

The archiving system of the Virgo antenna for gravitational wave detection

Fabrizio Barone

Università di Napoli "Federico II," Dipartimento di Scienze Fisiche, Mostra d'Oltremare Pad.19, I-80125 Napoli, Italy and Istituto Nazionale di Fisica Nucleare, sez. Napoli, Mostra d'Oltremare Pad.19, I-80125 Napoli, Italia

Fabio Garufi

Istituto Nazionale di Fisica Nucleare, sez. Napoli, Mostra d'Oltremare Pad.19, I-80125 Napoli, Italia

Leopoldo Milano

Università di Napoli "Federico II," Dipartimento di Scienze Fisiche, Mostra d'Oltremare Pad.19, I-80125 Napoli, Italy and Istituto Nazionale di Fisica Nucleare, sez. Napoli, Mostra d'Oltremare Pad.19, I-80125 Napoli, Italia

Benoit Mours

L.A.P.P. BP 110, F-74941 Annecy-le-Vieux CEDEX, France

(Received 11 June 1997; accepted for publication 30 June 1997)

In this article we describe the hardware and software architecture of the prototypes of the archiving system of the very long base line interferometric Virgo antenna for gravitational wave detection. The characteristics which make this system very interesting and innovative are not only its performances but also the modularity of the architecture which makes it easier to both follow the progress of the technology without dramatic changes of the hardware and software architecture and to match the requirements according to the Virgo needs. This prototype has been tested in the data acquisition of the 3 m pendularly suspended and evacuated Michelson interferometer prototype which is operational in Naples and in the monitoring of physical environmental quantities, like electromagnetic noise, acoustic noise, etc., necessary for the correlation of the output of the interferometer with the environment. © 1997 American Institute of Physics.

[S0034-6748(97)05309-4]

I. INTRODUCTION

Gravitational wave detection is one of the most important scientific goals for today's physics and for this task different types of detectors are already operational or are planned throughout the world (bar antennas, interferometric detectors, etc.) with different measurement band and sensitivities.^{1,2} Particular effort is being given in these years to the development and construction of very long base line (up to 3–4 km armlength) interferometric detectors on Earth because technology nowadays makes it possible to reach very high sensitivities ($h \approx 10^{-24}$ – 10^{-26} for integration times of the order of 10^7 s) and large measurement bands (from few Hz to many kHz). These characteristics make the interferometers suitable for the detection of gravitational waves from different classes of astrophysical objects, like, for example, radiation bursts from supernovae events occurring in the Virgo cluster, periodic signals from old and new pulsars, radiation from coalescing compact binaries, stochastic background gravitational radiation, quasi-normal modes of black holes, etc., opening a completely new channel of information on astrophysical objects.^{1,2} The construction of many of these large detectors (GEO600, LIGO, TAMA, VIRGO, etc.) has already started and we expect these detectors to be fully operational at the beginning of the next century.^{3–6}

The drawback related to having such large measurement bands is that the amount of data to be archived is very large. In fact, if we assume, for example, the scheduled sampling frequency of the output signal of the Virgo antenna ($f_s = 20$

kHz with 16 bit acquisition accuracy), then a data flow of at least 80 kbyte/s must be expected for this signal alone, including the time recording. On the other hand, the largest contribution to the data flow will be due to the acquisition of auxiliary signals both from the interferometer (data quality check against malfunctioning of the interferometer) and from the environment (cross correlation of the output of the interferometer with all the relevant environmental noise sources which can simulate a gravitational wave event) which lead to an expected continuous data flow spanning from 1 to 5 MByte/s, that is 30–150 TByte/year, but up to 10 MByte/s in burst mode.⁶

All the data are archived and distributed by the archiving system which is the Virgo subsystem which takes care of the on-line and off-line archive and distribution of Virgo data.⁶ Actually, such a large amount of expected data and the need of accessing totally or partially these data have suggested the specialization of this system into two subsystems, strongly related but different for structure and use, that are the raw data archiving system and the data distribution system, whose requirements are summarized in Tables I and II.

The block scheme of the Virgo data flow is shown in Fig. 1, where all the main logical links among the subsystems involved in data acquisition, archiving and data analysis are shown.⁶ As it is possible to see in this figure, all the data acquired by local readout systems, located within the Virgo buildings, are then organized in frames by the frame builder. Each frame is a structured unit of information containing all the raw data necessary for understanding the behavior of the interferometer over a finite time interval

TABLE I. Raw data archiving system requirements.

