Data Archiving and Distribution of the Virgo Antenna for Gravitational Wave Detection

F. Barone[1,2], A. Eleuteri[2], F. Garufi[2], L. Milano[1,2]

[1] Dipartimento di Scienze Fisiche dell'Università di Napoli "Federico II", Complesso Universitario di Monte S.Angelo, via Cintia, I-80126 Napoli, Italy

[2] Istituto Nazionale di Fisica Nucleare – Sezione di Napoli, Complesso Universitario di Monte S.Angelo, via Cintia, I-80126 Napoli, Italy

Abstract

We describe the architecture of the Data Archiving and Distribution of the Virgo antenna for gravitational wave detection. The main characteristic of this system is the modularity of the architecture. This solution allows system upgrades without dramatic changes of the hardware and software components. The main performances are:

1. maximum sustained data flow of 10 *Mbyte/s* on DLT tapes (35/70 *Gbyte*) for the raw data archiving;
2. up to 1 *Tbyte* data archiving capacity on disk at a maximum sustained data flow of 25 *Mbyte/s* for the on-line data distribution;
3. up to 10 *Mbyte/s* data retrieval flow for the on-line data distribution.

The basic architecture of the system consists of two sections: an acquisition and storage section and a data management section. The former is a LynxOS based system with the disks directly connected to the CPU slave boards, the latter is a DEC-Unix Alpha Server (Data Server), NFS mounting the LynxOS disks through a Fast Ethernet network.

I. INTRODUCTION

Gravitational Wave detection is certainly one of the most challenging goals for today's physics. In fact, it will be not only the first direct proof in favour of the Einstein's General Relativity description of phenomena related to the dynamics of gravitation, but also the opening of a completely new channel of information on astrophysical objects [1], [2]. For this task many detectors are already operational or are planned around the world with different measurement band and sensitivities, both at Earth (resonant mass detectors, laser interferometers, etc.) and in space (spacecraft tracking, space interferometers, etc.) [2].

Despite the great efforts produced in the last 30 years in this field, there is not yet direct evidence confirming the existence of gravitational waves (GW). The only credited indirect evidence of GW emission is the famous two-neutron star binary pulsar PSR1913+16 discovered in 1975 by Hulse and Taylor. This binary is decaying due to the loss of orbital energy emitted as gravitational waves exactly at the rate predicted by the General Relativity (better than 1% accuracy) [3].

A great improvement in the direction of direct GW detection will be given by the new generation of very long baseline (up to ~ 4 km armlength) Earth-based interferometric detectors. These detectors may reach very high sensitivities (from $h \sim 10^{-24}$ to $h \sim 10^{-26}$ for integration times of the order of 10^7s) and measurement bands (from few Hz to many kHz) [1], [2]. These characteristics make the interferometers suitable for GW detection from different classes of astrophysical objects (e.g. radiation bursts from supernovae, periodic signals from old and new pulsars, radiation from coalescing compact binaries, etc.) [1], [2]. The construction of many of these large detectors (GEO600 [4], LIGO [5], TAMA [6], VIRGO [7], etc.) has already started and we expect them to be fully operational at the beginning of the next century.

In particular, the VIRGO interferometer has been designed to reach a sensitivity of the order of $\tilde{h} = \Delta L / L \sim 10^{-21} / \sqrt{Hz}$ at 10 *Hz* (being L the armlength) and $\tilde{h} \sim 3 \cdot 10^{-21} / \sqrt{Hz}$ at 200 *Hz* and a detection band that spans from ~4 *Hz* up to 6 *kHz*.

The drawback of such a large measurement band is the large amount of data to be archived. In fact, such a large detection band led to the choice of a data sampling frequency of f_s=20 *kHz* (due to aliasing problems). Then, assuming 16 bit accuracy of the data and the above quoted sampling frequency of 20 *kHz* we get a minimum data flow of ~80 *kbyte/s* (including the time recording). This amount of data would be still easily manageable, but, unfortunately, the bare acquisition of the interferometer output would not be sufficient for a correct data analysis (and for the debugging of the Virgo antenna in the initial phase). In fact, due to its large sensitivity, the Virgo antenna will be largely sensitive both to environmental (temperature, seismic, acoustic, electromagnetic, etc.) and to internally generated noises (vacuum pumps, analog and digital controls, etc.). In addition, the general status of the interferometer has also to be monitored. The identification and correlation of these noises with the interferometer output requires the acquisition of many channels at the same sampling frequency of the signal. If we estimate the number of acquisition probes in ~50, then, a continuous data flow of 2 *Mbyte/s* has to be expected.

