LIGO SCIENTIFIC COLLABORATION
VIRGO COLLABORATION

Document Type	**LIGO-T0900389-01-Z** **VIR-040A-09**

The LSC-Virgo white paper on gravitational wave data analysis

Science goals, status and plans, priorities

(2009-2010 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 2009/08/14

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, and finalized by early fall. This is the document for 2009-2010. It is intended to facilitate:

- the understanding of the science that we are doing
- the identification of “holes” in our science plan and of tasks that demand more manpower
- the prioritization of our objectives
- the identification of areas when manpower should be shifted to and/or removed from
- the exploitation of synergies among the work carried out in different search groups
- an harmonious exploitation of common resources

In July 2009 begins the 6th science run of the of the LIGO detectors(S6) and the 2nd science run of the Virgo detector (VSR2). The sensitivity goal for S6/VSR2 is about a 2-fold increase in strain sensitivity with respect to the initial design sensitivity goal, which LIGO achieved during the S5 run. Although a detection is not guaranteed, it is not unexpected.

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. One of the new aspects of the work that we will carry out this year is to begin connecting to these other observatories. In some cases this requires the prompt production of GW-alerts,so to speak, to allow fast pointing of telescopes. In other cases an a-posteriori analysis of triggers suffices. 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 data and think ahead at the era of the advanced detectors, in which we expect routine detections and this type of work constitutes the main science output of our research.

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

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

2 The characterization of the data

2.1 LIGO Detector Characterization

2.1.1 Introduction

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

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

In practice, detector characterization is carried out at several different levels within the LSC and by a variety of scientists focused on different problems. Commissioners operate at the cutting edge of detector characterization, evaluating and updating interferometer noise budgets, as improvements are made between data runs. By the nature of commissioning, long-term stability is difficult to evaluate when such work is most intense. In the past, data runs have served as testing grounds for that stability, and there have been some unpleasant surprises. As experience has accumulated, as background monitoring tools have improved, and as more data have been collected in science mode, the rapidity of diagnostic feedback has improved dramatically. Feedback useful for mitigation and commissioning is now routine. Some investigations focus on interferometer-based detector characterization, such as investigation of line artifacts or environmental disturbances, while others focus on astrophysics-search-targeted artifacts, such as coherent or accidentally coincident glitches that could pollute inspiral and burst searches, or wandering line features that could mimic a pulsar.

As new artifacts are found and new characterization methods developed offline, there is a steady effort to migrate those improvements to the real-time online monitoring for more rapid detection of problems. This online monitoring includes programs run under the Data Monitoring Tool (DMT) environment[1], controls system software (EPICS), and a variety of customized tools written in C++ and Matlab. It also includes a human element, namely the attentiveness and active data exploration by interferometer operators and scientific monitors (scimons).

The LSC Detector Characterization (DetChar) community has a broad membership, including full-time commissioners, full-time data analysts, and many in between. In practice, the DetChar working groups have concentrated most on providing online characterization tools, e.g., DMT monitors and on providing characterization of interferometer data in science runs for astrophysical analysis. Every search working group has active DetChar members, and information flows in both directions, to mutual benefit. *The Det Char group has few members exclusively or even mostly dedicated to the group, which has proven a benefit for broadness but hampers the speed of progress, we hope to have more members committing a substantial* *amoutn of time to the group.*

In the past years, the DetChar efforts have included completing characterization of the S5 data already taken and support for the H2 Astrowatch running. In the next few years, the group wil implement improved diagnostics for the S6 run (with "Enhanced LIGO") that starts in 2009, and plan for Advanced LIGO diagnostics.

In the following subsections, we describe the software infrastructure, the storage of data quality information, and the different working groups working in DetChar and highlight the high priorities for each group.

2.1.2 Software Infrastructure

The interferometer controls system based on EPICS (Experimental Physics and Industrial Control Sytesm) software[2] is essential to operations. That software includes simple automated monitoring (e.g., alarms for values out of range) and the capability via customized microprocessor programs to carry out more sophisticated monitoring of interferometer state. This real-time controls system provides the first line of defense against wandering detector conditions and records literally thousands of data channels that permit later reconstruction of conditions, if needed. An online Data Viewer program permits engineers and scientists to view selected data channels in the style of an oscilloscope in either real-time or playback.

Closely coupled to the detector controls system is the Global Diagnostics System (GDS) software that includes both the interactive Diagnostic Test Tool (DTT)[3] and the background monitoring of the DMT. The DTT allows rapid exploration of data in the time and Fourier domains and includes user-selected filtering and extensive choice of data sources, real-time or stored in LIGO's distributed archives. The DTT also permits stimulation of detector channels for measuring transfer functions. The same underlying driver is used to inject sinusoidal "calibration lines" and simulated GW signals of various types into the interferometer hardware.

The DMT[3] offers an interactive ROOT-driven[4] environment for exploration and algorithm development and a background-process environment for continuous monitoring. Most DMT detector characterization is carried out via the 24/7 background monitors, which have been written by scientists from well over a dozen LSC institutions. As for the EPICS system, the DMT programs permanently record trended data channels, derived from the original interferometer and environmental data channels, in addition to providing real-time feedback to operators and scientists. That feedback comes in several forms: graphical displays on control room workstations (the most important of which are projected onto the walls), alarms (the most important of which are audible), and status web pages.

Offline detector characterization investigations are carried out using a variety of tools, ranging from offline DMT programs to Matlab, to LAL programs, to TCL scripts examining DMT trends, to simple interactive data viewing with the Data Viewer or ligo_viewer[5]. Many of these offline studies typically work with data products (e.g., trends or triggers) produced by programs upstream in a pipeline. The S5 science run has seen the development and widespread use of the Q-transform-based QScan tool[6] for examining interesting transient phenomenal, along with the use of an event display program, both using spectrograms and whitened time series.

These studies also benefit from the production of reduced data sets in which only selected raw data channels are included, some of which are downsampled for further reduction.

After many years of development, the suite of online DMT monitors is quite mature. Many of the existing DMT monitors have been tested enough that their output is used to produce online "data quality flags" (see **??**). On the other hand, the output of many other existing DMT monitors has not been studied in enough detail to deliver a final product that can be used for cleaning data or in vetoes for the astrophysical searches.

Existing programs monitor the controls state of the interferometers, servo unity-gain frequency, environmental noise (including seismic bands, overflying aircraft and liquid nitrogen dewar shifts), non-Gaussianity, spectral line contamination, glitchiness and non-stationarity, hardware/software overflows, faulty ADC's, timing stability, and spectral stability. In addition, several monitors produce astrophysically motivated figures of merit (FOM's) for display and archiving: sensitivity to inspiral mergers, sensitivity to bursts, sensitivity to stochastic background, and sensitivity to pulsars. For the inspiral search there was also a near-real-time display of results of single detector searches in the control room from template banks run on the observatory computing clusters, and a similar one will be used in S6.

Nonetheless, there is need for additional online monitoring, especially of monitors that produce data quality information ready to use n the astrophysical searches.

We highlight in the next paragraphs the short and intermediate term needs of the DetChar group related to software infrastructure. These activities will be coordinated with the Software Working Group, as described in Section 7.

Short term needs for S6 are:

- selection of channels for reduced data sets, to be used for detector characterization and for data analysis;
- continued upkeep of existing DMT infrastructure, especially what concerns the automated use of its output;
- reviewing results from existing monitors producing online flags and adapting them them to changing instruments if needed. In particular, this refers to DQ flag production for calibration quality, seismic noise, arm power fluctuations, non-stationarity of inspiral range, overflying aircrafts, and digital overflows in feedback loops;
- migration of monitor outputs proven useful to on-line production of data quality flags. Identified monitors relate to calibration lines, upconversion in GW channel, saturated control signals, magnetic transients, liquid nitrogen dewar shifts, artifacts from thermal compensation systems, noise at frequencies of known pulsars, overflows in suspension and tidal channels, and non-gaussianity affecting searches for unmodelled bursts.
- monitors trending and producing alarms for degrading hardware (i.e., optical levers and drifting suspensions);
- monitors trending the stationarity of the data, especially relating to search needs (derived from glitch rates, from variations in sensitivity and spectrograms);
- monitors identifying glitches in the signals from different photodiodes used for the GW signal, inconsistent with an astrophysical signal;
- design of display of data quality monitoring output (Figures of Merit);
- prompt analysis of all output produced by DMT monitors, delivering information that can be used for improving the sensitivity of astrophysical searches;
- systematic (and robust) archiving and retrieval of figures of merit, including spectral snapshots;
- faster real-time graphical display of ordinary or generalized spectrograms;
- standardized and simple interfacing of the real-time data streams to commonly used external interactive graphical programs, most notably Matlab.

2.1.3 Data Quality Storage

The DetChar group compiles information from its investigations to create a repository of data quality (DQ) information for each engineering and science run. The information is used to form time intervals which are flagged with “data quality flags”. For the S5 run an SQL database was created to store the DQ information, which has proved advantageous for automated inquiries. A similar database structure is being worked out for S6, coordinated by the Software Working Group, as described in Section 7. There is infrastructure ready for S6 that creates many DQ flags online, in a way that is used by online searches, as well the automatic insertion of those online DQ flags into the database. The database also allows the insertion of DQ flags created after the data is taken, from investigations in the DetChar group or from scimons and operators.

The high priority outstanding needs for the database infrastructure are :

- reviewing the robustness and functionality of the new database infrastructure;
- adapting the database structure to allow flagging used for auxiliary (i.e., non gravitational wave) channels, especially environmental channels, indicating the status of the data quality in those channels (as sometimes these channels can be disconnected or malfunctioning)
- thorough tracking of the creation and modification of DQ flags (which will be especially important for multi year runs where the data quality needs may be evolving);
- automated graphical tools to track the DQ online production, and their effect on reducing non-gw transients.

2.1.4 Glitch Investigations

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

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

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

These goals are pursued both online and offline. During the S5 science run, the Glitch Working Group reported weekly on found anomalies and investigations of them[18]; this experience was successful and we expect it to be repeated in S6, which will be a very high priority for the dedication of human resources in the DetChar group.

This rapid-feedback analysis is based on transients found in the gravitational wave channel and in auxiliary channels (e.g. KleineWelle and BlockNormal triggers) and of the output of DMT monitors such as BurstMon. This was accomplished, during S5 via multi-day 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.

Currently, there is an online offline DMT glitch-finding program, generating KleineWelle[20] triggers, and the DetChar group will use triggers for single detectors (i.e., not necessarily coincident in two or more detectors) generated by some of the search algorithms used in the Burst and the Compact Binary Coalescence groups. The KleineWelle triggers are also generated for many auxiliary and environmental channels, to help with the diagnosing of the transients in the gravitational wave channel.

The list of high priority activities for the Glitch group are:

- organize (i.e., guarantee the people to staff) the shifts during the S6 data taking run, with prompt delivery of conclusions to Laboratory staff and to LSC scientists;
- automated production of graphical visualization of the data products needed for evaluating data quality and identifying transients;
- automated *and* off-line identification of time intervals to be flagged as having uncertain or poor data quality;
- comparison between data quality flags and event candidates for burst and inspiral searches, identifying false alarms due to undiagnosed problems;
- determine when key auxiliary channels are disconnected or malfunctioning;
- high precision time-frequency characterization of glitches;
- grouping of glitches leading to their easier diagnosing;

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

2.1.5 Spectral features Group

Another working group of the LSC Detector Characterization group is charged with investigating spectral features of the gravitational wave spectral noise density, which is especially important for the searches of gravitational waves from rotating stars and stochastic background.

Many of the spectral features are due to environmental disturbances, including seismicity, high wind, acoustic noise, and electromagnetic interference. Some sources are natural, but many are anthropogenic, including sources from observatory infrastructure, e.g., nearby motors and HVAC systems. A wide variety of environmental channels have been commissioned and are monitored, but unusual artifacts typically require detailed on-site investigations and eventually mitigation, work carried out by scientists from the Observatories and from several LSC institutions, as part of commissioning. Acoustic mitigation has played an especially critical role in lowering interferometer noise floors[23]. The retrofitting of LLO vacuum chambers with feed-forward, hydraulic pre-actuators led to dramatic improvement in L1 duty cycle, allowing the interferometer to ride out the passage of trains without lock loss. Nonetheless, significant increase in gravitational wave channel noise is seen during such a passage and in general during high seismic noise times, due to not very well understood upconversion of the noise into the gravitational wave band (40Hz-6kHz).

Environmental disturbances may also, of course, be manifested through linear couplings to the interferometer as direct glitches or lines, for sources with characteristic frequencies in the LIGO band of sensitivity. There have been extensive studies during S5 to understand better the sources of steady-state environmental couplings, particularly lines.

The list of high priority activities related to characterizing spectral features in S6 are:

- produce a noise budget for the noise, investigating the origin of discrepancies with the actual noise spectrum and variations in time of the noise sources;
- measure, document and maintain a list of known environmental couplings to the GW channel;
- produce the coherence between PEM channels and DARM_ERR, averaged on a daily, weekly and monthly basis;
- automate the monitoring of spectral lines in the gravitational wave band, tracking and documenting the origin of known sources;
- monitor the coherence between PEM channels and DARM_ERR at the frequencies used for searches by the CW group;
- automate the measurements of correlations between the GW channel and auxiliary channels, documenting the results;
- systematically review the output of all the noise line search methods in order to identify, track, report and document the discovered noise lines.

In the longer term, a plan should be devised on deriving the requirements and methods for the diagnosing of spectral features for Advanced LIGO.

2.1.6 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.

There is a significant travel burden associated with LSC scientists taking scimon shifts, but it is the judgement of the collaboration that ensuring close monitoring of data quality outweighs the cost. That said, there is a potential cost benefit in stationing more LSC graduate students and postdocs longterm at the observatories which can naturally ensure more seasoned expertise among the scimons. The LSC has recently approved a policy for a 7-day minimum stay for people working as "scientific monitors"; we expect will result in more efficient and expert data monitoring.

Another important aspect of data run support is injection of simulated astrophysical signals into the detector hardware [27], to validate data analysis pipelines with high confidence. LIGO Laboratory and LSC scientists have provided the manpower to set up the injection infrastructure and carry out the injections during data runs. In addition, environmental signal injections of a wide variety have been carried out by Lab and LSC scientists. The sophistication and automation of signal and environmental injections has increased with each data run, and that steady improvement based on experience is expected to continue. In particular, the software should be enhanced so that it is able to inject distinct signals at the the sites which are consistent with an astrophysical signal arriving from a specified direction with specified polarization components. The system should also be made more robust against a few failure modes encountered during the S5 run. At the moment, only a handful of LSC scientists are expert in signal or environmental injections; increasing those numbers would be helpful and prudent.

High priority activities for the monitoring of data quality in the S6 science run are:

- prepare list of responsibilities for scientific monitors for S6 that helps the readiness from prompt analysis and discovery;

- prepare or upgrade software and hardware needed to make hardware injections on the detectors, including coherent blind injections;
- automate tools for monitoring the result of hardware injections, using them to track the instruments' sensitivity

Longer term goals include the strategizing for the scientific monitoring needed in a continuous data-taking run expected for Advanced LIGO, where detections are expected to happen regularly.

2.1.7 Calibrations

Calibration of the LIGO interferometers is a task critical to the success of the data analysis algorithms, and the confidence associated with their results. As such, the LSC created in its bylaws a Calibration committee, separate from the DetChar group, although there are still many common members and activities.

Calibration of a detector is a complex task that involves instrumental hardware measurements, detector modeling, computer programs, and extensive validation. The Calibration Committee responsible for this essential work includes LIGO Laboratory and other LSC scientists. A dedicated Calibration Review Committee provides advice and vetting of this work. The Calibration Committee's results are posted and documented on a web page[8] available to the Collaboration, as well as recorded in the electronic logs, software repositories, and LIGO documents[9].

The calibration procedure has evolved in sophistication since the S1 run, most notably in automation, modeling, and redundant validation methods, with calibration provided both in the frequency domain (a frequency-indexed response function to be applied to the Fourier transform of the gravitational wave channel) and in the time-domain (a derived digital time series, "h(t)", representing strain as a function of time, which LIGO started to generate in S4)[10, 11]. From the S6 run, the time domain data is going to be the main calibration product that stands upon the conventional frequency domain calibration analyses. Since 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, now the Calibration Committee is co-chaired by the time-domain chair and the 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 have been developed and implemented in the S5 run. Due to the deeper understanding of the measurement process and physics behind the techniques involved, the various methods now agree to within 10%. In the S6 run, we will have routine calibrations by the coil calibration and the photon calibration, and hopefully some other methods based on different physical principles, with agreement at 10% or better level.

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. The Calibration Review Committee advised to perform a series of hardware measurements to characterize the timing aspects of the LIGO interferometers, which was done by the timing experts, students and site staffs (see 2.1.8). Due to this effort, we are now confident about the S5 LIGO timing at 20 us level. However, an entirely new Advanced LIGO style CDS system was introduced to accomodate an output mode cleaner in Enhanced LIGO. Since there has been no long term study of the timing of this system, the Calibration Committee intends to monitor the timing throughout the S6 run.

Myungkee Sung has worked on the generation and analysis of the time-dependent calibration coefficients for S5. The coefficients were produced for the last version (V4) of the S5 calibration and 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, and Myungkee Sung is developing a DMT monitor for calibration based on LineMonitor. By constantly monitoring the calibration lines to make the above mentioned calibration coefficients, the new DMT monitor not only brings the already-established calibration studies into the DMT, but also provides new functionality like the model-less monitoring of the open loop transfer function at the calibration line frequencies, diagnosis of the photon calibrator, and comparison of the photon and coil calibration.

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 some communication channel is established during the S6 run.

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

2.1.8 Timing

Traceable and closely monitored timing performance of the LIGO detectors is mission critical for reliable interferometer operation and astrophysical data analysis. For example, (a) timing jitter of digitization of the GW signal could directly contribute to the noise level, i.e., 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, (d) in case of a coincident detection of GWs and astrophysical events, such as GRBs or supernovae, traceable timing information is a requirement.

The Timing Stability Working Group (TSWG) includes scientists from both the LSC and the LIGO Laboratory. The group is responsible for the timing diagnostics of the detectors and it also reports on the achieved timing accuracy. The TSWG also collaborates with the Calibration group [15] to ensure the accurate phase and timing calibration of the gravitational wave data stream (h(t), also see 2.1.7).

Beyond documenting the timing performance of the current detector, long term tests of the advLIGO systems [16] are also needed during the Enhanced LIGO era.

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

Based on past experiences, short term high priority needs are:

- Organization (coordination, training, supervision, publication input, etc.)
- Timing verification studies and documentation of timing stability during S6;
- Support of potential discoveries in S6;
- Further develop and test injection techniques to determine accurate timing through direct test mass excitations;

- If neccessary address emergencies and/or expand the capabilities of presently running data monitoring tools related to timing; [
- Support S6 data analysis publications relying on timing performance.

2.2 GEO Detector Characterization

2.2.1 Introduction

The GEO 600 detector is a significantly different design to the 3 LIGO detectors. As such, it has separate commissioning and characterisation teams. The vast majority of the characterisation work is carried out by the commissioning team with additional help by other scientists within the GEO collaboration. The characterisation work at GEO 600 can be considered as two threads of work with some overlap: characterisation of the detector for the purposes of guiding and informing the commissioning plan; and characterisation of the resulting data in order to prepare it for input into the various LSC search groups.

The characterisation of the detector focusses mainly on the areas of calibration, tracking of the detector sensitivity, studying the long-term stability of the detector, identification of limiting noise sources, and on studying the couplings of the many sub-systems that make up the entire instrument.

As GEO 600 approaches its predicted design sensitivity, the commissioning work depends more and more on an accurate understanding of the noise present in the main output channels. In the case of GEO 600, this requires two things: a reliable on-line calibrated strain sensitivity, and a detailed understanding of the limiting noise sources. With these two requirements satisfied, it is possible to track small changes in the sensitivity of the instrument which arise from any particular experiment or hardware change that takes place. In this way, the commissioning team can quickly and accurately evaluate the impact of their work.

Part of this sensitivity-tracking is achieved through the use of Summary Reports. For every 8 hours of data collected at GEO 600, every second of the calibrated output data stream is analysed in various ways, together with data from many other channels. The results from these various analyses are presented as web pages in the form of summary reports [28]. In addition, many on-line monitors are run to present the commissioners with an up-to-date view of the current detector sensitivity. These on-line monitors typically include:

- Continuously updating, low-latency spectrum of the calibrated $h(t)$ signal, with a background reference trace.
- Continuously updating, low-latency spectra of many interferometer and enviromental monitor signals, with background reference traces.
- Time-frequency map of detected glitches for the last N hours (typically 24 hours). The glitch detection algorithm used at GEO 600 is a modified version of the HACR algorithm [29, 30] and the monitor is called `hacrMon`.
- Range to 3 different optimally oriented binary systems which would produce an snr of 8 in the current $h(t)$ data stream. This monitor is called `inspiralMon`.
- The `burstMon` monitor determines the h_{rss} needed to detect 6 different sine-gaussian signals ($Q = 9$, central frequencies between 200 Hz and 2 kHz) at an snr of 8.

A lot of the software that carries out the on-line analysis is dedicated C-code. Such codes include the calibration code (that produces the on-line $h(t)$ data-stream), and four on-line monitors which analyse a large fraction of the recorded data (`hacrMon`, `inspiralSensMon`, `burstMon`, `saturationMon`). The output of these monitors is captured in a database where it is subsequently mined and presented in the summary reports. In addition, the contents of the database are routinely mined by scientists carrying out characterisation investigations.

Most of the daily on-site data analysis carried out by the commissioning team is performed in MATLAB using a suite of dedicated tools (`geomatapps`) designed to easily interact with the GEO 600 data servers and to allow easy access to typical signal processing routines.

Starting around July 2009, GEO will leave 'astrowatch'-mode and begin a program of hardware upgrades (known as the GEO-HF program). During this time, various changes to the instrument will be made which have a significant impact on how the detector characterisation is performed. In particular, the move to a dc-readout scheme will involve modification of the calibration scheme; some of the new hardware subsystems, like optical-squeezing, will be controlled by new digital control systems based on the Advanced LIGO CDS design. Bringing these digital control systems in to GEO requires that they are first studied and characterised, and then tightly integrated with the existing control and data acquisitions systems.

It is the intention of the GEO 600team to take science data with GEO 600in S6, whenever it seems appropriate. In particular that means taking data over nights and weekends when actual commissioning allows. Due to the fact that the intended upgrades are modular and mostly independent of each other (to some extend), we are planning to bring GEO 600back to some science data taking after each step of upgrades. This is essential for commissioning anyway, but may yield data to be used for follow-up analysis of one or the other LVC-S6/VSR2 candidate events.

2.2.2 Data Acquisition and timing

The data acquisition system used at GEO 600 uses the same analog-to-digital converter (ADC) cards as those used at the LIGO detectors. However, the infrastructure in which these cards are placed is somewhat different. There are three data-collecting units (DCUs) at GEO 600: one in each building. The main DCU, located in the central-building, contains one 32-channel ADC card sampling at 16384 Hz, and one 64-channel ADC card sampling at 512 Hz. The other two DCUS contain only a single 16-channel ADC board sampling at 16384 Hz. All cards are clocked by timing signals locked to a GPS reference clock. In addition, to ensure good time-stamping accuracy, the clock signals sent to the ADC cards are read-back by a dedicated circuit which counts the number of pulses in a defined period of time to ensure that the sample rate remains constant at 16384 Hz. More details of the GEO 600 data acquisition system can be found in [31] and [32].

To determine that the GPS references used to clock the ADC are performing properly, a second Rubidium atomic clock (locked to a GPS receiver for long term stability), is used to generate a periodic ramp signal. This ramp signal is recorded in the DAQ system and used to check that the ramp originates close to the beginning of each second of data. The measured offset of the ramp from the second boundary of the data-stream is monitored and an alarm is raised should this value exceed a certain threshold. Typically, this offset is measured to be around 6 μs with a standard deviation around 500 ns. This measured timing offset is also the basis for one the data quality flags generated at GEO 600 (see Section 2.2.5 for more details).

The measurement of the relative timing accuracy between GEO 600 and other LSC detectors is an ongoing area of research and discussion within the LSC.

With the impending integration of digital control systems based on the Advanced LIGO design, the current data acquisition system must be merged with these new systems to provide a single coherent set of data. This will involve changes to the data handling software and to the flow and organisation of the data products.

2.2.3 Strain calibration

GEO 600 is the first long-baseline GW detector to use Signal-Recycling. This, together with the heterodyne readout employed, results in GW signal being spread between the two demodulated quadratures of the main output photodiode signal. In other words, there is no demodulation phase that can be chosen that puts all the GW signal in one quadrature for all frequencies. These two output signals are typically referred to as P and Q, and both must be properly calibrated and then combined to recover the best estimate of the detected strain with the optimal snr. Because of this, estimating the sensitivity of the detector from the uncalibrated output signals is difficult, and as such an on-line time-domain calibration scheme was developed at GEO in order to give the commissioning team a single data-stream that properly reflects the strain sensitivity of the detector. This on-line calibration scheme is detailed in [33, 34, 35].

The calibration process is a complex one, relying on very detailed knowledge of the detector and supporting sub-systems. The accuracy of the calibration is paramount to the success of much of the data analysis that is performed within the LSC, particularly when more than one detector is used in a network.

The calibration accuracy can be thought of as two measures: the absolute calibration accuracy, how the voltage recorded at the two main output signals is scaled to be in units of GW strain; and the relative calibration accuracy, how the calibration accuracy varies as a function of frequency relative to some idealised, perfect calibration system. The relative calibration accuracy is reasonably easy to determine in GEO. The main actuators used to hold the Michelson interferometer at its dark-fringe condition follow a simple pendulum response: $1/f^2$ for $f \gg 0.6\,\mathrm{Hz}$ (main longitudinal pendulum resonance) and $f \ll 10\,\mathrm{kHz}$ (the first internal mode of the main test-masses). As such, we can induce known differential length changes and measure the predicted response to that given by the calibrated $h(t)$ signal. This typically gives results within 10% and 10 degrees across the detection band (50 Hz to 2 kHz).

The absolute calibration of GEO 600 is based ultimately on the length of the first Mode-cleaner (which is known to $<1\%$). By driving only one of the end-mirrors of the Michelson interferometer we induce some differential arm-length change, but also some (common-mode) length change of the Power-Recycling (PR) cavity. Since the laser frequency in GEO is stabilised to the length of the PR cavity, this length change ultimately results in a feedback signal which is applied to the control piezo of the master laser. The calibration of the master laser control piezo is then done by sweeping the frequency of the master laser across Free-Spectral-Ranges of the first mode-cleaner. We then relate the FSR of the mode-cleaner to its length. This measurement process is a difficult one, requiring many (careful) steps. Overall, the accuracy of this method is believed to be of the order 5%.

GEO 600 has been experimenting with a photon-pressure actuator for some time. The first aim was to confirm the absolute calibration achieved using the method described above. Until recently, the model used to convert photon-pressure (through momentum transfer) into displacement of the test-mass was thought to be simple, yielding a reliable alternative calibration method. The results of the experiments carried out at GEO and LIGO have shown that this is not the case. For example, it has become clear that the test-masses can not be considered as rigid bodies when applying localised forces. The effect of this on the simple model is particularly strong at high ($> 2\,\mathrm{kHz}$) frequencies, where it begins to dominate the simple pendulum response. Nevertheless, the initial calibrations derived using the photon-pressure actuator were within 30% of the calibration achieved using the routine method described above. The use of photon-pressure actuators is an active area of research within the LSC and VIRGO, and comparison of results between the different calibration teams is an essential step in properly characterising this method.

