Very low latency search pipeline for low mass compact binary coalescences in the LIGO S6 and Virgo VSR2 data

2010 Class. Quantum Grav. 27 194013

(http://iopscience.iop.org/0264-9381/27/19/194013)

Class. Quantum Grav. **27** (2010) 194013 (9pp) doi:10.1088/0264-9381/27/19/194013

Very low latency search pipeline for low mass compact binary coalescences in the LIGO S6 and Virgo VSR2 data

D Buskulic (on behalf of the LIGO Scientific Collaboration and the Virgo Collaboration[1])

LAPP, Université de Savoie, CNRS, Chemin de Bellevue, BP 110, F-74941 Annecy-le-Vieux Cedex, France

E-mail: buskulic@lapp.in2p3.fr

Received 2 May 2010, in final form 16 June 2010
Published 21 September 2010
Online at stacks.iop.org/CQG/27/194013

Abstract
A very low latency search pipeline has been developed for the LIGO S6 and Virgo VSR2 science runs, targeting signals from coalescing compact binary systems with total mass from 2 to 35 solar masses. The goal of this search is to provide both single-detector triggers and multi-detector coincident triggers with a latency of a few minutes, the former for online detector monitoring and the latter to allow searching for electromagnetic counterparts to possible gravitational wave candidates. The features and current performance of this low latency search pipeline are presented.

PACS numbers: 04.80.Nn, 07.05.Hd, 07.05.Kf

(Some figures in this article are in colour only in the electronic version)

1. Introduction

The two LIGO [1] and the Virgo [2] gravitational wave detectors performed a coincident run period from 7 July 2009 to 8 January 2010 during their S6 (for LIGO) and VSR2 (for Virgo) runs. The need for a very fast search of compact binary coalescence triggers appeared for two main reasons. First, the extraction of single-detector triggers in real time allows for some fast detector characterization, such as monitoring of the trigger rate and data quality studies. The second reason is the identification and localization in the sky of interesting candidates that deserve a follow-up by electromagnetic detectors, such as telescopes or gamma ray observatories. The speed of the identification is critical to ensure proper detection of an electromagnetic counterpart to the potential gravitational wave signal. The localization uses only candidates coincident in the three detectors.

[1] A list of members of the LIGO Scientific Collaboration and the VIRGO Collaboration can be found at the end of this issue.

We describe a low latency search pipeline and the first performance results obtained during the coincident run period.

When compared to the baseline search used for publications of detections or setting of upper limits [3], the desired very low latency of the search described hereafter imposes a streamlined and simple pipeline. The analysis threshold on the signal-to-noise ratio (SNR) of the single-detector triggers is raised from 5.5 in the baseline search to 6 in the low latency search. We restrict the search of coalescing binary signals to the so-called low-mass range from 1 to 34 $M_\odot$ for each component, with a total mass of the system between 2 and 35 $M_\odot$ and perform limited consistency tests. Because of the need to localize the event in the sky, we focus on 3-site coincidences for multi-detector analysis.

The cuts and thresholds described above are suboptimal when compared to the baseline search, leading to a loss of sensitivity. But the goal is not the same; we seek here solid events that could be sent to electromagnetic detectors for follow-up with a good level of confidence.

1.1. The case for electromagnetic follow-ups

The mergers of binary neutron star systems (NS-NS) or neutron star–black hole (NS-BH) are thought to be plausible progenitors of short, hard gamma ray bursts (GRBs) [4]. The possible observation of a gravitational wave signal coincident with a GRB would confirm the relation between a merger and a GRB and would give a great confidence in the gravitational wave detection itself. Furthermore, the GRB detection would bring additional information about the source that would help the gravitational wave search, like an accurate sky position, the identification of the host galaxy and its redshift.

Searches of gravitational wave events triggered by external short, hard GRB events are part of the LIGO-Virgo analyses [5, 6], but it is worth going in the opposite way, triggering an afterglow search on a gravitational wave trigger. Because GRBs are believed to result from collimated outflows, the resulting beaming factor reduces the chance of observation by gamma satellites. It is possible that some GRBs could be observed only through their afterglows (orphan afterglows) [7], giving more chances to see a coincidence between a gravitational waves signal and a GRB, even in the absence of a direct GRB observation. The time scale of afterglows, which is of the order of hours to days, is compatible with this approach.