Data flow rate	1-5	Mbyte/s	Continuous mode
	<10	Mbyte/s	Burst mode
Data transfer rate	<10	Mbyte/s	From the frame builder

which includes several samplings.⁷ Each frame, organized as a set of C structures described by a header, is basically made of three parts (see Fig. 2).⁶

(1) Structures filled by the frame builder, which contain all the raw data collected by detectors and probes.

(2) Structures filled by the on-line processing (or by the off-line reprocessing), which contain the reconstructed data ([t,h] pairs at the sampling frequency of 20 kHz) and the necessary auxiliary information.

(3) Structures filled by the simulation, which are necessary for comparison of the behavior of the interferometer with the modelled one.

All the frames (raw data) are then sent to the raw data archive which archives them on suitable standard media (tapes, according to the status of the technology). At the same time these data are sent to the on-line processing which processes them and adds other structures to the frames. These structures contain quantities related to the physics of the gravitational waves, like the couples [t,h] (t is the acquisition time and h is the reconstructed dimensionless amplitude), and all the available information coming from the data quality system and global control system necessary for the assignment of a quality coefficient to the data.⁶ These data are then analyzed by real-time parallel algorithms running on the on-line processing system, which selects all the frames which may contain a gravitational wave event. All the selected frames (but all the couples [t,h] for all the ones not selected) are then sent to the data distribution system which archives them on line (on disks) and off line [on data summary tapes (DST)], managing the data distribution to Virgo users.

In order to test the structure of the archiving system in connection with the Virgo requirements, we implemented in Naples a prototype of this system within the framework of the hardware and software architecture of the 3 m suspended Michelson interferometer prototype in vacuum which is operational in Naples and whose hardware and software architecture is very close to the Virgo one.^{8,9} In this way we checked well in advance the compatibility of the system with other systems of the Virgo antenna, collecting all the necessary information for the construction of the final system and for the development and the test of all the software necessary for the related data management.⁶

In this article, we will describe the architecture and the performances of the archiving system, showing how it can

TABLE II. Data distribution system requirements.

Data flow rate	<1	Mbyte/s	Continuous mode
	<10	Mbyte/s	Burst mode
Data transfer rate	<10	Mbyte/s	From the on-line processing
	<1	Mbyte/s	To users via FDDI line
Data on line	>500	Gbyte	1 month data

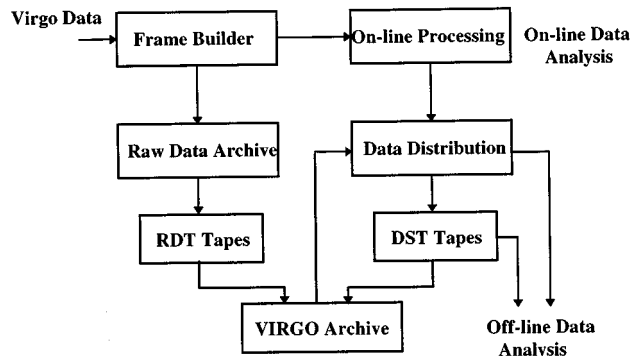


FIG. 1. Scheme of principle of the data distribution and archiving system and its connection with the frame builder and on-line processing system.

easily follow the progress of the technology, without the introduction of dramatic changes in its hardware and software architecture, and how the system can be further improved.

II. ARCHITECTURE OF THE ARCHIVING SYSTEM

Although the raw data archiving system and the data distribution system are structurally different, the technical solutions are designed according to the two following general criteria:

(1) System completely expandable: every further upgrade necessary for the improvement of the performances of the system can be obtained by simply adding components or upgrading part of them without changing the hardware and software structure of the system.

(2) System made of as many as possible standard and commercial hardware and software parts, in order to reduce the number of spare components.