Clearly it is highly desirable to have multiple methods for determining the absolute calibration accuracy. Within the LSC, various methods are employed at different detectors, and very important and fruitful exchanges of information and methods are carried out by members of the LSC Calibration Committee and the VIRGO calibration experts. It is essential that this collaboration continues in order to maximise the chances of finding more and better ways of calibrating these highly complex instruments.

Recently, a new absolute calibration method has been under investigation in GEO. This new method involved using a free-swinging Michelson configuration to determine the relationship between driving the electro-static actuators and the resulting motion in terms of full fringe deviations of the Michelson. Since the fringe amplitude is related to the (well-known) wavelength of the laser light, and since the number of fringes passed through in each swing of the pendulum relates to the amplitude of the drive signal fed to the electro-static actuators, we can obtain an absolute calibration of the actuators. LIGO have also used such a scheme in the past, and further work and development at GEO will be done to understand this method and compare it to the existing methods.

GEO will change to a dc-readout scheme in the near future, requiring a modification of the existing calibration scheme and software. In particular, new models of the instrument response will need to be developed and validated and then integrated in to the calibration software. Moving to a new optical readout scheme and calibration model will require a significant amount of characterisation and modelling to achieve a well calibrated output data stream.

2.2.4 Glitch studies

Due to the difference in sensitivity of the GEO detector compared to the LIGO detectors at low-frequency, data from GEO can have maximum scientific impact in the search for transient (burst) GW signals, particularly above 500 Hz where the sensitivity of GEO comes close to that of the LIGO detectors. As such, a significant amount of the detector and data characterisation at GEO 600 centers around the measurement, removal and characterisation of transient signals. As stated above, the HACR glitch detection algorithm is run continuously on many signals recorded at GEO. HACR characterises detected glitches (events) by a few parameters. For example, central frequency, central time, bandwidth, and duration.

Coincidence analysis is routinely performed between many instrumental/environmental signals and the $h(t)$ signal. Any signal that shows a significant population of coincident glitches is either targeted for repair, or (if repair is not possible), is studied further in the hope of creating a robust veto signal that can be used to reduce the final set of $h(t)$ events that will be considered as candidate GW bursts in subsequent (multi-detector) analysis.

One aspect of the new GEO-HF program is to increase the circulating power in the interferometer by roughly one order of magnitude. We have seen in the past that increased optical power can cause various non-stationary and transient effects. Glitch studies will likely play a significant role in the characterisation of the detector when it runs at these higher powers.

2.2.5 Data quality

A number of data quality flags are used at GEO in order to decide whether the produced $h(t)$ signal is considered to be science quality data. These data quality flags are automatically produced and checked during the calibration process so that a final Data Quality channel is produced which, for any value other than zero in the first 6 bits, indicates non-science data. For S5, this data quality channel is encoded in the bits of a 16-bit integer as described in Table 1. There are additional data quality flags encoded in bits 6 though 11. These represent a measure of the calibration quality. For S5, we build an additional data quality flag if bit 6 is 1 for more than 10 consecutive seconds.

It is highly likely that new data quality flags will be introduced that indicate, for example, the state of the optical squeezing and the power level in the interferometer. Understanding and implementing these data quality flags will be a necessary area of research as the new hardware systems are put in place.

Bit	Description	Science condition
0	Detector lock status.	0
1	Hardware maintenance on.	0
2	Software maintenance on.	0
3	Calibration lines missing.	0
4	Configuration re-read.	0
5	DAQ time-stamping not valid.	0
6	Calibration quality, $\chi^2 > 5$.	-
7	Calibration quality, $\chi^2 > 6$.	-
8	Calibration quality, $\chi^2 > 7$.	-
9	Calibration quality, $\chi^2 > 8$.	-
10	Calibration quality, $\chi^2 > 9$.	-
11	Calibration quality, $\chi^2 > 10$.	-
12-15	Not used.	-

Table 1: Data quality flags encoded in calibration data quality channel.

2.3 Virgo Detector Characterization

2.4 Introduction

The search for gravitational signals requires a careful characterization of the detector and of the noise sources. The response function have to be well estimated trough a suitable calibration, timing must be accurate and all noise sources, both internal ones and due to external environment, understood and reduced or controlled. In pratics the characterization of the interferometer is carried out at various levels. A first analysis is performed within the commissioning activities, and it is more devoted to the understanding of noises and to the interferometer calibration. Noise analysis is performed both to evaluate noise budget and reduce eventual noise sources. There studies are performed both on Science Mode Data, corresponding to best interferometer operation, and by exciting the interferometer with suitable noise sources to evaluate, typically, transfer function from environmental noise sources to dark fringe data. During the last year, of particular importance have been the studies of seismic and acoustic noises, that can generate spurious signals mostly by modulating the position of scattering light elements and of electromagnetic noise. Commissioning shifts are also dedicated to the careful calibration of the instruments, performed using the same actuator used to control the interferometer. The possibility of using an external actuator to produce the excitation signal, the photon calibrator, is under evaluation.

In a second level, which mainly concerns on line Data Monitoring , vetoes and on line-analysis, a community broader with respect to commissioning is involved. In particular data analysis is devoted to identify on-line possible interesting events and to communicate them immediately to people attending to the interferometer, and to provide eventual vetoes on the signals by analyzing the transients of suitable channels (like seismic, acoustic, electromagnetic...). Presently the work of DetChar community is devoted both to analise the VSR1 data and to create the appropriate up-grades of the tools in view of VSR2. In the following subsections the main activities within the DetChar will be summarized, with reference to papers or internal notes for the details.

2.5 Calibration

The calibration of the Virgo interferometer is necessary in order to perform precise data analysis. The 'standard' calibration has been automated and extended to have some redundant measurements during VSR1. It includes measurement of the absolute time of the Virgo data, measurement of the transfer function of the dark fringe photodiode readout electronics, measurement of the mirror and marionette actuation transfer functions and monitoring of the finesse of the arm cavities. The calibration output are then used in the frequency-domain calibration, resulting in the Virgo sensitivity curve, and in the time-domain calibration, resulting in the h(t) strain digital time series. The goal of the VSR2 calibration is to have everything measured before the start of the run to produce reliable online h(t) time series. The calibration parameters will then be monitored during the run.

An independent method, developed to calibrate the mirror actuation gain, using the radiation pressure of a secondary laser to push on the input mirrors, has been used (photon calibration) for VSR1. It confirmed the mirror actuation modulus measurements below 100 Hz, within the systematic errors of the order of 15/100 of this method. It was also used to set the sign of the h(t) strain in relation with the convention agreed with LIGO. This method is under improvement in order to reduce the systematic errors at low frequency below 10/100 before VSR2. As a second step, a different setup will be designed and tested in order to increase the measurement range up to a few kHz. Fruitful exchanges of ideas with the scientists performing the photon calibration for the LIGO detectors are under way.

The calibration output were used in the new version of the h(t) reconstruction for VSR1. Improvements of the reconstruction procedure were performed. They include more computation in the frequency domain, better high-pass filter not to distort the phase of the h(t) above 10 Hz, handling of the NE and WE marionette actuations to improve the h(t) accuracy below 50 Hz, corrections of the optical response variations (arm cavity finesse variations), use of calibrated time delays. The final reconstructed h(t) for VSR1 is valid from 10 Hz up to the Nyquist frequency of the channel used, i.e. up to 2048, 8192 or 10000Hz. In this validity range, the systematic error on the h(t) amplitude is 6/100. The systematic error on the h(t) absolute timing is 70 mrad below 1.9kHz and 6?s above.

The calibration data taking had been automated to monitor the parameters during VSR1. The procedures have to be updated but will mainly remain similar for VSR2. A full set of measurements is aimed to be done before the start of VSR2 in order to give precise inputs to the online h(t) reconstruction. This requires that the configurations of the Virgo+ mirror and marionette actuation are stable a few weeks before the start of the run. The new code of the h(t) reconstruction developed for VSR1 reprocessing will be used to produce the VSR2 h(t) online.

More information and the calibration and reconstruction notes can be found at: https://workarea.ego-gw.it/ego2/virgo/data-analysis/calibration-reconstruction

2.6 Environmental noise

The Environmental Noise group[42] within VIRGO is in charge of studying environmentally induced noise (of acoustic, seismic or electro-magnetic type) which affects the interferometer (ITF) sensitivity and operation. Noise sources include natural phenomena (earthquakes, sea microseism, wind, thunders and lightnings, ...) but also, of particular interest (as viable of mitigation) are the several electro-mechanical devices which serve the ITF operation: air conditioning sys, illumination sys, vacuum and water pumps, racks, electronic boards, electro-mechanical devices, power supplies...

A network of environmental probes (microphones, seismometers, magnetometers, power line sniffers, thermometers, pressure and humidity meters, weather station) is used to monitor ITF critical locations. These probes are used to asses the data quality by tagging potentially "bad" data periods associated to the presence of environmental disturbances, and by flagging transient events as possible veto candidates. In addition portable probes (a recent addition is one RF antenna) and tools (shakers, coils, loudspeakers, headphones) are used to characterize noisy devices, understand noise paths, and measure noise coupling by means of controlled noise injections.

The group task can be described as two-folds: (1) during the commissioning phase our main foucus is on noise couplings measurement and noise reduction, aiming to optimize sensitivity and ITF operation; 2) during data-taking (science mode run) we have to assure the good fuctioning of environmental sensors, if proves necessary add some sensors to monitor critical noise coupling locations. During data taking, data-analysis pipelines and on-line monitoring tools (coherence,lines, non-stationarity, glitches and transients monitors) provide fundamental information to the investigation of environmentally induced disturbances. If needed, noise mitigations can be carried out during maintenance or short commissioning breaks.

During the first phase of Virgo+ commissioning considerable effort has been dedicated to the study (measuring and modelling) and mitigation of upconverted noise due to light back-scattered from external optical benches, and in-vacuum absorbing glass baffles[43]. These were limiting the 20-200Hz sensitivity. Mitigations suceeded in reducing the amplitude of benches and baffles resonant motion, sufficiently to set the projected noise below the Virgo design curve. Other noise mitigation actions concerned: vacuum and water pumps, air conditioning and illumination systems. Other environmental related mitigation is the optimization of covers on optical benches and critical beam paths in order to reduce the air flows which induce noise in alignment sensors, and possibly in the GW channel. Bench cleaness is also a relevant issue, to be better addressed in next upgrades. Glitches associated to dust particles crossing output beams on detector bench have been studied and are being monitored. Magnetic noise coupling reduced after the replacement of mirror magnets with smaller ones and removal of unused lateral magnets. Residual noise stands at the level of present design sensitivity[43].

Residual environmental noises which are known or suspected to affect the initial VSR2 sensitivity are mainly two[43]: (1) upconversion of micro-seismic noise associated to diffused light at the WE, which in case of severe sea activity can deteriorate sensitivity below 50Hz; (2) beam jitter noise coupling to external injection bench noise which is affecting the 40-50Hz region.

In collaboration with the Noise and Data-Quality groups several monitors have been developed, although some need additional work and tuning. These are: monitors of environmental coherences, band limited RMS informing of sea activity, earthquakes, and anomalously high seismic and acoustic activity, transient detectors of airplanes and lightnings. The on-line detector monitor informs the shift crew if environmental data (as well as other fundamental ITF channels) are not acquired correctly (because of data acquisition servers problems). It is being added the capability to detect also sensors malfunctionings (noisy condition, saturation, faulty conditioning electronics). An activity which has just started, and will continue during VSR2 data-taking is the characterization of the laboratory site noise environment, like: (i) noise from road trafic (inside and outside lab.); (ii) the recently installed wind farm (6km N-E from the closest NE building); (iii) other electro-mechanical which do not have evidence to affect dark fringe, but produce some detectable environmental noise close to the ITF.

Forthcoming VIRGO upgrades (V+MS and AdV) necessitate improvements on the environmental noise side, which we are studying. At present, works of major relevance seems to be: (i) isolation of the external injection bench; (ii) improve environmental sensors network by adding probes, for example magnetic probes seem of help to monitor local disturbances in the vicinity of electronic components; (iii) implement a monitoring of the on/off status of several moving or electronic devices (such as drain pumps, air compressors, heaters and chillers, ...) which are potential sources of environmental noise.

2.7 Virgo Data Quality and vetoes

A working group in Virgo (VDQ) [41] is in charge of providing to the search groups the Data Quality (DQ) segment vetoes and event-by-event vetoes for each run (Commissioning and Science runs). The DQ segments are generated from the data quality information generated from different sources: some data quality information are produced online by the h(t) reconstruction algorithm, and by the online data monitoring linked to the online data acquisition. For VSR2, all DQ flag segments are generated online and stored in the Virgo DataBase (VDB) in order to provide fast veto feedback to real time analyses. 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 noise that can generate an increase of the glitch rate during few seconds or longer. Among the possible sources, there are the seismic activity (human activity, sea, wind, earthquake), interferometer misalignments, electronic problems, etc ... More precisely, the VDQ group is in charge of the following tasks:

- Identification of the glitches found in the GW channel, using different algorithms searching for transient or online monitoring tools: e.g. Q-online, OutlierMoni. Q-scan is also heavily used to find hints.
- Investigation of the origin of the glitches and the noise coupling. This activity is done in collaboration with the Noise group (see Section on Noise Monitors and Section on Environmental Noise). The group especially investigates the possible correlation between the GW channel and all the auxiliary channels. For that, we use the Kleine Welle (KW) algorithm to generate triggers for all the auxiliary channels recorded by Virgo ($\sim$ 200 channels).
- Generation of the DQ segments. It includes all cross-checks that are necessary to validate that the DQ segment lists are ready to be applied by the search groups.
- Study of the effect of the DQ vetoes on the burst and inspiral analyzes, using triggers that are generated online. Statistical information such that efficiency, use percentage and dead time of each DQ list is computed.
- The VDQ group, in collaboration with the search groups (search group members are involved in the VDQ group), assign category to each DQ that tells how the DQ vetoes should be applied in the analysis pipeline.
- Generation of the KW triggers for all auxiliary channels. Those triggers are then used by specific veto studies.
- Generation of the event by event vetoes based on the auxiliary channels KW triggers. That allows to get rid of the glitches that have not been suppressed by a DQ veto. Different algorithms are used to identify which auxiliary channels are interesting. Studies to understand the noise coupling help to assign a category to each event by event vetoes.
- Validate, on a monthly basis, all the DQ flags generated online. DQ segment lists can then be used by offline analyses.
- Validate the safety property of the DQ and event by event vetoes using the hardware injections

The VDQ group is making use of the Virgo logbook reporting incident or any noise disturbance during a run. That allows to define, offline, some DQ segment that flags bad quality data periods and/or to understand some mysterious noise increase.

All DQ segments are stored in Virgo Data Base (see Section on Virgo Data Base). This tools 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 allows provide useful facility to combine several lists. All DQ segments can be accessed from VDB by anyone. In addition, the VDQ group posts on a web page the combined DQ segment lists that should be used by the burst and inspiral offline analyzes.

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 deadtime, etc ...) will have to be done and published within a few hours after the data taking.

2.8 Noise Monitoring Tools: Data quality and 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. This information is also provided [qcmoni]. 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 [fdio] and the same configuration syntax [moni]. Moreover, the quality tests done within each detector monitoring algorithm can depend from the interferometer's state provided by the automation system. Among others are monitored the timing system, the detection photodiodes and detection system alignment, the seismic isolation positions and orientations, the environment status, the beams alignment in the north and west arms, the input bench alignment, the laser injection system and the vacuum system. 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 [qcmoni]. 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 behaviour 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 [dy] to provide an easy way to display data online or offline, allowing to follow in control room the changes made on the interferometer or to do some minimal data analysis. This software allows to combine signals and to do several types of plots: time plots, spectra, coherences, transfer functions, 1D or 2D distributions, time-frequency plots, etc. It runs under Unix-like operating systems and is widely used within the collaboration.

Web Monitoring Tools based on web display, VEGA scripts [vega] and matplotlib (python) scripts have been set up to provide periodically updated plots providing a monitoring, through the web, for the main channels of each detector's subsystems, for the data acquisition or for online processing and data analysis: http://wwwcascina.virgo.infn.it/MonitoringWeb and http://wwwcascina.virgo.infn.it/MonitoringWeb/Noise A specific monitoring process called SpectroMoni [spectro] provides data to create a set of spectrograms within those web pages.

Noise monitors The goal is to have an easy to use instrument to collect information for the scientist in control room and save the relevant event during the run. These set of monitors could be used also by commissioners during daily activities to check the status of data. all the monitors have been collected at the main web page http://wwwcascina.virgo.infn.it/MonitoringWeb/Noise

We produce web pages which are automatically update with the results of running algorithms. Some of these monitors are linked to on line chain and produce results in real time. Some other tools, which require to accumulate data to produce results, run on the last data written on disk and so produce results in in-time.

Each of this algorithms produce as outputs plots and web pages, which are sent directly to web server in such a way to have in real-time a snapshot of the noise even from remote labs. The algorithms produce also ascii files which will be archived. Some monitors produce outputs in mysql format in such a way to feed a Noise Data Base where we want to collect the relevant informations. The goal is also to produce a web or GUI interface to Noise Data Base to let easy querying to the database itself. The noise group will provide for each of this monitor the useful informations for scientist who will be in shift during VSR data taking in such a way to be aware of checks which must be done and reports to be written.

- Transient event
  - **VirgoHACR:** Based on the HACR algorithm developed in GEO, it computes time-frequency maps with short FFTs of the dark fringe signal and compares each time-frequency bin to a moving average and rms of the spectrum. Mainly, it does in time-frequency domain, what glitchMoni does only in time domain. Its results are not in the trend data. They are sent online to a php database. This database is inquired to produce a set of web pages showing the events versus time or versus frequency and various statistics about the events:
https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=11
  - **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 criterium 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
  - **outlierMoni:** It computes a moving averaged spectrum of the dark fringe and use it to whiten the data. Then, it search 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:
http://wwwcascina.virgo.infn.it/MonitoringWeb/Noise/outlierMoni
  - **glitchMoni:** It computes a moving average and a moving rms of the dark fringe signal over a given time length. It searches for samples more than 6 rms above the mean. This algorithm is dedicated to short glitches detection. Its results are in the trend data and a web page was developed for it but needs to be updated and improved:
http://wwwcascina.virgo.infn.it/MonitoringWeb/Noise/glitchMoni
  - **Omega-online:** The previous list of online noise monitoring algorithms does not include Omega because it is a burst search algorithm. It is nevertheless very useful for noise monitoring illustrating the coupling between monitoring tools and online analysis tools. Omega is running on one machine at Cascina, reading data from disk with a latency of several minutes. A web interface provides figure of merit based on the omega triggers. Another olnode machine is used to run omegascan over several auxiliary channels for the time of the loudest dark fringe trigger in the last 1mn, last hour and last day. Those results and time-frequency plots are available from: http://wwwcascina.virgo.infn.it/MonitoringWeb/Bursts. Omegascan results of all loud triggers are available in a centralized web page for further human investigation.

  - **Glitchness:** 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=40

- 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 bands RMS:
    http://wwwcascina.virgo.infn.it/MonitoringWeb/Noise/NonStatMoni
    and
    http://wwwcascina.virgo.infn.it/MonitoringWeb/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 seismicity on the site:
    http://wwwcascina.virgo.infn.it/MonitoringWeb/Environment

- Non linearity
  - **Bicoherence:** It will run offline to give hints on noise up-conversions in some frequency regions. To be implemented. https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=20
  - **Skewness-Kurtosis monitoring:** Skewness is a measure of symmetry , or more precisely, the lack of symmetry for data distribution.
    Kurtosis is a measure of whether the data are peaked or flat relative to a normal distribution.
    to be implemented https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=21
    https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=22

- Lines behavior
  - LineMonitor It computes a background spectrum of the dark fringe signal and looks for spectral lines above a given threshold. The results are only available in a php database and are used to update periodically a list of lines available on the web:
    https://pub3.virgo.infn.it/itf/linemonitor

- High-snr events follow up
  - WDF-FollowUp The last hour 5 events with the highest SNR values, are automatically analyzed to check their features in time and time-frequency domains. This is done on Dark Fringe signal and a set of auxiliary data.
    https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=34
  - Omega-scan OmegaScan scan in Time Frequency domain the most energetic events cross-checking with a given set of auxiliary data.
    http://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).

    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. The background color of each entry is more red the largest the coherence is. Again the name of each channel is a link to the corresponding plot.

    The page is expected to be updated once every 3 hours, using only 5 minutes of data.
    https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=15
  - Multicoherence This tool estimates the coherence of the Dark Fringe signal and a set of auxiliary channel, trying to understand which are the channels which contribute more to noise in the Dark Fringe signal.
    to be implemented https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=16

- Whitening We produce a snapshot of the last minutes of whitened data Noise distribution, before and after whitening
  https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=23
- Noise Budget The Noise Budget is regularly updated on the web. This tool using model, or measured data for transfer function, explains the different noise contribution to sensitivity curve
  https://pub3.virgo.infn.it/MonitoringWeb/Noise/html/index.php?callContent=28

2.8.1 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 the analysis of the Virgo data:

- 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 the official location of Virgo data quality segments lists and it is the main tool used by the Virgo dataquality group (VDQ) . The peculiarity if the VDB system is the implementation of a two level versioning, using the concept of stable and unstable list. This schema simplifies the access to the VDB resources by users, avoiding pollution and inconsistency on official data publication. In fact, all authorized people of scientific groups can upload DQ lists into VDB. These lists are tagged as unstable. At this point the VDQ community has time to check if these is concern and interest about such lists, taking then the decision about the ”stabilization” of the lists. After that the lists are official. The VDB system provides two main 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 perform several activities; such as show each single DQ list with its properties or combine together several segments lists using user defined logical expression or also access the online documentation about each list.

- The VDBtk_segments.exe - That is the command line tools for the DQ segments lists data. This extends the functionalities available in the Web UI, in particular it is used to upload DQ segments list into the VDB.

At the moment the VDB hosts more than 300 DQ segments lists. 42 stable Virgo DQ lists and the others are L1, H1, H2 DQ lists. LSC data are daily synchronized using Pyton scripts. In preparation to the VSR2, we are improving the VDB system in order to accept online DQ segments lists and share data directly with the LSC Databases.

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 [44]. 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 [44].

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 [45, 46, 47] as a function of system parameters (masses, etc.).
- Associate gravitational waves from binary coalescence with coincident observations using other astronomical detectors, including radio, optical, x-ray and gamma ray telescopes, and low-energy and high energy neutrino detectors. Implement prompt methods for the flow of external triggers both to and from gravitational wave observatories to extract the most science out of these associations. In the absence of a gravitational waves signal at the time of an electromagnetic or neutrino trigger, use this information to exclude certain progenitor models, as was done for GRB 070201 [48].
- 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.
- Establish the relationship between gamma ray bursts and compact binary coalescences involving neutron stars [48].
- 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 [49] and alternative theories of gravity, such as scalar-tensor theories which can result in modified phasing of the gravitational waves from binary inspiral [50].
- Bound the mass of the graviton by direct observation of the gravitational waves from binary inspiral [50].
- 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 [51, 52].

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 ideal candidates for gravitational-wave astronomy. These systems include binaries with masses as low as $1M_\odot$ up to several thousands of solar masses, composed of black holes, neutron stars, and perhaps other exotic compact objects with densities similar to or greater than neutron stars. The beauty of these systems for data analysis is that the waveforms can be computed within the context of General Relativity and/or alternative theories of gravity (e.g., [53, 54]). 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.

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 [53] 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 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, it is clear that post-Newtonian waveforms should be adequate to detect inspirals of systems with total mass as high as $20 - 30M_\odot$.

Waveform models that are accurate to a fraction of a cycle over many cycles — more than 1600 cycles over 25 seconds for neutron star binaries between 40 and 1500 Hz — allow optimal integration of the signal and thus 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, [55]). 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 [56]. 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 [55]. 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 [57, 58, 59, 60]. Analytical models, including the effective one body (EOB) [61, 62, 144, 64] framework, and hybrid or phenomenological approaches [65], 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. 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 [66, 67, 68]. These results will also allow for improved sensitivity of our searches.

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, and the 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 patern. 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 computaionally feasible to cover the full parameter space with a template bank [69, 70, 71]. This is an area of ongoing development and research [72, 73].

Eccentric orbits will also modify the gravitational waveform [74, 75]. 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 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 [313] 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 [77, 47], was developed as an batch mode analysis and is run offline to analyze blocks of data (currently of a week's length), and perform a detailed estimate of the noise background, the astrophysical sensitivity and significance of any event candidates. In the remainder of this section, we describe the components of the analysis procedure in greater detail.

3.3.1 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 [78]. 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$\%$ [79]. The template banks for the mass-space are well in hand [80, 81] for searches using Post-Newtonian approximations; work is ongoing to include spins in the most efficient manner [82].

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 [83]. 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 [84]. Other signal dependent discriminators are being used in current analyses [85, 86], and many others are being explored. General methods are under development to make more efficient use of information information about each event to discriminate signal and background.

3.3.2 Low latency pipeline

We will also be using *low-latency* pipelines to generate triggers for detector characterization, including near-realtime feedback to the control room via the InspiralMon figure of merit. Low-latency coincident triggers will be used, in conjunction with fast source localization algorithms, to generate external triggers for pointing EM telescopes. The low latency search will focus on the low mass range: total mass between 2 $M_\odot$ and 35 $M_\odot$ and component masses greater than 1 $M_\odot$. The baseline low-latency pipeline is based on the MBTA package[313] developed by the Virgo Collaboration.

MBTA splits the matched filtering of the data into two frequency bands for efficieny 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 will 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.

3.3.3 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 [77, 47]. The current pipeline 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. All data satisfying this minimal data quality cut is 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 [84], 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 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. 3.3.4) and for measuring the pipeline efficiency (Sec. 3.3.6).

This pipeline was developed and refined over the S3 and S4 analyses [47] has been used in the S5 analyses [87, 88]. 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 [56] (in S4 data) has been submitted for publication, and searches in S5 data for the combined inspiral-merger phase and for the ringdown phase using the full pipeline are currently in progress.

3.3.4 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 templae 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 order, by amounts that are long in comparison with typical signal durations, 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 allow 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 search for reliable background estimators for these coincident triggers is ongoing. 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 events observed events, more time slides will need to be performed. For externally triggered searches (eg, for short-hard GRBs [48]) 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 is also used in other searches for transients.

3.3.5 Instrumental and environmental vetoes

The complicated nature of interferometric gravitational-wave detectors means that instrumental and environmental effects produce non-gravitational-wave signals 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 auxilliary 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 auxilliary data channels is the subject of continued investigation. In coordination with the DetChar and Burst groups, these veto identification procedures are being automated for S6/VSR2. A number of new ideas are being tested and our ability to use this information should continue to improve.

3.3.6 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. To enhance the detectability of these signals, techniques using either fully phase coherence of inspiral-merger-ringdown signals or burst style searches in combination with the inspiral searches are under development. As more information becomes available from numerical simulations, it will be incorporated into the construction of template banks. The current analysis pipelines can be easily adapted to do this.

3.3.7 Follow-up of candidate events

Because of the non-Gaussian, non-stationary nature of the noise exhibited by the interferometric gravitational-wave detectors, residual false alarms can be found at the end of the analysis pipeline. A critical aspect of the search is then to assess our confidence for gravitational waves and to distinguish them from those false alarms. To this purpose the CBC group uses a detection checklist that is applied to all statistically significant candidate-events (i.e. the coincident triggers which have a low false alarm probability as deduced from the background estimation). As a sanity check, the CBC groups 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. The goal is to check for possible non-stationarities in the detectors or unusual excesses of noise which would translate into a higher rate of noise triggers and thus reduce our confidence in gravitational-wave candidates.

