

The Advanced LIGO timing system

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2010 Class. Quantum Grav. 27 084025

(<http://iopscience.iop.org/0264-9381/27/8/084025>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 131.215.220.185

The article was downloaded on 21/10/2010 at 16:16

Please note that [terms and conditions apply](#).

The Advanced LIGO timing system

Imre Bartos¹, Rolf Bork², Maxim Factourovich¹, Jay Heefner², Szabolcs Márka¹, Zsuzsa Márka¹, Zoltán Raics^{1,3}, Paul Schwinberg⁴ and Daniel Sigg⁴

¹ Department of Physics, Columbia University, New York, NY 10027, USA

² LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA

³ University of Debrecen, Debrecen 4034, Hungary

⁴ LIGO Hanford Observatory, Richland, WA 99352, USA

E-mail: ibartos@phys.columbia.edu

Received 25 November 2009, in final form 22 December 2009

Published 6 April 2010

Online at stacks.iop.org/CQG/27/084025

Abstract

Gravitational wave detection using a network of detectors relies upon the precise time stamping of gravitational wave signals. The relative arrival times between detectors are crucial, e.g. in recovering the source direction, an essential step in using gravitational waves for multi-messenger astronomy. Due to the large size of gravitational wave detectors, timing at different parts of a given detector also needs to be highly synchronized. In general, the requirement toward the precision of timing is determined such that, upon detection, the deduced (astro-) physical results should not be limited by the precision of timing. The Advanced LIGO optical timing distribution system is designed to provide UTC-synchronized timing information for the Advanced LIGO detectors that satisfies the above criterium. The Advanced LIGO timing system has modular structure, enabling quick and easy adaptation to the detector frame as well as possible changes or additions of components. It also includes a self-diagnostics system that enables the remote monitoring of the status of timing. After the description of the Advanced LIGO timing system, several tests are presented that demonstrate its precision and robustness.

PACS numbers: 95.55.-n, 95.55.Ym

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

1. Overview

The Laser Interferometer Gravitational Wave Observatory (LIGO) [1] is a network of interferometers aiming to directly observe gravitational waves. After finishing its initial-phase measurements (Initial LIGO), its enhanced version is currently in operation (Enhanced

LIGO). Its upgraded version, called Advanced LIGO [2], will consist of three detectors, each with 4 km arm length, two at the Hanford, Washington site, and one at Livingston, Louisiana. It is planned to begin operations in 2014. Advanced LIGO will exceed the sensitivity of Initial LIGO by more than an order of magnitude at its most sensitive frequency. While the sensitive bandwidth of Advanced LIGO will be increased compared to earlier versions, it will also include signal recycling cavities [3] that will enable the tuning of the sensitivity curve of the interferometers.

A network of detectors, besides lowering the noise floor compared to a single detector, provides increased sensitivity by being able to efficiently reject signal-like background events (called glitches) using the fact that detectors separated by large distances will have uncorrelated backgrounds, while real gravitational wave signals that are correlated appear coincidentally and have coherent waveforms in the detectors. Direction reconstruction, a crucial element of retrieving astrophysical source information from a measurement, relies on the accurate knowledge of the arrival times of gravitational wave signals to determine the source direction.

This paper introduces the Advanced LIGO timing system. Section 2 discusses the role of timing from the perspective of astrophysical measurements and data analysis. Section 3 introduces and describes the Advanced LIGO timing system that will provide timing for the Advanced LIGO detectors. We discuss the structure and specific components, the synchronization methods and system diagnostics. Finally, section 4 presents some performance tests of the Advanced LIGO timing system that demonstrate its precision in practice and its robustness against environmental and configurational changes.

2. Timing in gravitational wave astrophysics

Timing information of gravitational wave signals is vital for multiple purposes. First, coincident arrival times in different detectors separated by large distances is one of the main tools of background rejection. Another important purpose of precision timing is the reconstruction of source direction. Not only does it increase detection confidence by enabling comparison to possible celestial sources, but it is also essential for electromagnetic follow-up observations [4]. Furthermore, accurate direction reconstruction can greatly enhance multi-messenger, e.g., gravitational wave–high energy neutrino joint searches [5].

The accuracy of direction reconstruction largely depends upon the accuracy of timing (it further depends on the network geometry, the source direction compared to detector locations, uncertainty introduced by the analysis method and detector calibration) [4]. For instance if one precisely knows the arrival time of a signal for a network of three detectors, the signal direction can be precisely reconstructed with a degeneracy including the actual direction and its mirror image with respect to the plane of the three detectors (for some cases amplitude information can be used to break this degeneracy) [4].