1.2. Follow-up instruments

In order to achieve the electromagnetic follow-ups, collaborations were established between LIGO-Virgo and a few instruments. As an example, Swift [8] is offering target of opportunity (ToO) observations from the gravitational wave detectors. At a ToO request, the Swift satellite will look for afterglows in the x-ray, UV and optical domains. LIGO-Virgo expects to send about 3 requests during the Swift Cycle 6, starting in April 2010. These events will most likely be due to detector noise, but could plausibly contain a true signal. In order to probe the detection process, one of the triggering candidates could be a test, a so-called blind hardware injection.

Some collaborations have started with wide-field optical telescopes such as 'TAROT' [9] and 'QUEST' [10] to look for electromagnetic counterparts. Discussions have started with other wide-field telescope arrays ('Pi of the Sky' [11]).

2. The very low latency pipeline

2.1. The multi-band template analysis pipeline

The multi-band template analysis (MBTA) pipeline [12] is the primary detection engine of the low latency search. It implements efficiently a matched filtering process [13] over a

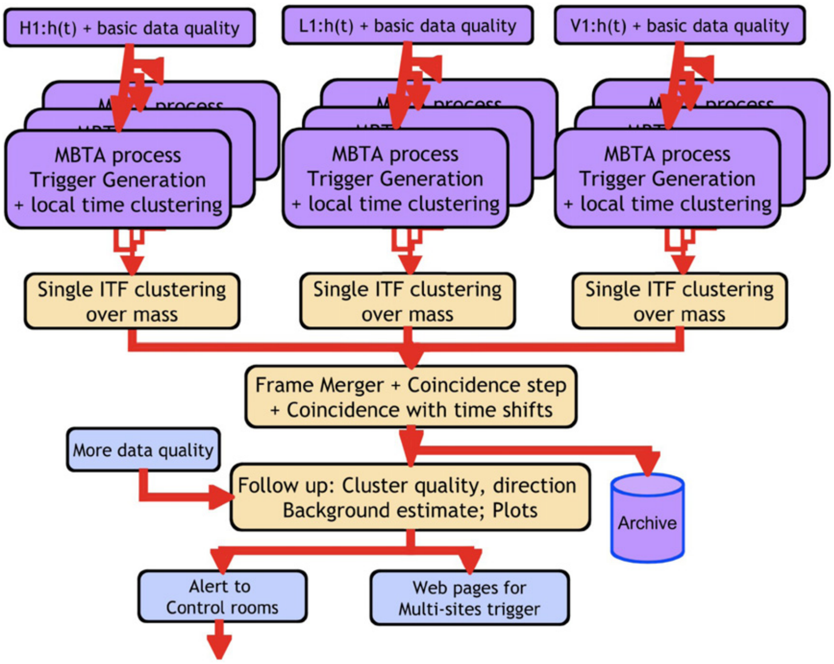

Figure 1. Diagram of the implementation of the very low latency pipeline.

family (bank) of templates. This technique consists of the comparison of the output signal of an interferometer with a family of expected theoretical waveforms, called templates. The comparison with one template is made through a Wiener filter [14] or matched filter, which is essentially an intercorrelation weighted with the inverse of the noise power spectral density. In the usual way of implementing a Wiener filtering, one builds a set (or bank) of templates which cover the parameter space of the search. In the multi-band technique, the matched filtering integral is calculated over two frequency bands, low and high, the output of which (SNR) is coherently added.

The templates of the banks are calculated in time domain using a second-order post-Newtonian approximation [15]. There are also adaptive mechanisms that follow detector non-stationarities.

The primary focus of the pipeline is low latency and speed, so a fast and simple pipeline was used. For example, the waveform consistency test is a two-band χ^2 test, with limited discriminating power but computationally inexpensive. Furthermore, while there are a few processes running in parallel to perform the matched filtering, there are no files involved in the information and data exchange between processes. We just use a TCP-IP-based protocol developed for the Virgo data acquisition system [16].

2.2. *Online implementation*

Figure 1 shows the online implementation of the very low latency pipeline. All the data are brought to the Virgo site where the processing is split by the interferometer and mass range.

The basic steps of the pipeline are as follows.