- *Environmental or intrumental causes*: We analyze the auxiliary channels of the interferometers, such as the environmental sensors or the signals involved in the mirror control loops, to check for the presence of possible noise transients. 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 different tools including time-frequency spectrograms of the data, time-series of the output of the match-filtering algorithm, and time-series of the multi-detector coherent SNR, and “null-stream”. The interpretation of these tests is based on comparative studies for simulated hardware and software injections and time shifted coincidences.

There is an ongoing effort to implement new tests in the candidate validation procedure and to automate this procedure into a “follow-up” pipeline.

3.3.8 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.

3.3.9 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 [89], 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, 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.

3.3.10 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; only a small number of binary systems containing neutron stars are known in our galaxy, and there are virtually no direct astrophysical constraints on systems containing black holes or extragalactic binary systems. Our group closely follows the astrophysics literature on such systems, especially rate predictions from population synthesis models [90, 91, 92]. The group maintains and regularly updates a document that summarizes the range of predictions for CBC rates accessible by LIGO and Virgo [93].

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 [94]. Therefore, we have chosen to quote upper limits on rates in terms of events per year per L_{10}, where L_{10} corresponds to the blue-light luminosity of 10^{10} suns (the Milky Way is approximately $1.7\ L_{10}$). In the absence of a detection, we evaluate the *cumulative luminosity* to which the search was sensitive in units of L_{10}. This is done by convolving the pipeline efficiency as a function of distance (and other parameters) with a list of source galaxies. The latter is drawn from a database of nearby galaxies (including sky location, distance, and blue light luminosity) assembled by group members [95], to which are added "fake" galaxies at larger distances (mostly beyond about 30 Mpc) so that a uniform density of $0.02\ L_{10}/\mathrm{Mpc}^3$ is obtained.

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 [45, 46], 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 early star formation in elliptical galaxies is expected to make additional contributions, requiring other tracers of CBC sources (e.g., galaxy mass). Additionally, 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). 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. Again, 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. In this case, however, the results of searches can be used to say something about the association of the external trigger with a coalescing compact binary progenitor (as was done with GRB 070201 [48]). We are also developing methods to make statistical statements about the presence of detectable gravitational waves from the entire population of GRBs examined by the searches.

3.4 CBC Searches and Results

3.4.1 Searches of 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 [96], S2 [97, 98], S3 and S4 [47] data
- A search for primordial black holes using post-Newtonian templates in S2 [99], S3 and S4 [47] data
- A search for stellar mass black holes (with component mass greater than $3M_\odot$) performed using phenomenological templates [100] in S2 [101], S3 and S4 [47] data
- A search for spinning binary black hole systems in the S3 data [72]
- A search for the ringdown of black holes formed after coalescence of binary systems, in S4 data [56].

3.4.2 Searches of the LIGO S5 and Virgo VSR1 data

The LIGO S5 run lasted for two years, the last six months of which coincided with the first Virgo science run (VSR1). The search of this data for CBC signals, using the methods described in the previous sections, is nearing completion. The search parameter space has been broken up somewhat differently than in the previous four science runs. The "low mass" CBC search, performed with post-Newtonian templates, covers binaries with a total mass between $1 M_\odot$ and $\sim 35 M_\odot$ and component masses not less than $1 M_\odot$. 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 [102] frequency domain templates are sufficiently accurate. The "high mass" search covers binaries with total mass between 25 and 100 $M_\odot$ (overlapping with the low-mass search in the mass range between 25 and 35 $M_\odot$). At these higher masses, the merger and ringdown of the signal are in the sensitive band and contribute a significant fraction of the signal-to-noise ratio. Therefore, developments in understanding the full waveform have a more significant impact on these higher mass signals. At even higher masses, a search for black hole ringdowns is being performed. In addition, a search for gravitational waves associated to short GRBs which occurred during S5-VSR1 has been done.

For both historical and technical reasons, the low mass search has been split into three epochs with three associated papers: the S5 first calendar year data (months 0-11), S5 months 12-18 data, and S5/VSR1 data from both LIGO and Virgo (S5 months 19-23). The LIGO data is filtered through templates with masses that extend up to 35 $M_\odot$, while the Virgo data have been filtered through binary neutron star templates only (between 1 and 6 $M_\odot$). The first two analyses and corresponding papers are complete; the LIGO-Virgo low mass search has been completed, but the final interpretation of the results and internal review of the search are ongoing.

There have been significant advances in understanding the waveforms during the last stages of inspiral, merger and ringdown, especially in the non-spinning case. This has been possible by combining information from numerical relativity, from extensions of PN theory such as the effective-one-body approach, and by doing phenomenological fits to the numerical relativity results. The S5 high mass search is using time-domain waveform templates known as EOB-NR, which 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.

During the S5 run, a short GRB (070201) whose sky location overlapped the Andromeda galaxy was observed. The analysis of the data at the time of this GRB has been completed. Since no gravitational waves were observed, we infer that the progenitor was not a coalescing binary in the Andromeda galaxy. The search for gravitational waves at the time of all short GRBs during the S5/VSR1 run has been completed, both the review and the writing of the paper are nearing completion.

A search for ringdown signals in S5 data is also being pursued along the same lines as the S4 ringdown search. The analysis is underway.

The completed analyses on the S5/VSR1 data are:

- Gravitational waves associated to Gamma Ray Burst 070201 whose allowable sky location range ovelapped the Andromeda galaxy [48]
- binary systems with total mass up to 35 $M_\odot$ in the first calendar year of S5 [87]
- binary systems with total mass up to 35 $M_\odot$ in months 12-18 of S5 [88]. This was the first search employing the *ihope* infrastructure which was developed in order to enable a more rapid turnaround for future analyses.

3.4.3 Recently completed analysis developments

In addition, several key analysis tools and requirements, common to many or all S5 analyses, are now either complete or in final stages of development and review:

- The definition and implementation of data quality vetoes for all of S5, in support of all S5 analyses, is now complete.
- The detection candidate follow-up procedures (section 3.3.7) have been developed and automated to the point where they are now in use by many group members, in support of the outstanding S5 searches listed below. The review of these procedures, and their continued development for S6, is ongoing.
- Improved detection statistics have been implemented for the S5/VSR1 LIGO-Virgo search and the S5 GRB search. Further development of search-specific detection statistics, improved signal based vetoes with automated tuning, and improved background estimation, is ongoing for S6.
- Two different families of NR-inspired IMR waveforms that incorporate all phases of binary coalescence have been implemented into the CBC software base, and can be used for both template bank construction and simulated signal injection for efficiency evaluation, for high mass searches.
- The Physical Template Family of spinning binary waveforms [82] has been implemented into the CBC software base, and can be used for both template bank construction and simulated signal injection for efficiency evaluation. Evaluation of the efficacy of a PTF-based search is ongoing.

3.5 CBC Group Priorities and Goals

The *highest priorities* for the CBC group during FY2010 are:

- Completion and publication of the remaining S5 searches: the S5 GRB search, the S5/VSR1 LIGO-Virgo low mass CBC search, the S5/VSR1 LIGO-Virgo high mass CBC search, and the S5/VSR1 LIGO-Virgo ringdown search. These are discussed in more detail in 3.6.
- Prompt and comprehensive search of the S6/VSR2, implementing and employing robust low-latency detection pipelines to search for low mass CBCs, high mass CBCs, CBCs associated with GRBs and other external triggers, and black hole ringdowns. S6/VSR2 plans are described in section 3.7.
- Continued development of key analysis tools, including: parameter estimation; improved separation of signal from background; multi-detector coherent analysis; detection confidence procedures; spinning binary search methods; and improved computational efficiency and efficacy. These are discussed in section 3.8.

In the longer term, the CBC group is committed to enhancing the search methods pioneered during the initial detector era to produce robust, low latency analysis of the advanced LIGO and Virgo data. Details of these longer term developments are described in 3.9.

3.6 Completion of S5 Searches

Our top priorities are to complete, and submit for publication, the following S5 searches by the end of calendar year 2009 (status as of July 2009):

3.6.1 S5/VSR1 GRB Search - CBC-S5-1

The S5 GRB search, covering 22 short-hard GRBs observed during S5, is now complete. No detection candidates were found. The review process is nearing completion and a mature paper draft is available. The S5 GRB paper should be completed by the end of the summer.

The search used templates for neutron star and neutron star–black hole binaries as these are one of the favoured models for the short hard GRBs. By concentrating on short stretches of time around the GRB and restricting attention to the sky location of the GRB, we are able to reduce the false rate of the search and thereby increase the sensitivity (as described in 3.3.9). For most GRBs, data from the two most sensitive LIGO detectors operating at the time were searched. For GRB070923, which was in a favorable position for Virgo, we searched for signals in H1, L1 and the Virgo detector. We searched for coincident triggers within 6 seconds of the GRB time, using nearby coincident data stretches to estimate the background. A detection statistic based on a likelihood ratio for signal versus background was developed and employed, and upper limits on the distance to the source as a function of mass of the neutron star companion. A statistical statement on the absence of detectable gravitational waves from the entire population of GRBs was also derived.

3.6.2 S5/VSR1 Low Mass CBC Search - CBC-S5-2

The S5/VSR2 LIGO-Virgo low mass CBC search is nearing completion. The data has been analyzed and no detection candidates were found. The interpretation of the result, and upper limit statement are being derived. The review has commenced. A paper draft is in preparation.

This analysis uses the final six months of the S5 search, which were joint with the Virgo VSR1, using data from all four detectors. It is the first joint LIGO-Virgo CBC analysis and made use of the *ihope* pipeline. The Virgo VSR1 run achieved noise performance comparable with the LIGO detectors at frequencies above several hundred Hz, but with considerable excess noise at lower frequencies. Thus, the Virgo detector is most powerful in detecting binary neutron star inspiral signals (which extend up to 1500 Hz), and the search only filtered the Virgo data with templates extending up to 6 $M_{\odot}$. As for the other low mass CBC searches in S5, the search was performed with non-spinning templates. We perform simulations of spinning waveforms into the data and observe good detection sensitivity to systems with spin, across a broad range of spin parameters. There were a variety of new challenges and innovations, including: the addition of data from the Virgo detector to the existing LSC CBC analysis; development of a new detection statistic in order to identify the most useful combinations of detectors and triggers; the development of tools to identify and implement data quality vetoes for the Virgo data; and the development of the coherent analysis code as part of the follow-up of detection candidates.

3.6.3 S5 high mass CBC search - CBC-S5-3

The S5 high mass CBC search, using full inspiral, merger and ringdown (IMR) waveform templates, is in advanced stages. The review of the analysis has begun and we aim for results and mature paper draft by the fall of 2009.

This search focuses on systems with total mass between 25 and 100 $M_{\odot}$, where the merger and ringdown of the signal are in the sensitive band and contribute a significant fraction of the signal-to-noise ratio, so that standard post-Newtonian inspiral waveforms are not very sensitive to the signals. The S5 high mass search is using time-domain waveform templates known as EOB-NR, which incorporate information from numerical relativity simulations in order to reliably model the inspiral, merger and ringdown phases (so far, without spinning components) of the binary coalescence. The high mass search is being performed using the same *ihope* pipeline as the low mass search. Because of the emphasis on low frequencies, the search is LIGO-only (ie, does not include Virgo data), covering all 24 months of S5 (including the overlap with VSR1). Much effort has gone into handling disconinuities in the templates and injected waveforms, re-tuning the signal-based vetoes, understanding the parameter estimation and coincidence windows, and developing appropriate data quality vetoes. Evaluation of the detection sensitivity for spinning waveforms is an ongoing challenge.

3.6.4 The S5 Ringdown Search - CBC-S5-4

The S5 ringdown search, focusing on black holes with masses between 75 and 750 $M_\odot$, is in progress. We aim to bring it to completion by the fall of 2009, and commence the review and paper.

For high mass binary coalescences ($M \gtrsim 100 M_\odot$) , the majority of the power received by the LIGO and Virgo detectors will be from the ringdown part of the coalescence. Therefore, the search for coalescencing binary systems can be done looking for the final ringdown which has known waveforms parametrized by mass and spin (determining frequency and damping time of the ringdown). The uncertainty in the theoretical predictions for how the inspiral, merger and ringdown phases couple into a single waveform governed by a single set of source parameters (masses and spins) leads us to pursue such a ringdown-only search. This search has been completed on the S4 data, and is now being run on the S5 data using the same *ihope* pipeline as the low mass search. Because of the emphasis on low frequencies, the search is LIGO-only (ie, does not include Virgo data), covering all 24 months of S5 (including the overlap with VSR1). The tuning of the signal-based vetoes, coincidence parameters, and other search parameters is essentially complete. An ongoing project is to understand and compare the sensitivity of the ringdown, high mass and unmodelled burst searches in regions of parameter space where they overlap.

3.7 S6/VSR2 goals

The LIGO and Virgo detectors are currently undergoing upgrades to enhanced LIGO and Virgo+ configurations. It is hoped that these will provide approximately a factor of two increase in sensitivity across a broad frequency range, and consequenty around an order of magnitude increase in the expected rate of compact binary coalescences over the S5/VSR1 sensitivity. The S6/VSR2 run began in July 2009. Two preparatory engineering runs (E13, E14) have already been completed. The CBC group has the following scientific goals for S6/VSR2:

1. **Low mass CBC**: A search for binaries with masses between $1 M_\odot$ and $35 M_\odot$ and component masses greater than $1 M_\odot$ which covers the binary neurton star, neutron star-black hole and lower mass binary black hole space. The merger and ringdown parts of these waveforms do not contribute significantly and therefore the post-Newtonian waveforms are appropriate. There is a significant benefit, both astrophysically and in terms of understanding the detectors if the analysis can be carried out promptly. Therefore, we will aim for a low mass search with a latency of two weeks at the beginning of S6/VSR2. As the run progresses, we will attempt to reduce the latency to one day or less.

2. **Electromagnetically Triggered Search**:

   There are various Gamma Ray Burst satellites which provide announcements. As in S5, we would like to perform a search of the data associated with GRB triggers. The search will follow the same mechanism as for S5, and therefore we should be able to search rapidly for signals associated to GRBs. We aim for a latency of days or less at the start of S6. Investigations for a more advanced background estimation for the time of the GRB trigger is also envisioned. Finally, we plan to extend set of external triggers used, to incorporate bursts observed in radio telescopes.

3. **High mass Inspiral–Merger–Ringdown CBC**: A search for binaries with mass $M \gtrsim 25 M_\odot$. Here, the best knowledge of the last cycles of inspiral, merger and ringdown are crucial. There have been important advances in building templates for inspiral, merger and ringdown by modeling analytically the non-perturbative information contained in numerical-relativity simulations. However, at the present time, templates calibrated to numerical-relativity simulations span a small region of the parameter space which does not include large mass-ratio and spinning, precessing black hole binaries.

The CBC group will be following progress on computing waveforms and, if necessary, will incorporate the new features into the search pipeline. In addition, the group will continue to investigate the relative merits of different search techniques, including: full IMR waveforms; ringdown only waveforms; inspiral-merger-ringdown coincidence analysis; parameter estimation techniques; searches for unmodelled waveforms. By doing this, we hope to find the best method of detecting and extracting the parameters of CBC waveforms, given the latest analytical and numerical insight.

4. **Spinning systems:** As discussed above, non-spinning templates capture most of the SNR of waveforms from spinning systems. Work is underway to identify regions of parameter space where spin waveforms differ significantly from the non-spinning templates. In parallel, search using a physical template family of spinning waveforms [71] is being developed. Depending upon the results of these studies, we may mount an S6 search using spinning templates — either as an independent search or a followup to candidates produced by the non-spinning waveforms. The issue of spin is equally significant, if not more so, in the high mass regime. For these masses, the post-Newtonian approximation does not adequately describe the revelant part of the waveform. Numerical results for spinning systems are being obtained, but a full exploration of the spin parameter space will take some time.

In order to achieve these science goals, the CBC group has identified several areas of development which are required by the start of the S6 analysis, or shortly thereafter:

- **Data infrastructure - CBC-S6-1**: Contribute to the data infrastructure work coordinated by DASWG, including the prompt generation of $h(t)$; the aggregation of $h(t)$ data to off-site clusters for analysis, and the aggregation of segment information — containing detector operation times, and times of various data quality flags — to off-site clusters for analysis.

- **Data Characterization - CBC-S6-2**: Contribute to the data characterization work coordinated by the DetChar group, including the prompt and automated identification and generation of data quality vetoes and KW vetoes, and the evaluation of the quality and accuracy of the calibration and the $h(t)$ data stream. A significant effort is required to: (i) identify auxiliary channels containing noise events which are correlated with inspiral triggers in the gravitational-wave channels; (ii) investigate the couplings between these channels and the gravitational-wave channel; (iii) understand and identify particular artifacts by the signature across all appropriate channels; (iv) develop methods to clean the data of such artifacts, e.g., by vetoing noisy times; and (v) automate these methods within the pipeline and the on-line analysis.

- **Hardware Injections - CBC-S6-3**: Coordinate coherent hardware injections in all detectors, and provide prompt results on whether they were "detected" using the low-latency and week-latency searches described below. Provide support for blind hardware injection.

- **Low Latency Analysis - CBC-S6-4**: Conduct a low-latency (minutes) filtering of the data through low-mass inspiral templates using the *MBTA* package. This will initially be used for control room display/feedback (via InspiralMon at the LIGO detectors), for daily glitch studies and rapid analysis of the hardware injections.

- **Weekly ihope Analysis - CBC-S6-5, CBC-S6-6**: Conduct searches for gravitational waves from: low mass binary coalescence; high mass binary coalescence; and black hole ringdowns. These searches will be performed with a cadence one week. We will use the *ihope* pipeline on off-site clusters. Backgrounds, efficiencies, upper limits, and detection confidence follow-up procedures will be similar to the S5 searches, with incremental improvements. Weekly search summaries will be produced, and potential detection candidates (which may include blind hardware injections) will be followed up on promptly.

- **Automated Externally Triggered Searches - CBC-S6-7**: Perform searches for binary coalescence in near-coincidence with externally triggered events. Initially the search will be focused on short GRBs but may later expand to radio bursts and other sources. This search will be automated to begin processing the data once an external trigger is received. Initial results are expected within a day or so of the trigger, and likelihood based full results ready within a week.
- **Followup of Candidates - CBC-S6-8**: Using the infrastructure developed during the S5/VSR1 analyses, perform a rapid followup analysis of candidate events, using the detection checklist. Feedback findings to the detector characterization group.

It is clear that a significant amount of effort has gone into reducing the latency of the analysis (from months to days or less). There are several motivations for this: (i) Looking at the data as it is taken helps provide an understand of the performance of the instruments, and particularly their non-stationarity. (ii) In the case of a candidate detection, the detectors will be in a similar state to when the candidate was observed, making followup investigations more feasible. (iii) We are working towards a low latency analysis and alert system so that gravitational wave candidates will be announced to astronomers for follow up electromagnetic observations.

At the time of writing, the majority of the items above are in place, and have been tested, at some level, during the E13 and E14. We envision that there will be significant investment required to ensure that the prototype systems that are in place are stable enough to operate over a year long science run. Also, despite the increasing automation of the analysis, there is still a significant investment required by the analysts in looking at the results, plots and figures of merit obtained and feeding back information to the operators, instrumentalists, detector characterization group and into providing pipeline enhancements.

3.8 High priority analysis development

In the previous section, we have described the state of the CBC analysis which will be in place at the start of the S6/VSR2 run. There are a number of developments which have the potential to significantly improve the searches being undertaken. These will be developed and deployed into the analyses as rapidly as is feasible.

- **Improved Background estimation - CBC-S6D-1**: An important aspect of the search for gravitational waves is to understand the background present in the experiment. This is usually addressed by re-analyzing the data through the pipeline while applying time slides to the data from each site. This method should provide an accurate estimate of the accidental background (from uncorrelated noise sources) for the experiment Experience with the use of time slides to estimate the background in gravitational wave experiments often shows unexpected results. In addition, it is difficult to perform enough time shifts to accurately evaluate the small false alarm rate that a detection candidate will surely require. A significant effort is required to: (i) study and understand the time-slide method in the context of large parameter space searches, (ii) study alternative methods of estimating the background (iii) investigate possible sources of bias in the background estimate from the environment.
- **Extension of externally triggered search - CBC-S6D-2**: Currently, the externally triggered search is triggered by announcements of short GRBs. There are arguments that binary coalescences will also produce signals in other frequency bands, such as radio and x-ray. The externally triggered search will be extended to include transients which are observed in these wavelengths.
- **Coherent Analysis - CBC-S6D-3**: In general relativity, the gravitational waveform contains only two polarizations. Therefore, in a network with more than two detectors, there is significant advantage to be gained from performing a fully coherent analysis of the data from all detectors to ensure

that the observed data is consistent with a gravitational waveform from a given location in the sky. Furthermore, for a given sky location, it is possible to construct a "null stream" which will contain no gravitational wave signal. If a candidate event is due to a noise glitch, it is unlikely to cancel in this null stream. Therefore, a coherent and null stream analysis provides a signficant benefit in both noise rejection and parameter extraction. There are currently several implementations of a coherent inspiral analyses. However, none of these are fully integrated into the analysis pipeline. Completing this task will provide a significant enhancement to our current searches, and should be implemented in all pipelines (low mass, high mass, GRB and ringdown).

- **Searches incorporating spin - CBC-S6D-4**: The components of an astrophysical, coalescing binary system are expected to be spinning. Depending on the spin magnitudes and orientations the emitted gravitational waveforms can be highly modified from the pure point-mass inspiral due to spin-orbit and spin-spin couplings and resulting precession of the orbital plane as the signal sweeps through the detectors' band. Extensive simulations have shown that for significant parts of the spin parameter space, there is only a small loss of efficiency (or "reach") from using non-spinning template waveforms. However, there are some regions of parameter space in which amplitude modulation due to spin-orbital coupling and phase modulation can lead to significant loss of SNR when using non-spinning waveforms. Work is ongoing to implement a search using a physical template family corresponding to single spin waveforms, and identifying regions of parameter space where such a search is warranted.

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

- **Low latency localization - CBC-S6D-6**: As the S6/VSR2 run progresses, we would like to be in a position to alert other astronomers of the time and sky location of gravitational wave candidate signals. Coincident triggers from *MBTA* will be used for low-latency (minutes) identification of potential gravitational wave detections, with rapid sky localization. These candidate detections will be made available for follow-up observation by telescopes, per external collaboration agreement.

- **Triggers for EM follow-up - CBC-S6D-7**: Externally triggered searches make use of triggers from EM (or neutrino) detectors to look deeper for GW signals. Conversely, GW detection candidates, well localized in time and sky location, can be used to trigger follow-ups by EM detectors (radio, optical, gamma ray telescopes). The Burst group has already developed methods, and relationships with many different external (EM telescope) groups, to realize this in S6 (with the full awareness that it is a long-shot). The CBC group will want to make use of these methods and relationships to trigger EM follow-ups of our most promising GW candidates.

- **Data quality, vetoes and followups - CBC-S6D-8** The CBC group applies a detection checklist to the candidate-events surviving to the end of the search. This checklist incorporates new tests on the events which help to discriminate background from signal. Current efforts are in progress to make the tests of the checklist more reliable, effective, and less subjective. As experience is gained, the group will refine the checklist in order to make the tests less subjective and more quantitative.

- **Improved Classification of Signal and Background - CBC-S6D-9**: 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, . . .). It is expected that by making use of more of this information, a better separation between signal and background would be 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:

  - **Multivariate Classification techniques**: Things like neural networks, decision bagging trees. These programs take in all the available information about a coincident trigger, and are "trained" to separate background coincident triggers from foreground signal.
  - **Likelihood Ratio techniques**: This approach calculates the one dimensional probability distributions of various parameters from both the foreground and background. Then, for a given trigger, it can produce a likelihood ranking based on these distributions, for use as a detection statistic.
  - **Bayesian Likelihood techniques**: By performing a Bayesian Likelihood calculation, marginalising over the various parameters associated to the signal, one can obtain a detection statistic that makes use of the extra information supplied in the Bayesian priors.

- **Overlap with Burst searches for high-mass systems - CBC-S6D-10**: 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_{sun}, 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.

  Indeed, burst searches may be more powerful, because: (a) burst searches typically make no specific assumptions about the nature of the waveform or its phase evolution, so that if the astrophysical waveforms differ significantly from the templates used in the CBC searches (eg, due to the presence of spin, or some failure of our NR matching methods), they be more robust in detection; (b) as of this writing, burst searches make use of the coherent detector network in their primary detection statistics, while the CBC group employs a coherent analysis only as a follow-up for loud detection candidates.

  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.

Publications that present the results of LIGO S5 and S6 searches for high mass CBCs and ringdowns may be prepared separately by the CBC and Burst groups, following the principle of distinct publications for distinct searches. Or, they may be prepared jointly by the two groups, on the principle of one publication for each distinct type of astrophysical source (as was done for GRB070201). This is an important question that requires resolution in the coming year.

- **Parameter estimation - CBC-S6D-11**: A good understanding of parameter estimation abilities of the detector network is critical for both the design of our analyses, and in extracting astrophysical parameters from gravitational wave observations. There are several areas of active development:
    - Understand the systematic effects of calibration errors on parameter estimation. This study must be carried out using real and simulated detector noise.
    - Use Markov-Chain-Monte-Carlo technique to determine posteriors on the parameters of a detected signal. At the same time, study alternative approaches and determine the relative merits of each.
    - Implement tools to include higher-order post-Newtonian corrections to the waveform amplitudes, and hence other harmonic structure. It is known that the other harmonics can significantly enhance our ability to detect gravitational waves from binary systems as the mass ratio decreases. These harmonics also break degeneracies between many of the parameters. By including this information in an observation, it may be possible to better measure the binary parameters.
    - Use Markov-Chain Monte-Carlo 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.
    - 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.

- **Thorough evaluation of our readiness for first detections - CBC-S6D-12**:

  Each GW detection candidate is assigned a false alarm probability (FAP) based on a procedure that is determined in advance of the box opening. this constitutes our primary, quantitative and unbiased discriminant for detection confidence. This is usually based on specific SNR thresholds, signal-based vetoes, data quality veto category criteria, background estimation methods, categorization of coincident triggers (eg, in regions of template bank space, coincidence criteria, date collection epoch), etc.

  The use of time-slide coincident events currently limits us to of order 100 independent samples, so that a FAP of less than 1/100 cannot be estimated. Work is in progress to make use of single-detector triggers to "extrapolate" to very low false alarm rates for accidental coincident triggers. In this way, we hope to be able to estimate false alarm probabilities as low as 0.1% or lower, for loud zero-lag coincident GW candidate events. Here again, all "lessons learned" from loud zero-lag coincident triggers must have no influence on the procedures for these estimations. This is a high-priority goal for the coming year.

  In addition, the CBC group makes use of many additional checks and tools to assess detection confidence, as part of the "follow-up pipeline" described above. These include the examination of coincident triggers satisfying different data quality category requirements or different coincidence criteria;

incorporating information from auxiliary channels; incorporating information from the coherent and null-stream analyses; information from Qscans; etc. We endeavor to make these tools as quantitative as possible, and roll them into the main analysis pipeline in the future. This is another high-priority goal for the coming year.

Until then, most of these detection confidence checks are at best semi-quantiative, and in many cases, are wholly qualitative. 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. But they lay no claim to be quantitative, unbiased, or statistically reliable estimators of the significance of a candidate event.

