LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -CALIFORNIA
INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Refereed paper	**LIGO-P010043-00-E**	19 December 2001

Network Data Analysis Server (NDAS) prototype development

(as published in *Classical and Quantum Gravity*)

Szabolcs Márka (LIGO/CIT), Benoit Mours (VIRGO/LAPP) and Roy Williams (CACR/CIT)

Distribution of this draft:
LIGO

Reference:

"Network Data Analysis Server (NDAS) prototype development" S. Márka, B. Mours and R. Williams, Proceedings of 2001 Amaldi Meeting, Perth, Australia; Classical and Quantum Gravity, (2002)

California Institute of Technology
LIGO Project - MS 18-34
Pasadena CA 91125
Phone (626) 395-2129
Fax (626) 304-9834
E-mail: info@ligo.caltech.edu

Massachusetts Institute of Technology
LIGO Project - MS 20B-145
Cambridge, MA 01239
Phone (617) 253-4824
Fax (617) 253-7014
E-mail: info@ligo.mit.edu

WWW: http://www.ligo.caltech.edu/

INSTITUTE OF PHYSICS PUBLISHING CLASSICAL AND QUANTUM GRAVITY
Class. Quantum Grav. **19** (2002) 1–4 PII: S0264-9381(02)29467-8

Network data analysis server (NDAS) prototype development

Szabolcs Márka[1], **Benoit Mours**[2,3] **and Roy Williams**[4]

[1] LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA
[2] LAPP Chemin de Bellevue – BP 110 74941 Annecy-le-Vieux Cedex, France
[3] LIGO Laboratory, California Institute of Technology, Pasadena, CA 91125, USA
[4] CACR, California Institute of Technology, Pasadena, CA 91125, USA

E-mail: smarka@ligo.caltech.edu, mours@lapp.in2p3.fr and roy@cacr.caltech.edu

Abstract
We have developed a simple and robust system based on standard UNIX tools and frame library code to transfer and merge data from multiple gravitational wave detectors distributed worldwide. The transfer and merger take place with less than 20 min delay and the output frames are available for all participants. Presently VIRGO and LIGO participate in the exchange and only environmental data is shared. The system is modular to allow future improvements and the use of new tools like grid.

1. Introduction

Interferometric gravitational wave detectors[5], capable of detecting various astrophysical sources, eventually at cosmological distances, are in advanced stages of their construction worldwide. The clear long-term goal of the international gravitational wave community is to operate these detectors in unison and share data freely and rapidly. Besides logistics, it is also extremely important from the viewpoint of achievable science: to extract maximum information from observations, to conduct efficient coincidence analysis, to determine source direction with the ultimate goal to provide this information to other detectors with pointing capability, etc. Sharing the large amount of sensitive data in near real time is not a trivial task even when considering only the technical difficulties and neglecting the organizational ones. It is mandatory to have a well-tested and robust prototype system by the time real gravitational

[5] LIGO: http://www.ligo.caltech.edu/
VIRGO: http://www.virgo.infn.it/
TAMA: http://tamago.mtk.nao.ac.jp/tama.html
GEO: http://www.geo600.uni-hannover.de/
ACIGA: http://www.anu.edu.au/Physics/ACIGA/

0264-9381/02/000001+04$30.00

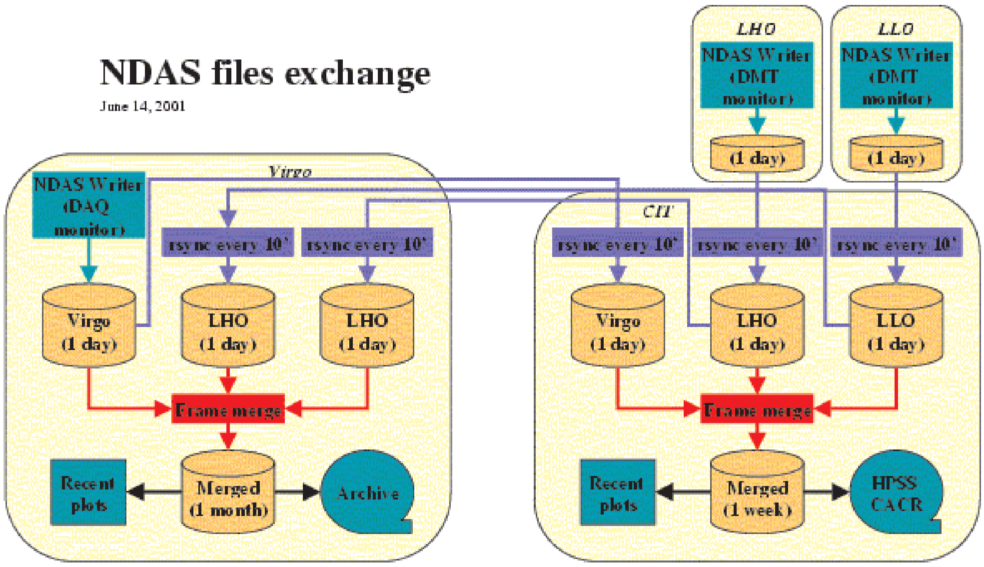

Figure 1. Anatomy of the network data analysis server prototype configured for the two LIGO and the VIRGO observatories.

wave data is available from multiple observatories. It is very likely that the evolution of this system will take a considerable amount of time and effort, therefore it is advisable to start this endeavour as soon as possible.