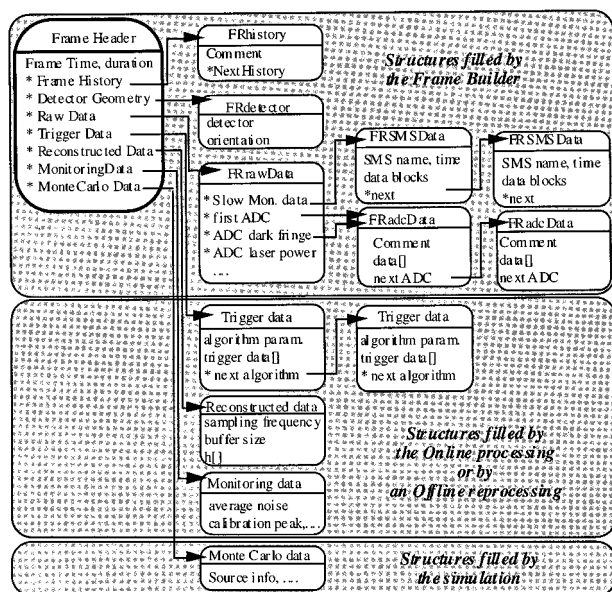


FIG. 2. Structure of a Virgo frame (see Ref. 6).

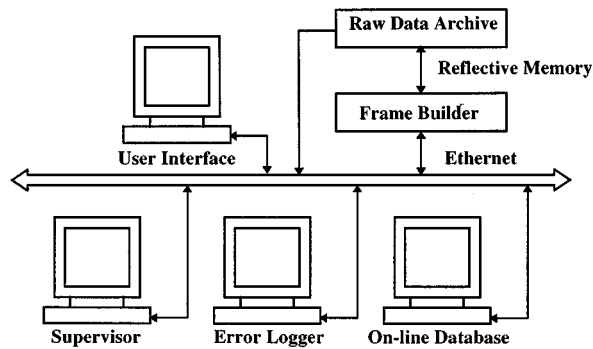


FIG. 3. Client/server architecture of the Virgo raw data archive.

A. Raw data archiving system

The raw data archiving system collects and stores on raw data tapes (RDT) 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 eventual data reprocessing starting from the original data. According to the requirements on the continuous data flow rate (see Table I), the amount of data to be archived spans from 86.4 Gbyte/day (1 Mbyte/s data flow rate) to 432 Gbyte/day (5 Mbyte/s data flow rate). Such a large amount of data cannot be maintained on line according to the actual status of the technology. Although an off-line archive can use both optical disks and magnetic tapes, actually the only economical affordable solution is that of using magnetic tapes. For this reason the raw data archiving system archives all the frames on RDTs, while the raw data retrieval and off-line data reprocessing is a data distribution system task.

At the moment the best candidate as media is the digital linear tape (DLT) which is becoming a defacto standard. These tapes have a good data storage capacity (20–40 Gbyte in compressed mode) and a maximum sustained transfer rate of the order of 3 Mbyte/s (writing mode) with a synchronous peak transfer rate of about 10 Mbyte/s and an average file access time less than 100 s. Therefore, assuming the maximum continuous data flow of 5 Mbyte/s and the maximum storage capacity of DLT (40 Gbyte), then about 4000 DLT/year/copy are necessary to maintain a raw data archive.

The DLT solution is much more efficient and reliable than a solution which would make use of the more diffused digital audio tape (DAT), which is also a defacto standard and has a good storage capacity (4–8 Gbyte in compressed mode) with a maximum sustained transfer rate of the order of 1.0 Mbyte/s (writing mode) and an average file access time of the order of 30 s. In this case, assuming again the maximum continuous data flow of 5 Mbyte/s and their maximum data storage capacity of 8 Gbyte, then at least 20 000 DAT/year/copy would have been necessary.

The architecture of the raw data archiving system is shown in Fig. 3. It is the typical client/server architecture, in which the raw data archive (server) is configured and controlled by its user interface (client). The configuration of the

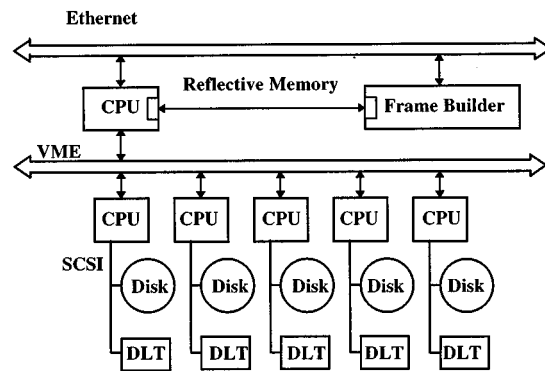


FIG. 4. Architecture of the Virgo raw data archiving server.