As a consequence, we may expect a continuous raw data flow ranging from 1 to 5 *Mbyte/s*, with peaks of 10 *Mbyte/s* in burst mode. The latter value takes into account the possibility of a dynamic change in the number of the acquisition channels. In fact, it may happen that some noise sources affect the output signal only if they overcome a certain threshold (e.g. the electromagnetic noise in presence of lightning, etc.). In this case a real-time acquisition of these noises must start and all these data must then be archived and distributed to study both the machine performances and

0018-9499/00$10.00 © 2000 IEEE

the effects of all the noises on the interferometer output. When the Virgo antenna is completely debugged, these data will be used for gravitational wave data analysis.

The global architecture of Virgo data acquisition and storage is described in the following. The raw data coming from the probes and the detectors are collected by Local Readout Systems and formatted in frames by the Frame Builder. These frames are then sent both to the Raw Data Archiving System, that stores them on DLT tapes (35/70 *Gbyte*), and to the On-line Processing system. This latter system processes the frames adding them the structures containing the reconstructed [*t*, *h*] pairs, and all the auxiliary information coming from the Data Quality and the Global Control Systems, necessary for the assignment of a quality coefficient to the acquired data [7]. The On-line Processing then provides a real-time data pre-processing and selects all the frames likely to contain a gravitational wave event. The selected frames are then sent to the Data Distribution system that archives them both on disks and on DLT tapes (Data Summary Tapes - DST), for distribution to Virgo users.

The block scheme of the Virgo data flow and all the logical links among the subsystems involved in data acquisition, archiving, distribution and data analysis are shown in Fig.1.

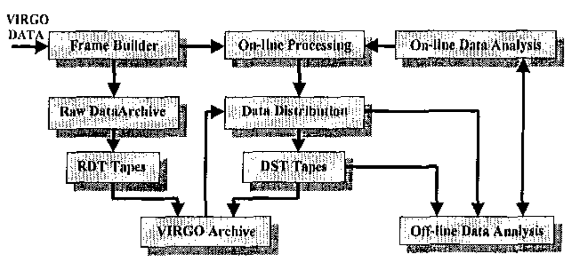

Fig.1: Scheme of principle of the Virgo Data Flow Archiving and Distribution.

A frame is a structured unit of information containing all the raw data necessary to understand the behaviour of the interferometer over a finite time interval including several samplings. Each frame can be divided into three main parts:

1. Structures filled by the Frame Builder, containing all the raw data collected by detectors and probes;
2. Structures filled by the on-line processing (or by off-line reprocessing), containing the reconstructed data ([*t*, *h*] pairs sampled at 20 *kHz*) and all the auxiliary information;
3. Structures filled by the simulation, useful to compare the interferometer behaviour with the models.

We have designed a special architecture for Data Archiving and Distribution. It is composed of two main subsystems: the Raw Data Archiving System (RDA) and the Data Distribution System (DD). The RDA stores all the raw data on tapes while the DD stores the pre-processed data on disks and on tapes (Data Summary Tapes–DST).The systems requirements are summarised in Table 1 and Table 2.

In this article, we will describe the architecture and the performances of both the RDA and DD.

Table 1
Raw Data Archiving system requirements.

Data flow rate (*Mbyte/s*)	**Mode**
1÷5	Continuous mode
<10	Burst mode
<10	From the Frame Builder

Table 2
Data Distribution system requirements.