- **Next Generation Low Latency Pipelines - CBC-S6D-13**: The Virgo MBTA pipeline is in place during S6 for the generation of CBC triggers for detector characterization, and for rapid identification of coincident triggers that could be GW detections. Work is in progress to rapidly locate sources in the sky associated with such triggers, for EM follow-up. In addition, work is in progress to develop new, stream-based low latency search and sky localization pipeline infrastructure for both S6 and for the long-duration chirps expected in Advanced LIGO.
- **GPU-enhanced CBC pipelines - CBC-S6D-14**: The optimal filtering and chisq computations that are at the heart of the CBC pipeline are very cpu-intensive. Modern Graphical Processor Units (GPUs), optimized for these kinds of computations, can greatly speed up, and/or greatly reduce the cost of, conducting a wide-parameter search for CBC (and also continuous wave) signals. The CUDA library greatly facilitates the conversion of our existing code to make use of GPUs. We are developing the necessary software modifications, and designing next-generation computing clusters to implement GPU-enhanced CBC search pipelines.

3.9 Advanced detector era

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

1. **Multi-band and Hierarchical Search Methods**: Advanced detectors will have greatly reduced noise at low frequencies, allowing for detection of gravitational waves down to 10 Hz. Inspiral waveforms can spend many minutes in the low frequency band. It will be necessary to restructure the existing pipeline, perhaps along the lines of the Virgo MBTA pipeline, in order to handle such long signals coherently. Hierarchical methods need to be implemented and tested. Several different hierarchical schemes exist which search over coarse grained parameters in an early part of a search, following up on triggers from this coarse grained search with finer grained parameters. An effort is required to: (i) develop code to translate triggers into hierarchical banks for injection and use by the filtering code; (ii) test and develop intuition for hierarchical searches using the current pipeline; (iii) develop tools to tune this hierarchical pipeline.
2. **New approaches to filtering**: The core of the analysis pipeline is the filtering of the data through bank(s) of templates, in the frequency domain (sections 3.3.1 and 3.3.3). This process is time consuming and (for large bank parameter spaces) computationally limited. This results in a variety of problems: analysis turn-around time is slow; securing adequate and sufficiently reliable computing resources is difficult; it is difficult to develop new (or tune existing) analysis features such as

signal-based vetoes; and there is an inherent latency in the pipeline, making it unsuitable for real-time searches or searches using very long-duration templates (necessary for Advanced LIGO data). There is exploratory work in progress to redesign the filtering stage to make it more modular, make use of multi-band hierarchical filtering, employ low-latency time-domain filtering, speeding up the processing (or reducing the number of computers) using GPUs, and other innovations.

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

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

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

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

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

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

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

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

4 Searches for general burst signals

4.1 Introduction

The mission of the LSC-Virgo Burst Analysis Working Group (also known as the "Burst Group") is to search as broadly as possible for gravitational-wave "bursts" (GWBs), *i.e.*, short-duration gravitational-wave signals. A wide variety of astrophysical sources are expected to produce such signals at some level; sophisticated models have been made in some cases, while in other cases the amplitude and waveform of the GW signal is highly uncertain. Thus, the Burst Group must utilize robust signal detection methods. Initially, the group focused on signals much shorter than 1 second in duration and with little or no assumption about their morphology [104]. Since then the Burst Group has evolved and today (mid-2009) is also pursuing longer signals, as well as incorporating knowledge of the signal waveform when available. It thus encompasses the broadest range of sources of gravitational radiation, seeking some specific targets while also providing coverage for yet unknown phenomena and mechanisms of emission and propagation. Section 4.2 summarizes recent GW burst search results, while section 4.3 describes many of the basic methodologies used in burst searches.

Sections 4.4 and 4.5 describe the goals and plans for GW burst searches with the LIGO, GEO and Virgo instruments over the next ∼3 years. During that time we will finish analyzing the data in hand from S5/VSR1, undertake a few selective searches with post-S5 "AstroWatch" data, and carry out a broad program of analyzing the data from the upcoming S6/VSR2 run. Many specific analyses and investigations are mentioned in this document as activities that current Burst Group members intend to carry out; other activities, not mentioned here, may also be valuable in furthering the general scientific goals of the group. The data analysis techniques and tools that we continue to refine or develop over the next few years will form the foundation of burst data analysis in the Advanced LIGO/Virgo era. (This document is not meant to address data analysis questions relevant *only* to the advanced detectors that will be taken up beyond the ∼3-year horizon.)

The Burst Group's ultimate goal in S6/VSR2 is the unequivocal detection and astrophysical interpretation of gravitational-wave burst-like signal(s). In the absence of any detection(s) the goal is to set physically significant upper limits on the gravitational-wave emission and astrophysical or cosmological mechanisms responsible for that. There are several approaches to the pursuit of these science goals. The broadest net is cast by the so-called untriggered burst search. This search "listens" to the whole sky at all possible times allowed by the operation of the instruments for statistically significant excursions of signal power within the sensitive frequency band of the instruments (sec. 4.4.1). In some cases, making astrophysically motivated assumptions about sources—including models of the GW waveform or the source population—points to ways to implement customized searches which are better than the general untriggered search in terms of sensitivity and ability to interpret the results (sec. 4.4.2). Another direction of interest is to look for gravitational-wave bursts starting with energetic transient astrophysical events observed in the electromagnetic or neutrino spectrum. The known sky position and generally short coincidence window associated with these searches together with the source phenomenology offer improvements in sensitivity and rich potential for interpretation within an astrophysical context. These are the so-called externally triggered burst searches (sec. 4.4.3). We also plan to symmetrize this relation by seeking electromagnetic counterparts starting with gravitational-wave burst candidates. For this we expect to form collaborations with observer groups and pursue Target of Opportunity (ToO) observations that will provide the possibily for gravitational-wave burst candidates to trigger non-gravitational-wave observatories (sec. 4.4.4).

The breadth of the burst searches is such that in some cases there is overlap with other LSC-Virgo groups in terms of the astrophysical sources that are being pursued as well as the methods invoked. Such cases include short-duration gamma-ray bursts and the merger and ringdown of binary compact objects, all of which are being pursued by both the Burst and CBC Groups. Longer duration transients (of order minutes or longer) may also be pursued using methods the stochastic and continuous waves groups have developed for their corresponding searches. The Burst Group will coordinate with the other working groups in areas of common interest to ensure that the best possible scientific results and corresponding publications are brought forward.

4.2 Recent observational results

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

Soft gamma repeaters (SGRs) are astrophysical objects which occasionally emit "flares", or bursts of soft X-rays. SGRs are believed to be highly magnetized neutron stars [105], whose vibrational modes may be excited during soft-gamma burst events. We have executed a triggered search for prompt GWBs associated with the SGR 1806–20 giant flare which occurred in December 2004, along with 190 lesser events from SGR 1806–20 and SGR 1900+14 which occurred during the first year of S5 [106]. This was the first search to target neutron star f-modes, usually considered the most efficient GW emitting modes. No evidence of GWBs associated with any SGR burst was found. Model-dependent upper limit estimates on the isotropic GW emission energies $E_{\rm GW}$ (at a nominal distance of 10 kpc) were set using f-mode ringdown waveforms predicted in [107] and band- and time-limited white noise bursts. Upper limits varied between 3×10^{45} erg and 9×10^{52} erg depending on waveform type, detector antenna factors and noise characteristics at the time of the SGR burst. These upper limits are within the theoretically predicted range of some SGR models, such as [108].

We previously had searched for quasiperiodic oscillations (QPOs) following the same giant flare, for frequencies and time intervals (a few minutes) matching what was observed in X-rays [109]. The X-ray QPOs were attributed to seismic modes of the progenitor neutron star [110], which could also emit GWs. No evidence of GW QPOs was observed in that search, and energy limits were set which were of the same order as the total (isotropic) energy emitted over the course of the giant flare in the electromagnetic spectrum.

The initial SGR GWB search was followed up with an externally-triggered "stacked search" for GWBs associated with the 2006 March 29 SGR 1900+14 storm [111]. This new search method stacked the GW data around the times of individual soft-gamma bursts in the storm to enhance sensitivity for models in which multiple bursts are accompanied by GW emission. GW excess power time-frequency tilings containing individual burst triggers were aligned to their corresponding electromagnetic emissions. Two GW emission models were used: a fluence-weighted model and a flat (unweighted) model for the most electromagnetically energetic bursts. No evidence of GWBs were found under either model. Model-dependent GW strain, isotropic GW emission energy $E_{\rm GW}$, and $\gamma \equiv E_{\rm GW}/E_{\rm EM}$ upper limits were estimated the same waveforms used in the original SGR search [106]. $E_{\rm GW}$ upper limit estimates (per burst, assuming a nominal distance of 10 kpc) were found in the range 2×10^{45} erg and 6×10^{50} erg depending on waveform type. These limits are an order of magnitude lower than the upper limits published previously for this storm and overlap with the range of electromagnetic energies emitted in SGR giant flares.

We have searched for a GWB associated with GRB 070201, a short-duration GRB whose electromagnetically determined sky position is coincident with the Andromeda galaxy (M31) [112], using a cross-correlation method previously applied to GRB 030329 [113]. This search was performed and published in conjunction with a templated search carried out by the Compact Binary Coalescence Working Group, as described in [112]. Possible progenitors of short GRBs include compact binary coalescences or SGR flares. No evidence of a GWB signal was found, and a compact binary progenitor with masses in the range 1 $M_\odot < m_1 < 3$ $M_\odot$ and 1 $M_\odot < m_2 < 40$ $M_\odot$, located in M31 was excluded at >99% confidence. If the GRB 070201 progenitor was not in M31, then we can exclude a binary neutron star merger progenitor with distance D $<$ 3.5 Mpc, assuming random inclination, at 90% confidence. An unmodeled GWB from GRB 070201 probably emitted less than 7.9×10^{50} erg in any 100 ms long period within the signal region if the source was in M31 and radiated isotropically at $\sim$150 Hz. Based in part on these findings, it is likely that the event was an SGR giant flare in M31.

During the S5/VSR1 run, over 130 GRBs were detected by satellite-based gamma-ray experiments during times when two or more GW detectors were operating; a joint LIGO-Virgo search for GW bursts associated with those GRBs has been carried out and the results will be made public in July.

Also in the past year, Virgo published the results of a search for GWs associated with GRB 050915a, which occurred during Virgo's C7 commissioning run [114].

Untriggered all-sky searches in the S5 LIGO science run were split into the first and second calendar year to enable analysis to move forward while more data was still being collected and calibrated. In the first calendar year, the all-sky burst search in the frequency range 64–2000 Hz [115] was run over a period of 268.6 days (combined triple-coincident and Hanford-only livetime). Root-sum-square (rss) strain limits were placed in the range $6 \times 10^{-22}\,\mathrm{Hz}^{-1/2}$ to a few times $10^{-21}\,\mathrm{Hz}^{-1/2}$ for a variety of waveform morphologies, and a frequentist upper limit of 3.6 events per year was placed on the rate of strong gravitational wave bursts at a 90% confidence level. These limits are the strictest ever placed on general gravitational wave burst emission. Using theoretical models of gravitational wave emission, the sensitivity of these analyses can also be expressed in terms of astrophysical reach. According to tests of models produced by the Goddard numerical relativity group [116], the binary merger of a system of two 10-solar-mass black holes would be detectable with 50% efficiency at a distance of 4 Mpc while a system of two 50-solar black holes would be detectable out to around 180 Mpc. Using models produced by Ott et al [117], supernova core collapse gravitational wave emission was estimated to be detectable at a range from 0.6 to 24 kpc, depending on the progenitor model.

Also during S5, all-sky "high frequency" burst searches were conducted for the first time. These analyses extended the frequency coverage of burst searches up to 6 kHz. The first calendar year high-frequency analysis [118] placed a frequentist upper limit of 5.4 events per year on the rate of strong gravitational wave bursts in LIGO triple coincident livetime.

Additionally, Virgo published the results of a search for GW bursts using data collected during Virgo's C7 commissioning run [119], which was conducted over the course of five days. The analysis placed root-sum-square strain limits around $\sim 10^{-20}$ on the emission and a frequentist upper limit on the rate of strong bursts of 1.1 events per day at a 90% confidence level.

In addition to untriggered analyses designed to be sensitive to a variety of transient waveforms, it is also possible to perform matched filter searches designed to look for specific burst signals. Such an analysis was conducted to look for gravitational wave bursts from cosmic strings in S4 data. The limits placed on the burst rate from strings were used to constrain the possible parameter space of cosmic string models, including string tension, reconnection probability and loop sizes [120].

4.3 Basic methodologies used in burst searches

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

4.3.1 Signal extraction methods

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

In a few special cases when an accurate signal model is available, such as for cosmic string cusps, a search can be done using matched filtering with a bank of templates. Otherwise, un-modeled bursts can be identified in the detector output data as excess-power events localized in time and frequency. Therefore, for better localization of burst events usually the analysis is performed in the time-frequency (TF) domain. To obtain a TF representation of data a number of transformations are used, including windowed Fourier transforms, discrete wavelet decompositions [121] (Symlets, Meyer wavelets) and continuous wavelet transforms [122] (Q-transform). These transformations have been actively used in GW burst search algorithms. At the same time, the Burst Group is exploring other interesting approaches such as the Hibert-Huang Transform (HHT) [123], an adaptive time-frequency decomposition which is able to resolve rapid signal frequency variations and thus can be used for more detailed study of the TF content of gravitational wave events.

A handful other burst search methods have been implemented which do not start with a TF representation. These include a change-point analysis of bandpassed data [124] (BlockNormal) and a cross-correlation analysis using pairs of detectors [125] (CorrPower).

Networks of GW detectors are very important for the detection of gravitational waves. Receipt of consistent signals by multiple instruments will increase confidence in an event, especially for a burst signal which otherwise may not be distinguished from the instrumental and environmental artifacts produced in the detectors. Also, with multiple detectors it is possible to disentangle the two gravitational wave polarizations and determine the direction of the source. For these reasons, the Burst Group has developed multi-detector search algorithms which can be approximately devided into two groups: incoherent and coherent. In the incoherent algorithms [124, 122, 127, 126], excess-power events in individual detectors are identified and then a time coincidence is required between the events in different detectors. The incoherent algorithms are particularly useful for characterization of individual detectors and they are actively used for studies of enviromental and instrumental artifacts (see Section 4.3.2). Coherent algorithms are based either on cross-correlating pairs of detectors [125, 106] or on a more general approach called coherent network analysis (CNA).

Coherent network analysis addresses the problem of detection and reconstraction of gravitational waves with networks of GW detectors. In coherent methods, a statistic is built as a coherent sum over the detector responses and, in general, it is more optimal (better sensitivity at the same false alarm rate) than the detection statistics of individual detectors. It is a high priority of the group to complete development of coherent network algorithms and apply them to burst searches. There are several CNA approaches employed by the Burst Group. The constrained likelihood method (coherent WaveBurst [128, 129, 130]) has been developed and used for the S4 and S5 all-sky searches. Also a likelihood method for triggered searches (X-pipeline [131]) was used during the S5/VSR1 run. The group is working on a Bayesian formulation [132] of coherent network analysis and on maximum entropy methods (MaxEnt [133]). All these methods are being upgraded or developed for use during the S6/VSR2 run. Also dedicated CNA algorithms which may use partial information about burst sources and incomplete source models are being investigated by the Burst Group (see Section 4.4.2).

Coherent algorithms enable not only the detection of gravitational waves, but also astrophysical analyses of GW signals and measurements of source properties, including reconstruction of the GW waveforms and source coordinates. These are necessary tools for the nascent field of GW astronomy. Prompt detection of GW signals and estimation of source coordinates enables coincident observations with other astronomical instruments, which can significantly increase the confidence of detection. Such measurements may not only aid the first detection of gravitational waves but also they will give us fundamentally new information about the GW sources and their population distribution. The Burst Group has made significant progress in the development of source localization methods. However this work is only in the beginning. The group will work to address the source localization problem and apply the coordinate reconstruction methods to the enhanced and advanced detector networks.

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

4.3.2 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. During the S5/VSR1 run, in a coordinated effort of Burst and CBC Group members (the Glitch Group, a part of the Detector Characterization effort described in section 2.1.4), we studied the correlation of the rate and strength of single detector transients to trends and transients in the sensing and control systems that maintain the interferometers in their operating condition, as well as monitors of the physical environment in which the interferometers sit: vibration, sound, magnetic and electric fields, power line voltages, and others. These studies led to the identification of times likely to be contaminated by non-GW effects, which the Burst Group uses to veto event candidates found by the GW channel analyses. Based on their duration and character, we distinguish between *data quality* vetoes and *event-by-event* vetoes.

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

- Category 1 vetoes define which data can be safely analyzed by the search algorithms; they remove features that could affect the power spectrum (for instance, detector instability prior to loss of lock or severe saturations).

- Category 2 vetoes define the *full* data set, where to search for detection candidates. They remove times when the detector is unambiguously misbehaving (examples are auxiliary channel saturations, and power main glitches that couple magnetically in the detector) with a well-understood physical coupling. They introduce a small dead time (fraction of a percent) and have high efficiency for removing single-detector outliers.

- Category 3 vetoes define the *clean* data set to be used to set an upper limit in the case of no detection. They identify times with an excess of single-detector triggers in the gravitational-wave channel, but the correlation is not unambiguous and they may introduce a large dead time, up to 10%. Typically, these vetoes are associated with high seismic activity. If a detection candidate is found at these times, the category 3 flag is taken into account in the event followup, as described in section 4.4.6.

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

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

Event-by-event vetoes (sometimes referred to simply as *vetoes*) are short intervals (typically 100 ms or shorter) that mark individual transients in an interferometer's output with a coincident transient in one or more diagnostic signals. They are identified with extensive statistical studies of coincidence between *KleineWelle* triggers in auxiliary channels and event candidates. Many possible veto conditions are considered in a hierarchical classification process in which the vetoes are ranked on the basis of the significance of their correlation with gravitational wave triggers, and their significance is re-evaluated in subsequent iterations, after each condition is applied. In the S5/VSR1 run, we identified a set of category 2 conditions, effective at removing environmental correlations at the Hanford site (H1-H2) and category 3 conditions, which include both interferometric and environmental signals.

This structure of data quality and event-by-event vetoes, with their categorization, will be extended ro S6/VSR2. In preparation for the run we are working to automate the DQ and veto categorization process. To ensure the success of the online analysis, the Burst Group is collaborating with the Detector Characterization group for a prompt and documented definition of DQ flags and veto criteria. In particular, Omega single-interferometer triggers will be produced online and time-frequency plots, as well as histograms and correlograms, will be available with a latency of 2-3 minutes. These plots will be available for commissioning and science monitoring, and used for the identification of DQ flags. Moreover, coherent WaveBurst coincidence triggers will also be available for science monitoring. These are multi-detector triggers and thus are relevant for understanding the background in the burst analysis. Finally, data quality categorization will be periodically revised, as the understanding of the detectors improves.

4.3.3 Accidental background estimation

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

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

In other cases where the trigger time of the possible GW source is known, the background estimation procedure can take advantage of this information by defining off-source samples even with a single detector.

The significance of any event candidate found in the on-source (unshifted) analysis can then 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. In order to have an objective, non controversial assessment of the FAR, we follow a *blind* statistical procedure, e.g. using the time-shifted data (without examining the unshifted data) to select the exact procedures to compute the test statistics and the thresholds applied to them.

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

4.3.4 Simulations

While tuning an analysis pipeline, and after finalizing all of the event selection criteria, the *detection efficiency* of the observation is assessed to interpret the results in terms of different signal models. This is done using the "mock data challenge" (MDC) technique, i.e. special runs of the analysis are done with simulated signals added to the actual detector data at pseudo-random times. The simulated signals are chosen to span the expected range of signal properties (frequencies, durations, etc.) but not to exhaustively test all plausible signals; the robustness of the signal extraction methods used allows extrapolating 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 else at the fixed sky position of an astrophysical event used to trigger the search.

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

Hardware signal injections

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

4.4 Science goals

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

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

4.4.1 Search as broadly as possible for GWs

There is strong astrophysical motivation [134] for searching for burst-like gravitational-wave signals using ground-based laser interferometers. The emphasis has historically been on astrophysical systems for which the resulting burst waveforms are either poorly modeled or remain unknown, including (but not limited to) binary compact star mergers and core-collapse supernovae. In recent years numerical relativity calculations have offered significant information on the waveform accompanying binary mergers [135] as well as important new information on the features of signals accompanying core-collapse [136]. Burst sources with well-modeled waveforms include emission from neutron star ringdowns following a pulsar glitch, black hole ringdowns or cosmic string cusps (sec. 4.4.2).

Typical signal durations predicted by these models are from less than a millisecond to hundreds of milliseconds and with signal power in the frequency range from 50 Hz and up to few kHz. Various models of GW emission from core-collapse supernovae [136] and gamma-ray burst progenitors [137] may result in signals lasting up to several seconds. Although the best sensitivity of ground-based interferometers is achieved around ∼150 Hz, the instruments remain sensitive to within a factor of ∼10 over most of their useable frequency band (60–6000 Hz), thus motivating searches to utilize their full capacity. 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 first approach in pursuing gravitational-wave bursts thus relies on generic search methods and on minimal assumptions about the signal morphology in order to provide detection capability for the widest possible range of sources. The search analyzes as much of the available data as possible from times when two or more detectors are running well, and no assumption is made about the direction and the source(s) of gravitational-wave bursts. This kind of untriggered burst search is often referred to as “all-times, all-sky”. The search is traditionally tuned in such a way that a very low number ($\ll 1$) of background events is expected over the duration of the observation. Thus, any signal events (foreground) resulting from this analysis constitute detection candidates for which further and exhaustive investigation is performed in order to understand their origin (sec. 4.4.6). The immediate result of this search is thus 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 specific 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 in a quantified way before being considered as part of the group-wide search for bursts.

Practical considerations have led “all-times, all-sky” burst searches to divide the frequency signal phase-space into broad categories covering “low” (∼50–2000 Hz) and “high” (∼2000–6000 Hz) frequency bursts. Also, the S5 data set has been split near the midpoint of the two-year run, with the second half overlapping at the end with the 5-month VSR1 run. The analysis of the first calendar year of S5 is now complete and two all-sky papers have been submitted for publication (as of June 2009). The analysis of the second year of S5, including the joint running time with VSR1, is in the final stages and remains the highest priority of the group. A single paper describing the search over the entire frequency regime is expected be submitted by early Fall 2009.

Continuing to pursue bursts over the full spectrum of frequencies allowed by the instrument and motivated by the source phenomenology will remain of the highest priority in the S6/VSR2 science run. Additionally, expanding coverage to seconds-long (if not longer) bursts is identified as something the group should give emphasis on. There are also questions surrounding the search for bursts that go beyond the all-sky "counting experiment" approach that the search has essentially taken so far. For example, can the distribution of weaker sources be revealed through the statistical analysis of the amplitude distribution of signal and background gravitational-wave burst candidates? Is the gravitational-wave burst sky uniform? How can the ability of the detector network to reconstruct event sky positions be used to perform sky map comparisons between foreground and background? Such questions can in principle improve sensitivity and provide hints for possible astrophysical models behind candidate signals. Moreover, they have routinely used by cosmic ray experiments in searching astrophysical sources of high energy photons and neutrinos [139, 140, 141, 138]. The availability of S5/VSR1 data provides an opportunity for the exploration and benchmarking of these search methods prior to applying them to data from the upcoming S6/VSR2 run.

4.4.2 Use knowledge of astrophysically plausible signals

Burst sources are un-modeled by definition. However, recent progress in astrophysics and numerical relativity provide knowledge about the intensity and morphology of expected gravitational wave signals from some sources; the most plausible being the coalescence of compact objects, over a broad parameter space, and core collapse supernovae. Predictions for binary black hole (BBH) coalescence rates are based on population synthesis models constrained by astronomical observations. Due to a number of assumptions and unknown model parameters, the rate estimates are still uncertain at present. According to estimates compiled from the literature [1], 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 [142]. 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 [136]. The Burst Group aims to exploit such knowledge for an astrophysical interpretation of its results.

Burst searches do not normally use these calculated waveforms as templates for matched filtering. However, available waveforms are used to test the efficiency of the detection pipelines and quantify the reach of the search. In particular, the coalescence of black holes, whose waveform only has a few cycles in the interferometers' sensitive band, is of particular interest for the Burst Group. The IMR (*Inspiral-Merger-Ringdown*) working group is a collaborative effort with the CBC Group, where the same set of phenomenological and EOBNR waveforms [143, 144, 145] is being analyzed by burst search methods, inspiral and ring-down matched filtering to compare detactability across the parameter space. This includes an event-by-event comparison of events found by each search (simulated and background) and a study of how to combine inspiral, burst and ringdown triggers in a single result. This study currently includes systems with 25-350 $M_\odot$ total mass and no spin. Preliminary studies based on numerical relativity waveforms show the magnitude and orientation of the spin of the coalescing black holes has a significant impact on the range of burst searches: this type of study will 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

---

[1] http://www.lsc-group.phys.uwm.edu/ligovirgo/cbc/protected/papers/rates/rates.pdf

to reconstruct the waveform, compare it to known models and extract source parameters. Some progress towards waveform reconstruction has already been made in this direction, via coherent techniques, but more progress is needed to compare a candidate to merger waveform parameters, using Bayesian [146] or MCMC [147] techniques (currently being explored by the CBC Group) or new techniques such as the Hilbert-Huang transform [148, 123].

For some targets, a customized search may be more appropriate than the all-sky flagship analysis of the Burst Group. Such searches will be pursued only after preliminary studies show them to be significantly more sensitive than the flagship analysis. In certain instances, the existing all-sky pipelines can be tuned to target a specific signature. For instance, elliptical constraints can be applied to *coherent WaveBurst* to search for elliptically polarized signals such as black hole mergers [149]. Or the configuration file for *Omega* could be modified to target high frequency, long waveforms from certain models for core collapse supernovae [150, 151, 152, 153, 154]. In other cases, a specific matched filtering search has been devised, such as that for cosmic string cusps [120] or a targeted matched filter search for supernovae, with specific models.

Customized burst analyses are also used to exploit the source sky locations. The group developed and implemented techniques to search for gravitational waves in coincidence with electromagnetic events such as GRBs, SGRs and optical supernovae, for which the sky location is known; see section 4.4.3 for a full discussion. Dedicated searches have been designed and will be used to target specific sky locations, including but not restricted to the galactic center or M31. In addition, the Burst Group will use directional information from burst candidate triggers to form sky map probabilities of their astrophysical origin. A comparison of such maps with background could unveil directional correlations with known or new point sources, as well as correlations with ensembles of galactic and extragalactic sources (as has been done by the Auger Collabration [138], for instance). This directional anisotropy search, in preparation for S6/VSR2, relies on triggers and direction information from the all-sky searches. It may ultimately be optimized for each possible arrival direction, like a *radiometer*.

4.4.3 Look for GW transients triggered by energetic astrophysical events

Many gravitational-wave sources should be observable in more traditional channels. Directly relevant astrophysical observations are in abundance, ranging from Earth-based astronomical data, through sensitive GRB/X-ray satellite detections, to neutrino signals. Knowledge of the time and/or sky position of the event increases the sensitivity of the GWB search compared to an untriggered all-sky search. Perhaps more importantly, the association with a known astrophysical event may be critical in increasing our confidence associated with a candidate GWB detection.

Here we summarize some of the promising externally-triggered (ExtTrig) transient searches: those associated with gamma-ray bursts (GRBs), soft-gamma repeaters (SGRs), neutrino bursts, optical supernovae, and pulsar glitches. A high priority for S6/VSR2 is to invoke in near-real-time the baseline GRB and SGR searches [112, 155, 156] used in S5/VSR1. The prompt availability of results from reviewed and published analysis method will allow relatively fast publications on exceptional astrophysical triggers.

