Laboratori Nazionali di Frascati

Submitted to N.I.M.

LNF-89/082(P)
13 Novembre 1989

L. Catani, G. Di Pirro, C. Milardi, L. Trasatti, F. Giacco, F. Serlini:

THE HIGH SPEED BUS SYSTEM FOR LISA

Presented at the International Conference on
"Accelerator and Large Experimental Physics Control System"
Vancouver, Canada, Nov 3, (1989)

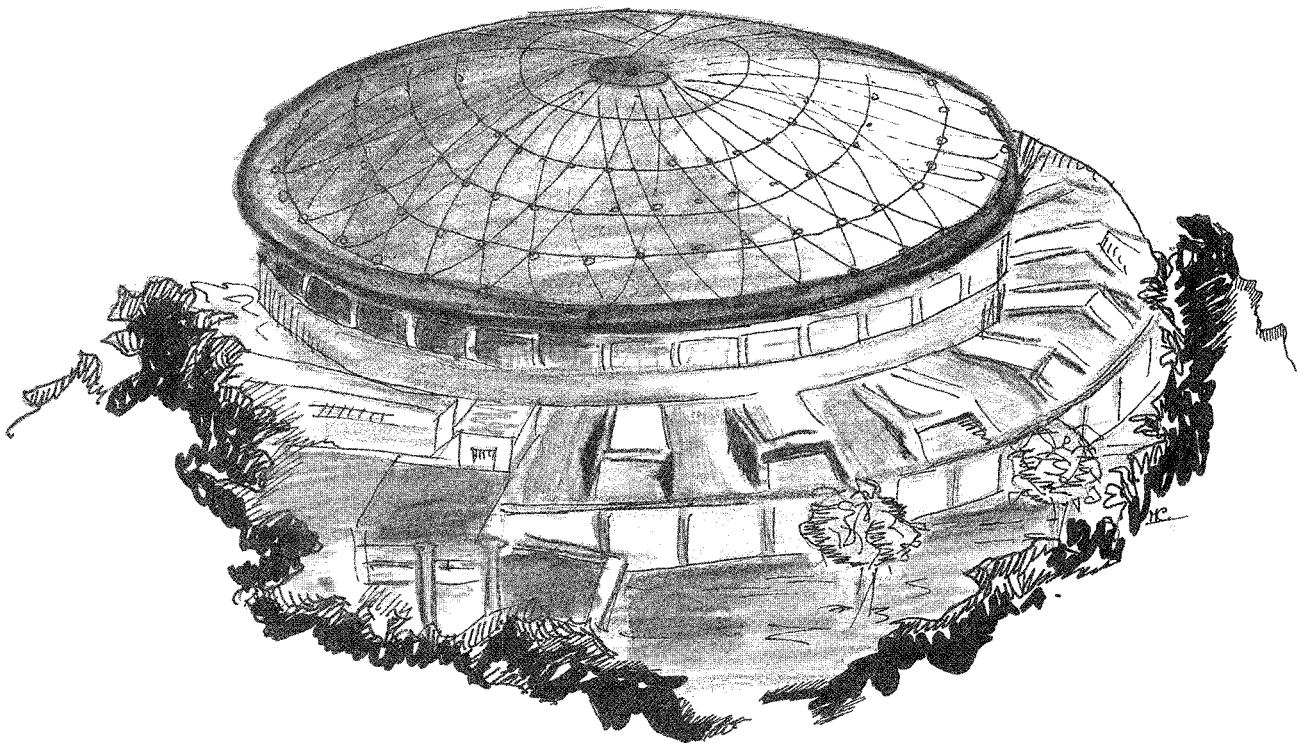

Servizio Documentazione
dei Laboratori Nazionali di Frascati
P.O. Box, 13 - 00044 Frascati (Italy)

INFN - Laboratori Nazionali di Frascati
Servizio Documentazione

LNF-89/082(P)
13 Novembre 1989

THE HIGH SPEED BUS SYSTEM FOR LISA

Luciano Catani, Giampiero Di Pirro, Catia Milardi, Luciano Trasatti
INFN, Laboratori Nazionali Di Frascati, P.O.BOX 13, I-00044- Frascati (Italy)

Ferdinando Giacco, Francesco Serlini
INFN - Sezione di Napoli and Universita' di Napoli, Dipartimento di Fisica, Pad. 20 Mostra d'Oltremare I-80125-Napoli (Italy)

ABSTRACT

The control system for LISA (INFN-LNF Frascati) is based on relatively short but very fast parallel buses. The high bandwidth available makes it possible to build a system where a central VME processor controls the whole data exchange in the two directions. Data are continuously and autonomously read by the central processor from the hardware interfaces and stored in a central memory. Commands are passed in the opposite direction under operator control. Measurements of data flow rate show this system to be much faster than the usual LAN communication systems and particularly suited for small size machines.