Data flow rate (*Mbyte/s*)	**Mode**
<1	Continuous mode
<10	Burst mode
Data transfer rate (*Mbyte/s*)	**Data flow**
<10	From the on-line processing
<1	To users via internal network

II ARCHITECTURE OF THE SYSTEM

A. Raw Data Archiving System

The Raw Data Archiving System stores on RDT (Raw Data Tapes) all the data acquired by the Virgo detectors and controls (raw data), collected by the distributed Local Readout processes and Slow Monitoring Systems and structured in frames by the Frame Builder. These data are the real output of the interferometer and are necessary for any data reprocessing starting from the original data. According to the requirements on the continuous data flow rate (see Table 1), the amount of data to be archived spans from 86.4 *Gbyte/day* (1 *Mbyte/s*) to 432 *Gbyte/day* (5 *Mbyte/s*). For this task we use DLT tapes (Digital Linear Tape – 35/70 *Gbyte* writing speed ~5 *Mbyte/s*). Therefore, assuming a maximum continuous data flow of 5 *Mbyte/s* and the *35 Gbyte* storage capacity of DLT, about 4500 DLT/year are necessary to maintain a Raw Data Archive.

The RDA is organised according to client/server architecture, as shown in Fig.2.

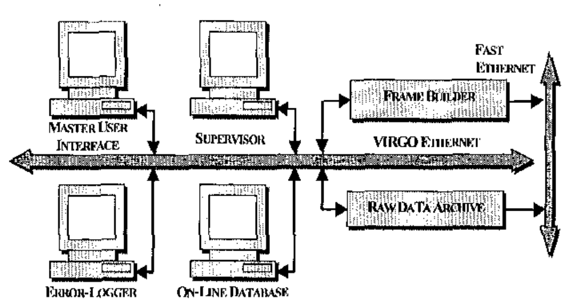

Fig.2: The Client/Server architecture of the Virgo Raw Data Archive.

The RDA server is configured and controlled by its User Interface (client). The system configuration is stored in an

On-line Database while an Error Logger records all the error messages generated the Raw Data Archive using the Cm package for communication [8]. A Supervisor may act as a User Interface [7].

In order to match the data flow to the DLT writing speed, we implemented a two-stage modular storage procedure consisting in the parallel staging of the data on disks (writing speed >10 *Mbyte/s* and storage capacity up to 18 *Gbyte*), and the subsequent copy of the data on DLT. The RDA server basic scheme is shown in Fig.3.

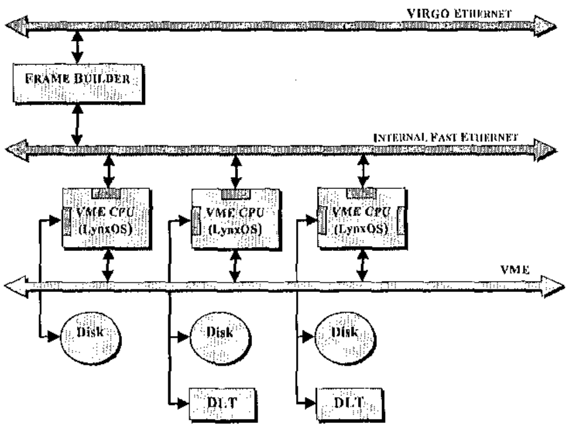

Fig.3: Architecture of the Virgo Raw Data Archiving Server.

The system is composed by a master CPU, running the LynxOS operating system, and two slave CPUs, each provided with disks and connected to a DLT autoloader, both housed in a dedicated VME crate. Through an internal Fast Ethernet network, the Frame Builder sends the frames to the master CPU that dispatches them through the VME bus to the first slave CPU enabling it to write on its disks. When these disks are full, the master CPU enables the second CPU to write on its disks, while the first one starts to download the acquired data on its DLT tapes. When it has finished, it changes the cartridge and waits until the master CPU enables it to write again. The procedure requires an adequate number of CPUs to match the data flow to the sustained transfer rate of the DLT drive.

This architecture easily allows reconfigurations and expansions of the system on the basis of the real data flow, and, if the disks and/or on-board random access memories (RAM) are used as buffer, to manage higher data flows for a limited amount of time.

In our case, with a data flow of 5 *Mbyte/s* and a sustained transfer rate on DLT of about 5 *Mbyte/s*, two CPUs are needed, each provided with 35 *Gbyte* of disks. In this configuration a full cycle lasts 4 hours. In fact a DLT cartridge (corresponding to ~2 *h* of data) takes 2 *h* to be filled and the data are written again on the same disk after 4 *h*. Each DLT autoloader houses 14 DLT cartridges. Therefore, with two autoloaders, each with a storage capacity of 490 *Gbyte* (980 *Gbyte* compressed mode), it is necessary to change the cartridge magazines at most every ~27 *h* (54 *h*).