Triggers from Gamma-ray bursts (GRBs) and Soft Gamma Repeaters (SGRs)

GRBs are intense flashes of gamma rays that are observed approximately once per day, isotropically distributed across the sky. GRBs are divided into two classes by their durations [157, 158]. Long ($\gtrsim$2 s) GRBs are associated with star-forming galaxies of redshifts of $z \lesssim 8.2$ and core-collapse supernovae [159, 160, 161, 162]. Short GRBs ($\lesssim$2 s) have been observed from distant galaxies of different types and from Galactic sources (distances from $\sim 1.5\,\mathrm{kpc}$ to $\sim 50\,\mathrm{kpc}$), such as SGR 1806$-$20 [164, 165, 163, 166]. Most short GRBs are believed to be due to the merger of neutron star or neutron star – black hole binaries [167, 168], while up to 15% may be due to SGRs [165, 169].

Each of these progenitors could be a source of gravitational waves. The central engine of both long GRBs and short GRBs produced by binary mergers (non-SGR GRBs) is likely a rapidly rotating black hole with accretion disk. Gravitational waves could be produced, for example, from fragmentation in the disk [170], or suspended accretion due to interactions between the BH spin and magnetic fields [171]. Binary mergers should, of course, also emit strong "inspiral" signals with relatively well-modeled amplitude and frequency evolution. The current LIGO-Virgo network is sensitive to GWB emission from these models to distances of $O(10)$ Mpc [172]. This compares favorably with estimates of the local rate density of low-luminosity GRBs (up to $10^3\,\mathrm{yr}^{-1}\mathrm{Gpc}^{-3}$ [173]), so that the enhanced and advanced detectors may well detect gravitational waves associated with GRBs or constrain these progenitor models.

SGRs emit short-duration X-ray and gamma-ray bursts at irregular intervals, and (rarely) giant gamma-ray flares with luminosities as high as 10^{47}erg/s [174]. Crustal deformations and catastrophic cracking [175, 176] and excitation of the star's nonradial modes [177, 178, 179] might lead to the emission of gravitational waves [178, 179, 180, 181, 182]. The pulsating X-ray tail of SGR1806$-$20 also revealed the presence of Quasi-Periodic Oscillations (QPOs) (e.g. see [183, 184, 185]) that were plausibly attributed to seismic modes of the neutron star, thus bringing up the possibility of associated gravitational wave emission [156, 186]. Enhanced and Advanced LIGO and Virgo shall probe the energetics of these Galactic sources several orders of magnitude below the typical electromagnetic energy output of giant SGR flares, constraining the unexplored and interesting region [179, 180, 181, 182].

Burst searches analyze all GRBs (both long and short) and SGR flares for which we have data from two or more detectors. The search algorithms look for a GWB coincident in time and sky direction with the external trigger, and are designed to be sensitive to generic GWBs. The CBC Group performs a dedicated matched-filter search for inspiral signals from short GRBs. The burst analysis of these same GRBs is complementary, providing the ability to detect or constrain any GW emission that does not precisely match the inspiral waveform. For SGRs, we also perform two specialized searches. "Stacking" searches look for the collective signature of weak GWBs from repeated flares from a single SGR source. "QPO" searches look for GWBs at frequencies matching those seen in the electromagnetic emission.

To date, we have completed the analysis of all S5/VSR1 GRBs with coincident LIGO/Virgo data. A results paper is in the late stages of review, and should be published in mid-2009. The analysis of AstroWatch GRBs using G1, H2, and V1 data is underway and will be completed in fall 2009. Looking ahead, the goal for S6/VSR2 is to perform a fully autonomous search in near-real-time for all GRBs. The complete analysis (including a detection statement or upper limit) is expected within approximately 24 hrs of receipt of the GRB trigger. A paper covering the full S6/VSR2 GRB set should be written within 3 months after the end of the run. Current research effort is focussed on building on the infrastructure used during S5/VSR1 (receiving and parsing GRB alerts, launching data analysis jobs, *etc.*) and making it robust for S6/VSR2. The search for GRB-GWB correlations can be tuned for specific assumptions of the GRB models (e.g. fragmentation of collapsar disks for long GRBs [170]); these options may be explored for S6/VSR2.

Two SGR searches on S5/VSR1 data have been published: the analysis of the first year, and the 2006 SGR 1900+14 storm. The next priority is the analysis of flares from the recently discovered SGR 0501+4516, thought to be at a distance of only 1.5 kpc. This will be followed by analysis of the remaining SGR flares from the second year of S5/VSR2 and the QPO search. Triggered searches for GWBs associated with SGR flares will also continue throughout S6/VSR2.

Low threshold gamma-ray triggers

An alternative approach to searches based on individual GRBs is the statistical analysis of a large ensemble of them. This search targets GRBs that were close-by, but whose jet axis points away from the Earth so that the received gamma-ray emission is weak. Since GRBs are believed to be "beamed", with a beaming factor of up to several hundred, there should be many more off-axis GRBs than observed GRBs.

Since gravitational waves are predicted to be much less directional than gamma rays for these objects, there is a possibility of nearby off-axis GRBs producing gravitational waves that can be detectable collectively. After the high-confidence GRB events are analyzed the search to harvest low-confidence GRB events will proceed.

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

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

Neutrinos as Astrophysical Triggers

Gravitational waves and high-energy neutrinos can originate from a common, very energetic source that exhibits burst activity. For example, GRBs are thought to be high energy neutrino emitters as implied by the internal shock model [187, 188, 189, 190, 191, 192, 193, 194, 195, 196, 197]. In the prompt and afterglow phases of long GRBs, high-energy neutrinos are expected to be generated by accelerated protons in relativistic shocks [187, 198, 199]. Short GRBs [168, 167, 200] are also thought to be associated with a detectable high-energy neutrino flux. "Low-luminosity" GRBs are associated with a particularly energetic population of core-collapse supernovae [162, 201, 160, 202, 203, 204, 205]. Typically discovered at relatively small distances (e.g. SN 1998bw at about 40 Mpc), the expected rate of these events in the local volume can be an order of magnitude larger than that of conventional long GRBs [206, 205] and strong neutrino emission is also expected from these sources [197, 207, 208]. A detectable number of neutrino events are additionally expected from the so-called "failed" GRBs [209, 210, 211, 212]. They can be hidden from conventional observational methods, however, neutrinos and gravitational waves detected together may reveal their properties. SGRs may also be considered as common sources for high-energy neutrinos and gravitational waves [179, 213, 214].

The IceCube [215] and ANTARES [216] neutrino observatories have extragalactic reach in the high energy regime and they can determine both the arrival time and the direction of individual neutrino events. The LIGO-Virgo detector network also provides extragalactic reach and relevant pointing information. The past and projected overlaps in the observation schedule of the collaborations ensures concurrent gravitational wave and high-energy neutrino observation time.

The feasibility and expected performance of a hypothetical IceCube-LIGO-Virgo network was demonstrated [217, 218], indicating that the false alarm rate for the combined detector network is expected to be very low (e.g. $\sim$(hundreds of years)$^{-1}$). The IceCube and ANTARES sky coverages are also complementary. A gravitational-wave – high-energy neutrino search can significantly expand the scientific reach of all experiments. The background rejection, arising from the combination of gravitational-wave and neutrino events, results in an increased sensitivity and recovery of cosmic signatures that could not be detected by individual detectors.

Supernova triggers from detectors sensitive to low energy (10s of MeV) neutrinos can be used to initiate a search for associated gravitational waves. For example, a core-collapse supernova near the galactic center is expected to produce a large flux of $\sim$ 8000 detected neutrinos in the Super-Kamiokande detector [219] with a pointing accuracy of 4°. Unlike electromagnetic signatures of supernova, neutrino bursts, like gravitational waves, mark the moment of core collapse, and are expected to be coincident in time to < 1 s. In fact, neutrinos and gravitational waves provide the only direct probes of the dynamics of core-collapse.

The expected strong neutrino signal for a galactic core-collapse supernova would provide excellent timing and good pointing, thus allowing an improved sensitivity gravitational-wave burst search, similar to that employed for GRBs. For extragalactic supernovae, the neutrino signature would in general provide timing but not pointing. At the distance of Andromeda the expected flux of detected neutrinos in Super-Kamiokande would fall off to $\mathcal{O}(1)$. In this case, joint neutrino-GW time-coincident searches would substantially increase detection probability and decrease the false-alarm rate. In general, a joint neutrino detection would boost confidence in a low-SNR GW detection.

Optical Supernovae as Astrophysical Triggers

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

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

Pulsar Glitches as Astrophysical Triggers

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

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

A search for a GWB associated with the August 2006 Vela pulsar glitch has been performed, using a new Bayesian model selection algorithm [220], and a paper is in preparation. This method will be applied to pulsar glitches of interest during S6/VSR2, and possibly to S5/VSR1 glitches from the X-ray pulsar J0537−6910.

Population study of astrophysical sources using data from other detectors

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

The first version of a population study method for GRB triggers was investigated in [221] and an improved one [222] was demonstrated using data from the S2, S3, and S4 runs of LIGO [223]. The population study method is complementary to the search for GW signals in coincidence with individual triggers [224]. An important feature of the method is the accumulation of signal-to-noise ratio as the sample size of triggers increases and a corresponding monotonic improvement in the astrophysical constraints obtained [221, 224].

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

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

The Rossi X-ray Timing Explorer (RXTE) mission has been extended to the end of 2010 and observations will therefore overlap with the S6/VSR2 run. Multi-messenger observations of Sco X-1 allow us the opportunity to search for correlations between the X-ray and potential GW emission from Sco X-1. Sco X-1 was the first extra-solar X-ray source to be discovered and there are still many unanswered questions regarding this object: What is the spin frequency? What mechanism is at work during low and high frequency quasi-periodic oscillations (QPOs) [225]? What is the limiting mechanism on the increase in spin frequency due to the accretion torque [226]? Recent papers have reported that some of these issues are related to gravitational radiation [227, 228, 229, 230]. 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. The coincidence analysis can improve the detection confidence and give a chance to set upper limits on GWs correlated with astrophysical events encoded in RXTE data.

Radio triggered searches

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

Models for prompt radio emission from compact binaries generally require that one of the compact objects is a magnetar. In the simplest of these model classes [231, 232], the orbital motion of the binary generates time dependent magnetic fields and consequently induced electric and magnetic fields. These fields then lead to the emission of radiation, which is predicted to be in the radio band.

A second, larger class of models similarly require a high magnetic field from one member of the binary; in these models the field either confines or otherwise interacts with a plasma. For example, the unmagnetized companion object can develop surface charge that can then be ejected from the surface of the star and subsequently undergo acceleration as it follows the magnetic field lines of the magnetar. Alternatively, [233] gravitational waves emitted during the inspiral and merger of a binary neutron star system may excite magnetohydrodynamic waves in the plasma (see also [234] for details), which then interact with charged particles in the plasma, inducing coherent radio emission.

A progenitor confirmation is still awaited for short hard GRBs (SHB). Results have been published on radio afterglows for SHB but the data shows only weak signals hours or days after the burst. There are a series of proposed models for radio afterglows for SHB that may be observed seconds to minutes after the burst. One of the models put forward by Usov and Katz [235] predicts that immediately after merger the rotational energy of the binary system is transferred to a highly magnetized, highly relativistic particle wind that interacts with the ambient warm gas and as a result EM radiation is emitted. The main bulk of radiation is below 1 MHz but its tail can reach 1–30 MHz. As a main prediction, according to the model we should detect an incoherent pulse of 1–30 MHz, duration of 1–100 s and fluence of a few percent of the total gamma-ray flux. The expected delay for such a pulse would be of around 10^4 s for a source placed at 3.2 Gpc.

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

4.4.4 Seek confirmation of GW candidates with follow-up observations

The previous section described many scenarios in which astrophysical systems are expected to emit electromagnetic (EM) radiation along with GW bursts, and the substantial benefits of joint observations. However, telescopes cannot cover the whole sky at all times with good sensitivity. Current EM survey projects typically have lengthy cadence times, or are considerably less sensitive than their more directed counterparts [236]; 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 [237, 238]. Thus there may be short-lived “orphan afterglows” from nearby energetic events which are going undetected [168]. If such an event is detected by the network of GW interferometers, then the reconstructed sky position can in principle be used to point one or more telescopes and catch the afterglow before it fades away.

Other possible sources of joint EM/GW emission include decaying neutron star matter ejected during merger [239, 240] and supernovas [241]. More details on these sources and their expected observational signatures can be found in [242, 236, 243].

In the event of the detection of a truly astrophysical signal, such observations will be greatly beneficial for two reasons. First, they will help establish the event as astrophysical in nature, effectively increasing the reach of the interferometer in the case of a low-signal-to-noise-ratio event that might not otherwise stand out from the background. Additionally, having an associated EM signal increases the astrophysical information that can be mined from a GW event. Several recent papers have stressed the progress that can be achieved only with a combination of EM and GW data [244, 245, 246]. This second motivation will remain strong even in the future advanced detector era, when GW signals are expected to be observed regularly.

During the S6/VSR2 run, we will have an ambitious program to seek confirmation of GW candidates with electromagnetic follow-up observations. This project is sometimes called “LoocUp”, which originally stood for “Locating and Observing Optical Counterparts to Unmodeled Pulses” [242], but the scope has evolved over time to include EM bands other than the optical and to be a joint effort with the CBC Group, who wish to follow up promising candidates from (well-modeled) inspiral searches.

There are several challenges that must be addressed to make this possible. One difficulty is the relatively large uncertainty expected for interferometer signal localization. For the LIGO-Virgo network, uncertainties of order a few degrees are expected, which is large compared to the field-of-view of many astronomical instruments. For this reason, it is natural to plan to work with wide-field instruments, with fields of view of several square degrees or more. We have had detailed discussions with a number of groups of astronomers who operate robotic wide-field optical telescopes around the globe—many of which are already being used to look for afterglows to GRBs—and expect to soon formalize arrangements for them to collect images following up on LIGO-Virgo GW triggers. Smaller field-of-view instruments may also be useful for secondary observations, including spectra, or in conjunction with a catalog of nearby galaxies to limit the region to be imaged, as outlined in [242]. To this end, we have made arrangements with the *Swift* project to plan for a limited number of target-of-opportunity (ToO) observations using that satellite's X-ray and UV/optical telescopes, and we are planning to apply for ToO observing time on a few large-aperture ground-based telescopes. Another approach being investigated is to "point" into normal survey data from Fermi, *Swift*, and/or RXTE afterwards to search for low-significance counterparts in these high-energy EM data sets.

The development for this project has been rapidly progressing. An "online", low-latency, GW data analysis pipeline is a key component of any follow-up plans. The Burst Group has made great strides towards making this a reality, with many components demonstrated during the pre-S6/VSR2 engineering runs and current analysis pipeline latencies on the order of 10 minutes. This online pipeline will continue to develop through S6/VSR2, along with LUMIN, a collection of specialized tools for identifying relevant events and passing them to automated telescopes for directed observation. We are also beginning to explore the details of telescope image processing, transient identification and false-coincidence rate estimation with our astronomer collaborators.

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

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

4.4.5 Be open to alternative theories of gravity

An intriguing possibility is that gravity may be better described by a theory other than General Relativity. The direct detection of gravitational waves could help determine the correct theory of gravity [247, 248, 249, 250]. 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 [251]. LIGO provides a unique opportunity to study the detailed properties of gravitational radiation.

Alternative theories of gravity may result in a difference between the speed of light and the speed of propagation of gravitational waves [252]. This has an impact on coordinated electromagnetic and gravitational wave searches that needs to be accounted for. If the propagation speed is much less than the speed of light, then standard GW burst searches themselves could also lose efficiency since they assume a certain maximum time delay between detector sites.

Alternative metric theories of gravity also predict extra polarization states, in addition to the +- and ×-polarization modes of Einstein gravity. Indeed, every other known viable metric theory of gravity predicts more than two polarization states [252], and the most general gravitational wave can have up to six polarization states [247, 248]. 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 [248], provided there are at least as many detector data streams as polarization modes. This is the case for externally triggered searches, and for searches where the gravitational wave triggers have sufficiently large signal-to-noise that triangulation can be used to pinpoint the sky-location. For all-sky searches new techniques need to be developed to separate out the expanded set of polarization states.

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

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

4.4.6 Be prepared to detect a GW signal with confidence

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

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

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

Beyond what we have been able to automate, it is necessary to check on a case-by-case basis any detection candidate that appears at the end of our pipelines. The "detection checklist" outlines a followup procedure for such events. It is meant to be applied to an event which has any chance of being identified as a gravitational wave. Currently the detection checklist consists of an 80-point list of items to be checked by hand covering the following categories:

- zero-level sanity: reports in detector logs, check hardware injections
- data integrity: frame file checksum, undocumented injections, check against raw frames
- state of the instrument: obvious disturbances reflected in auxiliary channels, verify coupling for any proposed veto, check by hand against known disturbances: dust, cosmic rays, power fluctuations, acoustic, etc.

- event properties: construct detailed spectrogram, reconstructed waveform and direction, compare background from various methods, check signal consistency across interferometers
- astrophysical interpretation: check for external EM or neutrino events, catalog sources consistent with reconstructed direction, compare waveform against simulations

The checklist serves several functions:

- provides a central resource for basic information about the event
- serves as a careful secondary review of the specifics of the end-to-end analysis that relate to the event
- provides an opportunity to check for obvious reasons to dismiss a candidate event (e.g. clear environmental cause) or increase our interest in an event (e.g. optical counterpart)
- outlines event details for each individual detector which otherwise could be buried in multi-detector statistics
- provides a test-bed for new ideas which have not been able to make their way into the quantitative detection statistic

The checklist comes after the blind analysis is complete and does not make any quantitative or objective end statement about the baseline statistical significance of an event (which is predetermined). However items which prove to be particularly useful at separating signal from background will have a strong motivation to be worked into the blind detection statistic for future analyses. We continue to work with members of the CBC group and others in the LIGO Scientific Collaboration and Virgo Collaboration, to improve our ability to use the followup checklist to hear the “ring of truth” in a gravitational wave detection.

4.5 First science targets for S6/VSR2

The upcoming S6/VSR2 science run should provide the best data yet with which to search for GW bursts, and we are planning to carry out many different searches. Nevertheless, there are a few topics that the Burst Group has identified as our *first* science targets for S6/VSR2. These are analyses of broad interest for which we have well-developed methods ready to apply to the new data, specifically:

- Rapid search for GW burst signals associated with exceptional astrophysical events such as a nearby GRB, an SGR giant flare, or a nearby supernova.
- All-sky search for arbitrary GW bursts (up to $\sim$1 second in duration) with frequency content in the range 40–6000 Hz.
- Search for GW bursts assocated with GRBs detected during the run.
- Search for GW bursts assocated with SGR flares observed during the run.

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

In order to meet these targets, we will need to be actively engaged in analyzing the S6/VSR2 data *during* the run. The calibration of the LIGO and Virgo detectors should be known to $\sim$10% in amplitude and $\sim 5^\circ$ in phase at the time it is collected, and the uncertainties must be known. The data will be analyzed with mature, well-understood analysis pipelines that we have already developed and reviewed (or will finish reviewing early on during the run). A major effort will be needed to study data quality and develop veto conditions on a daily and weekly basis as the data is collected, tracking changes that will inevitably occur.

At this time, the software and procedures for these first science targets are fairly advanced, but more work is needed to integrate the components. We expect these analyses to be carried out by teams of people working together, and will need to ensure that they have enough manpower to succeed.

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

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. (We use the term "neutron star" broadly, keeping in mind that some such stars may contain quark matter or other exotica.) These stars are expected to emit gravitational radiation through a variety of mechanisms, including elastic deformations [257, 258, 260], magnetic deformations [259, 263], unstable r-mode oscillations [261, 257, 262], and free precession [267], all of which operate differently in accreting and non-accreting stars. We present a review of these emission mechanisms in [256]. Indirect upper limits on gravitational wave emission inferred from photon astronomy are more optimistic for non-accreting stars, but according to present theories accreting neutron stars are more likely to be emitting at or near the indirect limits.

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

The breadth of investigation is fundamental to our search method. Given the large uncertainties in neutron star demographics (only $\sim$2000 of 10^8-10^9 neutron stars in the galaxy have been detected), evolution, and structure, we cannot confidently predict which type of source will provide our first continuous-wave discovery. Prudence demands an eyes-wide-open approach and enough flexibility to exploit unexpected waveform models. That said, however, we do adhere to certain priorities in allocating resources (scientists and computers) to different searches. Specifically, we place the highest priority on targeted searches for known pulsars (especially those for which the spindown limit is achievable – see below) and on all-sky searches for unknown isolated pulsars.

The recent merging of LSC and Virgo CW efforts has revealed strong and well-developed programmes from both groups. An ongoing task is to combine these effectively, maximising both the return on time already invested in developing codes and the science delivered by new joint ventures. Our state right now is one of transition and this is clearly reflected in the methodologies detailed below. Over the coming year we will carefully review the scientific justification and available manpower for these investigations, balancing the importance of redundancy with our current resources, and it is likely that some of these programmes will converge to common code. For the moment however, we will present LSC and Virgo investigations as distinct efforts within a common astrophysical programme.

5.1 Non-accreting pulsars

We include in this source type all objects for which pulsations are observed in radio, X-ray, or other electromagnetic radiation, with the exception of accreting millisecond pulsars. Photon astronomy can tell us precisely the sky positions, frequencies, and frequency changes of these objects, meaning that our analyses need search only a small parameter space and are not computationally limited (see section 5.1.1 below). Photon astronomy also sets an upper limit on the gravitational wave strain we could see from a known pulsar, assuming that all of the observed spin-down is due to gravitational waves. In terms of the distance D, gravitational wave frequency $f_{\rm gw}$ and its time derivative $\dot{f}_{\rm gw}$, this indirect limit is [256]

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

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 just a portion of the LIGO S5 data has already beaten this indirect "spin-down limit" by nearly a factor of 4.1 for the Crab pulsar (59.45 Hz), and could just beat it for PSRs J0537-6910 (123.95 Hz) and J1952+3252 (55.69 Hz) with another run at comparable sensitivity [264]. Other stars for which the spindown limit may be reached in S6/VSR2 include J0737-3039A (88.11 Hz) and J1913+1011 (50.59 Hz). At lower frequencies, VSR2 data may permit beating the spindown limit for Vela (22.38 Hz).

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

While this form of gravitational wave emission is the simplest and most plausible, it is by no means the only possible wave generation mechanism. Other emission mechanisms include free precession, excited modes of oscillation of the fluid, and the spin-down of a multi-component star. The astrophysical returns from detection of such wave generation could be considerable, potentially giving information on asymmetries in the inertia tensor, the viscosity of the fluid, and the internal structure of the star. However, the observational challenge is correspondingly greater, as the gravitational wave emission no longer occurs at twice the spin frequency. This means that searches for such waves require careful thought in order to pick out a range of parameter space (i.e., the wave frequency and its time derivatives) that is both astrophysically reasonable and computationally feasible to search over. As described below (5.1.2), such a search has already been carried out for the Crab pulsar, concentrating on a small patch of parameter space centred on (twice) the observed spin frequency. Clearly, a more comprehensive search over an enlarged parameter space and for more pulsars is needed to fully exploit the science potential of targeted searches.

5.1.1 Time domain Bayesian method

This method has been developed to tackle targeted searches, that is searches for gravitational wave emission from pulsars of known position, rotation frequency, spin-down rate, and binary orbital parameters where necessary. This additional information greatly reduces the size of the parameter space over which we must search, and allows us to perform a fully coherent search for signals over all the available data. Timing accuracy is sufficient to maintain coherence both during and between science runs, and the relative phasing of the interferometers is also sufficiently well determined for us to be able to combine data from all runs and all detectors coherently, resulting in the lowest signal sensitivities achievable by LIGO. This method has been applied successfully to data from the first four LSC science runs [253, 254, 272] and is currently being used to produce results from the full data set from the fifth science run. A detailed discussion of the method can be found in [270], with the implementation of the inclusion of binary system parameters in [273].

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 a strong collaboration with the radio group at Jodrell Bank, UK, who have generated timing solutions for our pulsar candidates over the LIGO observing runs, and checked that no glitches have occurred in these pulsars. We also now have collaborations with groups observing with the Green Bank and Parkes telescopes, giving access to a larger number of pulsars, for example the many millisecond pulsars observed in the Terzan 5 and M28 globular clusters. These collaborations have provided new and useful data for the S5 searches and should continue to 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 J0537−6910, which is another pulsar for which we could soon beat the spin-down limit.

There are just over 200 pulsars within the sensitive band of the LIGO interferometers, that is with spin frequencies greater than 20 Hz. For all these we are able to perform the first stage of the search heterodyne process, but for S5 we only have 116 with radio/X-ray obsverations overlapping the times of the run. For all pulsars, the radio timing data give uncertainties in the pulsars' parameters. In the past we have used these uncertainties, without taking account of covariances between source parameters, as an estimate on whether a single template search is valid for the pulsar. This has meant discarding some pulsars from our analysis. For the majority of pulsars we now have covariances for the parameters and can use these to make a better estimate of whether a single template search is sufficient. This has lead to the inclusion of 12 pulsars in the current search that would have been vetoed previously. We are also no longer restricted to a single template as we have now extended our parameter estimation techniques to use Markov Chain Monte Carlo techniques to search over the parameter uncertainties, whilst also taking into account the covariances between the phase parameters. In the cases where no signal is seen this will marginalise over the uncertainties and fold them into our upper limit in a natural way. For the majority of pulsars a signal template (i.e. not added extra phase parameters into the MCMC search) would still be sufficient, but for at least one pulsar the additional search space would be required to properly recover the signal. In addition, we now account for glitches in pulsars by adapting the timing model to allow for step changes in rotational phase at these points. For the S5 search this has been applied for three pulsars that were seen to glitch during the run: the Crab pulsar, J0537−6910 and B1951+32. We note that Virgo's design sensitivity at lower frequencies offers the opportunity to search for sources inaccessible to LIGO (see section 5.1.3).

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 [272, 273]. 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 [271] (also see §5.1.2.) This search enabled us to beat the spin-down limit by a factor of 4.1 using uniform priors over the unknown parameters. For the first time astrophysically motivated priors on the inclination and polarisation angles have been used, allowing us to further beat the spin-down limit by an extra factor of 1.3. This result has allowed us to constrain the amount of the pulsar spin-down power budget released by gravitational radiation to be less than 4%. However, we have since discovered that our template for this search was incorrect and would have missed a signal had one been present. Reanalysis has been performed with the correct template and updated calibrations and has found that the upper limit on gravitational wave amplitude increases by a factor of about one and a half times.

A paper containing the results of the search for 116 known pulsars using all the LIGO S5 data (expanded from the 78 searched for in S3 and S4) is close to completion and will be submitted by summer 2009.

5.1.2 Wide Parameter Searches for Known Pulsars

We know of several physical mechanisms that could cause the frequency of a neutron star's emitted gravitational waves to differ slightly from the typically assumed $2f_{\rm rot}$, with $f_{\rm rot}$ being the rotation frequency. We also know of emission mechanisms which can cause gravitational wave emission at other harmonics, such as $(4/3)f_{\rm rot}$ for r-modes. In our search we consider the cases of free precession and a two-component model of the neutron star's crust and core, and work out how much difference might occur between the true gravitational frequency and twice the rotation frequency. We also consider the uncertainity in frequency associated with r-mode emission when searching around $(4/3)f_{\rm rot}$.

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

We have searched a smaller parameter space for just the Crab pulsar with nine months of data from the S5 run, up until a timing glitch in the pulsar [271]. By using known information on the orientation of the Crab pulsar, the search placed 95% confidence strain upper limit of 1.2×10^{-24}, beating the indirect spin-down limit of 1.4×10^{-24} across the entire parameter space searched. It is a less stringent upper limit than placed by the time domain search because of the use of 3×10^{7} more templates and the statistics of the noise. This search will continued using the larger amount of data available after the glitch, which means the search should see over a 20% improvement in sensitivity for the Crab pulsar.

A paper describing these wider searches in S5 data is expected to be submitted for publication in late 2009.

