thunderstorms etc). We expect such applications to intensify over the 2011-2014 period, and will likely lead to additional publications. Substantial code development will be needed, for example to make the pipeline useful in control rooms, which will likely be synchronized with the Software Engineering Runs that are planned in this time period.

6.4.4 Updates to the Stochastic Pipeline

In addition to the planned searches discussed above, the Stochastic Group is considering improvements to the pipeline that should take place during 2011-2014, before the advanced detectors start providing sensitive strain data. The most important upgrade under consideration is moving the code platform away from Matlab. While Matlab platform certainly has many advantages, such as availability of various analysis tools, ease of troubleshooting problems etc, it also has limitations. For example, Matlab platform is not well suited for parallel processing, and it is not clear whether such features will continue to be supported in the future. Consequently, the Group is considering moving the stochastic code suite (including isotropic and anisotropic codes, STAMP codes, pre- and post-processing codes) to C, C++, Python, or some other platform. This move would require a substantial amount of effort, but it may be necessary to enable smooth operation of the Group in the advanced detector era. While the move should ideally happen before 2014, its timing will be highly dependent on the available manpower in the Group, as well as on balancing against the ongoing searches using S5 and S6 data. The details of this effort, including the choice of the new platform, ordering of code translations, testing the new code, and potential use of the software engineering runs to guide the development, will also be dependent on the available manpower and expertise.

7 LSC 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 [3] 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.

[3] `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 certiïňĄcates 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 [4] 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 technoligies 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.

[4] `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`).

8 Virgo Computing and Software

The Virgo computing model [439] relies on 3 main computing centers: Cascina where Virgo data are acquired, the IN2P3 computing center (CCIN2P3 in Lyon, France) and the INFN computing center (CNAF in Bologna, Italy). The CCIN2P3 and CNAF computing centers provide permanent data storage facilities and large computing resources for data analysis, while Cascina computing resources are mainly used for the online searches and detector characterisation applications that require an immediate access to the raw data in order to provide a fast feedback to the commissioning activities.

Virgo delivers about 1 TB per day of data. The data acquisition volume is mainly composed (99%) of raw data. Gravitational wave strain, and also reduced data set (trend, 50 Hz and RDS) are also generated online. A storage buffer (disk) of about 6 months of data is available at Cascina. Permanent storage is provided in CCIN2P3 and CNAF. Virgo has made the choice to provide users with access to the full raw data, and not only to the reduced data sets. This choice puts strigent constraints on the quality of data access at the computing centers.

Since the beginning of Virgo, no logical separation between online and offline software has been chosen in order to avoid the duplication of libraries. Many GW search developments have been designed to run online as soon as possible. One drawback of this strategy is the still limited number of Virgo pipelines running offline in CCIN2P3 and CNAF.

8.1 Organization

The Virgo data analysis software (VDAS) group has been set up in 2010 to help and support data analysis activities in Virgo and to manage the computing resources. It includes software developments support, data transfer monitoring, data access support and management of the relations with the LSC and EGO. The EGO computing department is an important actor which provides manpower and support. They are in charge of the installation and the maintenance of all computing equipment on site. They also contribute to Virgo software developments for the control of the detector and its characterization. Data transfer from Cascina to CCIN2P3 and CNAF is carried out by the EGO computing department.

Every year a "Virgo computing needs" document is prepared for discussions with the compiting centers and the EGO directions. This gathers all requests from data analysis groups.

8.2 Current status

The Virgo computing model defined more than 10 years ago requires to be updated. An effort to improve the use of computing centers is obvious. The available computing resources are used by a very reduced number of Virgo users. One of the reasons is the lack of interoperability of the LSC and Virgo software. Common LSC/Virgo software development requires now to choose to run workflows either on LSC clusters or on Virgo clusters exclusively. At CCIN2P3 and CNAF jobs submission is performed through BQS or LSF batch scheduler. A Virgo GRID Virtual Organisation (VO) is existing since few years, allowing users to submit jobs from any labs, including LSC clusters. LSC clusters batch scheduler is mainly *Condor* and all LSC workflows are embedded with Condor. That prevents to run directly these workflows at CCIN2P3 or CNAF. An attempt to submit Condor jobs to CNAF using GRID pilot job paradigm has been tested. This is a promising solution, but is not used for the moment.

Most of the data analysis software developed by Virgo are under the responsability of individuals, but few packages are developed under the umbrella of the VDAS group.

8.2.1 Computing at the EGO/Virgo site

DAQ Virgo acquires about 11 MByte/s of raw data, which are stored on a hierarchy of circular buffers, ultimately allowing to keep about 6 months of data at the site. A separation is in place, at network and disk level, between the computers and applications critical for Virgo operation and monitoring, and the rest of the computing/user environment.

Detector control and monitoring A number of real time machines and workstations are dedicated to run the control and data monitoring algorithms, with an high level of automation and web reporting. Most of the actions are either automatic or run via an interface which also takes care of tracking and reporting the actions in logfiles.

Online processing A number of real time processes take care of formatting the acquired data and of performing a first processing, including the production of online hrec and the basic data quality assessment. More information about this and the previous activities is available at `http://wwwcascina.virgo.infn.it/WebGUI.htm`

Data logging and transfer Data are automatically indexed, producing Frame File Lists, and transferred at the computing centers of CNAF (Bologna) and CCIN2P3 (Lyon) for permanent storage, using about half of the 500 Mbit/s link to the GARR network.

Online/offline analysis A cluster of about 100 nodes including 32 dual processors and 64 dual-core dual processor machines, is available for online analysis as well as for offline studies on recent data. These machines receive either data from the online processing, thus with a very small latency (a few seconds), or read data from the raw data buffer, with a larger latency (a few minutes).

General purpose workstation A small farm of workstations is available for general purposes; both these machines and the user's workstations access all the recent interferometer data.

8.2.2 Computing at CCIN2P3 (Lyon, France)

The CCIN2P3 (http://cc.in2p3.fr/) is the computing center of the Institut National de Physique Nucleaire et des Particules funded by the CNRS, and is used by the Virgo Collaboration since 2000.

It serves as official repository of the Virgo data since 2000. All data are stored in high mass storage system (HPSS) that provides quasi unlimited storage capacity (in 2011 more than 13 PByte of data are stored in HPSS. Virgo data amounts to 716 TByte). A large enough cache disk space and a software interface is used to get a transparent data access. As soon as data are produced or transfered to Lyon they are declared in the SRB (Storage Resource Broker) catalogue that allows a remote access to the data from the laboratories of the Virgo Collaboration.

The CCIN2P3 also provides large computing facilities, based on about 7000 Linux processors whose efficiency use is higher than 80%. The CCIN2P3 operates an LHC Computing Grid (LCG) Tier-1 center.

Jobs can be submitted either via a batch queue system (BQS) or via LCG Grid. Virgo runs at CCIN2P3 GW burst and pulsar searches. Data quality investigations and the h(t) reprocessing are also performed at CCIN2P3.