The implemented Virgo RDA system consists of a 21 slot ELMA VME crate with a master CPU (VMPC4a PowerPC604e from Cetia) running LynxOS v3.0 operating system and linking the system to the Fast Ethernet network, two slave CPUs (same model) each provided with two 18 *Gbyte* SCSI disks (Cheetah from Seagate) and two 14 cartridge DLT autoloaders (DLTstor7114 from Quantum). This configuration leaves the autoloaders enough dead time to rewind and to load the new tapes.

B. Data Distribution System

The DD archives all the frames selected by the on-line data analysis algorithms (i.e. the frames containing also the structures filled by the on-line processing and by the simulation). These frames are the most likely to contain the gravitational wave events, thus needing a deeper analysis. To this task, besides the [*t, h*] pairs, the DD stores the necessary auxiliary information on all the frames at 20 *kHz* sampling rate. Unfortunately, again, the large amount of data limits the number of frames that can be maintained on-line. In fact, assuming a mean data flow of 150 *kbyte/s*, for the selected frames, we get 12 *Gbyte/day* and 10 *Gbyte/day* for the [*t, h*], including all the auxiliary information for data analysis. For this reason the DD has been split into two parts: on-line and off-line section.

The on-line section allows the retrieval of all the data produced by Virgo in the last month of run, through the network. This task requires to maintain at least 500 Gbytes on line (for a data flow of 2 *Mbyte/s*). At the same time, the content of the on-line data distribution is written on data summary tapes (DST) and distributed on request to the Virgo laboratories. To optimise the data distribution and analysis, special DST tapes are foreseen. In this way no limitation is imposed to the off-line data analysis, since whenever a raw data reprocessing is needed, the RDT can be loaded and reproduced by the data distribution system (off-line reprocessing) using the same tools used for the on-line processing. To this task, a list of contents of the data stored on RDT and DST is maintained on line on a data distribution database together with an archive of all the environmental parameters, necessary to check the Virgo environmental status.

In synthesis the DD collects the data produced by the on-line and stores them on disks (on-line DD) and on DLTs (off-line DD - DST). These data contain all the information useful for the off-line analysis, i.e:

1. Reconstructed *h* values ([*t,h*] pairs with quality coefficients and all the necessary auxiliary information-max data flow ~100 k*byte/s* at 20 *kHz)*.
2. Slow monitoring records. These data are organised in an on-line database (historical monitoring) and can be accessed at any time for checks and analyses of the environmental quantities, with a maximum continuous data flow of 5 *kbyte/s*.

3. Frames selected by the Virgo real-time analysis algorithms, running in parallel on the on–line data processing and containing all the raw data required for a full analysis of the signal candidates (max. continuous data flow of 150 *kbyte/s* for an event selection of 10 days over a data flow of 5 *Mbyte/s*).
4. Raw data, retrieved from the archived RDTs, to be reprocessed by the DD, in order to reconstruct the full frame, or data retrieved from DSTs.

The data in DD are hierarchically organised in a tree structure as shown in Fig.4; the frame directory, the slow monitoring variables and the auxiliary information are stored in a database.

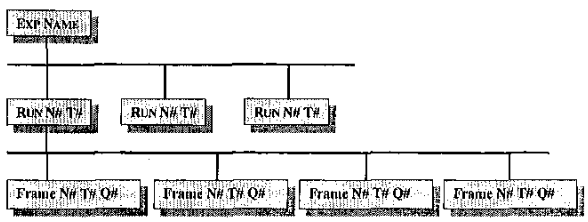

Fig.4: Structure of the archived data in the data distribution system; N# is the number, T# the time and Q# the quality factor.

Only the last month of run data is stored in ~500 Gbytes of disks (30×18 *Gbyte* fast & wide SCSI disks), while all the data are stored on DSTs. The total amount of these data can be calculated as follows:

1. For the reconstructed *h* data and slow monitoring records at 20 *kHz*, ~3.6 *Tbyte/year* are required;
2. The frames with slow monitoring records and [*t*, *h*] pairs, resampled at 2 *kHz*, require a maximum amount of ~360 *Gbyte/year*;
3. The frames produced by the real-time analysis algorithms running on the on-line processing system require a maximum amount of ~4.4 *Tbyte/year*.