Enhanced LIGO is expected to a see factor of 2 improvement in sensitivity above 100 Hz during the S6/VSR2 science run. However, since the Crab pulsar rests on the wings of the 60 Hz line noise peak, the sensitivity of the Crab pulsar search will be most affected by any improvements to this particular noise source. A year long wide parameter search without a break due to pulsar glitches would improve sensitivity by 30% over the current upper limits, assuming current 60 Hz noise levels.

5.1.3 Virgo targeted searches

Two different procedures are being used for targeted searches. Both start from the short FFT database (SFDB)and have in common the initial time domain data cleaning, based on an auto-regressive method which allows to efficiently remove short time domain disturbances [278]. The first method is based on the computation of the analytical signal [274]. 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 and the correction for the Doppler effect, spin-down or any other frequency variation is applied. 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 more noisy periods. We are now working on the final stage which consists in applying a matched filter on the Fourier transform of the data to correct for the signal power spread due to the amplitude modulation produced by detector radiation pattern and the variation of the polarization angle due to the Earth rotation.

In the second method we shall search for the signal using coherent matched-filtering. 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 [297]. If the computed value of the $\mathcal{F}$-statistic is not significant we shall impose an upper limit on the gravitational wave signal. From current observations of Vela nebula the polarization and inclination angles can be estimated to a very good accuracy [298]. We shall use this knowledge to improve our search. This requires rederivation of the optimum statistic. In our analysis we shall take into account non-Gaussianity and non-stationarity of the noise.

We will apply these methods to target various interesting known pulsars. In particular they are being applied to the analysis of VSR1 data for the Vela pulsar. For this analysis we are using updated ephemeris provided by radio-astronomers (Aidan Hotan, Jim Palfreyman - Hobart Radiotelescope). With the current sensitivity (May, 2009) we expect to beat the spin-down limit for the Vela in 3 months. If the sensitivity of VSR2 run will touch the design one at the Vela frequency, 3 weeks would be enough. 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 design sensitivity. The search for emission at $f_{\rm rot}$ will be also performed.

The status of the search for Vela in VSR1 data, with preliminary results, will be described at the next Amaldi Conference. Final results of the analyses are expected to be published in a paper in the second half of 2009.

5.2 Non-pulsing non-accreting neutron stars

This type includes point sources, such as central compact objects in supernova remnants, as well as highly localized regions, such as the innermost parsec of the galactic center. Photon astronomy can provide sky positions for these objects, but since no pulses are observed, the external measurements cannot provide us with frequencies or spin-down parameters. Since we must search over many frequencies and spin-down parameters, sky-survey positional errors (such as from ROSAT) are too large: we require arcminute accuracy or better to keep down the computational cost of a long integration time and thus a deep search. Although no f and $\dot{f}$ are observed for these objects, we can still define an indirect limit we need to beat. If we assume an object has spun down significantly from its original frequency and that this spin-down has been dominated by gravitational wave emission, we can rewrite Eq. (1) as

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

in terms of the age a, which can be inferred in various ways.

Initial LIGO can beat this upper limit for several objects of this type, including the youngest – the object in supernova remnant Cas A ($\tau_{\rm sd} = 326\,{\rm yr}$, $h_{\rm IL} = 1.2 \times 10^{-24}$) – and the closest, Vela Junior ($D > 200\,{\rm pc}$, though the precise value is uncertain). Several more objects have indirect limits attainable with advanced LIGO, including the remnant of Supernova 1987A ($h_{\rm IL} = 3.2 \times 10^{-25}$). However this putative neutron star is only 20 years old and would require a search over a large parameter space including six or seven frequency derivatives. Searches over small sky areas (single "pixels") are computationally the same as searches for known point sources, and for several of these (such as the galactic center) even initial LIGO could beat indirect limits. We are presently searching Cassiopeia A (see section 5.2.1) and shortly will start searching the central parsec or so of the galactic center (see section 5.3.4) and at least one globular cluster (see section 5.3.5). 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 5.2.5).

The first search for a source with no timing (Cas A below) used the $\mathcal{F}$-statistic code with a single integration of $\mathcal{O}(10)$ d. Our estimate of computational cost and sensitivity [299] 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 will also investigate the potential of resampling methods to reduce not only the computational cost for a given search, but also the cost's scaling with integration time. This, combined with hierarchical methods, will allow us to search a significant fraction of the S5 data set (and future sets of comparable length) rather than $\mathcal{O}(10)$ d.

5.2.1 Cassiopeia A

This search is in its final stages, after substantial development work on template generation using optimal lattices and careful treatment of the boundaries of the parameter space. Preliminary results have been produced, and the search pipeline review is nearly complete. A paper draft is in circulation within the CW Group. The search pipeline developed for Cassiopeia A will also be applied to similar targets in S6 data that require only one sky position . The sensitivity, computational cost, and indirect limits to aim at are described in [299].

5.2.2 Other coherent directed searches

The coherent search for Cas A will be extended to top targets from each category of directed search (see Sec. 5.2.5), 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.

5.2.3 Supernova 1987A using the Cross-correlation technique

As described elsewhere, the semi-coherent excess power methods are more robust than the full 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 [305], can be found in [290]. 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 $T_{\text{obs}} \gg T_{\text{coh}-\text{min}} \gg T_{\text{sft}}$. This is however a computationally inefficient way of calculating the coherent statistic, for which it is better to use the existing $\mathcal{F}$-statistic, so we expect that the cross-correlation is useful only with $T_{\text{coh}-\text{min}}$ either comparable or lesser than T_{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 $T_{\text{coh}-\text{min}}$ will be tuned so that the search can be completed on a time scale of 1-2 weeks which leads to $T_{\text{coh}-\text{min}} \approx 1$ hr. In searching for such a young object, searching over frequency derivatives can be prohibitive because one would need to search over higher derivatives as well. It turns out that the search space can be narrowed by using a physical model for the frequency evolution: $\dot{n}u = Q_1\nu^5 + Q_2\nu^n$. The first term is the usual term due to gravitational wave emission while the second term represents all other effects (ideally, for electromagnetic braking, one would expect a braking index of $n = 3$ but this is not observed in practice). With this model, and using $T_{\text{coh}-\text{min}} \approx 1$ hr, it turns out that the computational cost becomes manageable.

5.2.4 Virgo semi-targeted search

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 our pipeline. 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) [276]. The removal (or the duplication) of the samples takes place typically each a few seconds. The list of samples to be removed or duplicated (named *mask*) is thus not huge and can be easily computed by simple geometrical consideration. As detailed in [276] the mask corresponding to a given direction is provided by the times when the antenna crosses one of the equi-phase parallel planes fixed in the space, perpendicular to the wave vector and each at a distance c/f_s from the next one. Each "crossing time" is computed by the scalar product of the antenna velocity and the wave direction (in practice by a few operations each second of data). The maximum phase error due to the non perfect synchronization is given by $2\pi f_0/f_s$ where f_0 is the signal expected frequency and f_s is the sampling one. As a reminder, a phase error around a few tenths of rad is small enough to guarantee that almost all the signal energy is recovered around the main frequency. It is thus important to resample the data working at the Virgo data acquisition frequency (20 kHz) in order to use the method effectively up to several hundred Hz. This frequency independence makes the method very appealing for sources where the direction is well fixed, but the emission frequency is uncertain (semi-targeted search). The pulsar spin-down is taken into account by properly shifting the equi-phase target plane during the acquisition time. As a consequence, a single mask requires specifying both the direction and the spin-down value of the source. The Einstein delay and the Shapiro effect can be also easily computed without any significant additional computational cost.

We are now developing an analysis pipeline to be applied to VSR1 Virgo 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 VSR1 period with a sampling time of a few seconds (enough to approximate Earth motion with enough accuracy).

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

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

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

Since the Doppler correction (computation of masks and their parallel application in decimation of the filtered data) is negligible, the optimization strategy for the semi-targeted search is straightforward. We need only choose the width of the pass-band filter ("slice"). Indeed this choice determines the downsampling factor, thus the length of the buffers governing the FFT computation time. Finally we must multiply the time required for the previous operation (filtering and parallel FFTs) times the number of slices required to cover all of the interesting detection band. The optimization of the pass-band filter width, obtained minimizing the total computation time, depends on the analysis to be performed.

Once this method is mature and has been validated, it is hoped that its computational speed will prove useful in S6/VSR2 searches leading to publication.

5.2.5 Other targets

We are collaborating with several astronomers on constructing lists of interesting targets for further directed searches, i.e., targets where LIGO and Virgo can beat the indirect limits on gravitational-wave emission. Apart from Cas A there are nearly ten central compact objects in supernova remnants with no observed pulsations and indirect limits on h_0 high enough to beat with S5 or S6/VSR2 coherent and semi-coherent searches. "Calvera" [300] is not associated with a remnant, but may be the closest observed neutron star. There are also several small, young supernova remnants (such as SN 1987A) and pulsar wind nebulae where the neutron star is not seen. Other small sky regions (further discussed below) also can be targets of this type of search. Examples include regions of massive star formation such as the galactic center and massive young clusters containing magnetars such as Westerlund 1. Globular clusters are not likely to contain young neutron stars, but some old neutron stars are known to possess planets and debris disks. Frequent perturbations in the dense environment of a cluster core could trigger bombardment episodes, and a star with an impact-related deformation counts as rejuvenated for purposes of a gravitational-wave search. Considering interaction timescales, most of the best targets are nearby, dense clusters such as NGC 6544. However 47 Tuc's interaction timescale is short enough to make it an attractive target even though it is further away; and furthermore the first GLAST/Fermi results show considerable high-energy diffuse emission which is likely related to neutron star activity in the relatively recent past.

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

5.3 Previously unknown non-accreting objects

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

LIGO or Virgo. It has been argued, based on the observed supernova rate and inferred population of neutron stars in the Galaxy, that the indirect limit on the strongest signal from this population is no more than

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

where τ is the mean time between supernovae in the Galaxy. The latest and most detailed derivation of this upper limit is given in [256]. Note, however, that a more recent simulation analysis finds significantly lower expectations that depend on the assumed source frequency and ellipticity [289].

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. Four methods are now used for carrying out "blind" searches: 1) a quick-look semi-coherent method known as PowerFlux using incoherent sums of strain spectral power from many 30-minute "Short Fourier Transforms" (SFTs), 2) a multi-interfereometer Hough search 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 Hough or StackSlide step (see below), and 4) a hierarchical method, developed in Virgo, based on the alternation of coherent and incoherent steps. Even with semi-coherent methods, computational savings due to efficient parameter space coverage and choice of parameter ranges are helpful as they increase the coherent integration time and thus the sensitivity of the search. Recent analysis [290] indicates that a cross-correlation method 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.

As indicated, an all-sky, full-band search is computationally formidable. If one focuses, however, upon small but intrinsically interesting patches of the sky, such as globular clusters or the galactic center, then the challenge is reduced, permitting a more favorable tradeoff in sensitivity, either by searching deeper with longer coherence times, or by exploiting the smaller number of search templates to lower SNR thresholds for detection candidate followup.

5.3.1 Semicoherent searches

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 [265], Hough [284, 255], and PowerFlux [285]. 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 [286] in Physical Review D.

In short, PowerFlux computes from many thousands of 30-minute SFTs an average strain power corrected for antenna pattern and with weighting based on detector noise and antenna pattern. Its signal estimator is direct excess strain power for an assumed source direction and polarization, using one circular polarization projection and four linear polarization projections. PowerFlux, like Hough or StackSlide, corrects explicitly for Doppler modulations of apparent source frequency due to the Earth's rotation and its orbital motion around the Solar System Barycenter. Source frequencies are searched with ∼0.56 mHz spacing and limits presented separately for 0.25 Hz bands.

A short publication based on an improved PowerFlux search over the first 8 months of S5 data was published in Physical Review Letters in early 2009 [287]. 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.

Over the past six months a series of improvements to computational efficiency were made to facilitate a PowerFlux run over the full S5 data while keeping memory requirements within 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. A production run over 50-400 Hz recently completed and will soon be followed by a run up to 800 Hz and beyond. Further refinements, including a coherent IFO-sum option for each SFT to gain further SNR [288], are planned for the PowerFlux follow-up analysis of high-SNR candidates. Recent multi-threading of the PowerFlux program and the availability of powerful, multi-core processors in the Atlas computing cluster offer the prospect of more efficient and rapid follow-up analysis. Full-S5 PowerFlux results are expected to be ready for publication in late 2009.

A fully automated coherent followup pipeline is also planned for PowerFlux. This pipeline will take the highest SNR outliers and utilize the ComputeFStatistic code to perform a coherent followup of each outlier. This followup will be performed in steps, with each step zooming in on the loudest event(s) from the previous step, and using around 3 to 4 times more data than the previous step. The followup step will be more sensitive than the PowerFlux, allowing it to better determine if the outlier was due to noise or an actual gravitational wave. It is expected that 1-2 thousand outliers could be examined using a few weeks of cluster time. This pipeline could also be applied to the outliers of other semi-coherent searches.

As in the PowerFlux method, the weighted Hough method (used already to analyze the S4 data [286]) takes into account the detector antenna pattern functions and the non-stationarities in the noise. This algorithm allows to combine data from different interferometers, doing 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 [291] to reduce the number of candidates and consequently increase the sensitivity of the search. A preliminary analysis of the first calendar of S5 data was carried out in 2008, and an analysis of the second year of S5 is underway, with coincidence of outliers between years 1 and 2 to be imposed in the follow-up. Power sums from H1, H2, and L1 are used in this search.

Full-S5 Hough results are expected to be ready for publication in mid to late 2009. A combined publication is planned in late 2009 that will include the all-sky results obtained by PowerFlux and multi-IFO Hough.

5.3.2 Cross-correlation searches

The cross-correlation method has been decsribed 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.

5.3.3 Hierarchical Searches and `Einstein@Home`

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 $\sim$ 160 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 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 [295]. A similar search was run on S5 data (S5R1), and the paper presenting these results is on the arXiv [296], and is submitted to PRD.

S5 R2/R3 postprocessing: The second S5 analysis was based on a greatly improved search pipeline, which eliminates the main problem limiting the sensitivity in previous `Einstein@Home` searches: this search was based on a Hierarchical search pipeline, consisting of individual $\mathcal{F}$-statistic searches on $N_{\rm stack} = 84$ data "stacks" spanning no more than $T_{\rm stack} = 25\,$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 improving sensitivity by a factor of ~ 6 with respect to the previous search method. After an initial shorter "test" run ("S5R2") with this pipeline, lasting for about 3 months, a further improved setup was launched as the 3rd `Einstein@Home` run on S5 data (codename "S5R3"), designed to run for about a year. This run has finished and the work on postprocessing of these results has begun. A paper describing this search of the first year of S5 data is expected to be submitted for publication by the end of 2009.

Current Hierarchical search, S5R5: (The S5R3 run analysed only data from roughly the first year of S5. A new `Einstein@Home` run (S5R5) was launched in Jan 2009, covering a frequency range up to 1kHz. This search is expected to finish in about Oct. 2009, after this a second installment of this search will cover the frequency range [1kHz, 1.2kHz], and is expected to run for another 6 months. These runs use a mostly identical pipeline to the previous (S5R3) run, but include 121 new segments of data from the second half of S5. The search setup also takes account of the speedup of the science application by nearly a factor of 2, and an increase of more than 30% in participating hosts. A paper describing this definitive E@H search of the full S5 data is expected to be submitted for publication sometime in 2010.

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

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

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

Relation to the "Grid":

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

5.3.4 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 present in the galactic center. The dispersion and scattering of the signal by interstellar material between potential sources and the Earth significantly reduces the depth of such observations. The current estimate of the total number of pulsars present in the galactic center (the inner $2\,\mathrm{deg} \times 0.8\,\mathrm{deg}$ of the galaxy) is 10^6 (Muno et al). Some of those objects could be promising sources of CW gravitational wave signals. Similar to searches in globular clusters, searching in the galactic center involves searching over a small sky-region but over large frequency and spin-down ranges. An important parameter is again the minimum spin- down age we wish to search for, which will determine the number of spin-down parameters required and also the maximum amplitude of such a signal.

We plan a search for traditional CW sources that last for the duration of the entire S5 run using the hierarchical search code. While `Einstein@Home` is carrying out an all sky-search using the hierarchical Hough code, this search is meant to be more sensitive over a more limited region of parameter space. Thus, it should search a larger range of spin-downs and also have a longer coherent integration time and reach younger spindown ages.

The implementation of the search focusing on the galactic center began by merging the Resampling code written by Patel into the hierarchical Hough code. The current implementation of the Hough search calculates the $\mathcal{F}$-statistic [297] directly in the frequency domain. This is, however, computationally inefficient for long observation times (i.e. ~ 1 yr). The computational cost, in this case, is proportional to T^2, where T is the time of analysis. A more efficient algorithm involves using the method of resampling, which calculates the $\mathcal{F}$-statistic in the time domain from the detector time series resampled to the barycenter reference frame. This reduces the computational cost of such a search to $T \log T$. Some progress has been made in this direction, and preliminary tests using resampling in lalapps_HierarchicalSearch have shown improvements in speed of the search by approximately 60 %.

Ongoing studies are investigating different search algorithms to determine which one best balances computational costs with search sensitivity. Possible algorithms include: a fully coherent method using the $\mathcal{F}$-statistic, the hierarchical Hough/StackSlide algorithm, an SFT-based semi-coherent method, and the cross-correlation algorithm.

Additional ongoing work is focused on determining the parameter space region and template grid layout. Important parameters of the search are the sky position, a single template pointed directly to the galactic center; the frequency band, covering the most sensitive region of the LIGO band, 50 Hz to 1 kHz; the data will be from S5, using only H1 and L1 detectors; and the search will cover a young spindown age of a few 100 years. 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 amount of data from S5 used and the computation time the search will run on the cluster has not yet been determined.

A paper describing an S5 search for large-spin-down sources near the galactic center is expected to be submitted for publication mid- to late-2010.

5.3.5 Searches for sources in globular clusters

Promising sources of gravitational waves are isolated spinning neutron stars lying within globular clusters. Due to the dense environments which are common in globular clusters, some debris may hit some of these spinning neutron stars, raising a significant quadrupole moment. These stars are then potentially detectable sources of gravitational wave emission. The effective spin-down age τ is the time elapsed after a collision event. Sources with a lower spin-down age would emit with a higher amplitude, and thus a search over the lowest possible spin-down age is worth conducting.

A search targeted at globular clusters will have the advantage of having to search over a smaller sky parameter space. For searches that are conducted over an integration time spanning a couple of weeks, the number of sky locations to search for a likely cluster (see section 5.2.5) is at most of order 10, and in the best cases just 1. Thus the limiting factor to this kind of a search is the number of spin-downs that can be afforded for a given computational time. The number of spin-downs required are a function of the spin-down age τ, that one is looking over and also the integration time. A search using 300 computers for about a month will be conducted on a suitable candidate, lying in the equatorial plane for greater discrimination from stationary instrumental lines, will be conducted in the near future. 47 Tuc is the target for the first round of this search. It is about 4 kpc distant from the Sun and has a core radius of about 0.40 arcmins. A lower bound on the effective spin-down age τ as described above is 300 years.

A paper describing an S5 search for sources in 47 Tuc is expected to be submitted for publication in the beginning of 2010. The search will reuse parts of the Cas A search code, which has already been reviewed. The search will also use the resampling code developed by Patel as described in the Galactic Center section above.

5.3.6 Semi-targeted searches for “transient CW signals”

This project aims at developing an efficient search method and pipeline to scan long stretches of data (of length $\sim$1 year) for *transient* quasi-monochromatic signals that only last for timscales 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 on the search method is underway, trying 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. If this study concludes that an astrophysically significant sensitivity can be achieved, such a search will be performed, first targeting known pulsars, then increasingly relaxing the parameter-space as far as possible given the required computing cost.

5.3.7 Followup-searches to confirm or veto CW signal candidates

Better theoretical understanding and the development of softwar 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.

5.3.8 Virgo blind searches

Two kind of pipelines have been developed and are being applied to VSR1 data: a hierarchical procedure using the Hough transform as the incoherent step and a coherent pipeline on short data segments based on the $\mathcal{F}$-statistic.

This hierarchical pipeline consists of a series of steps [277, 278, 279]. 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. It goes through the computation of the analytical signal (see 5.1.3) followed by a spectral filtering stage. This consists in a matched filter on the power spectrum for which of the order of 10 templates are enough. A Matlab version of the software for the spectral filtering procedure has been developed and tested [281]. The Hough transform step, which is the heaviest from a computational point of view, is performed on the *INFN Production Grid*, which allows us to transparently access a large set of geographically distributed resources. A supervisor program has been developed to automatically manage the analysis on the grid.

In the past, we have applied the procedure, apart from the final coherent step, to data of both commissioning runs (C6 and C7) and WSR runs (WSR8,9,10), also testing the detection efficiency, and accuracy in source parameter estimation, through software injections [280]. The analysis of a small portion of LIGO S5 data has also been done, already providing some interesting information on the data itself and hints on how to improve the analysis procedure.

We are now analyzing data of the first Virgo scientific run VSR1. The production of Hough candidates has been completed for the first half of the run, and coincidences will be done with the set of candidates coming from the analysis of the second half. The search is all-sky, over the frequency range $[50, 1200]$ Hz and with a minimal spin-down time of 20 000 yr. The adaptive version of the Hough transform stage, which takes into account the detector antenna pattern and noise non-stationarities, has been developed and implemented [279] and will be used in the next analyses. The Hough code at the moment performs a transformation between the time-frequency plane (peak maps) and the source ecliptic coordinates. A new Hough procedure, based on the transformation between the time-frequency plane and the source frequency/spin-down plane, has been developed. Using simulated data, it has been shown to have an efficiency 15% better respect to the standard Hough, which means, at fixed efficiency, a gain of a factor 5 in computing time [283]. It should also be more robust with respect to the presence of spectral disturbances and non-stationarities. The new Hough is now being studied on portion of VSR1 data before starting to use it as part of the hierarchical procedure (it should be used for the analysis of VSR2). The new frequency-Hough should also be very efficient in finding and removing lines of instrumental origin from the peak maps and, to this aim, a procedure, based on the inverse transform from the Hough plane to the peak map plane, is being studied.

The other analysis method for blind searches is based matched filtering using the $\mathcal{F}$-statistic [297]. It consists of a coherent stage over 2-day data segments, each covering 1 Hz bandwidth plus a follow up analysis of candidates in a 4-day data segment. We shall assume the minimal spin-down time of 1000 yr. We shall use a constrained optimal grid with minimal match MM = $\sqrt{3}/2$. The constraints are such that we need to resample the data only once for each sky position and such that we can use the FFT to calculate the F-statistic. With an available time of half a year for such computation we expect to analyze coherently 1000 data segments. Our guiding principle in the choice of a segment to analyze will be the quality of the data. The data for analysis will be narrow-banded and cleaned using the procedures described in 4.1.3. We shall set a low threshold of 30 for twice the value of the $\mathcal{F}$-statistic above which we shall register the parameters of the candidates. We shall verify the candidates by coincidence test among the candidates from different segments and by the F-test for the $\mathcal{F}$-statistic value gain when we increase the observation time twice. We shall collaborate on the coincidence analysis with the LSC E@HEinstein@Home team. We also plan to do search using the above method in the subspace of the parameter space defined by the spin down parameter equal to 0. For this subspace we plan to analyze 25000 2-day 1Hz band sequences.

A paper describing these all-sky VSR1 searches is expected to be submitted for publication in the second half of 2009.

5.4 Accreting neutron stars

For this class of source the gravitational radiation is thought to be powered by accretion onto the neutron star and not, as is the case for isolated neutron starts, by its own rotation. In this scenario, first proposed by [268], 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 [256]. The resulting indirect limit can be put in terms of X-ray flux F_x and spin frequency $f_{\rm rot}$ as

$$h_{\rm IL} = 5\times10^{-27}\left(\frac{300\,{\rm Hz}}{f_{\rm rot}}\right)^{1/2}\left(\frac{F_x}{10^{-8}\,{\rm erg\,cm^{-2}\,s^{-1}}}\right)^{1/2}. \tag{4}$$

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 completed a single coherent analysis on the accreting neutron star in the LMXB Sco X-1 using S2 data [256]. 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 [305].

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 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 [256], and a more sensitive search than the incoherent cross-correlation search done in [305].

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 S5 data, with an eye towards having the pipeline fully developed by the end of S6/VSR2.

A paper describing an S5 search for Sco X-1 using this method is expected to be submitted for publication in 2010.

5.4.2 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 maximise 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 semimajor 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 [293, 294]. This is equivalent to convolving the detection statistic frequency series with a "comb" of unit amplitude spikes separated in frequency by the inverse of the orbital period. The incoherent summing makes this a non-optimal strategy, but one that can have greater sensitivity to GW signals than a matched–filter approach because its observation length is not computationally limited. When using this approach, the parameter space resolution (and hence the number of search templates) is significantly reduced. It should also be noted that the sensitivity of this search to GWs scales with $T^{1/2}$, as with a coherent search (and unlike other incoherent searches), however, the sensitivity also scales as $M^{-1/4}$ (M is the number of sidebands) and hence high frequency, large orbit sources will be harder to detect with this method.

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

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

As written above, the method assumes constant frequency over the observation. But it can be extended to the case of changing frequency, e.g. due to fluctuating accretion rate, with semi-coherent methods. A natural choice to investigate in this context is the stack-slide method, which could use coherent integration lengths of order two weeks [265].

A paper describing an S6/VSR2 search for known binaries using the simplest version of the method is expected to be submitted for publication sometime in 2010.

5.4.3 TwoSpect search for unknown binary systems

The TwoSpect search is a hierarchical method under development for detecting unknown continuous wave sources from binary systems. The goal of the TwoSpect search is to probe regions of the large parameter space of pulsars in binary systems without exhausting the existing computational resources available. We plan to complete the search pipeline and begin production running within the next 8 months. 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 are also exploring the use of 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. Developing the pipeline will require an assessment of the tradeoffs in sensitivity vs speed in this multi-stage hierarchical method.

A paper describing an S5 (perhaps early-S6/VSR2) search for unknown binaries using this method is expected to be submitted for publication in 2010.

5.4.4 Quadratic search for unknown binary systems

As discussed above, searches for unknown binaries present formidable computing challenges. The orbital movement of the neutron star around the center of gravity of the binary system may induce large and rapidly changing frequency modulations of the gravitational wave. The frequency $f_{\rm ssb}$ detected in the solar system barycenter may be modeled as

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

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

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

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

In order to accommodate such large frequency shifts, a new search algorithm is developed. An extension of the coherent search method with extra parameters to describe the orbital motion of the neutron star is not computationally feasible (for coherence times in the order of 1 h, the extra amount of parameters needed to cover all likely Keplerian orbits exceed a factor of 10^9). A hierarchical search method like the stack-slide or Hough transform methods as discussed in Ref. [265] 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 $10^{-6}\mathrm{Hz/s}^2$.

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 draft of this latter document will be made available for the collaboration before July first. Analysis of S6 data will not commence before the code review is completed.

5.4.5 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 [301] and the Gregory-Loredo method [302]. We have begun to explore the use of the Gregory-Loredo method, which has been used previously in radio, X-ray and gamma ray astronomy to search for non-harmonic periodicity in time-series. It is designed to be efficient in detecting pulsating signals in sparse-sampled data. We will study the tradeoffs in detection efficiency *vs.* computational cost for non-sinusoidal pulses when applied to the high-duty-factor LIGO and Virgo data. Initial exploratory studies are being carried out in the Matlab environment. If the method proves to be promising, it will likely be implemented in the offline DMT environment, to increase computational efficiency.