- After applying a set of basic data quality vetoes, the MBTA processes are run and generate the primary triggers on each individual interferometer flow of data. A time clustering is applied at this step; all triggers separated by less than 0.1 s are grouped together.
- A few templates in a bank may respond to each event in the data; hence, a clustering is applied over the masses of the triggers on each individual interferometer.
- A clustering and merging of the triggers from the individual interferometers is then performed, with a coincidence step. The time window between the LIGO–Hanford and LIGO–Livingston interferometers is set to 20 ms, while between any LIGO interferometer and Virgo, the time window is 40 ms. At this step, one may apply time shifts of the data for estimating the background. At the end of the merger, the pipeline provides triggers consisting of single-site, 2-site or 3-site coincidences.
- The last step before sending an alert to the control room is a follow-up procedure (that we will call internal follow-up to distinguish it from the electromagnetic final follow-up) looking at the cluster quality, the direction of the potential source, estimating the background and generating a set of plots on web pages.

The generated coincident trigger is then recorded into a database (named 'GraceDB'). Information coming out of the follow-up procedure is also recorded. The database is connected to an alert system which is common with the pipelines searching for burst events. The alert system was used by the burst searches while it was not yet enabled for coalescing binary searches at the time of this writing.

The telescopes that will do the electromagnetic follow-up need to know the direction of the source for as accurate as possible a pointing. During the internal follow-up, a sky localization procedure is activated.

2.3. Sky localization

The localization of the potential event in the sky is performed by triangulation based on the time of flight between the LIGO–Hanford (H1), LIGO–Livingston (L1) and Virgo (V1) detectors [17]. Only events presenting a 3-site coincidence are used. The procedure consists of a scan of the sky identifying the points the signal is most likely to come from (see figure 2).

Instead of using the end time of the templates coming out naturally from the MBTA pipeline, one uses the time when the signal crosses a determined reference frequency, of the order of 150 Hz [18]. This allows for better accuracy in the final localization result. Furthermore, a symmetry ambiguity in the localization arises when using only three detectors. An effective distance [19] measured at each detector helps to lift this ambiguity.

The performance of the sky localization procedure will be assessed by the use of hardware and software injections. First tests show a modest pointing accuracy of around 10° for signals at the detection threshold.

3. Performance of the pipeline

3.1. Latency

The pipeline was built with low latency in mind. Figure 3 shows a set of distributions corresponding to various steps of the pipeline. The final latency from the generation of data to the availability of coincident triggers for the follow-up procedure is less than 3 min.

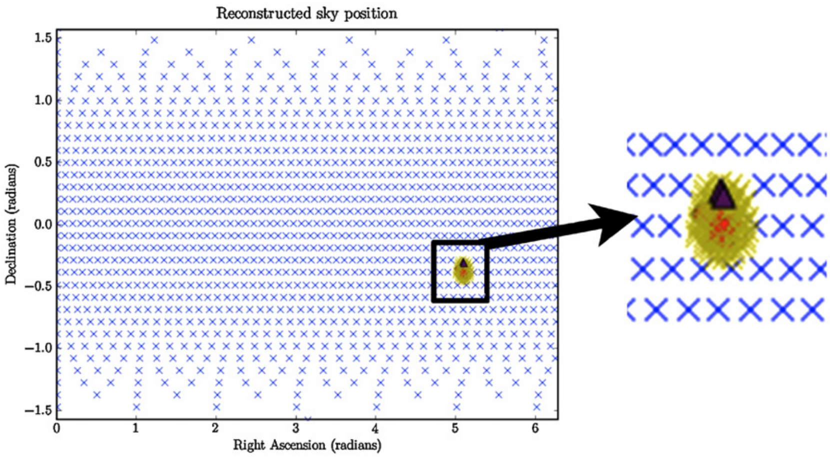

Figure 2. Given a 3-site coincident trigger (software injection), example of a scan of the sky searching for the direction having the highest probability. The gray area corresponds to a 90% confidence level. The dark triangle represents the injected position.