The CCIN2P3 is linked at high speed with CERN (thanks to a dedicated 10 Gbit/s optical link) and with the USA.

8.2.3 Computing at CNAF (Bologna, Italy)

The CNAF (http://www.cnaf.infn.it) is the main computing center of the Istituto Nazionale di Fisica Nucleare and serves as repository of two year's worth of the most recent Virgo data. It also provides a large computing facility, consisting of a collection of Linux workstations (about 10^3 bi-processor nodes) accessing about 430 TByte of disk space and 1000 TByte of tape space. Jobs can be launched either via GRID or via a standard batch scheduler system.

The CNAF is linked at 10 Gbit/s with CERN, and at 1 GBit/s with various other locations in Italy and Europe.

In 2011, Virgo data have been migrated to a new storage system that keeps on spinning media the most recent data or most recently accessed data and migrates the older data on tapes automatically. Virgo runs at CNAF mostly pulsar searches, and up to now has used only a small fraction of the total CNAF power. It plans to run in Bologna also the other offline searches.

8.3 Virgo software and services

The Virgo data analysis software is organized in a hierarchy of libraries, from the very basic ones which allow to format and to access frame data, to very specialized ones for the searches, with, as already underlined a strong connection to the online control of Virgo. While referring to http://wwwcascina.virgo.infn.it/sDoc/ for more information, we give here just a brief overview, which should not be taken as complete.

Basic software Data are formatted and accessed with the `Frame` library. Data visualization is possible using the `dataDisplay` software, which is capable of reading data from the online and offline, also from geographically distance places. Simulation of the detector and the sources is possible with the `siesta` software, which can be programmed with "cards" which describe objects and their relation, and outputs data in frame format. Interactive work with frame data is possible with the `vega` software, which is `ROOT` linked with various Virgo libraries. `Cm` is a task-to-task communication protocol based on TCP-IP used by all online running processes in Virgo. It is also used to transfer frames with low latency to LSC clusters. Many other packages, devloped for the online control of Virgo are also used in some Virgo GW search workflows.

Software management In Virgo, all packages are managed, configured and build by CMT (http://www.cmt.org). A simple requirements file describe the build options, and more importantly the dependences with other packages. Each Virgo C or C++ software project is made of a package that comprises at minimum 2 folders:

- src: contains all source and include files.
- cmt: contains the requirements file.

All Virgo packages are archived in the Cascina CVS repository accessible by any LSC/Virgo users through a pserver whose address is :pserver:<user>@cvs.ego-gw.it:/cvsroot

All packages can be installed using CMT directly from the CVS archive. The most stable version of data analysis packages are also available from the standard Virgo Common Software distribution available at http://wwwcascina.virgo.infn.it/sDoc/VirgoReleases/current.html.

GRID The use of GRID has been promoted in Virgo since many years. A Virgo Virtual Organization has been defined in the LHC Computing Grid (LCG) organization (http://lcg.web.cern.ch/lcg/). The Virgo pulsar group is routinely running jobs over the GRID, inside the Virgo VO, both at the two TIER1 computing centers of CNAF and CCIN2P3, and at several TIER2 sites, fully belonging to Virgo or not.

Data transfer Virgo transfers data at the Computing Centers using the BBFTP software (a multi-threaded FTP client/server system) complemented with a set of PERL scripts and a DB database between Cascina and CNAF and using SRB between Cascina and CCIN2P3 While satisfactory for point-to-point data transfer, this solution is not optimal for distributing data on request, hence Virgo is developing a GRID-compatible solution based on the FTS architecture, and plans to exploit more fully the range of middleware available by the GRID development.

In the medium-long term, every effort has to be devoted to make this GRID architecture compatible with the one adopted by the LSC, also operating towards a greater interoperability of the EU and US grids. Data transfer between Cascina and LSC clusters (Hannover) is performed using LDR tools. A server at Cascina has been equipped with all required LSC software.

Noise analysis software A C++ package called NAP is dedicated to the noise analysis, from the computation of basic statistics to more sophisticated approaches, like AR and ARMA modeling, as well as multi-coherence and non-linear analysis. More information is available at http://wwwcascina.virgo.infn.it/DataAnalysis/Noise/nap_index.html.

Search software The Burst group has developed a comprehensive C++ BuL library which hosts most of the search algorithms, as well as utilities and service routines which allow to build pipelines running on real or simulated data. Burst pipeline software also includes matlab based post-processing functions.

The Coalescing Binaries group uses the inspiral library for templates, signals and template grids. It uses MBTA for running the searches on real or simulated data.

Both Burst and CBC software is available in the standard Virgo Common Software distribution, available at http://wwwcascina.virgo.infn.it/sDoc/VirgoReleases/current.html

The Pulsar group has developed a comprehensive package, mostly based on Matlab or Matlab compiled routines. More information is available at http://grwavsf.roma1.infn.it/pss/

The Stochastic Background Group has developed a specific search and simulation library SB, documented at http://wwwcascina.virgo.infn.it/DataAnalysis/Stochastic/software/SB/index.html which leverages also on the noise library NAP

VDB The Virgo Database (VDB) system is based on a mysql v5 server and on a set of user and administration tools. The aim of this project is to store and manage different kinds of information that are important for data analysis:

- bookeeping: Virgo frame files geographical position (SITE, PATHNAME, FILENAME, ...)
- metadata information: data about frame data (science mode, data quality and ITF status)
- segments information: ITF specific segments (e.g. science mode) and user defined segments
- events: inspiral, burst and others
- triggers and veto: ITF specific or user defined

VDB is one of the main tool used by the Virgo Data Quality group (VDQ). It is the official location of Virgo data quality segments lists. One peculiarity of the VDB system is, on one hand the possibility to upload online DQ segments as well as Scimon or offline DQ segments, on the other hand the fact that it provides interesting tools to upload, manage and download segments lists:

Virgo problem report system The SPR system to report problems for both software and hardware has been recently upgraded. It's now based on MantisBT. That especially allows to use the Cascina Active Directory authentification system. Any user can submit problems at `http://sprserver.ego-gw.it/mantisbt/my_view_page.php`.

8.4 Preparing for the advanced detector era

Plans are under discussion.

References

[1] LSC Detector Characterization Group home page:
https://wiki.ligo.org/foswiki/bin/view/DetChar/WebHome

[2] 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/

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

[4] Home page for the omega pipeline:
https://trac.ligo.caltech.edu/omega/

[5] 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.

[6] "Reducing the effect of seismic noise in LIGO searches by targeted veto generation", Duncan M. Macleod, Stephen Fairhurst, Brennan Hughey, Andrew P. Lundgren, Larne Pekowsky, Jameson Rollins, Joshua R. Smith, Submitted to Classical and Quantum Gravity.
http://arxiv.org/abs/1108.0312

[7] 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).