Timing jitter, another crucial factor that can affect measurement precision, has special importance due to the intrinsic sensitivity curve of the Advanced LIGO detectors. Timing jitter introduced in the analog to digital converter (ADC) of a detector will result in the sampling process creating higher harmonics to the existing noise. Due to the steep negative slope of the Advanced LIGO noise curve (see, e.g., [2]) at low frequencies, the background at low frequencies below the interferometer's most sensitive frequency regime will be much higher than the background at the frequencies of its higher harmonics. This will result in the higher harmonics being much more emphasized compared to the case of uniform noise spectrum. It is therefore critical to reduce the appearance of higher harmonics by minimizing timing jitter⁵.

⁵ A detailed study of signal contamination due to timing jitter will be the topic of a follow-up article.

As an example to the precision of arrival time reconstruction, assume that a gravitational wave signal of frequency f_0 is incident on a detector. The uncertainty of the measured arrival time can be approximated as [6, 7]

$$\Delta t \simeq \frac{1}{2\pi f_0} \frac{1}{\text{SNR}} \quad (1)$$

where SNR is the signal to noise ratio. We will assume a signal of higher characteristic frequency of $f_0 \approx 1$ kHz, that can be produced by several sources interesting to Advanced LIGO, such as core-collapse supernovae, long-soft gamma-ray bursts, or rotating black holes [8]. With a signal to noise ratio of $\text{SNR} \approx 10$, we obtain $\Delta t \approx 15 \mu\text{s}$.

The Advanced LIGO timing system is required to provide timing information with uncertainty below $1 \mu\text{s}$. This uncertainty limit is safely below the requirements of arrival time or direction reconstruction. This requirement also allows for a possible 1% amplitude and 1° phase calibration—for the case of, e.g., a 1 kHz signal, this would mean $2.8 \mu\text{s}$ maximum allowed uncertainty. We note here that, between the actual timing of gravitational wave data and the timing signal from the Advanced LIGO timing system, there are other data acquisition subsystems, whose accuracy might also introduce timing/phase errors that will be added in the total timing error of gravitational wave data. It is therefore important that the precision of the Advanced LIGO timing system stays below the limit required by the considerations described above to ‘leave room’ for sources of uncertainty from other subsystems.

The provided timing signals are required to be synchronized to coordinated universal time (UTC) provided by the GPS satellite system. We note that the nominal accuracy of the GPS system provided by the US government is $\sim O(100 \text{ ns})$.

3. The Advanced LIGO timing system

Timing information for gravitational wave detectors requires different levels of synchronization. First, one needs to provide synchronous timing to sub-systems of a given detector, a challenging task due to the large sizes of the Advanced LIGO interferometers. Furthermore, timing needs to be synchronized between different detectors, as well as synchronized to an absolute time measure so that results can be compared between each detector as well as with non-gravitational wave measurements.

The Advanced LIGO optical timing distribution system (Advanced LIGO timing system) [9, 10] provides timing information to detector sub-systems. Timing is synchronized to coordinated universal time (UTC), this way insuring both relative and absolute synchronization⁶. The system distributes timing information in a modular structure. A central component, called the *master module*, synchronizes its internal clock to UTC, and distributes this synchronized timing through optical fibers to peripheral components, called *slave modules*. Slave modules synchronize their internal clocks to the received timing signal, and provide timing information based on their internal clock signal to external devices. The distribution of timing information is aided using intermediate components, so called *fan-out modules*. Fan-out modules synchronize their internal clock to the received timing from an upstream source, and distribute it to up to 16 destinations (either other fan-out or slave modules), enabling a highly flexible structure of arbitrary size. We note here that another

⁶ Optionally, common view GPS can be used to confirm timing precision if requirements make it necessary.