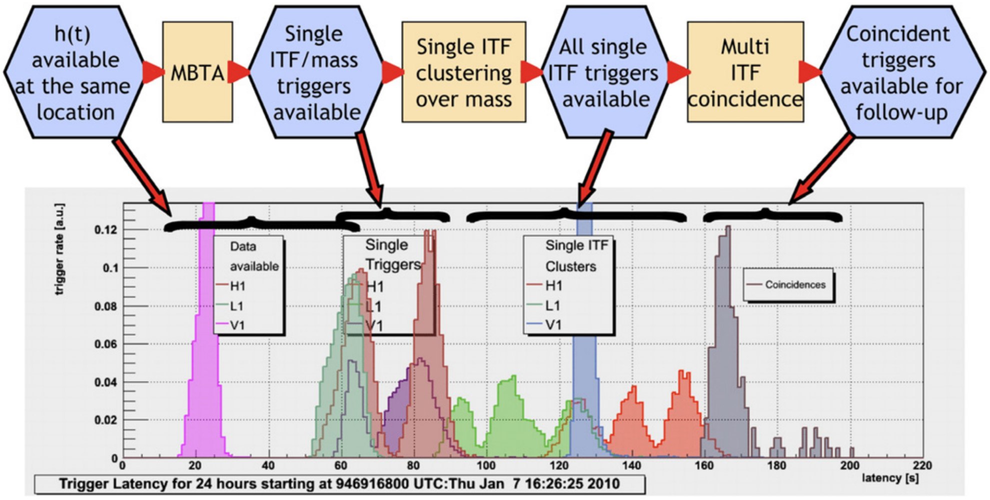

Figure 3. Latency distribution of triggers at various stages of the very low latency pipeline. The overall latency is less than 3 min.

3.2. Controlling the trigger rate

Latency is not the only variable to be considered for the electromagnetic follow-ups. The trigger rate is equally important in that it should not be too high. The goal is to tune thresholds for of the order of 1 trigger per month to be considered for possible follow-up by electromagnetic detectors.

The first obvious cut is on the time of flight between the three sites. The single-detector trigger time resolution was measured to be $\sigma_{t1}(\mathrm{SNR}) = \frac{5\ \mathrm{ms}}{\mathrm{SNR}}$. Taking into account the time of flight, the single-detector resolution at 3σ with one detector at each end of the line of flight,

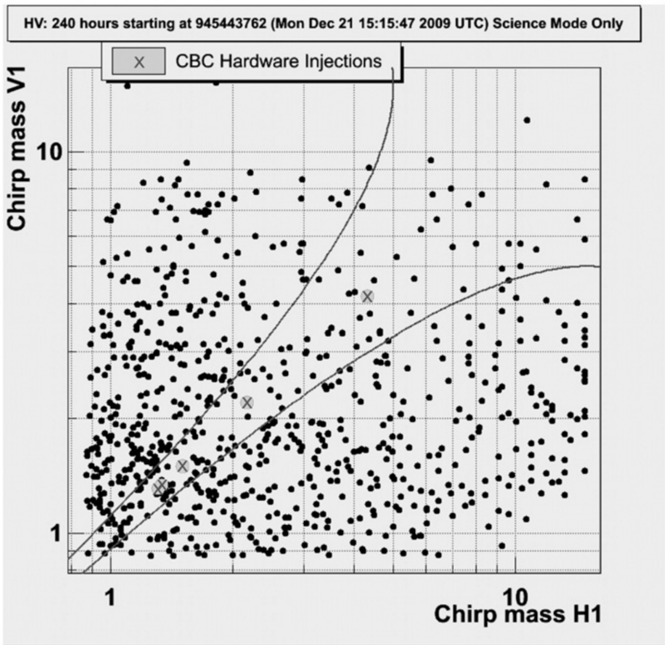

Figure 4. Example of the distribution of the chirp masses seen in each interferometer for double interferometer triggers, here for H1 and V1. The true signals should be consistent between the various detectors. This is shown as the injected simulated signals lay on the diagonal. The lines show the cuts applied in the procedure.

giving a factor $\sqrt{2}\,(3 \times \sqrt{2} \times \sigma_{t1}(6))$ and a safety margin, we understand why the coincidence windows are 20 ms between the triggers coming from the two LIGO sites, while it is 40 ms between the Virgo and LIGO detectors. The efficiency and performance of the pipeline were checked with software injections.

Starting from an observed single interferometer trigger rate of 0.1 Hz, it is possible to estimate the coincidence rates, and those rates are consistent with what is observed. Before applying any other cuts, the observed rates are $\sim$1.5 events h^{-1} for H1-L1 coincidences and $\sim$3 events h^{-1} for H1-V1 or L1-V1 coincidences.

With such double-coincident rates, the triple-coincident rate would be far too high. Thus, a consistency test was applied on the chirp masses detected for each trigger in the three interferometers. An example is shown in figure 4.

When applying the chirp mass cuts, the expected and observed triple-coincident H1-L1-V1 trigger rate is less than one coincident event per month. This rate is very low for estimating the background directly from triple coincidences. The background is estimated from single-trigger rates and the double-coincidence rate is used to check the consistency of this estimation.

3.3. *Figures of merit*

A few figure of merit plots are produced online and automatically reported on web pages. As an example, figure 5 shows the single-trigger rates for the three LIGO-Virgo interferometers. It can be seen that the average rate is around 0.1 Hz for L1 and V1 interferometers, while it is a little lower, around 0.06 for H1.