[8] "A hierarchical method for vetoing noise transients in gravitational-wave detectors", J. R. Smith, T. Abbott, E. Hirose, N. Leroy, D. MacLeod, J. McIver, P. Saulson, P. Shawhan, accepted by Classical and Quantum Gravity July 2011.
http://arxiv.org/abs/1107.2948.

[9] "Used percentage veto for LIGO and virgo binary inspiral searches", Tomoki Isogai and the Ligo Scientific Collaboration and the Virgo Collaboration, Journal of Physics: Conference Series, 243 2010.
http://stacks.iop.org/1742-6596/243/i=1/a=012005

[10] ligoDV home page:
https://www.lsc-group.phys.uwm.edu/daswg/projects/ligodv.html

[11] mDV is available from the 40m svn.

[12] S6 glitch group homepage:
https://wiki.ligo.org/DetChar/GlitchStudies

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

[14] 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).

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

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

[17] 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).

[18] 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).

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

[20] Y. Aso et al., "Accurate measurement of the time delay in the response of the LIGO gravitational wave detectors", Class. Quantum Grav. 26 (2009) 055010.

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

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

[23] H. Grote and H. Lück, "The GEO-HF upgrade program", LIGO Document L1000195, 2010.

[24] J. R. Smith *et al.*, "A hierarchical method for vetoing noise transients in gravitational-wave detectors", gr-qc arXiv:1107.2948, 2011.

[25] Tomoki Isogai, the Ligo Scientific Collaboration, and the Virgo Collaboration, "Used percentage veto for LIGO and virgo binary inspiral searches", *Journal of Physics: Conference Series* 243(1):012005, 2010.

[26] A. Parameswaran, "Vetoes for transient gravitational-wave triggers using instrumental coupling models", LIGO Document G1100437, 2011.

[27] K. Hodge, "Multivariate Statistical Classification for Detector Characterization and Data Quality - Special Report to June 30 2011 DetChar Call", LIGO Document G1100790, 2011.

[28] S. Klimenko, "WaveMon and Burst FOMs - LSC Meeting, March 15 - 18th, Livingston, LA", LIGO Document G040173, 2004.

[29] C. Costa and C. Torres, "The Critical Coupling Likelihood Method: A new approach for seamless integration of environmental and operating conditions of gravitational wavedetectors into gravitational wave searches", LIGO Document P1000187, 2011.

[30] P. J. Sutton, "S3 Performance of the LIGO Interferometers as Measured by SenseMonitor - LSC Meeting, March 15 - 18th, Livingston, LA", LIGO Document G040102, 2004

[31] M. Wąs, F. Robinet, and N. Leroy, “GRB/SN range, a good burst figure of merit?”, LIGO Document G1100565, 2011.

[32] J. R. Smith *et al.*, “Linear projection of technical noise for interferometric gravitational-wave detectors”, *Classical and Quantum Gravity* 23(2):527, 2006.

[33] J. R. Leong *et al.*, “Absolute amplitude calibration of GEO600”, LIGO Document P1100083, 2011.

[34] Data analysis report for the STAC and EGO council (June 2011). The Virgo Collaboration, Virgo note, VIR-0290A-11, June 2011.

[35] Stability of the timing system during VSR2. L. Rolland, N. Letendre, A. Masserot and B. Mours, Virgo note,VIR-0354A-10, June 2010.

[36] VSR2 mirror and marionette actuator calibration. L. Rolland, Virgo note, VIR-0076B-10, February 2010.

[37] Calibration status in September 2009. L. Rolland, Virgo note, VIR-NOT-00576A-09, October 2009.

[38] Calibration and sensitivity of the Virgo detector during its second science run. T. Accadia et al., (Virgo collaboration), Class. Quantum Grav. 28 (2011) 025005.

[39] Preliminary VSR3 calibration (July 2010). L. Rolland, Virgo note, VIR-0477A-10, August 2010.

[40] Virgo calibration during VSR3. L. Rolland, Virgo note, VIR-0610A-10, November 2010.

[41] h(t) reconstruction for VSR2. B. Mours and L. Rolland, Virgo note, VIR-0340A-10, May 2010.

[42] h(t) reconstruction for VSR3. B. Mours and L. Rolland, Virgo note, VIR-0056A-11, February 2011.

[43] Power and timing calibration of the photon calibrator for VSR2. T. Accadia, L. Rolland and B. Mours, Virgo note, VIR-0404A-10, June 2010.

[44] B.Canuel and E.Genin, VIR-0375A-10 Determination of back scattering and direct reflection recoupling from single optics: application to the End Benches.

[45] Saccorotti G, Piccinini D, Cauchie L and Fiori I (2011) Seismic Noise by Wind Farms: A Case Study from the Virgo Gravitational Wave Observatory Italy. Bull. Seism. Soc. Amer. 101, 2, doi: 10.1785/0120100203.

[46] Web page of the Virgo Environmental Monitoring: https://wwwcascina.virgo.infn.it/EnvMon/

[47] Home page for the Virgo Data Quality group: https://wwwcascina.virgo.infn.it/DataAnalysis/VDBdoc/index.html

[48] Data quality and vetoes for the gravitational-wave burst and inspiral analyses in Virgo’s first science run GWDAW13 proceedings - Class. Quantum Grav. 26 (2009) 204007

[49] Data quality in gravitational wave bursts and inspiral searches in the second Virgo Science Run Florent Robinet et al - Class. Quantum Grav. 27 (2010) 194012

[50] VSR1/VSR2/VSR3 Detector Characterization paper in preparation

[51] QcMoni Documentation: virgo note VIR-075A-08

[52] Virgo Online Frame Distribution: virgo note VIR-087B-08

[53] Detector Monitoring Library : virgo note VIR-074A-08

[54] Data Display Users Manual: virgo note VIR-008A-08

[55] ROOT based visual environment for GW data analysis: lappweb.in2p3.fr/virgo/vega

[56] Virgo Monitoring web pages: https://wwwcascina.virgo.infn.it/MonitoringWeb

[57] Virgo SpectroMoni documentation: virgo note VIR-076A-08

[58] web interface to Noise DataBase (Virgo internal notes VIR-0227A-11 and VIR-0226B-11)

[59] The NMAPI to interface to database directly via web (Virgo internal note VIR-0226B-11)

[60] https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php

[61] Lines Database Web Interface Software Requirements, VIR-0227A-11

[62] 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.

[63] I. W. Harry, S. Fairhurst, "A targeted coherent search for gravitational waves from compact binary coalescences," Phys. Rev. **D83**, 084002 (2011). [arXiv:1012.4939 [gr-qc]].

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

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

[66] 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.

[67] 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.

[68] 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.

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

[70] 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.

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

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

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

[74] 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.

[75] 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.

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

[77] 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.

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

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

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

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

[82] 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.

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

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

[85] 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.

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

[87] 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.

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

[89] 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).

[90] 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).

[91] 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.

[92] 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).

[93] 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.