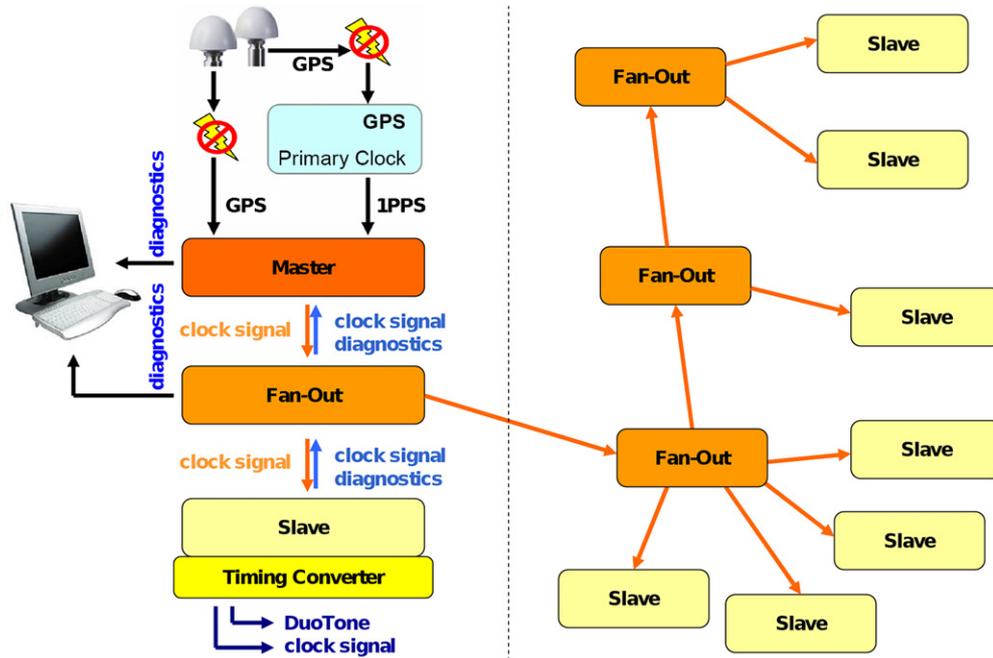


Figure 1. (Left) schematic structural diagram of the Advanced LIGO timing system. The master module is connected to a GPS receiver module that receives timing from a GPS antenna. Another, failover GPS antenna is connected directly to the master module. GPS antennae are connected to the GPS receiver and the master module through lightning arresters. The master module provides timing information to a fan-out module, while the fan-out module sends back diagnostics information. The fan-out module provides timing information to a slave module, while the slave module sends back diagnostics information (including the clock signal that is used to compensate for the time delay in the fiber). The slave module, together with a timing converter interface outputs timing information to an external device. The master and fan-out modules output diagnostics information to an external computer through the internet. (Right) example organization of Advanced LIGO timing modules in a star-like structure. Practically any number of fan-out modules can be serially connected, while one fan-out or master module can provide timing information to up to 16 fan-out or slave modules.

major element of providing flexibility to the Advanced LIGO timing system is the use of *field-programmable gate arrays* (FPGAs, type: Xilinx Spartan-3E) in each module. FPGAs are integrated circuits designed such that they can be easily (re-) programmed by the user or developer on site, without the need to re-manufacture the modules.

Below we describe the structure of the Advanced LIGO timing system by following the path of timing information throughout the system. The structure of the system as well as an example configuration are depicted in figure 1.

We start with the UTC time stamp that is provided by the GPS satellite network. UTC time is received by the master module through two paths. An external GPS receiver is used to receive and pre-process UTC timing and provide it as the *primary input* of the master module in the form of one pulse per second (1PPS) signal. The master module also has its internal GPS receiver sub-module (type: Oncore) that, upon being connected to a GPS antenna, will provide the master module's secondary, *failover* 1PPS timing input that is used if 1PPS is not received through the primary input.

3.1. Master module

The master module [9] uses the received 1PPS signal to which it synchronizes its internal clock. The core of the master module's internal clock is an *oven-Controlled crystal oscillator* (OCXO, type: Wenzel) that has excellent phase-noise characteristics and naturally operates on a 2^N Hz base frequency. The OCXO is synchronized to the external 1PPS through a phase-locked loop with a bandwidth of several millihertz, that filters out higher frequency fluctuations in GPS timing.

The master module provides synchronized timing information through optical fibers based on its internal clock signal, to up to 16 fan-out or slave modules. Timing information is transferred in the form of a 2^{23} Hz (~ 8.4 MHz) square-wave signal. The positive edges of this square-wave signal are spaced equally, and are used as the clock signal. The spacing of the negative edges is used to transfer additional information between the modules, such as the 1PPS signal, or diagnostics information (pulse-width-based binary encoding, see also section 3.5).