system is stored in the on-line database while an error logger records all the error messages generated by the raw data archive using a Cm (communication message) package for communication.^{10,11} The supervisor may act as a user interface.⁶

The matching of the sustained transfer rate (writing mode) of the DLT with the requirement on the maximum continuous data flow and with the restriction of a sequential writing on DLT, which makes easy both the data distribution and retrieval, is obtained with the implementation of a two stage storage procedure. A modular solution to this problem, which can be easily integrated in Virgo, consists in the parallel staging of the data on disks (which can sustain a transfer data flow higher than 6 MByte/s in writing mode and can have a storage capacity up to 4–9 Gbyte) and then copying the data on DLT.

The scheme of principle of the raw data archive server structure which implements this solution is shown in Fig. 4. The system is made of a dedicated VME (Versa Module Europe) crate controlled by a master CPU, running the operating system LynxOS. In the same crate VME slave CPUs are housed, each one provided with disks and with a DLT autoloader. The master CPU controls the frame acquisition from the frame builder via reflective memory enabling the first slave CPU which writes the frames on its disks. When these disks are full, then the master CPU enables the second slave CPU to write the frames on its disks, while the first one starts to download the content of the disks on the first DLT. When it finishes, it changes the tape and waits until the master CPU enables it to write again. The procedure requires the number of CPUs necessary to match the data flow to the sustained transfer rate of the DLT. This solution also makes easy the reconfiguration of the system on the basis of the actual data flow (also in the case in which the data flow may result higher than the requirements) and to have higher data flows for short times if the disks and/or the random access memory (RAM) on board are used as buffers.

The temporal course of the storage procedure with DLT is shown in the scheme of Fig. 5 for a data flow equal to 5 Mbyte/s and a sustained transfer rate of 2.5 Mbyte/s on DLT. Three slave CPUs are necessary, each provided with 40 Gbyte total disk memory and autoloader with seven media. In this configuration a full cycle lasts 24 ks (approx 6.7 h),




FIG. 5. Time course of the raw data archiving server.

that is new data are written on the same disk after 24 ks. Now, assuming that the autoloader houses up to 7 DLTs, a total memory of 840 Gbyte (in compressed mode) is foreseen and it is necessary to change the cartridge magazine of each slave CPU approximately every 48 h. In order to increase this time it is necessary to provide each CPU with larger autoloaders or with more DLT autoloaders. This configuration has also the advantage of leaving the autoloaders enough dead time to rewind a DLT and to load another one on the same unit. Finally, this configuration is completely open and expandable and can easily sustain higher data flow rates, also writing on disks in parallel. Moreover, the hardware can be easily upgraded without changing substantially the software architecture.

Due to the importance of these data, which represent the real output of the Virgo interferometer, the raw data archive system will be provided with a shadow companion in order to render the system fault tolerant. In this way, the system writes two copies of the archive, which is located in two different places, for safety.

On the basis of the above defined structure, the raw data archiving system server will be composed of a dedicated VME crate controlled by a master CPU (PowerPC604 - 64 Mbyte RAM Memory), running the operating system LynxOS and linking the raw data archiving system to the network (Ethernet). In this crate three slave CPUs [PowerPC604 with 64 Mbyte RAM Memory - 3 PMC (PCI - Mezzanine Card) slots on board], each provided with 40 GByte fast disks and a DLT unit with autoloader (with seven or more media), are housed. The data transfer among the frame builder and the raw data archiving server is performed using reflective memory PMC boards which allow a data transfer rate up to 40 MByte/s. This configuration also has the advantage of leaving the autoloaders enough dead time to rewind a DLT and to load another one on the same unit.