[94] 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.

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

[96] 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.

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

[98] 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.

[99] 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.

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

[101] 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.

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

[103] 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.

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

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

[106] 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.

[107] 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.

[108] 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. Phys. Rev. D 82 (2010) 102001.

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

[110] 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.

[111] 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.

[112] 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.

[113] 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).

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

[115] 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.

[116] T. Bulik, K. Belcyznki, and A. Prestwich. IC10 X-1/NGC300 X-1: The Very Immediate Progenitors of BH-BH Binaries. *Astrophysical Journal*, 730:140, 2011.

[117] 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.

[118] 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.

[119] 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.

[120] 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.

[121] 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).

[122] 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.

[123] 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.

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

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

[126] 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).

[127] 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.

[128] Citation for coherent methods in CBC.

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

[130] gstlal: GW analysis routines with gstreamer elements. https://www.lsc-group.phys.uwm.edu/daswg/projects/gstlal.html

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

[132] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Implementation and testing of the first prompt search for electromagnetic counterparts to gravitational wave transients (AKA the LUMIN paper). In preparation for submission to Astron. & Astroph. (2011). LIGO Document P1000061.

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

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

[135] LVAlert: https://www.lsc-group.phys.uwm.edu/daswg/projects/lvalert.html .

[136] 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`.

[137] CBC group S6a quarterly report: LIGO Document G1000037. CBC group S6b quarterly report: LIGO Document G1000649. CBC group S6c quarterly report: LIGO Document G1001141.

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

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

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

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

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

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

[144] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Search for gravitational waves from binary black hole inspiral, merger and ringdown. Phys. Rev. D 83 (2011) 122005.

[145] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Implications For The Origin Of GRB 051103 From LIGO Observations. In preparation for submission to Ap.J. LIGO Document P1000097.

[146] S. Privitera for the LSC and Virgo. Sensitivity to Gravitational Waves from Compact Binary Coalescences Achieved during LIGO's Fifth and Virgo's First Science Run. LIGO Document T0900499 (2010).

[147] S. Privitera for the LSC and Virgo. Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO's Sixth and Virgo's Second and Third Science Runs. LIGO Document T1100338 (2011).

[148] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3. LIGO Document P1100034, in preparation for submission to Phys. Rev. D (2011).

[149] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Parameter Estimation for GW100916 (Blind Injection!) LIGO Document P1100030, unpublished (2011).

[150] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). Evidence for the Direct Detection of Gravitational Waves from a Black Hole Binary Coalescence (Blind Injection!) LIGO Document P1000146, unpublished (2011).

[151] F. Marion *et al.*. Detection Committee report. LIGO Document T1100114 (2011).

[152] J. Abadie et al. (The LIGO Scientific Collaboration and Virgo Collaboration). The Low-Latency Search for Binary Inspirals and their Electromagnetic Counterparts in LIGO S6 and Virgo VSR3. In preparation for submission to Astron. & Astroph. LIGO Document P1100065.

[153] S. Anderson and R. Williams, LIGO Data Management Plan. LIGO Document M1000066 (2011). See also: A. Weinstein *et al.*, An Astrophysical Metric for LIGO Open Data Release. LIGO Document T1000414 (2010).

[154] The NR-AR project wiki: `https://www.ninja-project.org/doku.php?id=nrar:home`.

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

[156] 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.

[157] 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.

[158] 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.

[159] A. Corsi and B. J. Owen. Maximum gravitational-wave energy emissible in magnetar flares. *Phys. Rev. D.*, 83(10):104014, May 2011.

[160] Y. Levin and M. van Hoven. On the excitation of f-modes and torsional modes by magnetar giant flares. *arXiv:1103.0880*, March 2011.

[161] N. Andersson and K. D. Kokkotas. Towards gravitational wave asteroseismology. MNRAS, 299:1059–1068, October 1998.

[162] J. A. de Freitas Pacheco. Do soft gamma repeaters emit gravitational waves? *Astronomy and Astrophysics*, 336:397–401, August 1998.

[163] K. Ioka. Magnetic deformation of magnetars for the giant flares of the soft gamma-ray repeaters. *MNRAS*, 327:639–662, October 2001.

[164] B. J. Owen. Maximum Elastic Deformations of Compact Stars with Exotic Equations of State. *Physical Review Letters*, 95(21):211101–+, November 2005.

[165] 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.

[166] B. Zink, P. D. Lasky, and K. D. Kokkotas. Are gravitational waves from giant magnetar flares observable? *ArXiv e-prints*, July 2011.

[167] 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.

[168] 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.

[169] 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.

[170] 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.

[171] R. Khan. Searching for gravitational wave fingerprints of SGR QPOs, poster at 2008 APS April .

[172] J. Abadie et al. Search for gravitational wave bursts from six magnetars. *The Astrophysical Journal Letters*, 734(2):L35, 2011.

[173] R. X. Xu. Solid quark stars? *The Astrophysical Journal Letters*, 596(1):L59–L62, 2003.

[174] 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.

[175] N. Gehrels et al. A new γ-ray burst classification scheme from GRB060614. *Nature*, 444:1044–1046, December 2006.

[176] S. Campana et al. The shock break-out of grb 060218/sn 2006aj. *Nature*, 442:1008–1010, 2006.

[177] 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.

[178] Jens Hjorth et al. A very energetic supernova associated with the gamma-ray burst of 29 march 2003. *Nature*, 423:847–850, 2003.

[179] T. J. Galama et al. Discovery of the peculiar supernova 1998bw in the error box of grb980425. *Nature*, 395:670, 1998.

[180] 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.

[181] E. Nakar. Short-hard gamma-ray bursts. *Physics Reports*, 442:166–236, April 2007.

[182] 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.

[183] 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.

[184] 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.

[185] C. Pankow et al. A burst search for gravitational waves from binary black holes. *Classical and Quantum Gravity*, 26:204004, 2009.

[186] 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.

[187] E. Troja, S. Rosswog, and N. Gehrels. Precursors of Short Gamma-ray Bursts. ApJ, 723:1711, November 2010.

[188] B. M. S. Hansen and M. Lyutikov. Radio and X-ray signatures of merging neutron stars. MNRAS, 322:695, April 2001.

[189] V. M. Lipunov and I. E. Panchenko. Pulsars revived by gravitational waves. A&A, 312:937, August 1996.

[190] D. A. Kann et al. The Afterglows of Swift-era Gamma-ray Bursts. I. Comparing pre-Swift and Swift-era Long/Soft (Type II) GRB Optical Afterglows. ApJ, 720:1513, September 2010.

[191] D. A. Kann et al. The Afterglows of Swift-era Gamma-Ray Bursts. II. Type I GRB versus Type II GRB Optical Afterglows. ApJ, 734:96, June 2011.