The time delay due to signal transfer in the optical fibers is automatically compensated for. The master module sends out a 1PPS signal to the downstream (fan-out or slave) module that, upon reception, immediately sends it back to the master module. The master module uses the measured time delay between these two signals to determine and compensate for the time delay due to signal transmission. This synchronization is done upon the connection of a new module, therefore its value does not change during proper operation. The 8.4 MHz clock signal is also adjusted to compensate for the delay due to the cable.

3.2. Fan-out module

Synchronized (and delay compensated) timing information is transmitted from the master module to a fan-out module [9]. Fan-out modules use the received timing information to which they synchronize their internal clocks. The core of a fan-out module's internal clock is a voltage-controlled crystal oscillator (VCXO, type: Vectron) whose oscillation frequency is synchronized to the received 8.4 MHz clock signal. The 1PPS of the fan-out module is determined based on the incoming 1PPS signal.

We note here that the fan-out and master modules are mostly identical, the only difference being the OCXO and GPS receivers that the master module has but the fan-out module does not.

The fan-out module, similarly to the master module, provides synchronized timing information of a 8.4 MHz clock signal and 1PPS signal, based on its internal clock signal, to up to 16 fan-out or slave modules. It also compensates for time delay in the optical fibers similarly to the master module.

3.3. Slave module

Synchronized timing information from the master module, after being transferred through an arbitrary number of fan-out modules, will be received by a slave module [9]. Slave modules, similarly to fan-out modules, use the received timing information to which they synchronize their internal clocks. Their internal clock is also identical to those of the fan-out modules.

Slave modules provide timing information to external devices, i.e. detector sub-components. They convert the received 8.4 MHz clock signal and 1PPS signal into the format required by the specific device they communicate with. Slave modules use *operation specific interfaces*: daughter boards that are placed onto the slave modules. While these interfaces provide the physical connection between the slave module and external devices,

operation specific timing information is generated within the (also operation specific) FPGA code programmed onto the slave module's FPGA chip.

Operation-specific interfaces include the timing converter interface described below, that will be used to synchronize the analog-to-digital and digital-to-analog converters, RF units and some of the CPU clocks of Advanced LIGO. Another interface is the *timing comparator interface*, that is used to (i) output UTC synchronized 1PPS signals from the slave module through multiple BNC ports and (ii) compare 1PPS signals from several different external devices to high precision. Further interfaces include the IRIG-B (inter-range instrumentation group B) interface, that will be used to communicate with devices that require IRIG-B input, such as some computer interfaces.

3.4. DuoTone generation

Advanced LIGO analog-to-digital and digital-to-analog converters require a two-tone signal to monitor their state of synchronization. The two-tone signal generated by the Advanced LIGO timing system, called DuoTone, is the sum of two sinusoids of 960 and 961 Hz. The frequency difference $\Delta f = 1$ Hz between its two components enables the precise determination of every 1 s mark by using the phase difference between the two sinusoids. While, for a 1PPS signal, one uses a few microseconds of data each second to determine the precise position of the 1 s mark (e.g. the rising edge of 1PPS), for a DuoTone signal, one can use the full 1 s long data stream for the same purpose. The narrow Fourier band width of a DuoTone signal is another important advantage that decreases possible cross-channel contamination. The LIGO gravitational wave channels can have significant narrow band noise at the 60 Hz power line harmonics. There can be a crosstalk between ADC channels; therefore, it was safe to place the frequencies of the timing signals on top of a high frequency line harmonics, where the crosstalk cannot cause an additional loss of bandwidth.

Advanced LIGO sub-components that require DuoTone signal, such as the analog-to-digital and digital-to-analog converters, will be connected to the timing system through a so called *timing converter interface* [9]. Besides a DuoTone signal, they also require other timing signals, which are shown and explained in figure 2.

3.5. Diagnostics

To monitor the operation of the Advanced LIGO timing system, diagnostics information is collected from each master module, fan-out and slave module, as well as about the structure and configuration of the system as a whole. Diagnostics information is transferred every second through the optical fibers connecting the modules, along with timing information (see figure 1) [11]. Data transfer is accomplished using pulse-width-based binary encoding. It utilizes the fact that only the positive edges of the 2^{23} Hz square-wave clock signal are used for synchronization purposes, leaving the negative edges available to carry additional information. Information bits are therefore encoded by shifting the negative edges by a quarter-cycle. Negative and positive shifts are both used to represent '1' bits, while no shift represents '0' bits. To account for the fiber receiving circuit's ac-coupling and to preserve dc-balance, positive and negative shifts are used one after the other for '1' bits. The only exception is the 1 s mark, which is differentiated from other information by having two consecutive '1' bit signs with the same shift (see [12] for further information on data encoding).