B. Data distribution system

The data distribution system archives all the frames selected by the on-line data analysis algorithms (i.e., the frames which also include the structures filled by the on-line processing and by the simulation). These frames should contain the gravitational wave events and, for this reason, they

need a more refined analysis. Actually, this on-line frame selection well applies to impulsive sources, but may be inadequate for continuous or quasi-continuous sources, which may need very refined off-line analyses. For this reason, the data distribution system also stores all the $[t, h]$ pairs and the necessary auxiliary information of all the frames at the real sampling rate (20 kHz). Unfortunately, also in this case the large amount of data limits the quantity of data which can be maintained on line because of the costs of the disks and of the CPUs. In fact, about 12 Gbyte/day for the storage of the selected frames (i.e., 150 Kbytes/s as mean data flow assuming 10 days of events over a year of acquisition at 5 Mbyte/s data flow) and 10 Gbyte/day for the $[t, h]$ pairs at the sampling rate of 20 kHz (including quality factor coefficients and the auxiliary information necessary for an off-line data analysis) would have been the necessary disk space for an on-line storage. For this reason, we have split the data distribution system into two parts: on-line data distribution and off-line data distribution. The on-line data distribution allows the Virgo user to get all the data produced by Virgo in the last month via computer network. This task requires to maintain at least 500 Gbyte on line (assuming a mean data flow of only about 2 Mbyte/s). At the same time all the content of the on-line archive is written on data summary tapes (DST) and distributed on request to the Virgo laboratories. In order to optimize the data distribution and analysis, special DST are also foreseen according to the data analysis needs.

These solutions put no restrictions to the off-line data analyses made by Virgo users. In fact, every time a raw data reprocessing is necessary (i.e., for making correlation with environmental data, with other antennas, etc.) the RDT can be loaded and reprocessed by the data distribution system (off-line reprocessing) using the same tools implemented for the on-line processing. For this reason a list of the contents of the data stored in RDT and in DST is maintained on line on a data distribution system database. This database also includes an on-line archive of the environmental parameters necessary for a check of the Virgo environmental status.

The data distribution system collects the data produced by the on-line processing system and stores them on disks (on-line data distribution) and on DLTs (and on DATs) (off-line data distribution - DST). These data contain all the useful information for the off-line data analysis, that are:

(1) Reconstructed h values ($[t, h]$ pairs with quality coefficients and all the necessary auxiliary information) (maximum continuous data flow 120 kByte/s at 20 kHz sampling rate).

(2) Slow monitoring records for monitoring the interferometer environment. These data are organized in an on-line database (historical monitoring) and can be accessed at any time for checking and analysis of the environmental variables (maximum continuous data flow 5 kByte/s).

(3) Frames selected by the Virgo real-time data analysis algorithms running in parallel on the on-line processing. These frames contain all the raw data required to perform a full analysis of the signal candidates (maximum equivalent continuous data flow 150 kbyte/s assuming 10 days of events at 5 Mbyte/s data flow).

(4) Raw data, retrieved from the RDTs of the archive, to




FIG. 6. Structure of the archived data in the data distribution system.

be reprocessed by the data distribution system in order to reconstruct the full frame, or data retrieved from DSTs.

The organization of the data in the data distribution system is hierarchical. The frame directory, the slow monitoring variables and all the derived data and auxiliary information are stored in classic database (a commercial database is used for a full compatibility with other scientific databases), while the frames selected by the on-line processing system are stored according to a tree structure, as shown in Fig. 6.

Actually, such a large amount of data (300 kbyte/s mean value) cannot be archived on line for a long time by the data distribution system, because it would make the archive too expensive. For this reason we limit the content of the on-line archive to the data produced by Virgo in the last month, which requires to maintain about 500 Gbyte on line (about 80 disks, assuming 9 Gbyte capacity and a global mean data flow of 2 Mbyte/s of raw data). At the same time, all the content of the on-line archive is also written on DSTs. Although the content of special DSTs can be defined according to the data analysis needs, the standard DSTs contains the following information:

(1) All the frames with reconstructed h data and slow monitoring records at the sampling rate (20 kHz) for performing a full data analysis for all the sources on the whole band of the interferometer. The storage of these data requires a maximum amount of 3600 Gbyte/year.

(2) All the frames with slow monitoring records and reconstructed $[t, h]$ pairs data resampled at lower frequency (2 kHz) for performing a specific search on continuous sources (i.e., pulsar search). The storage of these data requires a maximum amount of 360 Gbyte/year.

(3) Frames selected by the parallel real-time data analysis algorithms running on the on-line processing system for performing a specific research on bursts or coalescing binaries. Assuming ten days of these data, then their storage requires a maximum amount of 4400 Gbyte/year.

On the basis of the assumptions made, then the storage of data on DSTs requires a maximum amount of 8400 Gbyte/year on tape, that corresponds to 210 DLT/year.