[192] B. D. Metzger, G. Martínez-Pinedo, S. Darbha, E. Quataert, A. Arcones, D. Kasen, R. Thomas, P. Nugent, I. V. Panov, and N. T. Zinner. Electromagnetic counterparts of compact object mergers powered by the radioactive decay of r-process nuclei. MNRAS, 406:2650, August 2010.

[193] V. V. Usov and J. I. Katz. Low frequency radio pulses from gamma-ray bursts? A&A, 364:655, December 2000.

[194] W. H. Lee and E. Ramirez-Ruiz. The progenitors of short gamma-ray bursts. *New Journal of Physics*, 9:17–+, January 2007.

[195] S. E. Woosley and J. S. Bloom. The Supernova Gamma-Ray Burst Connection. *Ann. Rev. Astron. Astrophys.*, 44:507, September 2006.

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

[197] C. L. Fryer, D. E. Holz, and S. A. Hughes. *Astrophys. J.*, 565:430, 2002.

[198] 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.

[199] Anthony L. Piro and Eric Pfahl. Fragmentation of collapsar disks and the production of gravitational waves. *The Astrophysical Journal*, 658(2):1173–1176, 2007.

[200] 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.

[201] 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.

[202] E. Waxman and J. Bahcall. High Energy Neutrinos from Cosmological Gamma-Ray Burst Fireballs. *Physical Review Letters*, 78:2292–2295, March 1997.

[203] J. P. Rachen and P. Mészáros. Photohadronic neutrinos from transients in astrophysical sources. *Phys. Rev. D*, 58(12):123005–+, December 1998.

[204] 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.

[205] E. Waxman and J. N. Bahcall. Neutrino Afterglow from Gamma-Ray Bursts: 10^{18} eV. *Astrophys. Journal*, 541:707, October 2000.

[206] P. Mészáros and E. Waxman. TeV Neutrinos from Successful and Choked Gamma-Ray Bursts. *Phys. Rev. Lett.*, 87(17):171102–+, October 2001.

[207] S. Razzaque, P. Mészáros, and E. Waxman. High Energy Neutrinos from Gamma-Ray Bursts with Precursor Supernovae. *Phys. Rev. Lett.*, 90(24):241103, June 2003.

[208] 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.

[209] C. D. Dermer and A. Atoyan. High-Energy Neutrinos from Gamma Ray Bursts. *Phys. Rev. Lett.*, 91(7):071102, August 2003.

[210] 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.

[211] 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, November 2006.

[212] S. Ando and J. F. Beacom. Revealing the Supernova Gamma-Ray Burst Connection with TeV Neutrinos. *Phys. Rev. Lett.*, 95(6):061103, August 2005.

[213] 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.

[214] B. Baret, I. Bartos, B. Bouhou, A. Corsi, I. D. Palma, C. Donzaud, V. V. Elewyck, C. Finley, G. Jones, A. Kouchner, S. Márka, Z. Márka, L. Moscoso, E. Chassande-Mottin, M. A. Papa, T. Pradier, P. Raffai, J. Rollins, and P. Sutton. Bounding the time delay between high-energy neutrinos and gravitational-wave transients from gamma-ray bursts. *Astroparticle Physics*, 35:1–7, August 2011.

[215] K. Ioka. Magnetic deformation of magnetars for the giant flares of the soft gamma-ray repeaters. MNRAS, 327:639–662, October 2001.

[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] Eric Chassande-Mottin and et al. Multimessenger astronomy with the Einstein Telescope. *Gen. Relat. Gravit.*, pages 1–28, 2010. 10.1007/s10714-010-1019-z.

[219] V. van Elewyck and et al. Joint Searches Between Gravitational-Wave Interferometers and High-Energy Neutrino Telescopes:. Science Reach and Analysis Strategies. *Int. J. Mod. Phys. D*, 18:1655–1659, 2009.

[220] H.-T. Janka, K. Langanke, A. Marek, G. Martínez-Pinedo, and B. Müller. Theory of core-collapse supernovae. Phys. Rep., 442:38, 2007.

[221] C. D. Ott. Probing the core-collapse supernova mechanism with gravitational waves. *Class. Quant. Grav.*, 26(20):204015, October 2009.

[222] 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.

[223] C. L. Fryer and K. C. B. New. Gravitational Waves from Gravitational Collapse. *Living Reviews in Relativity*, 14:1–+, January 2011.

[224] S. van den Bergh and G. A. Tammann. Galactic and extragalactic supernova rates. ARA&A, 29:363, 1991.

[225] G. A. Tammann, W. Loeffler, and A. Schroeder. The Galactic supernova rate. ApJS, 92:487, June 1994.

[226] M. D. Kistler, H. Yüksel, S. Ando, J. F. Beacom, and Y. Suzuki. Core-collapse astrophysics with a five-megaton neutrino detector. Phys. Rev. D, 83(12):123008, June 2011.

[227] 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.

[228] G. Pagliaroli, F. Vissani, E. Coccia, and W. Fulgione. Neutrinos from Supernovae as a Trigger for Gravitational Wave Search. *Phys. Rev. Lett.*, 103(3):031102, July 2009.

[229] M. Ikeda et al. Search for Supernova Neutrino Bursts at Super-Kamiokande. ApJ, 669:519, November 2007.

[230] K. Scholberg. Future supernova neutrino detectors. *J. Phys. Conf. Ser.*, 203(1):012079, January 2010.

[231] I. Leonor, L. Cadonati, E. Coccia, S. D'Antonio, A. Di Credico, V. Fafone, R. Frey, W. Fulgione, E. Katsavounidis, C. D. Ott, G. Pagliaroli, K. Scholberg, E. Thrane, and F. Vissani. Searching for prompt signatures of nearby core-collapse supernovae by a joint analysis of neutrino and gravitational wave data. *Class. Quant. Grav.*, 27(8):084019, April 2010.

[232] C. S. Kochanek, J. F. Beacom, M. D. Kistler, J. L. Prieto, K. Z. Stanek, T. A. Thompson, and H. Yüksel. A Survey About Nothing: Monitoring a Million Supergiants for Failed Supernovae. ApJ, 684:1336, September 2008.

[233] S. Horiuchi, J. F. Beacom, C. S. Kochanek, J. L. Prieto, K. Z. Stanek, and T. A. Thompson. The Cosmic Core-collapse Supernova Rate does not match the Massive-Star Formation Rate. *ArXiv:1102.1977*, February 2011.

[234] 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.

[235] 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.

[236] 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.

[237] 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.

[238] N. Andersson, G. L. Comer, and R. Prix. The superfluid two-stream instability. 354:101–110, October 2004.

[239] A. Melatos, C. Peralta, and J. S. B. Wyithe. Avalanche Dynamics of Radio Pulsar Glitches. ApJ, 672:1103–1118, January 2008.

[240] Y. Levin and G. Ushomirsky. Crust-core coupling and r-mode damping in neutron stars: a toy model. MNRAS, 324:917, July 2001.