Data are collected by each master and fan-out module from all downstream components, and is sent to a computer using RS422 or ethernet connection (see figure 1). Assigned addresses of modules are used to construct a map of the physical fiber connections between modules.

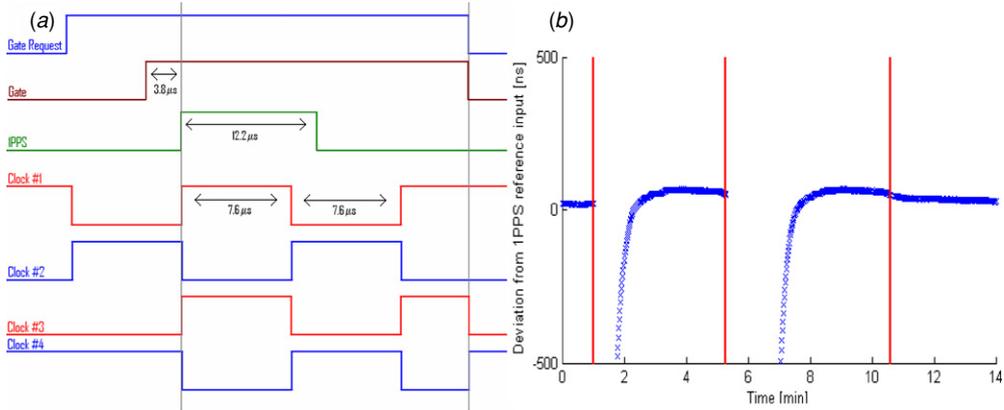


Figure 2. (a) Input/output signals of the timing converter interface (the DuoTone signal excluded), compared to 1PPS. The timing converter continuously outputs clock 1 and clock 2, two 2^{17} Hz clock signals which are the inverse of each other and are synchronized to 1PPS. The external device (i.e. an analog-to-digital converter), upon starting conversion at an arbitrary time, sends a gate request signal to the timing converter. The timing converter outputs a gate signal after gate request is turned on, starting right before the 1PPS signal such that it is closer to the 1PPS signal than the half cycle of clock 1. The timing converter also outputs two additional clock signals, clock 3 and clock 4, that are the logical products of gate and clock 1 and clock 2, respectively. (b) On/off test of the Advanced LIGO timing system: deviation between the input and output 1PPS signals of a master-fan-out-slave module chain as a function of time. The (red) vertical lines show the times of power recycle for the master module (first and second line) and the slave (third line) modules. The master module and fan-out modules were connected through a 4 km long fiber. The system restores itself and acquires precise timing within 3 min after the master module is restarted, and within 20 s after the slave module is restarted. Note that the fan-out module uses the same synchronizing mechanism as the slave module.

Diagnostics information further includes status and configuration information for each module. If needed, one can also send information to each module from a computer interface. Note that status information, besides signaling whether the system is up and running, also includes information on every error, e.g. temporary loss of signal, the time compensation applied for each connection due to delay in fiber, etc. For further information on data structure and processing, see [11].

4. Performance and tests

As discussed above, the Advanced LIGO timing system is required to provide sub-microsecond timing synchronized throughout a detector, as well as between different detectors. Besides the internal precision of the system, this sub-microsecond precision needs to include the uncertainty introduced by communication with external devices, i.e. timing reconstruction using the DuoTone signal output. Of further importance is the timing system's robustness against environmental and structural changes. In the following we present several test results that show the system's performance.

The Advanced LIGO timing system transfers timing information using a modular system, each module is synchronized to its upstream module. To demonstrate proper synchronization of the entire system, we tested the synchronization between the two ends of a typical modular configuration including all elements. The configuration took its input timing 1PPS from a UTC GPS receiver. A master module was connected and synchronized to this GPS 1PPS.

A fan-out module was connected to the master module, and a slave module to the fan-out module. The master and fan-out modules were connected with a 4 km long optical fiber, that allows for testing long-distance synchronization, as well as being the realistic setup for Advanced LIGO. We tested the synchronization of this GPS-master-fan-out-slave module network by comparing the 1PPS output of the slave module to the GPS 1PPS. Conducting a 24 h long measurement, we obtained a median of 14.3 ns and a 99% confidence interval of 6.2 ns (with maximum of 24 ns and minimum of 10.5 ns), showing that the uncertainty of internal synchronization is a small part of the total allowed error budget of 1 μ s.