In synthesis a maximum amount of ~8.4 *Tbyte/year* of data on DSTs is required, corresponding to 240 DLT/year (120).

All the archived data are available to the authorised users, via the standard network or directly on DST, for the off-line analysis. They can retrieve them by means of the Software Tools – Data Distribution [7], using the list of contents of DST and RDT maintained on the DD.

Also the DD is organised according to a client-server architecture, as shown in Fig.5. A master user interface (client) configures and controls the DD server using the configuration stored in the on-line database, checks the errors (recorded by the error logger system), accesses the data to change their structure, deletes or moves files, and performs all the operations relevant to the management of the system. The Virgo users can access the stored data in read-only mode, using the Software Tools – Data Distribution clients [7]. The history of each stored quantity can be presented to the users by using the historical monitoring software.

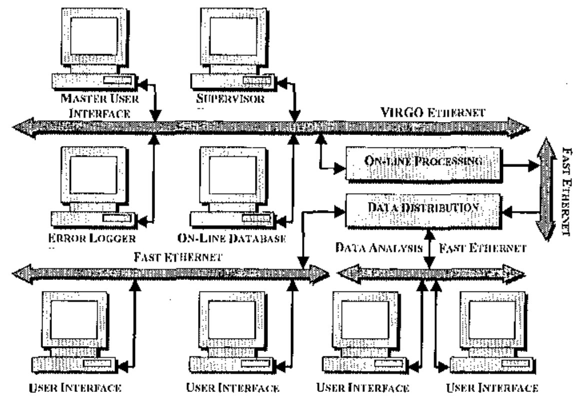

Fig.5: Client/Server architecture of the Virgo Data Distribution System.

Due to the large amount of data and to avoid any interference between the Virgo data collection and the data distribution, the DD uses separate network lines (see Fig.6).

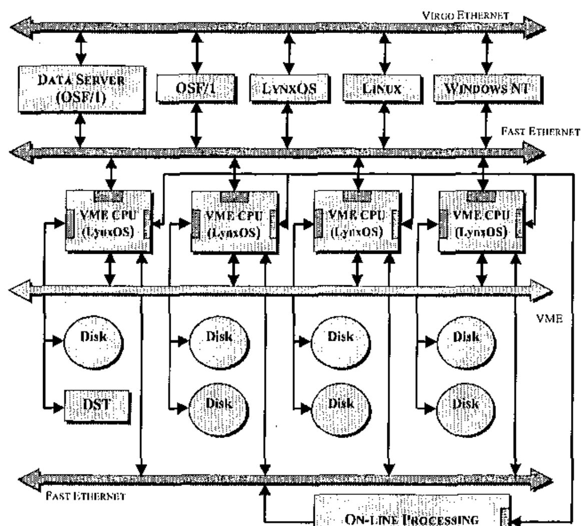

Fig.6: Architecture of the Virgo Data Distribution Server

In synthesis, the DD has two main tasks:

1. Data acquisition and storage from on-line processing;
2. Data distribution.

The DD has been designed in order to perform these two tasks in a nearly completely independent way, giving the highest priority to the data acquisition and storage tasks when access conflicts occur.

A sketch of the software architecture of the DD distribution part is shown in Fig.7, where the different operating systems run by potential clients are shown. The

Siesta block in the figure refers to the official Virgo simulation package.

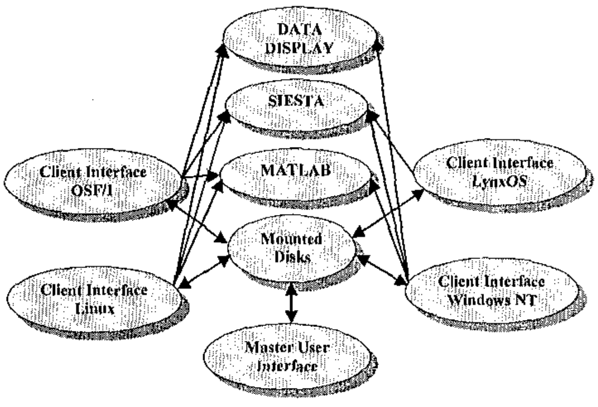

Fig.7: Software architecture of the Virgo Data Distribution section.