All the data archived by the data distribution system are available to the authorized Virgo users via a standard network (fast Ethernet) or directly on DST for the off-line analysis. They can get them using the software tools - data distribution,⁶ choosing them within the list of contents of the data stored in RDTs and in DSTs, which is maintained on line on the data distribution system.

The architecture of the data distribution system is shown in Fig. 7. A user interface (client) configures and controls the

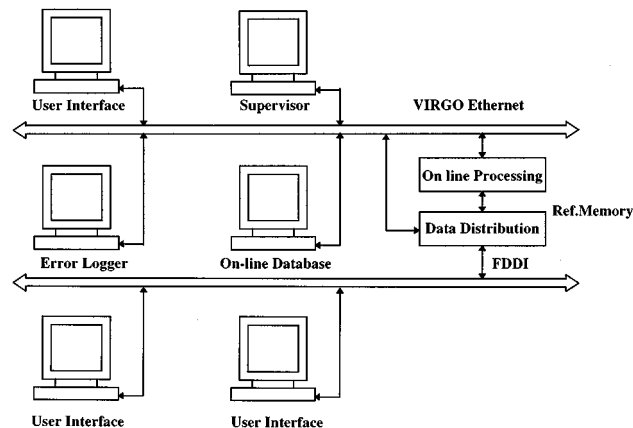


FIG. 7. Client/server architecture of the Virgo data distribution system.

data distribution server according to a typical client/server architecture, setting the configuration (stored in the on-line database system), checking the errors (recorded by the error logger system), accessing the stored data for changing their structure, for deleting and moving files and for all the relevant operations of management of the system. Standard Virgo user interfaces access the stored data in read only mode using the software tools - data distribution.⁶ In particular, the history of each stored quantity can be displayed on these workstations using the historical monitoring software. Due to the large amount of data and to avoid any possible interference between the Virgo data collection and the data distribution, the data distribution system uses separate standard network lines. In synthesis, the data distribution system has two main tasks:

- (1) Data acquisition and storage from the on-line processing.
- (2) Data distribution.

The data distribution system has been designed in order to perform these two tasks in a nearly completely independent way following two different philosophies, giving full priority to the data acquisition and storage tasks if access conflicts exist. In particular, the data acquisition and storage section is designed according to the real-time VME acquisition techniques, while the data distribution section follows

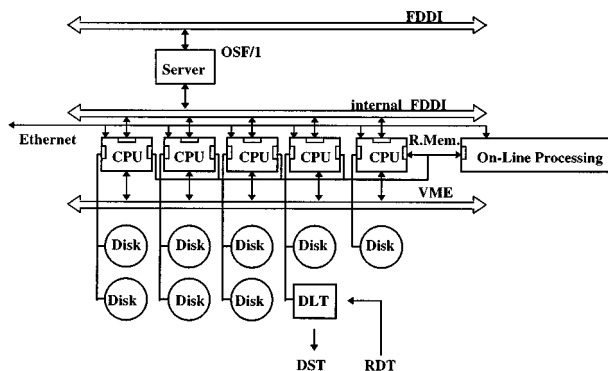


FIG. 8. Architecture of the Virgo data distribution server.

the standard computer networking techniques, as we will show in the following (see Fig. 8).

C. Data acquisition and storage from the on-line processing

As we said before, 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 (RDT) and data writing output (DST). It is a VME based system in which a master CPU (all the data acquisition CPUs are provided with up to 3 PMC bus interfaces on board) runs the operating systems LynxOS. On the same VME bus slave CPUs are housed, each one handling a number of disks connected to its native small computer system interface (SCSI) (wide and fast) bus or to added PMC-SCSI interfaces. The master CPU acquires the data from the on-line processing via reflective memories and distributes them in turn in a sequential way to a slave CPU (via VME bus or via reflective memory PMC boards if the slave CPU is not housed in the same crate of the master), which stores them on the disk. At the same time the master CPU sends to the data distribution server the part of the frame relative to the environment monitoring parameters together with the information of the location and the content of the stored frames, necessary for any data retrieval and content list making. All this information is stored in a dedicated database.

The procedure of sequential archiving frames on different disks has the great advantage of greatly increasing the writing speed of the system. In fact, using this technique the writing speed is limited by the number of slave CPUs used in the crate, being the writing speed on each disk multiplied by the number of slave CPUs used, provided that each CPU has enough memory on board in order to contain a full frame. Therefore, the real limit of this architecture is given by the block data transfer on VME from the master CPU to the slave one (which is of the order of tens of MByte/s) or by the reflective memory data transfer if this technique is used (which may be at most one order of magnitude higher).