2. Overview

Inspired by the above arguments we developed a very simple, portable, secure, robust, stable and scalable system based on freely available UNIX tools. Due to the standardized nature of the software components, this system can function under all flavours of UNIX and these different implementations can be mixed and can still work well together. The system incorporates multiple buffering to eliminate possible data loss due to network problems. Since none of the observatories will provide high duty cycle gravitational-wave data on the short term, we decided to use environmental data for our prototype system. Besides allowing a first look at correlated noises, environmental data has several advantages at the moment. It is available from most of the gravitational-wave observatories worldwide. The data is produced by well-characterized commercial sensors. Modest data rate is acceptable and will spoil neither the initial bandwidth available for the observatories nor the value of the scientific analysis. Environmental data is also less sensitive and can be shared internationally much easier than the gravity wave data. These arguments strongly support our decision to limit the amount of shared data to 10 Kb s^{-1} for each node (observatory) and only share the scientifically interesting channels. Presently we are exchanging power line monitor, seismic and magnetometer data, since they are the most likely channels to indicate real or accidental intersite coincidence of environmental artefacts.

3. System anatomy

Figure 1 shows the block diagram of the NDAS system when configured for communication between the LIGO and VIRGO observatories.

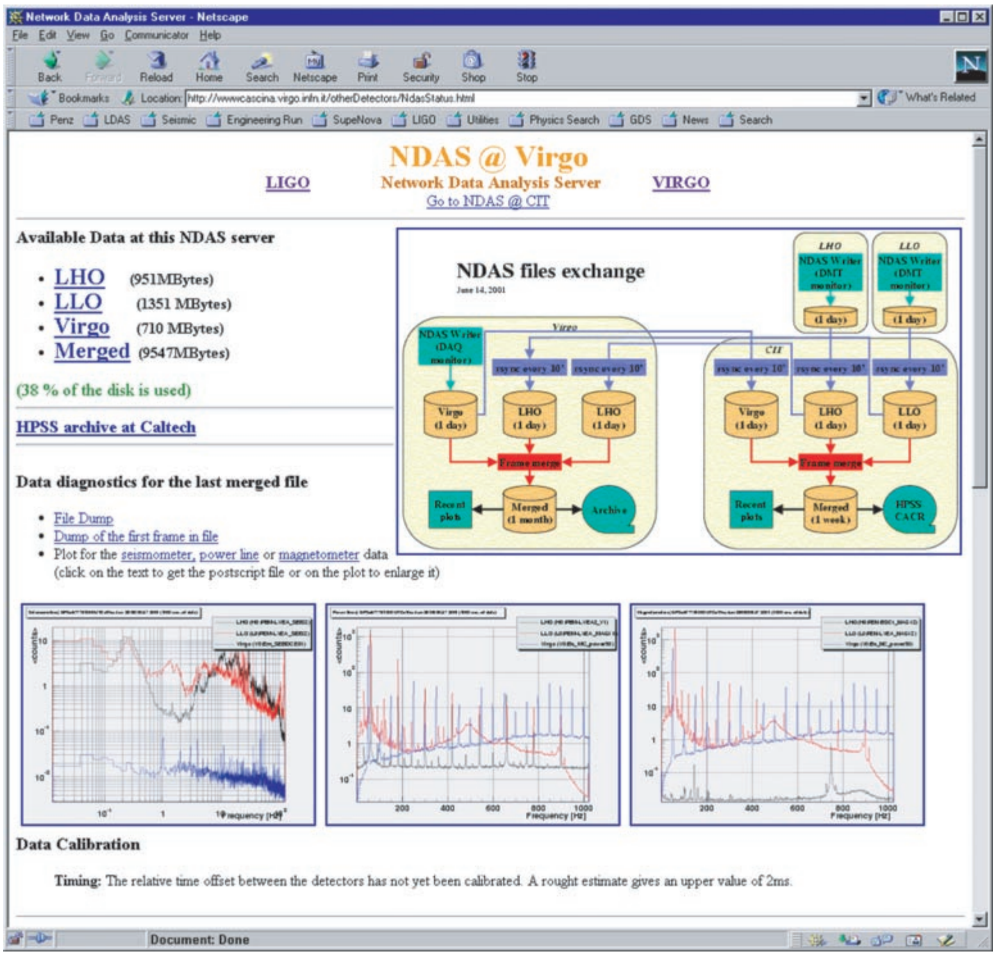

Figure 2. Screenshot of a typical moment of the NDAS diagnostic web page.