1. INTRODUCTION

LISA is the new 25 MeV superconducting electron linac [1] under construction at the INFN Frascati National Laboratories.

The machine is intended mainly as a workbench to study several new components, like efficient injection, ricirculation and to install a high peak gain FEL.

Such an experimental machine requires a very flexible control system, capable to adapt to new and largely unforeseeable requests of the experimenters.

We are implementing a control system based on Macintosh II computers as consoles, and on the VME standard for the peripheral apparatus [2]. The peripheral CPU's are Motorola MVME 133. The software environment is Apple's MPW3.

2. CONTROL SYSTEM OVERVIEW

Due to the limited physical size of LISA, it has been possible to design the interconnection between the separate parts of the control system using high speed, direct memory access parallel buses (MacVee [3] and VMV [4]), instead of a network system. This in turn has led to a hierarchical structure where a single central processor is responsible for data storage and communication control.

This structure greatly simplifies the communication protocol: since you have direct access to the memory containing the mailbox, a message is written and immediately read back to test its integrity. No acknowledge is necessary, since delivery of the messages is guaranteed.

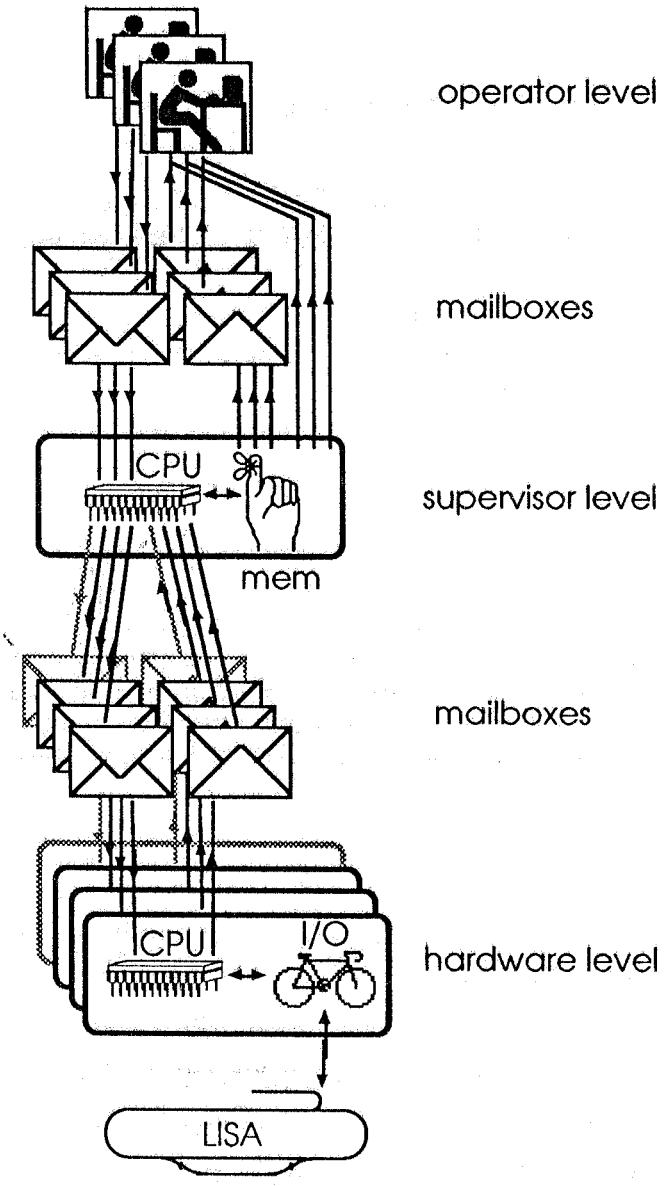

FIG. 1 - LISA mailbox scheme.

The system architecture is based on three processor levels (see Fig. 1):

1) Operator interface (three consoles)

2) Supervisor (one central VME crate)

3) Hardware interface (several VME crates).

The first level is the user interface, implemented using three Macintosh IIx personal computers. The three consoles are all equivalent, i.e. every one of them can control the whole system and no privileged functions on specific machine subsystem are provided for any of them.