D. Data distribution

The data distribution section, which manages the requests of the Virgo users, follows a standard networking procedure. For this task, a fully expandable server (Alpha server) or more than one if necessary, which can accommodate many Alpha CPUs running the operating systems (OSF/1) (Open Software Foundation), memory boards, peripheral component interconnect (PCI) peripheral buses for the input-output (I/O) which supports several SCSI buses to which it is possible to attach disks and tapes, is used. The link with the VME CPUs is made via an internal FDDI link (but the architecture is open to the introduction of new and faster data transfer protocol standards) on which all the disks are mounted and can be readable. The data retrieval is done using a fiber distributed data interface (FDDI) internal network on optical fiber. For this reason, each slave VME CPU is provided with a FDDI-PMC interface which is used by the server for accessing the data stored in the disks.

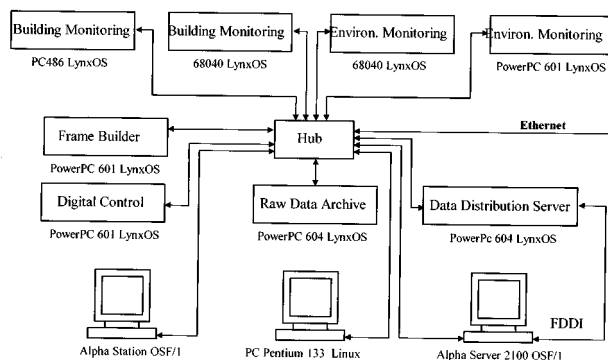


FIG. 9. Scheme of the network for data acquisition, archive and control of the 3 m prototype suspended Michelson interferometer for gravitational wave detection which is operational in the Naples Virgo lab.

III. RESULTS AND DISCUSSION

A raw data archiving prototype and a data distribution prototype have been implemented and are operational in the Naples Virgo lab within the framework of the architecture of the digital system for the control and the data acquisition of the 3 m suspended Michelson interferometer in vacuum,⁸⁻⁹ whose architecture is modular (see Fig. 9).

In fact, every system is made of many independent VME-bus crates in which commercial and prototype VME boards are housed (subsystems). The VME subsystems are linked together through a dedicated Ethernet network, which is isolated from the outer world by means of a bridge [Digital Equipment Corporation (DEC) bridge 90FL from Digital Equipment]. Each VME subsystem has a standard basic hardware configuration: a CPU board (68040 or PowerPC) which handles the operating system LynxOS and a local hard disk. The architecture of the remaining part of the subsystem (CPUs, ADCs, digital-to-analog converter (DACs), memory boards, etc.) depends on its purpose and may greatly differ from one system to the other. The architecture of the user interfaces and of the raw data archive and data distribution are different. In particular the user interfaces are basically workstations (alpha station) which handle the operating system OSF/1 or PC which handle the operating system Linux.

On the other hand, the prototype of raw data archiving system is made of a seven-slot VME crate from ELMA, in which a master CPU [VTMR2a from Cetia-PowerPc 604-100 Mhz-64 Mbyte random access memory (RAM)], provided with a 1 Gbyte disk (from Quantum) and running the operating system LynxOS (v2.3) is housed. In the same crate, two other CPUs of the same type but with 32 MByte RAM on board and slave for VME, each provided with a 4 Gbyte fast disk (Barracuda from Seagate) and a DAT auto-loader (from HP) with 6 media for a total capacity of 48 Gbyte are also housed. On this system we developed the whole hardware and software architecture of the Virgo raw data archive. In fact, the archive has been implemented in this classic client/server architecture, with the server running on the master CPU of the archive, while the user interface (client) can run on a standard workstation (OSF/1, Unix, LynxOS, Linux, etc.). The results of the tests show that this architecture works according to the specifications. In fact,

tests on the writing speed of the slave CPUs on the buffer disks of 4 GByte have reached a maximum writing speed of more than 6 MByte/s that is more than the Virgo requirements.