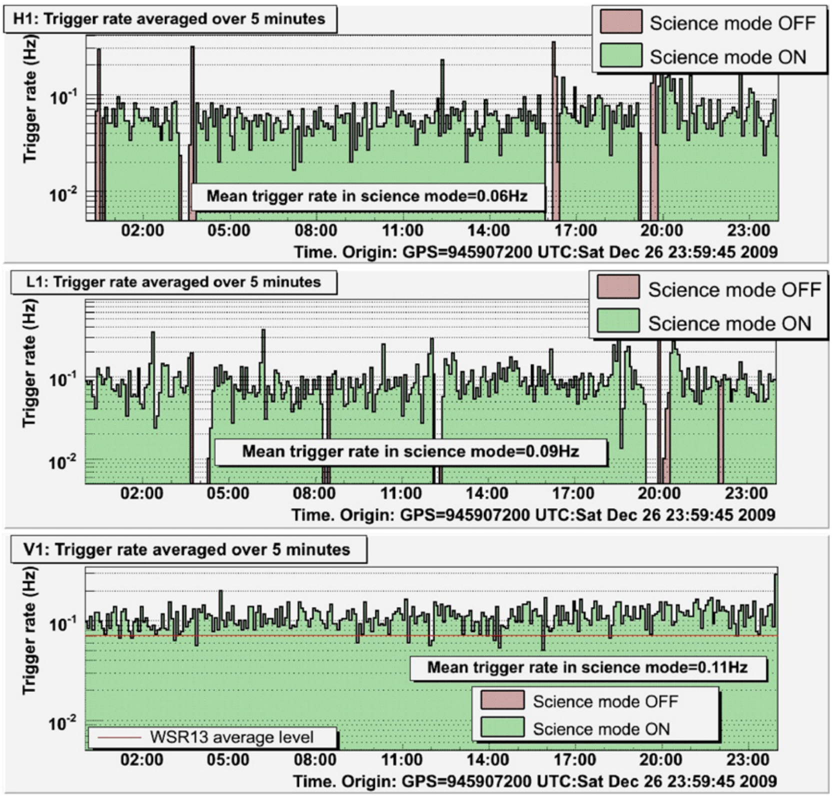

Figure 5. Example of figures of merit plots produced online by the low latency search pipeline. Rate distribution of single interferometers triggers as a function of time for each of the three detectors. H1 (top), L1 (middle), V1 (bottom).

Table 1. MBTA duty cycle for the first 2 months of running.

	H1	L1	V1	H1-L1	H1-L1-V1
MBTA duty cycle	94%	96%	97%	93%	92%

Figure 6 shows the predicted and observed triple-coincident rate as a function of time. During the period shown, there was one hardware injection of a simulated coalescing binary signal. The lower plot shows the combined H1-L1-V1 range, defined as the distance of a 1.4–1.4 $M_\odot$ binary coalescence seen with an SNR of 6 and averaged over the orientation of the system. The range was between 8 and 9 Mpc.

Another interesting figure of merit of the overall pipeline is the duty cycle. MBTA should produce triggers whenever it receives data corresponding to at least one detector in science mode. There is however a small dead time due to maintenance of the pipeline or data transfer inefficiencies between the sites. Table 1 shows the MBTA duty cycle for the first 2 months of running when the three single interferometers are in science mode as well as for the double H1-L1 and triple H1-L1-V1 combinations running in coincidence.