Each node observatory has a semi-custom process (NDASWriter) running on a computer with real time access to the data. The desired channels are resampled to a 2^n rate to ensure compatibility among the sites and saved in a standard frame format on a local disk. The local disk serves as a circular buffer and it is large enough to store data at least for a day. Therefore, short network outages and transient bandwidth problems will not lead to data loss. At the same time, each local buffer is accessible to a single NDAS process running on a CIT (CalTech) computer. This process reaches out and transfers the recently created data frames to its local disk. This disk also serves as a circular buffer and caters out each data frame to participating observatories. The VIRGO and CIT machines presently have the capability to store individual frames for one day. These machines also create a new merged dataset for each time segment, where data is available from all participating observatories. These merged datasets can be stored locally for an interval ranging from couple of days to a month. After that the merged sets must be archived. Following the merging process, a detailed state of health and diagnostic information is generated, which is posted on the web in real time. Presently, CIT and VIRGO feature such an active

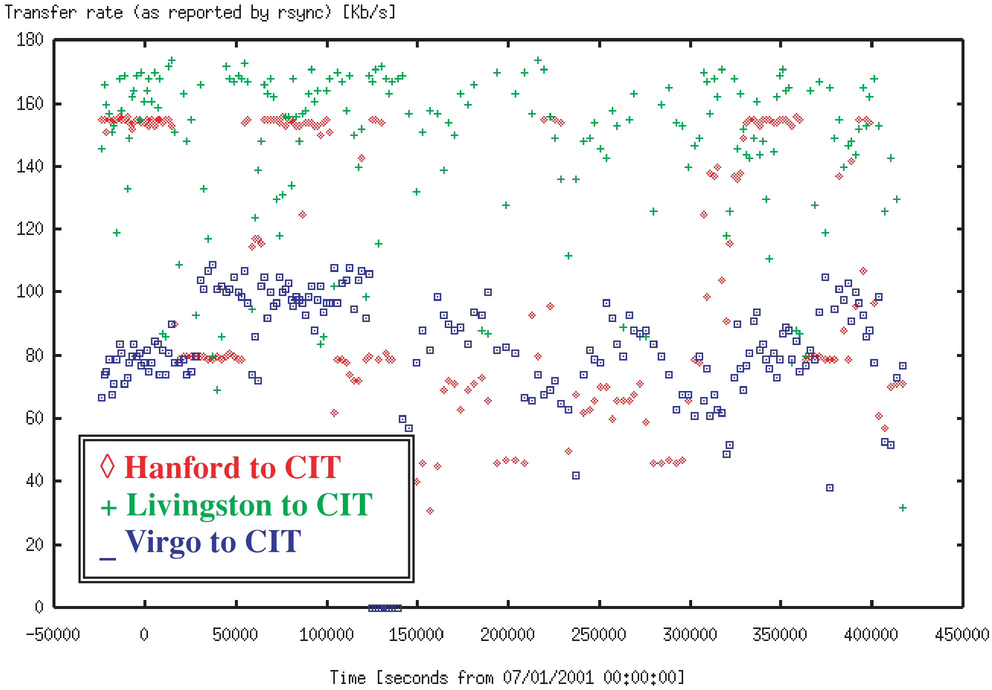

Figure 3. Transfer rate log for a typical period. Note the large daily variation due to uneven load of the network.

NDAS diagnostic web page, which is available at http://www.ligo.caltech.edu/∼ndas/ and http://wwwcascina.virgo.infn.it/otherDetectors/NdasStatus.html. The screenshot of a page taken at a typical moment is shown in figure 2. The transfer rate log for a typical period is shown in figure 3.

4. Summary

We developed a very simple, but robust and scalable system based on freely available UNIX tools to transfer and merge data from multiple gravitational wave detectors in near-real time. Our system is up and running and we already have significant amount of shared and merged data from the two LIGO and the Italian–French VIRGO observatories. The expansion of the network is under way with GEO joining in at the moment and interest is shown from our Australian colleagues. We plan to survey the advantages and implement prototype systems using grid technology in the near future.

Acknowledgments

We would like to acknowledge the support of Albert Lazzarini and Barry Barish of LIGO Laboratory, the LIGO laboratory and the VIRGO Collaboration, which made it possible to create this system. LIGO is funded by the National Science Foundation under the Cooperative Agreement PHY-9210038.