[241] J. Abadie, B. P. Abbott, R. Abbott, T. Accadia, F. Acernese, R. Adhikari, P. Ajith, B. Allen, G. Allen, E. Amador Ceron, and et al. Search for Gravitational-wave Inspiral Signals Associated with Short Gamma-ray Bursts During LIGO's Fifth and Virgo's First Science Run. ApJ, 715:1453–1461, June 2010.

[242] S. Zwart and S.L.W. McMillan. The Runaway Growth of Intermediate-Mass Black Holes in Dense Star Clusters. *Astrophys. J.*, 576:899, 2002.

[243] Tod E. Strohmayer and Richard F. Mushotzky. Evidence for an Intermediate Mass Black Hole in NGC 5408 X-1. *Astrophys. J.*, 703:1386, 2009.

[244] L. Zampieri et al. In X-Ray Astronomy 2009, Present Status. *Multiwavelength Approach and Future Perspectives, edited by L. A. A. Comastri, M. Cappi (AIP)*, 2010.

[245] P.Casella et al. Weighing the black holes in ultraluminous X-ray sources through timing. *Mon. Not. Roy. Astron. Soc.*, 387(4):5, 2008.

[246] Alessandro Patruno and Luca Zampieri. The black hole in NGC 1313 X-2: constraints on the mass from optical observations. *Mon. Not. Roy. Astron. Soc.*, 403:L69–L73, 2010.

[247] M. Safonova and P. Shastri. *Astrophysics and Space Science*, 325:47, 2010.

[248] K. Gultekin, M. C. Miller, and D. P. Hamilton. *Astrophys. J.*, 616:221, 2004.

[249] S. Umbreit, J. M. Fregeau, and F. A. Rasio. *Astrophys. J.*, 719:915, 2010.

[250] Kocsis O'Leary and Loeb. Gravitational waves from scattering of stellar-mass black holes in galactic nuclei. *Mon. Not. R. Astron. Soc.*, 395:2127âĂŞ2146, 2009.

[251] 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.

[252] A. Vilenkin and E. Shellard. *Cosmic strings and other Topological Defects*. Cambridge University Press, 2000.

[253] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cosmic strings. *Phys. Rev. Lett.*, 85(18):3761–3764, Oct 2000.

[254] Thibault Damour and Alexander Vilenkin. Gravitational wave bursts from cusps and kinks on cosmic strings. *Phys. Rev. D*, 64(6):064008, Aug 2001.

[255] Thibault Damour and Alexander Vilenkin. Gravitational radiation from cosmic (super)strings: Bursts, stochastic background, and observational windows. *Phys. Rev. D*, 71(6):063510, Mar 2005.

[256] AurÃl'lien A Fraisse. Limits on defects formation and hybrid inflationary models with three-year wmap observations. *Journal of Cosmology and Astroparticle Physics*, 2007(03):008, 2007.

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

[258] P. Binétruy, A. Bohé, T. Hertog, and D. A. Steer. Gravitational wave signatures from kink proliferation on cosmic (super-) strings. *Phys. Rev. D*, 82(12):126007, Dec 2010.

[259] J. Abadie et al. Search for gravitational waves associated with the august 2006 timing glitch of the vela pulsar. *Phys. Rev. D*, 83(4):042001, Feb 2011.

[260] 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.

[261] 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.

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

[263] B. Abbott et al. First joint search for gravitational-wave bursts in LIGO and GEO 600 data. *Class. Quantum Grav.*, 25:245008, 2008.

[264] 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.

[265] Omega Pipeline documentation and wiki, https://trac.ligo.caltech.edu/omega/.

[266] D Buskulic, the LIGO Scientific Collaboration, and the Virgo Collaboration. Very low latency search pipeline for low mass compact binary coalescences in the ligo s6 and virgo vsr2 data. *Classical and Quantum Gravity*, 27(19):194013, 2010.

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

[268] 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.

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

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

[271] 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.

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

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

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

[275] 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.

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

[277] 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.

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

[279] Antony C. Searle, Patrick J. Sutton, and Massimo Tinto. Bayesian detection of unmodeled bursts of gravitational waves. *Classical and Quantum Gravity*, 26(15):155017–+, August 2009.

[280] 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.

[281] JG Rollins. *Multimessenger Astronomy with Low-Latency Searches for Transient Gravitational Waves*. PhD thesis, Columbia University, 2011.

[282] L Blackburn, L Cadonati, S Caride, S Caudill, S Chatterji, N Christensen, J Dalrymple, S Desai, A Di Credico, G Ely, J Garofoli, L Goggin, G Gonzalez, R Gouaty, C Gray, A Gretarsson, D Hoak, T Isogai, E Katsavounidis, J Kissel, S Klimenko, R A Mercer, S Mohapatra, S Mukherjee, F Raab, K Riles, P Saulson, R Schofield, P Shawhan, J Slutsky, J R Smith, R Stone, C Vorvick, M Zanolin, N Zotov, and J Zweizig. The lsc glitch group: monitoring noise transients during the fifth ligo science run. *Classical and Quantum Gravity*, 25(18):184004, 2008.

[283] J Slutsky, L Blackburn, D A Brown, L Cadonati, J Cain, M Cavaglia, S Chatterji, N Christensen, M Coughlin, S Desai, G Gonzalez, T Isogai, E Katsavounidis, B Rankins, T Reed, K Riles, P Shawhan, J R Smith, N Zotov, and J Zweizig. Methods for reducing false alarms in searches for compact binary coalescences in ligo data. *Classical and Quantum Gravity*, 27(16):165023, 2010.

[284] S Chatterji, L Blackburn, G Martin, and E Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21(20):S1809, 2004.

[285] Joshua R Smith et al. A hierarchical method for vetoing noise transients in gravitational-wave detectors. 2011.

[286] 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.

[287] 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.

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

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

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

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

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

[293] B. Abbott et al. Search for gravitational waves associated with the gamma ray burst grb030329 using the ligo detectors. Phys. Rev. D, 72(4):042002, 2005.

[294] LIGO Scientific Collaboration and K. Hurley. Implications for the Origin of GRB 070201 from LIGO Observations. ApJ, 681:1419–1430, July 2008.

[295] F Acernese et al. Gravitational waves by gamma-ray bursts and the virgo detector: the case of grb 050915a. *Classical and Quantum Gravity*, 24(19):S671–S679, 2007.

[296] F Acernese et al. Search for gravitational waves associated with grb 050915a using the virgo detector. Class. Quant. Grav., 25(22):225001, 2008.

[297] K. Hurley et al. The status and future of the third interplanetary network. In Charles Meegan, Chryssa Kouveliotou, and Neil Gehrels, editors, *GAMMA-RAY BURST: Sixth Huntsville Symposium*, volume 1133 of *American Institute of Physics Conference Series*, pages 55–57, May 2009.

[298] 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.

[299] 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.

[300] S. D. Mohanty. Population study of gamma ray bursts: detection sensitivity and upper limits. *Classical and Quantum Gravity*, 23:723–+, October 2006.