1) Data acquisition and storage from the on-line processing

As previously said, the main tasks of the data acquisition and storage section are the data acquisition from the on-line processing system, the data retrieval from raw data and the creation of the data summary tapes (DST) using DLT tapes.

This section is a VME based system composed by a master CPU (VMPC4a from Cetia) running LynxOS and slaves CPUs (same model) handling each 5 disks (18 *Gbyte* Chetaah from Seagate) connected with their fast & wide SCSI native interface and/or with a PCI/SCSI mezzanine card interfaces. The master CPU acquires the data from the on-line processing via a Fast Ethernet link and distributes them sequentially to the slave CPUs (via VME bus) that store them on the disks. At the same time the master CPU sends the part of the frame relative to the environmental monitoring parameters to the DD server together with the information on the location and the content of the stored frames. The latter information is necessary for any data retrieval and the construction of a list of contents. This information is stored in a dedicated database.

The procedure of sequentially archiving the frames on different disks has the advantage of greatly increasing the writing speed of the system. In fact, provided that each slave CPU has enough memory to hold a full frame, the writing speed on each disk must be multiplied by the number of CPUs. Thus the actual bottleneck of this system is the transfer rate from the master to the slave CPU through the VME bus (~25 *Mbyte/s*).

2) Data Distribution

The distribution section, that manages the user requests, is designed following a standard networking procedure. As a fully expandable data server we used a DEC Alpha server 4100, housing up to four Alpha CPUs, running DEC/Unix and supporting several SCSI buses to connect disks. The link with the VME CPUs is ensured by a Fast Ethernet interface, but both the server and the CPUs may also support FDDI. The whole architecture is open to faster data transfer protocols (e.g. FPDP), also for the data acquisition section.

The data retrieval is done by NFS mounting the disks connected to the VME CPUs on the Alpha server. In order to minimise the network traffic, the disks are automounted only on users' request.

III RESULTS AND DISCUSSION

The Data Archiving and Distribution Systems are now available in the Napoli Virgo Lab and ready to be installed in the Virgo experimental site in the next months. The performances of these systems are fully satisfactory and well tuned to the present archiving and distribution needs of the project. The modularity of this architecture now allows further expansions and improvements of the performances.

Due to the modularity and the adaptability of our systems to different uses, this architecture is well suited to be used also in other experiments. In particular it has been chosen by the italian-chinese experiment, ARGO [10] for astroparticle physics, that is now being built in Yanbajing (Tibet).

Finally, in order to improve these systems we are studying the possibility of integrating different and faster transmission protocols and of increasing the writing speed and the storage capacity.

IV. REFERENCES

[1] D.G.Blair, "The Detection of Gravitational Waves", Cambridge, England, University Press, 1991.

[2] P.R.Saulson, "Fundamentals of Interferometric Gravitational Wave Detectors", New Jersey, World Scientific, 1994.

[3] J.H.Taylor, *Rev. Mod. Phys.*, Vol.66, p.252, 1994.

[4] J.Hough et al., "Proposal for a joint German-British Interferometric Gravitational Wave detector", Max Planck Institut Report No.MPQ147, 1989 (unpublished).

[5] R.E.Vogt, R.W.Drever, F.J.Raab, and K.S.Thorne, "Proposal for the Construction of a Large Interferometric Detector of Gravitational Waves", Proposal to the National Science Foundation, California Institute of Technology, 1989.

[6] TAMA (unpublished)

[7] "The Virgo Project: Final Design of the Italian-French large base interferometric antenna Virgo for gravitational wave detection", INFN, Italy and CNRS, France, 1989, 1992, 1995 (unpublished).

[8] C.Arnault, and P.Massart, "Cm: a multitask Communication Package. Version 6.1", LAL-Orsay, France, June 12, 1995 (unpublished).

[9] F.Barone, F.Garufi, L.Milano, and B.Mours, "The Archiving System of the Virgo Antenna for Gravitational Wave Detection", *Rev. Sci. Instr.*, Vol.68 (10), pp.3907-3913, October 1997.

[10] The Argo Collaboration "Astroparticle Physics with ARGO", INFN Italy, June 1995 (unpublished)