Some external devices require a DuoTone signal for sub-second time stamping or diagnostics. To determine the precision of this time stamping, we compared the 1PPS signal reconstructed from the DuoTone signal and the actual 1PPS signal. We conducted an \sim 9 days long measurement, calculating the difference between the reconstructed and actual 1PPS signals once every hour. The obtained time difference was $\Delta t = -263 \pm 2$ ns. This result demonstrates that the DuoTone signal can be used to precisely reconstruct the timing information.

Besides timing precision, the Advanced LIGO timing system is also capable of quickly re-synchronizing itself if one changes its configuration by disconnecting or connecting some of its modules. To test this property, we used the configuration used for testing internal synchronization, i.e. a master-(4 km fiber)-fan-out-slave module chain. After allowing the system to synchronize, we power cycled the master module (twice) and then the slave module. The results can be seen in figure 2(b). After restarting the master module, the system re-acquired precise timing in about 3 min. This time frame is due to the low pass filter applied in synchronizing the master module's OCXO. After restarting the slave module, the system re-synchronized in about 20 s which is the time it takes for the FPGA chip on the board to reboot itself and the time of measuring the signal time delay in the fiber that connects the slave module to the rest of the timing system. We note that the same reboot durations apply for components that are newly connected to the system. This measurement demonstrated that the system is robust against power cycling or disconnection of its components, which provides additional safety and easy access to the system, enabling the quick addition and/or removal of elements.

5. Summary

The Advanced LIGO optical timing distribution system provides UTC synchronized time-stamping throughout the Advanced LIGO detector, as well as to a network of detectors to sub-microsecond precision in hardware. This accuracy guarantees that uncertainties in arrival time and direction reconstruction will be unaffected by synchronization uncertainties. Sub-microsecond precision also enables a possible 1% amplitude and 1 $^\circ$ phase calibration of the Advanced LIGO detectors. The self-diagnostics system of the Advanced LIGO timing system allows for remotely monitoring status and structural information that enables the traceability of the modular structure, as well as possible errors. We demonstrated that the Advanced LIGO timing system robustly recovers after power cycling or after the (dis-) connection of components.

Acknowledgments

The authors would like to thank Keith Thorne, David Strom and Raymond Frey for their valuable comments.

References

- [1] Abbott B P *et al* 2009 LIGO: The Laser Interferometer Gravitational-Wave Observatory *Rep. Prog. Phys.* **72** 076901
- [2] Fritschel P 2003 Second generation instruments for the Laser Interferometer Gravitational Wave Observatory (LIGO) **4856** 282 (arXiv:[gr-qc/0308090v1](https://arxiv.org/abs/gr-qc/0308090v1))
- [3] Meers B J 1988 Recycling in laser-interferometric gravitational-wave detectors *Phys. Rev. D* **38** 2317–26
- [4] Fairhurst S 2009 Triangulation of gravitational wave sources with a network of detectors *New J. Phys.* **11** 123006
- [5] Aso Y, Marka Z, Finley C, Dwyer J, Kotake K and Marka S 2008 Search method for coincident events from ligo and icecube detectors *Class. Quantum Grav.* **25** 114039
- [6] Sigg D 2003 Strain calibration in LIGO *Technical Note LIGO-T970101-B-D*
- [7] Cavalier F, Barsuglia M, Bizouard M-A, Brisson V, Clapson A-C, Davier M, Hello P, Kreckelbergh S, Leroy N and Varvella M 2006 Reconstruction of source location in a network of gravitational wave interferometric detectors *Phys. Rev. D* **74** 082004
- [8] Abbott B P *et al* 2009 Search for high frequency gravitational-wave bursts in the first calendar year of LIGO's fifth science run *Phys. Rev. D* **80** 102002
- [9] Bartos I, Bork R, Factourovich M, Heefner J, Márka S, Raics Z, Schwinberg P and Sigg D 2009 Timing system document map Advanced LIGO *Technical Note LIGO-E090003*
- [10] Bartos I, Bork R, Factourovich M, Heefner J, Márka S, Raics Z, Schwinberg P and Sigg D 2008 Quick start guide optical timing distribution system Advanced LIGO *Technical Note LIGO-E080541*
- [11] Bartos I and Sigg D 2009 Diagnostics information for the Advanced LIGO optical timing distribution system *Technical Note LIGO-E0900036*
- [12] Factourovich M 2007 Timing synchronization for Advanced LIGO timing system *Technical Note LIGO-T070218*