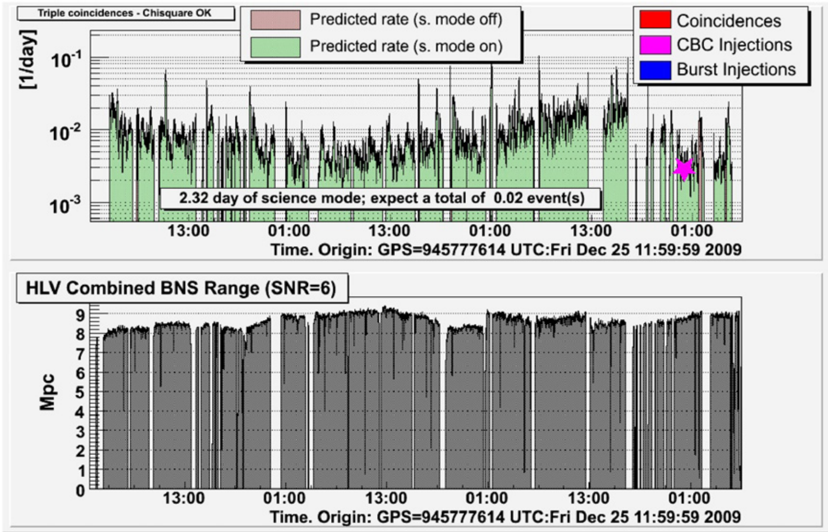

Figure 6. Top: rate of triple coincident events predicted from single-detector events. The star indicates a hardware injection of a coalescing binary simulated signal. Bottom: range reachable for the combined H1-L1-V1 network for a binary neutron star system seen with a SNR of 6 (average orientation).

4. Conclusions

Compact coalescing binaries involving a neutron star are potentially observable also as GRBs or as their afterglow. A joint observation in the gravitational waves and electromagnetic domain would be highly desirable. Gravitational wave candidates could trigger searches for an electromagnetic counterpart. A very low latency search may be a key point to achieve this goal.

A very low latency pipeline named MBTA was run during the common part of the S6 (LIGO) and VSR2 (Virgo) runs. A latency lower than 3 min until the availability of a trigger was achieved. The duty cycle of the pipeline for three detector coincidences was higher than 90% over the full run. The analysis pipeline was designed to obtain a rate of less than one triple-coincident event per month.

Our goal is to enable the trigger submission for electromagnetic follow-ups for the data taking period foreseen to start in the course of July 2010, at the restart of the Virgo detector.

Acknowledgments

The authors gratefully acknowledge the support of the United States National Science Foundation for the construction and operation of the LIGO Laboratory, the Science and Technology Facilities Council of the United Kingdom, the Max-Planck-Society and the State of Niedersachsen/Germany for support of the construction and operation of the GEO600 detector, and the Italian Istituto Nazionale di Fisica Nucleare and the French Centre National de la Recherche Scientifique for the construction and operation of the Virgo detector. The authors also gratefully acknowledge the support of the research by these agencies and by the Australian Research Council, the Council of Scientific and Industrial Research of

India, the Istituto Nazionale di Fisica Nucleare of Italy, the Spanish Ministerio de Educación y Ciencia, the Conselleria d'Economia Hisenda i Innovació of the Govern de les Illes Balears, the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research, the Polish Ministry of Science and Higher Education, the FOCUS Programme of Foundation for Polish Science, the Royal Society, the Scottish Funding Council, the Scottish Universities Physics Alliance, The National Aeronautics and Space Administration, the Carnegie Trust, the Leverhulme Trust, the David and Lucile Packard Foundation, the Research Corporation and the Alfred P Sloan Foundation.

References

[1] http://www.ligo.org
[2] Bradaschia C *et al* 1990 *Nucl. Instrum. Methods Phys. Res.* A **289** 518–25 (http://www.virgo.infn.it)
[3] Abbott B *et al* (LIGO Scientific Collaboration) 2009 *Phys. Rev.* D **79** 122001
[4] Nakar E 2007 *Phys. Rep.* **442** 166–236 and references therein
[5] Abbott B P *et al* 2010 *Astrophys. J.* **715** 1453
[6] Abbott B P *et al* 2010 *Astrophys. J.* **715** 1438
[7] Rhoads J E 2001 *Astrophys. J.* **557** 943
[8] Gehrels N *et al* 2004 *Astrophys. J.* **611** 1005
[9] Bringer M, Boër M, Peignot C, Fontan G and Merce C 1999 *Astron. Astrophys. Suppl.* **138** 581
[10] http://hepwww.physics.yale.edu/lasillaquest/
[11] http://grb.fuw.edu.pl/index.html
[12] Marion F (for the Virgo Collaboration) 2004 Multi-band search of coalescing binaries applied to CITF data *Proc. Rencontres de Moriond on Gravitational Waves and Experimental Gravity 2003*
[13] Owen B J 1996 *Phys. Rev.* D **53** 6749–61
[14] Wiener N 1949 *The Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications* (New York: Wiley)
[15] Droz S and Poisson E 1997 *Phys. Rev.* D **56** 4449
[16] Verkindt D 2007 *15th IEEE NPSS Real Time Conference*
[17] Fairhurst S 2009 *New J. Phys.* **11** 123006
[18] Acerneze F *et al* 2007 *Class. Quantum Grav.* **24** S617–25
[19] Abbott B *et al* 2004 *Phys. Rev.* D **69** 122001