[301] 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.

[302] B. P. Abbott et al. Stacked Search for Gravitational Waves from the 2006 SGR 1900+14 Storm. ApJ Lett., 701:L68–L74, August 2009.

[303] L Matone and S MÃąrka. Search algorithm for the detection of long-duration narrow-band transients in gw interferometers. *Classical and Quantum Gravity*, 24(19):S649, 2007.

[304] B. Abbott et al. Search for gravitational wave radiation associated with the pulsating tail of the SGR 1806–20 hyperflare of 27 December 2004 using LIGO. Phys. Rev. D, 76:062003, 2007.

[305] IceCube. http://icecube.wisc.edu.

[306] ANTARES. http://antares.in2p3.fr/.

[307] 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.

[308] 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.

[309] 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.

[310] K. Scholberg. Supernova neutrino detection. *ArXiv Astrophysics e-prints*, January 2007.

[311] 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.

[312] D. R. Lorimer. Blind surveys for radio pulsars and transients. *AIP Conference Proceedings*, 1357, 2011.

[313] J. M. Cordes et al. Arecibo pulsar survey using alfa. i. survey strategy and first discoveries. ApJ, 637:446–455, 2006.

[314] Tanmay Vachaspati. Cosmic sparks from superconducting strings. Phys. Rev. Lett., 101(14):141301, Sep 2008.

[315] M. van der Sluys, V. Raymond, I. Mandel, C. Röver, N. Christensen, V. Kalogera, R. Meyer, and A. Vecchio. Parameter estimation of spinning binary inspirals using Markov chain Monte Carlo. *Classical and Quantum Gravity*, 25(18):184011–+, September 2008.

[316] T. Littenberg and N. Cornish. Separating Gravitational Wave Signals from Instrument Artifacts. *APS Meeting Abstracts*, pages 14004–+, February 2010.

[317] 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.

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

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

[320] 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.

[321] 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.

[322] John Veitch and Alberto Vecchio. Assigning confidence to inspiral gravitational wave candidates with Bayesian model selection. *Class. Quant. Grav.*, 25:184010, 2008.

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

[324] 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.

[325] 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.

[326] 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.

[327] 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.

[328] 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.

[329] C. W. Stubbs. *Class. Quantum Grav.*, 25:184033, 2008.

[330] D. B. Fox et al. *Nature*, 437:845–50, 2005.

[331] N. Gehrels, J. K. Cannizzo, and J. P. Norris. *New Journal of Physics*, 9, 2007.

[332] L. Li and B. Pacyznski. *Astrophys. J.*, 507:L59–62, 1998.

[333] S. Kulkarni. Modeling supernova-like explosions associated with gamma-ray bursts with short durations. 2005.

[334] J. Kanner, T. L. Huard, S. Marka, D. C. Murphy, J. Piscionere, M. Reed, and P. Shawhan. *Class. Quantum Grav.*, 25:184034, 2008.

[335] J. Sylvestre. *Astrophys. J.*, 591:1152–6, 2003.

[336] J. S. Bloom et al. Astro2010 decadal survey whitepaper: Coordinated science in the gravitational and electromagnetic skies. 2009.

[337] S. Desai, E. O. Kahya, and R. P. Woodard. *Phys. Rev. D*, 77:124041, 2008.

[338] N. Dalal, D. E. Holz, S. Hughes, and B. Jain. *Phys. Rev. D*, 74:063006, 2006.

[339] 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.

[340] 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.

[341] Teviet Creighton, Fredrick A. Jenet, and Richard H. Price. Pulsar timing and spacetime curvature. 2008.

[342] 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.

[343] Stephon Alexander, Lee Samuel Finn, and Nicolas Yunes. A gravitational-wave probe of effective quantum gravity. *Phys. Rev.*, D78:066005, 2008.

[344] Clifford M. Will. The confrontation between general relativity and experiment: A 1998 update. 1998.

[345] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004).

[346] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. Lett. **94**, 181103 (2005).

[347] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 102004 (2005).

[348] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 082001 (2007)

[349] L. Bildsten, Astrophys. J. **501**, L89 (1998).

[350] G. Ushomirsky, C. Cutler and L. Bildsten, Mon. Not. Roy. Astron. Soc. **319**, 902 (2000).

[351] C. Cutler, Phys. Rev. D **66**, 084025 (2002).

[352] B. J. Owen, Phys. Rev. Lett. **95**, 211101 (2005).

[353] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio and N. Andersson, Phys. Rev. D **58**, 084020 (1998).

[354] N. Andersson, K. D. Kokkotas and N. Stergioulas, Astrophys. J. **516**, 307 (1999).

[355] A. Melatos and D. J. B. Payne, Astrophys. J. **623**, 1044 (2005).

[356] B. J. Owen, Class. Quant. Grav. **23**, S1 (2006).

[357] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).

[358] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).

[359] D. I. Jones and N. Andersson, Mon. Not. Roy. Astron. Soc. **331**, 203 (2002).

[360] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).

[361] K. S. Thorne, in *Three hundred years of gravitation,* eds. S. W. Hawking and W. Israel (Cambridge University Press: 1987).

[362] R.J Dupuis, Ph.D. Thesis, University of Glasgow (2004); `LIGO-P050014-00-R`.

[363] B Abbott *et al.* (The LIGO Scientific Collaboration), Ap J. Lett. **683** 45 (2008).

[364] B Abbott *et al.* (The LIGO Scientific Collaboration), Ap. J. **713** 671 (2010).

[365] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D, **76**, 042001 (2007).

[366] M. Pitkin and G. Woan, Phys. Rev. D, **76**, 042006 (2007).

[367] P. Astone *et al.*, Phys. Rev. D **65**, 022001 (2001).

[368] S. Frasca and C. Palomba, Class. Quant. Grav. **21**, S1645 (2004).

[369] D. Passuello and S. Braccini, Virgo Note (VIR-046A-07).

[370] S. Frasca, P. Astone and C. Palomba, Class. Quant. Grav. **22**, S1013 (2005).

[371] P. Astone, S. Frasca and C. Palomba, Class. Quant. Grav. **22**, S1197 (2005).

[372] C. Palomba, P. Astone and S. Frasca, Class. Quant. Grav. **22**, S1255 (2005).

[373] F. Acernese *et al.*, Class. Quant. Grav. **24**, S491 (2007).

[374] S. D'Antonio, S. Frasca, C. Palomba, submitted to CQG (2009).

[375] S. Frasca and C. La Posta, Il Nuovo Cimento, **14 C**, (1991).

[376] F. Antonucci, P. Astone, S. D'Antonio, S. Frasca, C. Palomba, CQG **25**, 184015 (2008).

[377] P. Jaranowski and Andrzej Królak, submitted to CQQ (2010), arXiv:1004.0324.

[378] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca and C. Palomba, Phys. Rev. D **70**, 082001 (2004).