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

5.5 Needs and plans for S6/VSR2

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

A potential exception is the possibility of searching for “transient continuous wave” sources, e.g., newborn neutron stars that are rapidly spinning down, based on external triggers from photon astronomy, using high-spin-down algorithms, such as described in section 5.3.4 for a galactic center search. An ideal example would be a supernova in our own galaxy, but there may be violent extragalactic events that warrant a search too, such as a nearby long gamma ray burst. Such triggers could warrant retrospective searches in older, less sensitive data. The ongoing development of high-spin-down algorithms should permit relatively fast followup of a detected nearby neutron star birth, if such a happy event occurs.

In addition, we will want to carry out “fast-track” searches of the early S6/VSR2 data, as was done for the early S5 data cf. targeted Crab and all-sky PowerFlux searches).

In looking ahead to S6/VSR2 data taking itself, the Pulsar Group has a strong interest in monitoring (and mitigating) instrumental spectral lines (see the detector characterization section) with low latency. We will want to continue the online daily line-finding machinery that ran successfully during S5. In addition, the offline auxiliary-channel coherence studies used in following up on S4 and S5 pulsar candidate outliers will be important to repeat for S6/VSR2 analyses. Special attention will be given to narrow spectral bands around a half dozen known pulsars for which the spindown limit may be reached in S6/VSR2.

In addition, we plan to exploit new "F-Scan" infrastructure developed during S5 for generating high-resolution spectrograms to detect wandering lines visually. We will complete the infrastructure needed to allow for daily inspection by scientists on shift in the observatory control rooms (with periodic review by pulsar analysts). LIGO Lab support will be needed for automated generation and storage of the SFTs used for F-Scans. We estimate approximately 20 Tbytes will be needed per year in S6/VSR2 to store the LIGO SFTs and F-Scan figures of merit.

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) in nature.

One of the searches performed by the Stochastic Background Working Group targets an isotropic gravitational-wave background. The isotropic background is predicted by different models, and it can be completely described in terms of dimensionless $\Omega_{\rm GW}(f)$, the gravitational-wave energy density per unit logarithmic frequency (in units of the closure density of the Universe). Different models predict different spectral shapes in the LIGO frequency band (roughly $50 - 150$ Hz), although they typically follow a power-law form. Hence, the group performs the stochastic background search for different power-law forms of $\Omega_{\rm GW}(f)$. The increasing sensitivity of LIGO interferometers has allowed the group to start exploring the implications of the stochastic background searches for various models. In particular, the most recent result of the isotropic background search, based on the LIGO S4 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 isotropic searches at the free-spectral range frequency (37.5 kHz), at which the strain sensitivity of the 4-km interferometers is similar to that at 100 Hz.

Another search performed by the group is non-isotropic in nature, and it is designed to be sensitive to possible localized foreground (astrophysical) sources of stochastic gravitational waves. Analogous to the CMB, such sources would be localized with a distribution that follows the local matter distribution in our galactic neighborhood. Potential sources that could fall into this category are the low-mass X-ray binaries, rotating neutron stars etc. In addition to this search, which focusses on localized sources, the group is also investigating ways of searching for other patterns of stochastic gravitational-wave background across the sky, and potentially for correlations between different directions on the sky.

6.2 Stochastic Search Method

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

6.2.1 All-Sky Search

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

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

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

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

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

The raw correlation depends on the (one-sided) power spectral density $S_{\mathrm{gw}}(f)$ 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)[306], which can be written as

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

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

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

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

We deploy a cross-correlation method to search for the stochastic GW background, following [307]. In particular, we define the following cross-correlation estimator:

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

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

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

where $P_i(f)$ are the one-sided power spectral densities (PSDs) of the two interferometers, and T is the measurement time. Optimization of the signal-to-noise ratio leads to the following form of the optimal filter [307]:

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

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

In order to handle gaps in the data, data non-stationarity, and for purposes of computational feasibility, the data for an interferometer pair are divided into many intervals of equal duration (typically 1-3 minutes), and Y_I and σ_{Y_I} are calculated for each interval I. The loss in duty-cycle due to the finite interval size is of order 1 minute for each analyzable data segment (which is typically several hours). The data in each interval are decimated from 16384 Hz to 1024 Hz and high-passed filtered with a 40 Hz cut-off. They are also Hann-windowed to avoid spectral leakage from strong lines present in the data. Since Hann-windowing effectively reduces the interval length by 50%, the data intervals are overlapped by 50% to recover the original signal-to-noise ratio. The effects of windowing are taken into account as discussed in [309].

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 [310], one can write:

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

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

6.2.3 Mapping

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

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

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

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

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

one can show that its expectation value is given by

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

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

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

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

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

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

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

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

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

where

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

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

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

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

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

6.2.4 Multi-baseline: LSC/VIRGO joint search

As shown in [307], 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 [308] the inclusion of the LIGO-Virgo pairs can enhance the sensitivity of the global GW detector network to an isotropic background of gravitational waves, particularly at frequencies above 200 Hz. Furthermore, the addition of a third instrument with comparable live time and sensitivity improves both the resolution and sensitivity of the mapping algorithm, effectively simplifying the regularization problem mentioned in section 6.2.3.

6.2.5 H1H2 All-Sky Search

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

One approach relies on the coherence, defined as

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

where P_{XY} is the cross-power spectrum between channels X and Y, and P_{XX} and P_{YY} are the two power spectra. As discussed in [311], it is possible to estimate the instrumental correlations between interferometers 1 and 2 by

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

where Z_i are the numerous environmental channels, including microphones, seismometers, accelerometers, power-line monitors etc. As discussed in [311], this method can be used to identify frequency bands in which the instrumental correlations between two interferometers are large. These bands could then be removed from the all-sky stochastic search. Moreover, the method can be used to estimate the residual contamination in the ”good” frequency bands.

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

6.3 Results and Plans

6.3.1 Status of S5 Searches

Isotropic Search The stochastic group is repeating the isotropic search with LHO-LLO interferometer pairs using the S5 data. Preliminary results were obtained using the first year of S5 data and the 4-km interferometer pair H1-L1: $\Omega_0 < 9.0 \times 10^{-6}$ for a frequency independent spectrum ($\alpha = 0$) in the band 41-178 Hz. This result is 7 times more sensitive than the published S4 result [303], 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 announced at the AAS meeting in January 2008. The group has finalized the isotropic search using LHO-LLO pairs in the entire S5 run, and has evaluated the implications of the new result for the models of early-universe cosmology, for cosmic (super)string models and for pre-big-bang models. A paper describing these results and implications has been written and is submitted for publication in the Nature journal.

Radiometer Search The stochastic group is repeating the radiometer search using the S5 data of the LHO-LLO 4-km interferometer pair (and using the new pipeline). This analysis is expected to produce $\sim 10\times$ more sensitive maps of the GW sky than those produced using the S4 data [305], and it will apply an algorithm for deconvolution of the antenna pattern from the maps. This search will also produce a second isotropic measurement, with similar sensitivity improvement over the S4 all-sky result.

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

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

The stochastic group is developing two methods to estimate and suppress the instrumental correlations, as discussed above in more detail. The group has applied these methods to the first 5 months of the S5 run, and the preliminary results indicate that the PEM-coherence and the time-shift approaches identify well the grossly contaminated frequency bands, which can then be removed from the analysis. Moreover, the PEM-coherence approach can be used to estimate the residual contamination in the ”good” frequency bands. More effort is needed to understand the systematic errors of the applied techniques. The group plans to apply these methods using the rest of the S5 H1H2 data and produce an isotropic measurement of the stochastic GW background.

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

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

Finally, the group is considering the possibility of performing directional searches with the LIGO-VIRGO pairs. Although currently less sensitive than LIGO-LIGO pairs, in the long run such directional searches could lead to a better angular resolution of the estimates of the stochastic gravitational-wave sky.

Non-Gaussian Search The group is exploring the possibility of searching for non-Gaussian stochastic gravitational-wave background, also known as the ”popcorn noise”. Much of the formalism for this search has been developed, although future modifications are likely. Preliminary tests of the formalism have been successfully performed, and runs on parts of the S5 data are being made. The goal of the group is to perform the non-Gaussian search using all of S5 data, thereby improving on the sensitivity to non-Gaussian stochastic signals as compared to the standard isotropic search.

Pipeline Upgrade The stochastic group has completed an upgrade of the analysis pipeline. In particular, since several searches rely on similar quantities (such as strain cross and power spectral densities of different interferometers), the group has decided to produce Stochastic Intermediate Data (SID), stored in the frame format, and containing the commonly used quantities calculated for segments throughout the S5 run. In addition to simplifying the standard stochastic searches, the SID frames will also find use in detector characterization studies and can also be used for searches for GW transients on minute or hour time-scales.

While different segment durations are of interest, the first set of SID frames was produced for 52-sec long segments. This segment duration allows relatively simple averaging of the intermediate data, producing 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.

Searches for Long-Lasting Transients The stochastic group is initiating an effort to search for transient GW signals on the time scales of minutes or hours. For this purpose, the SID frames will be converted into time-frequency maps of the cross-correlation between two interferometers with time resolution of order 1 minute and with the frequency resolution of order 0.1-0.25 Hz. The time-frequency map will then be parsed in search for different types of GW signals. Several directions will be pursued, including searches for relatively short broad-band signals and for relatively long narrow-band signals. The group will draw from experiences of other groups (such as the burst group) performing searches for short duration (one second scale) transients.

6.3.2 Plans for S6 Searches

Online Searches The stochastic group performed an online H1-L1 search during the S5 run. Although this search was not used to produce a measurement of the stochastic gravitational-wave background, it was very useful for tracking the evolution of the sensitivity of the experiment to the stochastic gravitational-wave background. Hence, the group plans to repeat this search during S6. Moreover, the group plans to add the coherence calculation to the S6 online search, which will allow early identification of any features correlated between two detectors, while not affecting the blindness of S6 stochastic searches (to be performed after the run).

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

The strain sensitivity of the 4-km interferometers (H1 and L1) during S6 are expected to be improved by a factor of ~ 2 over most of the relevant frequency band, as compared to S5. In addition, the sensitivity of the VIRGO interferometer during S6 is expected to be similar to those of LIGO interferometers at low frequencies (around 100 Hz). Hence, we expect the S6 stochastic searches to produce up to a factor of 10 more sensitive measurements of the stochastic gravitational-wave background. Moreover, due to the longer baseline between the LIGO and VIRGO sites, adding VIRGO interferometer is expected to improve the angular resolution of the directional search.

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

Similarly, the SID would simplify directional searches to be performed for very narrow frequency bands. An example of such a search was performed as a part of the S4 radiometer analysis, targeting Sco-X1. However, repeating the analysis for every direction in the sky and for every frequency bin, becomes expensive in terms of both memory and processing time. The SID would remove the steps of loading and manipulating the time-series data, thereby reducing the processing time.

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

7 Computing and Software

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

The LSC has converged on commodity computer clusters as the solution that meets its computational needs most cost effectively. LIGO has super-computer class requirements and data that can be handled efficiently in the simple parallel environment of clusters. In recent years the LSC has migrated to the grid concept of geographically distributed computing with clusters located at several sites. This approach has the advantage that it puts resources close to the university researchers who are analyzing the data. Grid middleware allows for relatively easy access to data and computing power. If local resources are inadequate or a poor match, a researcher can access additional resources on the grid.

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

7.1 Current status

The LIGO Data Grid (LDG) is the combination of computational resources, data storage, 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, 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, and so we are committed to a modular approach which allows us to remain light on our feet and implement solutions to 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, such a rapidly developing operating system. These are the traditional system-administration roles independent of grid activities.

---

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

2. **Grid Middleware Administration** Each local cluster must maintain a rapidly evolving set of middleware in as stable a fashion as possible. The primary means to do this is the LDG Server software, discussed in Section **??**. 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, Section **??**) 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. **Certificate Authority and User Accounts** LIGO uses X.509 certificates for authentication and authorization of users. LIGO uses the DOEGrids CA to provide certificates for authentication. RA agents are required to verify certificate requests and then approve them. User accounts are requested by a centralized web site; approvals are sent to each site where local admins add the accounts.

5. **LIGO Data Grid Client/Server Bundles** LSC staff leveraged experience with the VDT and built upon the 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. This bundle is maintained and released on a cycle similar to the VDT.

6. **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 effecient 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 experts, 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.

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

7.3 Data Analysis Software Development Activities

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

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

2. **GLUE** The Grid LSC User Environment (GLUE) provides workflow creation tools and metadata services, written in Python, which allow LSC scientists to efficiently use grid computing resources within and external to the LIGO Data Grid. GLUE leverages grid middleware to provide secure authentication mechanisms (Globus), job submission and control on clusters (Condor) and workflow planning for grids (Pegasus). 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 across the OSG. The pipeline module of GLUE provides this facility and is used in all large-scale inspiral and burst searches, and is being adopted by the pulsar and stochastic groups.

3. **LSC Algorithm Library**

   The LSC Algorithm Library (LAL) is a library of C language routines that form the engine of the computationally-intensive data analysis programs. LAL 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. LAL 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 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. As a result, there is considerable duplication within the current repository. The next step is to develop a set of atomic functions, together with associated documentation and automated test suites, that address the most common functions and evolve the existing analysis tools to use these new functions. Candidates for this factorization include a single MATLAB Application Programming Interfaces (API) for I/O to the several gravitational wave data file storage formats; data calibration functions; and higher level functions that are identified as commonly used and useful. By developing and migrating the existing code base to use these APIs and atomic functions the collaboration will significantly increase the verifiability and maintainability of this analysis software, while simultaneously reducing the barrier to the development of analysis software by individual researchers, educators and students.

5. **Support and Release** Release 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 OS's decided by the Comp Comm and DASC. Explore possible build and test harnesses for this software bundle with a view to activating it by summer 2008. Documentation relating to software installation will be updated and enhanced where appropriate.

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.

1. **Problem Tracking and Security** 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.

2. **Authentication and Authorization** The LSC relies on the Grid Security Infrastructure (GSI) from the Globus toolkit as part of its public key infrastructure. There is currently no centralized, fine grained authorization control within GSI; system administrators at each site map use gridmap files to map credentials to standard unix accounts. This approach will 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.

   A sub-committee of the CompComm has been exploring issues of authentication and authorization on LDG and other Collaboration resources. This sub-committee has provided a report to the CompComm outlining a strategic approach to authentication and authorization appropriate to meet the

Collaboration's varied needs. The Authentication and Authorization working group is tasked to develop and deploy a solution that addresses the strategic approach developed by the AuthComm and approved by the CompComm. This is a multi-year effort and will take a staged approach to deployment which allows all collaborators to continue to effectively use the LDG resources for the purposes of gravitational-wave astronomy throughout the process.

There is ongoing planning and design of Membership Services - the backend database that drives the Directory Services project and will drive the LAMS and Collaborator Services project - to continue to insure project extensibility with all other administrative interfaces that the Collaboration devises. Refine the requirements for author list generation and define requirements for demographics capture and support of elections and votes through Membership Services. Extend planning and design of the Directory Services interface and Membership Services database backend in order to better improve the current Attachment Z infrastructure, provide increased management capability for Collaboration Leadership (such as the ability for the Spokesperson to create new committees, etc.), as well as add the capability for Collaboration members to join opt-in mailing lists.

3. **Monitoring Services** Another type of metadata inherent to grid computing models describes the status of clusters and their processes and services, the status of jobs on the cluster, and the status of connectivity between clusters. In order to maximize the harvest of cluster cycle time, users and job submitting agents need to be aware of these metadata. 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 one or more available tools and collating the information collected by these to provide a consolidated Grid friendly interface is an essential step to improve efficiency.

4. **Enhancing LDAS for the LDG** The Globus Toolkit provides Application Programming Interfaces (API) to several popular computer languages such as C and Java. But Tcl/Tk is not one of the available API interfaces. LDAS software technology is based on the principle of extending the Tcl/Tk scripting language with highly efficient C++ Tcl packages (software libraries). Since the LDAS control level is in Tcl, LDAS requires an extension to the Tcl language to bring Grid technologies into its infrastructure. This is achieved using the TclGlobus package being developed at Caltech. The TclGlobus project is using SWIG to expose Globus Toolkit functions to the Tcl/Tk scripting language. The continued use of the Tcl/Tk scripting language within LIGO along with the evolutionary migration towards Grid based computational technologies by LIGO make TclGlobus a critical component for continued support and maintenance. LIGO is in the midst of enhancing LDAS to incorporate several of the Grid components found in the Globus Toolkit.

   In addition to the LDAS system, many of LIGO's client side tools, collectively called LIGOtools, are based on the Tcl/Tk scripting language. These client side tools are also slated to be upgraded with Grid technologies using the TclGlobus package, providing a common Grid enabled interface for users of LIGOtools.

5. **LIGO Data Replicator** The LIGO Data Replicator (LDR) replicates in bulk interferometer data files to LSC computing sites. LDR also provides a metadata catalog for LSC files that in conjunction with other tools allows LSC scientists and their codes to discover data and other files within the LIGO Data Grid. 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 S5 science run the LDR metadata catalog is expected to contain metadata information on more then 25 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. With at least seven LDR installations in the LIGO Data Grid the RLS network is expected to serve between 100 and 350 million mappings, making the LDR deployment in the LIGO Data Grid the single largest deployment of the Globus RLS. Intense investigations and testing are needed to insure that this solution continues to scale through the operation of initial and enhanced LIGO. Moreover the LDR metadata catalogs will need to be integrated into the LIGO Data Grid metadata service to provide a unified look-and-feel to the user community.

6. **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 Globus (which provides authentication and data location services) and a relational database (MySQL). This framework is designed to be extensible; the first application using it is the interferometer data quality service described. Develop for low latency discovery and archival of Science and DQ segments for S6 online and offline analysis. Complete the implementation of prototype of the S6 infrastructure at Syracuse. Develop robust monitoring and fail-over solutions for the segment database. Integrate data quality information generated by DMT monitors into the database. Deploy the new infrastructure at LIGO-Caltech, LLO and LHO. Collaborate with L. Bossi (Virgo) to exchange data quality information with Virgo. Provide support for the new service during the S6 run. Extend this metadata framework further enable scientists and applications to efficiently submit and retrieve instrument data quality information from a metadata database. Integrate the service with the Monitoring Project to provide robust monitoring solutions;

7. **Event Database and Archival Project** As the LSC moves into the era of gravitational-wave astronomy, there is a need for easy access to archived results from various searches. Draw up a preliminary design document and implement a prototype system. The system will be designed to use mult-site archival to insure data integrity and easy failover. The data will be generally read-only. An authenticated server will provide ability of designated individuals to move the data into the archive. As part of this project, database solutions which allow event storage and annotation will be explored; integration with the Virtual Observatory projects will be considered a high priority for this effort.

8. **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 focues 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.

A second and complimentary approach being investigated for planning and scheduling LIGO workflows across multiple sites leverages the Pegasus workflow mapping engine. LIGO workflows described using abstract directed acyclic graphs (DAX) are planned by Pegasus across both LDG and non-LDG sites and during the planning phase the abstract DAX representation of the workflow is converted to a concrete graph (DAG) with instructions for running parts of the workflow across specific (multiple) sites. The actual management of the workflow is handled by Condor DAGMan. Pegasus can 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 sites most efficiently.

7.5 Long-term goals

In the long term, the LIGO Data Grid must evolve into a robust system for processing gravitational-wave data in real-time, delivering easy access to compute resources and catalogs of gravitational-wave observations. Here are some of the points that must be addressed to reach this goal:

- **Hardware** The current LDG infrastructure is the best value for money solution available at the present time. Hardware is evolving rapidly and it is important that plans for the LIGO Data Grid in the era of advanced detectors and gravitational-wave astronomy should take account of this. During the initial LIGO era, institutions (other than Caltech and MIT) from the LIGO Scientific Collaboration provided approximately half of the computing resources needed for the analysis of LIGO data. The computing hardware plan for the next few years and continuing into the Advanced LIGO era remains the same, i.e. half of the computing hardware will be provided by LIGO Lab and the other half will be provided by other institutions in the LIGO Scientific Collaboration.
- **Software** The evolution of software technology continues to play an important role in developing tools for scientific data analysis. The open-source software development effort has been of great benefit to the LSC. The LSC will continue to assess its policies and procedures for software development in a manner that allows rapid response to new developments without compromising the quality of the scientific products delivered by the Collaboration.
- **Services** The nature of the services which the collaboration will deliver to members and to the broader community will continue to evolve. In the future, reduced data products will have more value to some scientists than the raw data. Nevertheless, the spectrum of scientific investigations will require a similar spectrum of services to deliver data and compute cycles, and to allow easy sharing of information with the whole scientific community.
- **International partners** The ultimate vision of gravitationalwave astronomy involves a world network of gravitational-wave detectors. The software and computing models of individual collaborations are likely to be quite different. The development of protocols and tools for collaboration will be critical to the rapid success of these efforts in gravitational-wave astronomy.

It is also important to continuously evaluate emerging technologies, having the potential to enable increased efficiency in data analysis.

7.6 Virgo computing and software

This section does not aim at providing a full overview of the Virgo computing and software, but only to provide a few entry points for further information. While in fact Virgo and the LSC keep independent computing resources, and their software developments are partially independent, still in the long term it can be expected an higher degree of interoperability and integration, hence it is important to start providing some more information in this document.

Concerning the computing resources, the Virgo Collaboration is using the resources of 3 computing centers: EGO/Virgo site mainly used to run online GW searches and detector characterization software that requires an immediate access to the raw data in order to provide a fast feedback to the commissioning activities. The offline GW searches are mainly run at the IN2P3 computing center (CCIN2P3 in Lyon, France) and the INFN computing center (CNAF, in Bologna, Italy) where large computing and storage resources are available. The CCIN2P3 and CNAF computing centers are the main Virgo data repository. It is worth underlining that both at CNAF and at CCIN2P3 the Virgo Collaboration uses resources which are shared with the High Energy Physics community: this means on one hand that the actual resources are negotiated year by year, on the other hand it means that Virgo leverages on the experience in user support and middleware development developed by the more mature HEP community.

7.6.1 Computing at the EGO/Virgo site

At EGO, the Virgo collaboration runs not only the Data AcQuisition, but also a relevant computing center. Several tasks are run here:

DAQ Virgo acquires about 6 MByte/s of raw data, which are stored on a hierarchy of circular buffers, ultimately allowing to keep about 3 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, saturating about half of the 100 MBit/s link to the GARR network.

Online/offline analysis A cluster of about 100 nodes of the latest generation, 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.

7.6.2 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 2009 5.4 PByte of data are stored in HPSS. Virgo data amounts to 176 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, pulsar and CBC searches. Data quality investigations and the h(t) reprocessing are also performed at CCIN2P3. In 2008, Virgo used 7% of the computing resources available at CCIN2P3.

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

7.6.3 Computing at CNAF (Bologna, Italy)

The CNAF http://www.cnaf.infn.it/joomla/index.php?lang=en 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.

At CNAF recent Virgo data are stored on spinning media, on large, data contiguous GPFS file systems. Less recent data are stored on CASTOR, a mass-storage system. 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, most notably the CBC searches.

7.6.4 Virgo software

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.

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.

Data transfer and GRID 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. 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.

The Virgo pulsar group is routinely running jobs over the GRID, inside the VIRGO Virtual Organization, both at the two TIER1 computing centers of CNAF and CCIN2P3, and at several TIER2 sites, fully belonging to Virgo or not.

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` and `Merlino` 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`

References

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

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

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

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

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

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

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

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

[9] R. Adhikari *et al.*, "Calibration of the LIGO Detectors for the First LIGO Scientific Run", LIGO-T030097 (2003);
G. Gonzalez *et al.*, "Calibration of the LIGO Detectors for S2", LIGO-T040060 (2004);
G. Gonzalez *et al.*, "Calibration of the LIGO Detectors for S3", LIGO-T050059 (2005);
A. Dietz *et al.*, "Calibration of the LIGO detectors for S4", LIGO-T050262-01-D (2006).

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

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

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

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

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

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

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

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

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

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

[20] 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.

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

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

[23] R. Schofield, “Acoustic Mitigation Status for S3”, LIGO-G030640 (2003).

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

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

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

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

[28] *GEO Summary Reports*, http://www.geo600.uni-hannover.de/georeports/.

[29] *Results from the first burst hardware injections performed on GEO 600*, R Balasubramanian, H Grote, I S Heng, M Hewitson, H Lück, J R Smith, K A Strain, H Ward and B Willke, Class. Quantum Grav. **22** No 14 (21 July 2005), 3015-3028.

[30] *First steps towards characterizing the hierarchical algorithm for curves and ridges pipeline*, I S Heng, R Balasubramanian, B S Sathyaprakash and B F Schutz, Class. Quantum Grav. **21** No 5 (7 March 2004) S821-S826.

[31] *On aspects of characterising and calibrating the interferometric gravitational wave detector,* GEO 600, M Hewitson, PhD Thesis, University of Glasgow, 2004.

[32] *Data Acquisition and Data Analysis for the Gravitational-Wave Detector* GEO 600, K Kötter, PhD Thesis, University of Hannover, 2004.

[33] *Principles of calibrating the dual-recycled GEO 600*, M Hewitson, G Heinzel, J R Smith, K A Strain, and H Ward, Rev. Sci. Instrum. 75, 4702 (2004).

[34] *Calibration of the dual-recycled GEO 600 detector for the S3 science run*, M Hewitson, S Babak, R Balasubramanian, K Danzmann, H Grote, G Heinzel, J Hough, H Lück, M A Papa, J R Smith, K A Strain, H Ward, B Willke and G Woan, Class. Quantum Grav. **21** No 20 (21 October 2004) S1711-S1722.

[35] *Optimal time-domain combination of the two calibrated output quadratures of GEO 600*, M Hewitson, H Grote, S Hild, H Lück, P Ajith, J R Smith, K A Strain, B Willke, G Woan, Class. Quantum Grav. 22 (2005) 4253-4261.

[36] *Using the null-stream of GEO 600 to veto transient events in the detector output*, M Hewitson and P Ajith, Class. Quantum Grav. **22** No 22 (21 November 2005) 4903-4912.

[37] *Linear projection of technical noise for interferometric gravitational-wave detectors*, J R Smith, P Ajith, H Grote, M Hewitson, S Hild, H Lück, K A Strain, B Willke, J Hough and K Danzmann, Class. Quantum Grav. **23** No 2 (21 January 2006) 527-537.

[38] *Robust vetoes for gravitational-wave burst triggers using known instrumental couplings*, P Ajith, M Hewitson, J R Smith and K A Strain, Class. Quantum Grav. **23** No 20 (21 October 2006) 5825-5837.

[39] *Physical instrumental vetoes for gravitational-wave burst triggers*, P. Ajith, M. Hewitson, J. R. Smith, H. Grote, S. Hild, K. A. Strain, submitted to Phys. Rev. D.

[40] some acronyms:
$[qcmoni]$: QcMoni Documentation: virgo note VIR-075A-08
$[fdio]$: Virgo Online Frame Distribution: virgo note VIR-087B-08
$[moni]$: Detector Monitoring Library : virgo note VIR-074A-08
$[dy]$: Data Display Users Manual: virgo note VIR-008A-08
$[vega]$: ROOT based visual environment for GW data analysis: lappweb.in2p3.fr/virgo/vega
$[spectro]$: Virgo SpectroMoni documentation: virgo note VIR-076A-08

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

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

[43] Commissioning progress report for the STAC and the EGO council, May 2009, virgo note VIR-034A-09

[44] 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.

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

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

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

[48] LIGO Scientific Collaboration and K Hurley. Implications for the origin of grb 070201 from ligo observations. *Astroph. Journal*, 681:1419–1430, 2008.

[49] 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.

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

[51] 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.

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

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

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

[55] 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.

[56] : B. Abbott. Search for gravitational wave ringdowns from perturbed black holes in LIGO S4 data. 2009.