This requirement is very important because the storage on the buffer disks is the real specification for the data flow rate. In fact, the writing speed on the DAT tapes is less relevant because it specifies only the number of slave CPUs, buffer disks and DAT autoloader units that is necessary to add to the system. In particular, we obtained a maximum sustained writing speed on DAT tapes equal to 0.7 Mbyte/s. Therefore, our system in Naples is able to sustain a raw data archiving data rate of about 1.4 MByte/s. Of course, with a maximum requirement of 5 MByte/s this would require about 8 CPUs with relative disks and autoloaders, but the solution is surely feasible. At the moment, in any case, we are moving toward the solution which uses digital linear tapes (DLTs). We made some preliminary tests on an autoloader unit (from ADIC - up to 2.5 MByte/s writing speed) with seven media. Therefore, with these media the system will be made of at most three slave CPUs and a total of 420 Gbyte on line (not working in compressed mode). Finally, the data retrieval from RDT has also been developed and tested.

The prototype of data distribution system is made of a 2 VME crates (from Schroff), in which a master CPU (CVME604 from Cetia - PPC604 - 100 MHz - 64 Mbyte RAM) with three PMC slots on board, running the operating system LynxOS (v2.3) and provided with a 1 Gbyte disk, is housed. In the other crate two slave CPUs of the same type, but with only 32 Mbyte RAM on board, are housed with two 4 Gbyte fast disks (Hawk type from Seagate) together with a DAT unit (from HP). The data arriving to the master CPU of the first crate are sent to the slave CPU in the second crate for the storage on the disks via reflective memory. For this task VME reflective memory boards (VSA100 from Lextel with 4 Mbyte memory on board) are housed in each crate in order to provide and to test this link. Although the use of this configuration surely makes the data transfer slower among the master CPUs and the slave ones (we have reached a maximum of 4 Mbyte/s data transfer) we tested in this way a configuration of the system which is quite close to the final one, in which if there is the need of a reflective memory link among the master and the slave CPUs, then each CPU must be provided with PMC reflective memory on board for this fast data transfer. At the same time, we also tested the direct

VME connection putting the master CPU and the slaves on the same bus VME, obtaining a data transfer of 25 Mbyte/s, which is well beyond the scheduled Virgo requirements.

While the architecture described above is for the data archive on disk, for the data retrieval we tested the FDDI connection. In fact, we provided each CPU with a PMC-FDDI optical interface which allows an Ethernet independent connection of each CPU with an external workstation. All the FDDI nodes are connected to a concentrator (DEC concentrator 9000 from Digital) which form an internal FDDI network to which two Alpha workstations (an Alpha 2100 server VME and Alpha standard workstation) running the operating systems OSF/1 are connected. The disks of the data distribution system are mounted in such a way they can be read from the workstation without any perturbation of the writing procedure because of the different priorities assigned to the two operations.

The structure described above has the great advantage of the modularity: the architecture of the system does not need to be changed, allowing easy upgrades according to the progress of the technology. As an example, probably the final data distribution system for Virgo will be made of nine or more GByte disks using fibre channel for data transfer and, in future, optical disks when the capacity and the writing-reading speed on these disks will become comparable with the magnetic ones.

¹D. G. Blair, *The Detection of Gravitational Waves* (Cambridge University Press, Cambridge, 1991).

²P. R. Saulson, *Fundamentals of Interferometric Gravitational Wave Detectors* (World Scientific, River Edge, NJ, 1994).

³J. Hough *et al.*, Report No. MPQ147, Max Planck Institut für Quantenoptik, Munich (1987).

⁴R. 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).

⁵TAMA (unpublished).

⁶The Virgo Project, final design of the Italian-French large base interferometric antenna Virgo for gravitational wave detection of which the authors are proponents and in whose construction the authors are collaborating (INFN, Italy, and CNRS, France, 1989, 1992, 1995).

⁷Mours, Virgo note (1996).

⁸F. Barone, E. Calloni, L. Di Fiore, A. Grado, L. Milano, and G. Russo, Virgo note No. PJT94-020, May 1994.

⁹F. Barone, E. Calloni, L. Di Fiore, A. Grado, L. Milano, and G. Russo, *Rev. Sci. Instrum.* **67**, 4353 (1996).

¹⁰C. Arnault, Virgo note No. CTLSW-001, January 1994.

¹¹C. Arnault and P. Massart, "Cm: a multitask Communication Package. Version 6.1," LAL-Orsay, France, June 12, 1995 (unpublished).