[379] V. Dergachev, “Description of PowerFlux Algorithms and Implementation”, *LIGO Technical Document* LIGO-T050186 (2005), available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[380] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008).

[381] 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).

[382] 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/`.

[383] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008).

[384] V. Dergachev, “On blind searches for noise dominated signals: a loosely coherent approach”, arXiv:1003.2178, March 2010.

[385] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Phys. Rev. D **77**, 082001 (2008).

[386] 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)

[387] 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

[388] L. Sancho de la Jordana, A.M. Sintes, Class. Quant. Grav. **25** 184014 (2008).

[389] Ransom, S. M., Cordes, J. M., & Eikenberry, S. S. 2003, Astrophys. J. **589**, 911

[390] C. Messenger and G. Woan, Class. Quant. Grav., **24**, 469 (2007).

[391] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009).

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

[393] H. Pletch and B. Allen, Phys. Rev. Lett. **103** 181102 (2009).

[394] P. Jaranowski, A. Krolak, and B. F. Schutz, “Data analysis of gravitational-wave signals from pulsars. I. The signal and its detection”, Phys. Rev. D **58**, 063001-24, (1998).

[395] C.-Y. Ng and R. W. Romani, “Fitting Pulsar Wind Tori. II. Error Analysis and Applications”, arXiv.:0710.4168v1

[396] J. Abadie *et al.* (The LIGO Scientific Collaboration), “First search for gravitational waves from the youngest known neutron star,” Astrophys. J. **722**, 1504–1513 (2010).

[397] K. Wette *et al.* [LIGO Collaboration], Class. Quant. Grav. **25** 235011 (2008).

[398] S. Zane *et al.*, to appear in MNRAS, arXiv:1009.0209 (October 2009).

[399] P. Patel *et al.* “Implementation of barycentric resampling for continous wave searches in gravitational wave data”, LIGO P0900301, January 2010.

[400] R. E. Rutledge, D. B. Fox and A. H. Shevchuk, arXiv:0705.1011 [astro-ph].

[401] S. Braccini *et al.*, Phys. Rev. D **83** 044033 (2011).

[402] R. Prix, S. Giampanis and C. Messenger, to appear in Phys. Rev. D, arXiv:1104.1704 (April 2011).

[403] E. Goetz and K. Riles, submitted to Class. Quant. Grav., arXiv:1103.1301 (March 2011).

[404] J.H. Taylor and G.R. Huguenin, Nature **221**, 816 (1969).

[405] P.C. Gregory and T.J. Loredo, ApJ. bf 398, 146 (1992).

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

[407] D. I. Jones, MNRAS **402**, 4, pp. 2503-2519 (2010).

[408] J. Abadie *et al.* (The LIGO Scientific Collaboratio and The Virgo Collaboration), to appear in ApJ, arXiv:1104.2712 (2011).

[409] B. Abbott *et al.* (The LIGO Scientific Collaboration), Ap. J. **659**, 918 (2007).

[410] B. Abbott *et al.* (The LIGO Scientific Collaboration and the ALLEGRO Collaboration), Phys. Rev. D **76**, 022001 (2007).

[411] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007)

[412] Flanagan É É 1993 *Phys. Rev.* **D48** 2389; astro-ph/9305029

[413] Maggiore, M 2008 *Gravitational waves* Oxford: Oxford University Press

[414] Allen B and Romano J D 1999 *Phys. Rev.* **D59** 102001; gr-qc/9710117.

[415] G. Cella *et al.*, Class. Quant. Grav. **24**, S639 (2007)

[416] B. Abbott et al., "Analysis of first LIGO science data for stochastic gravitational waves", Phys. Rev. **D69**, 122004 (2004).

[417] Ballmer S W 2006 "A radiometer for stochastic gravitational waves", *Class. Quant. Grav.* **23** S179; gr-qc/0510096

[418] S. Mitra et al., "Gravitational wave radiometry: Mapping a stochastic gravitational wave background", *Phys. Rev.* **D77**, 042002 (2008).

[419] N.V. Fotopoulos, "Identifying Correlated Environmental Noise in Co-Located Interferometers with Application to Stochastic Gravitational Wave Analysis", in preparation for GWDAW 10 proceedings.

[420] B. Abbott et al, "An upper limit on the stochastic gravitational-wave background of cosmological origin", Nature **460**, 990 (2009).

[421] 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).

[422] http://wwwcascina.virgo.infn.it/DataAnalysis/Noise/nap_index.html

[423] T. Regimbau and J.A. de Freitas Pacheco, "Gravitational wave background from magnetars", Astron. and Astrophys. 447, 1 (2006).

[424] T. Regimbau and J.A. de Freitas Pacheco, ”Cosmic background of gravitational waves from rotating neutron stars”, Astron. and Astrophys. 376, 381 (2001).

[425] T. Regimbau and V. Mandic, ”Astrophysical sources of a stochastic gravitational-wave background”, Class. Quantum Grav. 25, 184018 (2008).

[426] X. Siemens, V. Mandic, and J. Creighton, ”Gravitational-Wave Stochastic Background from Cosmic Strings”, Phys. Rev. Lett. 98, 111101 (2007).

[427] V. Mandic and A. Buonanno, ”Accessibility of the Pre-Big Bang Models to LIGO”, Phys. Rev. D 73, 063008 (2006).

[428] E. Thrane et al, ”Long gravitational-wave transients and associated detection strategies for a network of terrestrial interferometers”, Phys. Rev. D 83, 083004 (2011).

[429] A.L. Piro and E. Pfahl, Astrop. J. 658, 1173 (2007).

[430] M.H.P.M. van Putten, Astrop. J. Lett., 575, 71 (2002).

[431] B.J. Owen et al, ”Gravitational waves from hot young rapidly rotating neutron stars”, Phys. Rev. D 58, 084020, (1998)

[432] L. Dessart et al, "Multidimensional Radiation/Hydrodynamic Simulations of Proto-Neutron Star Convection", Astrop. J. 645, 534 (2006).

[433] M. Shibata et al, ”Three-dimensional simulations of stellar core collapse in full general relativity: Nonaxisymmetric dynamical instabilities”, Phys. Rev. D 71, 024014 (2005).

[434] A. Corsi and P. Mészáros, Class. Quant. Grav. 26, 204016 (2009)

[435] D.A. Frail et al, Astrop. J. Lett. 562, L55 (2001).

[436] E. Thrane, https://wiki.ligo.org/Main/RoadmapForNewSearch, (2011).

[437] R.E. Morganstern and H.-Y. Chiu, ”Scalar Gravitational Fields in Pulsating Stars”, Phys. Rev. 157, 1228 (1967).

[438] D.C. Robinson and J. Winicour, ”An Observable Peculiarity of the Brans-Dicke Radiation Zone”, Phys. Rev. Lett. 22, 198 (1969).

[439] The Virgo Computing Model, VIR-TRE-DIR-7100-102, May 2000.