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

[58] 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.

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

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

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

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

[63] 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.

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

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

[66] 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.

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

[68] 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.

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

[70] 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).

[71] 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).

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

[73] D Fazi.

[74] 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.

[75] 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.

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

[77] 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.

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

[79] 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.

[80] 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.

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

[82] 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.

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

[84] 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.

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

[86] The LIGO Scientific Collaboration. Tuning matched filter searches for compact binary coalescence. Technical Report LIGO-T070109-01, 2007.

[87] B. Abbott et al. Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO's S5 Data. 2009.

[88] B. Abbott et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO's fifth Science Run. 2009.

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

[90] 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.

[91] 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.

[92] 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.

[93] CBC Group. The rate of compact binary coalescences.

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

[95] 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.

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

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

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

[99] B. Abbott et al. Search for gravitational waves from primordial black hole binary coalescences in the galactic halo. *Phys. Rev. D*, 72:082002, 2005.

[100] 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).

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

[102] 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.

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

[104] B. Abbott et al. (LIGO Scientific Collaboration), Phys. Rev. D **69**, 102001 (2004).

[105] R. C. Duncan and C. Thompson, *ApJ Lett.* **392**, L9 (1992).

[106] B. Abbott et al., *Phys. Rev. Lett.* **101**, 211102 (2008).

[107] O. Benhar, V. Ferrari, and L. Gualtieri, *Phys. Rev. D* **70**, 124015 (2004), `arXiv:astro-ph/0407529`.

[108] K. Ioka, *MNRAS* **327**, 639 (2001a), `arXiv:astro-ph/0009327`.

[109] B. Abbott et al. (LIGO Scientific Collaboration), *Phys. Rev. D* **76**, 062003 (2007), `arXiv:astro-ph/0703419`.

[110] A. L. Watts and T. E. Strohmayer, *ApJ Lett.* **637**, L117 (2006a), `arXiv:astro-ph/0512630`.

[111] B. Abbott et al., ArXiv e-prints (2009a), to appear in Astrophys. J. Lett., `arXiv:0905.0005`.

[112] LIGO Scientific Collaboration and K. Hurley, *ApJ* **681**, 1419 (2008).

[113] B. Abbott et al. (LIGO Scientific Collaboration), *Phys. Rev. D* **72**, 042002 (pages 17) (2005).

[114] F. Acernese et al. (Virgo Collaboration), *Class. Quant. Grav.* **25**, 225001 (2008a).

[115] B. Abbott et al. (LIGO Scientific Collaboration) (2009b), `arXiv:0905.0020[gr-qc]`.

[116] J. Baker et al., Phys. Rev. D **73**, 104002 (2006).

[117] C. D. Ott, A. Burrows, L. Dessart, and E. Livne, Phys. Rev. Lett. **96**, 201102 (2006a).

[118] B. Abbott et al. (LIGO Scientific Collaboration) (2009c), `arXiv:0904.4910[gr-qc]`.

[119] F. Acernese et al. (Virgo Collaboration), *Class. Quant. Grav.* **26**, 085009 (2008b).

[120] B. Abbott et al. (LIGO Scientific Collaboration) (2009d), submitted to Phys. Rev. D, `arXiv:0904.4718[gr-qc]`.

[121] S. Klimenko and G. Mitselmakher, Class. Quant. Grav. **21**, S1819 (2004).

[122] S. Chatterji et al., Class. Quant. Grav. **21**, S1809 (2004).

[123] J. Camp et al., Phys. Rev. D **75**, 061101(R) (2007).

[124] Shantanu Desai, Sam Finn, John McNabb, Amber Stuver, Tiffany Summerscales and Keith Thorne, Class. Quantum Grav. **21**, S1705 (2004).

[125] C. L and M. S, Class. Quantum Grav. **22**, S1159 (2005), `arXiv:0809.2809`.

[126] L. Blackburn et al., LIGO-G050158-00-Z (2005), URL `http://www.ligo.caltech.edu/docs/G/G050158-00`.

[127] S. Klimenko et al., Class. Quant. Grav. **21**, S1685 (2004).

[128] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, Phys. Rev. D **72**, 122002 (2005), `arXiv:gr-qc/0508068`.

[129] S. Klimenko et al., Class. Quantum Grav. **25**, 114029 (2008), `arXiv:0802.3232`.

[130] R. A. Mercer and S. Klimenko, Class. Quantum Grav. **23**, 184025 (2008).

[131] P. J. Sutton, S. Chatterji, G. Jones, S. Poprocki, J. Rollins, A. Searle, L. Stein, M. Tinto, and M. Was, in preparation (2009), `https://geco.phys.columbia.edu/xpipeline/browser/trunk/docs/T080088/T080088-02.pdf`.

[132] A. C. Searle, P. J. Sutton, and M. Tinto, Class. Quantum Grav. (to appear) (2009), `arXiv:0809.2809`.

[133] T. Z. Summerscales, A. Burrows, L. S. Finn, and C. D. Ott, Astrophys. J. **678**, 1142 (2008).

[134] C. Culter and K. Thorne, in *General Relativity and Gravitation*, edited by N. T. Bishop and D. M. Sunil (2002), p. 72, `arXiv:gr-qc/0204090`.

[135] J. G. Baker, M. Campanelli, F. Pretorius, and Y. Zlochower, Classical and Quantum Gravity **24**, 25 (2007), `arXiv:gr-qc/0608133`.

[136] C. D. Ott, Classical and Quantum Gravity **26**, 063001 (2009), `arXiv:0809.0695`.

[137] M. H. P. M. van Putten, *ApJ* **583**, 374 (2003), `arXiv:astro-ph/0209518`.

[138] The Pierre Auger Collaboration, Science **318**, 938 (2007).

[139] T. M. Collaboration, *ApJ* **546**, 1038 (2001).

[140] T. I. Collaboration, *Phys. Rev. D* **79**, 062001 (2009).

[141] The Super-Kamiokande Collaboration, *ApJ* **652**, 198 (2006).

[142] S. Ando, J. F. Beacom, and H. Yüksel, Physical Review Letters **95**, 171101 (2005), `arXiv:astro-ph/0503321`.

[143] P. Ajith et al., Phys. Rev. **D77**, 104017 (2008), `arXiv:0710.2335`.

[144] A. Buonanno et al., Phys. Rev. D **76**, 104049 (2007), `arXiv:0706.3732`.

[145] Y. Pan et al., Phys. Rev. D **77**, 024014 (2008), `arXiv:0704.1964`.

[146] J. Veitch and A. Vecchio, Class. Quant. Grav. **25**, 184010 (2008), `arXiv:0807.4483`.

[147] M. van der Sluys et al., Class. Quant. Grav. **25**, 184011 (2008), `arXiv:0805.1689`.

[148] N. Huang et al., Proc R Soc Lond A **454**, 903 (1998).

[149] C. Pankow et al. (2009), `arXiv:0905.3120`.

[150] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, B. Zink, I. Hawke, and E. Schnetter, Classical and Quantum Gravity **24**, 139 (2007a), `arXiv:astro-ph/0612638`.

[151] C. D. Ott, H. Dimmelmeier, A. Marek, H.-T. Janka, I. Hawke, B. Zink, and E. Schnetter, Physical Review Letters **98**, 261101 (2007b), `arXiv:astro-ph/0609819`.

[152] C. D. Ott, A. Burrows, L. Dessart, and E. Livne, Physical Review Letters **96**, 201102 (2006b), `arXiv:astro-ph/0605493`.

[153] K. Kotake, W. Iwakami, N. Ohnishi, and S. Yamada, *ApJ Lett.* **697**, L133 (2009), `arXiv:0904.4300`.

[154] A. Marek, H.-T. Janka, and E. Müller, *A&A* **496**, 475 (2009), `arXiv:0808.4136`.

[155] Abbott, B. and others, *Phys. Rev. Lett.* **101**, 211102 (2008).

[156] Abbott, B. and others, Phys. Rev. D **76**, 062003 (2007).

[157] C. Kouveliotou, C. A. Meegan, G. J. Fishman, N. P. Bhat, M. S. Briggs, T. M. Koshut, W. S. Paciesas, and G. N. Pendleton, Astrophysical Journal Letters **413**, L101 (1993).

[158] N. Gehrels et al., Nature **444**, 1044 (2006).

[159] S. Campana et al., Nature **442**, 1008 (2006), `arXiv:astro-ph/0603279`.

[160] D. Malesani et al., Astrophysical Journal Letters **609**, L5 (2004), `arXiv:astro-ph/0405449`.

[161] J. Hjorth et al., Nature **423**, 847 (2003), `arXiv:astro-ph/0306347`.

[162] T. J. Galama et al., Nature **395**, 670 (1998), `arXiv:astro-ph/9806175`.

[163] K. Hurley et al., Nature **434**, 1098 (2005), `arXiv:astro-ph/0502329`.

[164] K. Hurley, C. Kouveliotou, T. Cline, E. Mazets, S. Golenetskii, D. D. Frederiks, and J. van Paradijs, Astrophysical Journal Letters **523**, L37 (1999), `arXiv:astro-ph/9906020`.

[165] E. Nakar, A. Gal-Yam, T. Piran, and D. B. Fox, Astrophysical Journal **640**, 849 (2006).

[166] D. M. Palmer et al., Nature **434**, 1107 (2005), `arXiv:astro-ph/0503030`.

[167] J. S. Bloom et al., Astrophysical Journal **654**, 878 (2007), `arXiv:astro-ph/0607223`.

[168] E. Nakar, Physics Reports **442**, 166 (2007), `arXiv:astro-ph/0701748`.

[169] R. Chapman, R. S. Priddey, and N. R. Tanvir, [arXiv:0709.4640] **709** (2007), `arXiv:0709.4640`.

[170] A. L. Piro and E. Pfahl, The Astrophysical Journal **658**, 1173 (2007).

[171] M. H. van Putten, A. Levinson, H. K. Lee, T. Regimbau, M. Punturo, and G. M. Harry, *Phys. Rev. D* **69**, 044007 (2004), `arXiv:gr-qc/0308016`.

[172] B. P. Abbott et al., in preparation (2009e).

[173] S. Chatterji et al., Phys. Rev. D **74**, 082005 (2006).

[174] P. M. Woods, C. Kouveliotou, M. H. Finger, E. Göğüş, C. A. Wilson, S. K. Patel, K. Hurley, and J. H. Swank, *ApJ* **654**, 470 (2007), `arXiv:astro-ph/0602402`.

[175] C. Thompson and R. C. Duncan, Mon. Not. R. Astron. Soc. **275**, 255 (1995).

[176] 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, *ApJ Lett.* **627**, L129 (2005), `arXiv:astro-ph/0504056`.

[177] N. Andersson and K. D. Kokkotas, *MNRAS* **299**, 1059 (1998), `arXiv:gr-qc/9711088`.

[178] J. A. de Freitas Pacheco, Astronomy and Astrophysics **336**, 397 (1998), `arXiv:astro-ph/9805321`.

[179] K. Ioka, MNRAS **327**, 639 (2001b), `arXiv:astro-ph/0009327`.

[180] R. X. Xu, The Astrophysical Journal Letters **596**, L59 (2003).

[181] B. J. Owen, Physical Review Letters **95**, 211101 (2005), `arXiv:astro-ph/0503399`.

[182] J. E. Horvath, Modern Physics Letters A **20**, 2799 (2005).

[183] T. E. Strohmayer and A. L. Watts, *ApJ* **653**, 593 (2006), `arXiv:astro-ph/0608463`.

[184] A. L. Watts and T. E. Strohmayer, *ApJ Lett.* **637**, L117 (2006b), `arXiv:astro-ph/0512630`.

[185] T. E. Strohmayer and A. L. Watts, *ApJ Lett.* **632**, L111 (2005), `arXiv:astro-ph/0508206`.

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

[187] E. Waxman and J. Bahcall, Physical Review Letters **78**, 2292 (1997), `arXiv:astro-ph/9701231`.

[188] J. P. Rachen and P. Mészáros, Phys. Rev. D **58**, 123005 (1998), `arXiv:astro-ph/9802280`.

[189] J. Alvarez-Muñiz, F. Halzen, and D. W. Hooper, Phys. Rev. D **62**, 093015 (2000), `arXiv:astro-ph/0006027`.

[190] E. Waxman and J. N. Bahcall, Astrophys. Journal **541**, 707 (2000), `arXiv:hep-ph/9909286`.

[191] P. Mészáros and E. Waxman, Physical Review Letters **87**, 171102 (2001), `arXiv:astro-ph/0103275`.

[192] D. Guetta and J. Granot, Physical Review Letters **90**, 201103 (2003), `arXiv:astro-ph/0212045`.

[193] S. Razzaque, P. Mészáros, and E. Waxman, Physical Review Letters **90**, 241103 (2003a), `arXiv:astro-ph/0212536`.

[194] S. Razzaque, P. Mészáros, and E. Waxman, Phys. Rev. D **68**, 083001 (2003b), `arXiv:astro-ph/0303505`.

[195] C. D. Dermer and A. Atoyan, Physical Review Letters **91**, 071102 (2003), `arXiv:astro-ph/0301030`.

[196] K. Murase and S. Nagataki, Physical Review Letters **97**, 051101 (2006), `arXiv:astro-ph/0604437`.

[197] K. Murase, K. Ioka, S. Nagataki, and T. Nakamura, Astrophys. Journal **651**, L5 (2006), `arXiv:astro-ph/0607104`.

[198] M. Vietri, Physical Review Letters **80**, 3690 (1998), `arXiv:astro-ph/9802241`.

[199] E. Waxman, Physica Scripta Volume T **85**, 117 (2000), `arXiv:astro-ph/9911395`.

[200] W. H. Lee and E. Ramirez-Ruiz, New Journal of Physics **9**, 17 (2007), `arXiv:astro-ph/0701874`.

[201] S. R. Kulkarni et al., Nature **395**, 663 (1998).

[202] A. M. Soderberg et al., Nature **430**, 648 (2004a), `arXiv:astro-ph/0408096`.

[203] B. E. Cobb, C. D. Bailyn, P. G. van Dokkum, and P. Natarajan, Astrophys. Journal **645**, L113 (2006), `arXiv:astro-ph/0603832`.

[204] E. Pian et al., Nature **442**, 1011 (2006), `arXiv:astro-ph/0603530`.

[205] A. M. Soderberg et al., Nature **442**, 1014 (2006), `arXiv:astro-ph/0604389`.

[206] E. Liang, B. Zhang, F. Virgili, and Z. G. Dai, Astrophys. Journal **662**, 1111 (2007), `arXiv:astro-ph/0605200`.

[207] N. Gupta and B. Zhang, Astroparticle Physics **27**, 386 (2007), `arXiv:astro-ph/0606744`.

[208] X.-Y. Wang, S. Razzaque, P. Mészáros, and Z.-G. Dai, Phys. Rev. D **76**, 083009 (2007), `arXiv:0705.0027`.

[209] S. Ando and J. F. Beacom, Physical Review Letters **95**, 061103 (2005), `arXiv:astro-ph/0502521`.

[210] A. M. Soderberg, D. A. Frail, and M. H. Wieringa, Astrophys. Journal **607**, L13 (2004b), `arXiv:astro-ph/0402163`.

[211] J. Granot and E. Ramirez-Ruiz, Astrophys. Journal **609**, L9 (2004), `arXiv:astro-ph/0403421`.

[212] P. A. Mazzali et al., Science **308**, 1284 (2005), `arXiv:astro-ph/0505199`.

[213] K. Ioka, S. Razzaque, S. Kobayashi, and P. Mészáros, Astrophys. Journal **633**, 1013 (2005), `arXiv:astro-ph/0503279`.

[214] A. Achterberg et al., Physical Review Letters **97**, 221101 (2006).

[215] IceCube, http://icecube.wisc.edu.

[216] ANTARES, http://antares.in2p3.fr/.

[217] Y. Aso, Z. Márka, C. Finley, J. Dwyer, K. Kotake, and S. Márka, Classical and Quantum Gravity **25**, 114039 (2008), `arXiv:0711.0107`.

[218] 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).

[219] K. Scholberg, ArXiv Astrophysics e-prints (2007), `arXiv:astro-ph/0701081`.

[220] J. Clark, I. S. Heng, M. Pitkin, and G. Woan, Physical Review D (Particles, Fields, Gravitation, and Cosmology) **76**, 043003 (pages 13) (2007).

[221] L. S. Finn, S. D. Mohanty, and J. D. Romano, Phys. Rev. D. **60**, 121101 (1999), `arXiv:gr-qc/9903101`.

[222] S. D. Mohanty, Classical and Quantum Gravity **22**, S1349 (2005).

[223] B. Abbott et al. (LIGO Scientific Collaboration), *Phys. Rev. D* **77**, 062004 (pages 22) (2008), URL `http://link.aps.org/abstract/PRD/v77/e062004`.

[224] S. D. Mohanty, Classical and Quantum Gravity **23**, 723 (2006).

[225] M. van der Klis, J. H. Swank, W. Zhang, K. Jahoda, E. H. Morgan, W. H. G. Lewin, B. Vaughan, and J. van Paradijs, *ApJ Lett.* **469**, L1+ (1996), `arXiv:astro-ph/9607047`.

[226] D. Chakrabarty, E. H. Morgan, M. P. Muno, D. K. Galloway, R. Wijnands, M. van der Klis, and C. B. Markwardt, *Nature* **424**, 42 (2003), `arXiv:astro-ph/0307029`.

[227] R. V. Wagoner, *ApJ* **278**, 345 (1984).

[228] N. Andersson, *ApJ* **502**, 708 (1998), `arXiv:gr-qc/9706075`.

[229] L. Bildsten, *ApJ Lett.* **501**, L89+ (1998), `arXiv:astro-ph/9804325`.

[230] J. Jernigan, AIPC **586**, 805 (2001).

[231] V. M. Lipunov and I. E. Panchenko, *A&A* **312**, 937 (1996), `arXiv:astro-ph/9608155`.

[232] B. M. S. Hansen and M. Lyutikov, *MNRAS* **322**, 695 (2001), `arXiv:astro-ph/0003218`.

[233] J. Moortgat and J. Kuijpers, ArXiv General Relativity and Quantum Cosmology e-prints (2005), `arXiv:gr-qc/0503074`.

[234] M. D. Duez, Y. T. Liu, S. L. Shapiro, and B. C. Stephens, *Phys. Rev. D* **72**, 024029 (2005), `arXiv:astro-ph/0503421`.

[235] V. V. Usov and J. I. Katz, *A&A* **364**, 655 (2000), `arXiv:astro-ph/0002278`.

[236] C. W. Stubbs, Class. Quantum Grav. **25**, 184033 (2008).

[237] D. B. Fox et al., Nature **437**, 845 (2005).

[238] N. Gehrels, J. K. Cannizzo, and J. P. Norris, New Journal of Physics **9** (2007).

[239] L. Li and B. Pacyznski, Astrophys. J. **507**, L59 (1998).

[240] S. Kulkarni (2005), `arXiv:astro-ph/0510256`.

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

[242] J. Kanner, T. L. Huard, S. Marka, D. C. Murphy, J. Piscionere, M. Reed, and P. Shawhan, Class. Quantum Grav. **25**, 184034 (2008).

[243] J. Sylvestre, Astrophys. J. **591**, 1152 (2003).

[244] J. S. Bloom et al. (2009), `arXiv:0902.1527`.

[245] S. Desai, E. O. Kahya, and R. P. Woodard, Phys. Rev. D **77**, 124041 (2008).

[246] N. Dalal, D. E. Holz, S. Hughes, and B. Jain, Phys. Rev. D **74**, 063006 (2006).

[247] D. M. Eardley, D. L. Lee, A. P. Lightman, R. V. Wagoner, and C. M. Will, Phys. Rev. Lett. **30**, 884 (1973a).

[248] D. M. Eardley, D. L. Lee, and A. P. Lightman, Phys. Rev. D **8**, 3308 (1973b).

[249] T. Creighton, F. A. Jenet, and R. H. Price (2008), `arXiv:0812.3941`.

[250] A. Nishizawa, A. Taruya, K. Hayama, S. Kawamura, and M. Sakagami (2009), `arXiv:0903.0528`.

[251] S. Alexander, L. S. Finn, and N. Yunes, Phys. Rev. **D78**, 066005 (2008), `arXiv:0712.2542`.

[252] C. M. Will (1998), `arXiv:gr-qc/9811036`.

[253] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **69**, 082004 (2004).

[254] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. Lett. **94**, 181103 (2005).

[255] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 102004 (2005).

[256] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **72**, 082001 (2007)

[257] L. Bildsten, Astrophys. J. **501**, L89 (1998).

[258] G. Ushomirsky, C. Cutler and L. Bildsten, Mon. Not. Roy. Astron. Soc. **319**, 902 (2000).

[259] C. Cutler, Phys. Rev. D **66**, 084025 (2002).

[260] B. J. Owen, Phys. Rev. Lett. **95**, 211101 (2005).

[261] B. J. Owen, L. Lindblom, C. Cutler, B. F. Schutz, A. Vecchio and N. Andersson, Phys. Rev. D **58**, 084020 (1998).

[262] N. Andersson, K. D. Kokkotas and N. Stergioulas, Astrophys. J. **516**, 307 (1999).

[263] A. Melatos and D. J. B. Payne, Astrophys. J. **623**, 1044 (2005).

[264] B. J. Owen, Class. Quant. Grav. **23**, S1 (2006).

[265] P. Brady and T. Creighton, Phys. Rev. D **61**, 082001 (2000).

[266] C. Cutler, I. Gholami, and B. Krishnan, Phys. Rev. D **72**, 042004 (2005).

[267] D. I. Jones and N. Andersson, Mon. Not. Roy. Astron. Soc. **331**, 203 (2002).

[268] R. V. Wagoner, Astrophys. J. **278**, 345 (1984).

[269] K. S. Thorne, in *Three hundred years of gravitation,* eds. S. W. Hawking and W. Israel (Cambridge University Press: 1987).

[270] R.J Dupuis, Ph.D. Thesis, University of Glasgow (2004); `LIGO-P050014-00-R`.

[271] B Abbott *et al.* (The LIGO Scientific Collaboration), to appear in Ap. J. Lett., arXiv:0805.4758.

[272] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D, **76**, 042001 (2007).

[273] M. Pitkin and G. Woan, Phys. Rev. D, **76**, 042006 (2007).

[274] P. Astone *et al.*, Phys. Rev. D **65**, 022001 (2001).

[275] S. Frasca and C. Palomba, Class. Quant. Grav. **21**, S1645 (2004).

[276] D. Passuello and S. Braccini, Virgo Note (VIR-046A-07).

[277] S. Frasca, P. Astone and C. Palomba, Class. Quant. Grav. **22**, S1013 (2005).

[278] P. Astone, S. Frasca and C. Palomba, Class. Quant. Grav. **22**, S1197 (2005).

[279] C. Palomba, P. Astone and S. Frasca, Class. Quant. Grav. **22**, S1255 (2005).

[280] F. Acernese *et al.*, Class. Quant. Grav. **24**, S491 (2007).

[281] S. D'Antonio, S. Frasca, C. Palomba, submitted to CQG (2009).

[282] S. Frasca and C. La Posta, Il Nuovo Cimento, **14 C**, (1991).

[283] F. Antonucci, P. Astone, S. D'Antonio, S. Frasca, C. Palomba, CQG **25**, 184015 (2008).

[284] B. Krishnan, A. M. Sintes, M. A. Papa, B. F. Schutz, S. Frasca and C. Palomba, Phys. Rev. D **70**, 082001 (2004).

[285] V. Dergachev, "Description of PowerFlux Algorithms and Implementation", *LIGO Technical Document* LIGO-T050186 (2005), available in `http://admdbsrv.ligo.caltech.edu/dcc/`.

[286] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **77**, 022001 (2008).

[287] 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).

[288] 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/`.

[289] B. Knispel and B. Allen, Phys. Rev. D **78**, 044031 (2008).

[290] S. Dhurandhar, B. Krishnan, H. Mukhopadhyay, and J. T. Whelan, Phys. Rev. D **77**, 082001 (2008).

[291] L. Sancho de la Jordana, A.M. Sintes, preprint arXiv:0804.1007 [gr-qc] (April 2008). To appear in Class. Quant. Grav.

[292] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007)

[293] Ransom, S. M., Cordes, J. M., & Eikenberry, S. S. 2003, Astrophys. J. **589**, 911

[294] Messenger, C., & Woan, G. 2007, Class. Quant. Grav., **24**, 469

[295] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **79**, 022001 (2009).

[296] B. Abbott *et al.* (The LIGO Scientific Collaboration), submitted to Phys. Rev. D, arXiv:0905.1705 (2009)

[297] P. Jaranowski, A. Krolak, and B. F. Schutz, "Data analysis of gravitational-wave signals from pulsars. I. The signal and its detection", Physical Review D **58**, 063001-24, (1998).

[298] C.-Y. Ng and R. W. Romani, "Fitting Pulsar Wind Tori. II. Error Analysis and Applications", arXiv.:0710.4168v1

[299] K. Wette *et al.* [LIGO Collaboration], arXiv:0802.3332 [gr-qc].

[300] R. E. Rutledge, D. B. Fox and A. H. Shevchuk, arXiv:0705.1011 [astro-ph].

[301] J.H. Taylor and G.R. Huguenin, Nature **221**, 816 (1969).

[302] P.C. Gregory and T.J. Loredo, ApJ. bf 398, 146 (1992).

[303] B. Abbott *et al.* (The LIGO Scientific Collaboration), Ap. J. **659**, 918 (2007).

[304] B. Abbott *et al.* (The LIGO Scientific Collaboration and the ALLEGRO Collaboration), Phys. Rev. D **76**, 022001 (2007).

[305] B. Abbott *et al.* (The LIGO Scientific Collaboration), Phys. Rev. D **76**, 082003 (2007)

[306] Flanagan É É 1993 *Phys. Rev.* **D48** 2389; astro-ph/9305029

[307] Allen B and Romano J D 1999 *Phys. Rev.* **D59** 102001; gr-qc/9710117.

[308] G. Cella *et al.*, Class. Quant. Grav. **24**, S639 (2007)

[309] B. Abbott et al., "Analysis of first LIGO science data for stochastic gravitational waves", Phys. Rev. D **69**, 122004 (2004).

[310] Ballmer S W 2006 *Class. Quant. Grav.* **23** S179; gr-qc/0510096

[311] N.V. Fotopoulos, "Identifying Correlated Environmental Noise in Co-Located Interferometers with Application to Stochastic Gravitational Wave Analysis", in preparation for GWDAW 10 proceedings.

[312] http://wwwcascina.virgo.infn.it/DataAnalysis/Noise/nap_index.html

[313] F. Beauville *et al.* [LIGO / Virgo working group], "Detailed comparison of LIGO and Virgo inspiral pipelines in preparation for a joint search," arXiv:gr-qc/0701027, submitted to PRD.

[314] F. Beauville *et al.* [LIGO-Virgo working group], "A comparison of methods for gravitational wave burst searches from LIGO and Virgo,", arXiv:gr-qc/0701026, submitted to PRD.

[315] F. Beauville *et al.*, "Benefits of joint LIGO - Virgo coincidence searches for burst and inspiral signals," J. Phys. Conf. Ser. **32**, 212 (2006) [arXiv:gr-qc/0509041].

[316] L. Blackburn *et al.*, "A first comparison of search methods for gravitational wave bursts using LIGO and Virgo simulated data," Class. Quant. Grav. **22**, S1293 (2005) [arXiv:gr-qc/0504060].

[317] L. Blackburn *et al.* [Joint LIGO / Virgo working group], "A first comparison between LIGO and Virgo inspiral search pipelines," Class. Quant. Grav. **22**, S1149 (2005) [arXiv:gr-qc/0504050].