The second level, which consists of a single VME crate, is the supervisor. A single CPU continuously monitors, using a round robin mechanism, the lower level subsystems. All of the data describing the machine status are immediately transferred to a RAM memory located in the supervisory crate, but accessible by the consoles. At the same time, commands or series of commands from the consoles are interpreted and transferred to the appropriate third level units for execution (e.g. emittance measurement, beam steering etc.). A log of commands sent is kept in the central memory.

The third and last level is made up of at least one CPU per crate, performing on-line control and/or monitoring functions on the accelerator hardware equipment. The third level CPU's must carry out three separate tasks: reading and executing all commands coming from the consoles through the second level supervisor; relaying to the upper levels only the relevant changes in the hardware status and the warning and fault messages, and finally continuously controlling the hardware.

The EPCS Control Protocol for power converters [5] has been implemented as appropriate.

This hierarchical system, where a single supervisory CPU controls both data acquisition and command dispatching is very appealing for the simplicity and coherence of the whole setup.

3. MAILBOX SYSTEM TESTS

The mailbox system is organized as follows (see Fig. 1):

I level to II level	Commands from the consoles to the supervisor;
II level to III level	Commands relayed to the hardware level;
III level to II level	Machine status measurements and warning-alarm messages to the supervisor, which uses them to update the central status memory. All messages from the supervisor to the consoles are relayed through the main status memory. A level II to level I mailbox for transmission of error and warning messages to the consoles is located in the same RAM.

The three level system has been implemented in the laboratory, and performances have been measured in a series of tests (see Fig. 2).

In the first test, we have measured the throughput of the command forwarding scheme: a Macintosh console continuously sent commands to the second level supervisory CPU, which relayed them to a third level CPU. The commands were constituted by 40 character strings. The third level CPU monitored the number of commands transferred per second. The total rate was 1700 messages per second.

The most critical part of our test was the measurement of the scanning speed when the supervisory CPU continuously monitors third level processors to check for messages waiting to be transferred to the central, second level RAM.

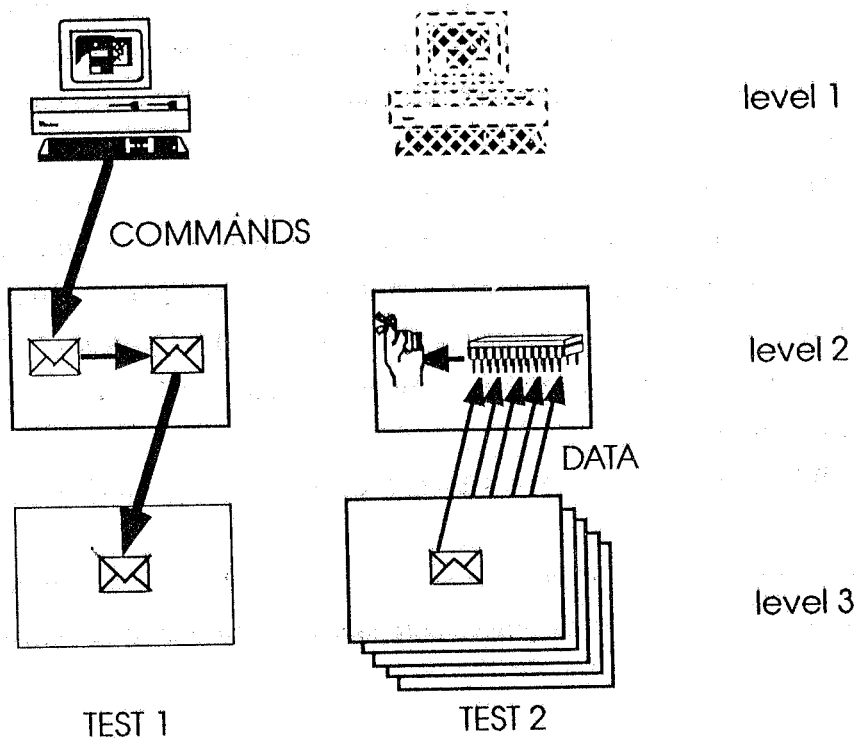

FIG. 2 - Bus test schematic.

Two tests were performed: in the first one ten CPU's of the third level had no data to transfer, giving the time response of the system to answer a message. The measured time was 100 μsec.

In the second test, all ten CPU's required to transfer a large buffer of memory (40 KBytes) to the central RAM. The transfer rate in this case was 400 KByte/sec.

Obviously some price must be paid for these high throughput: such buses, made up of two cables containing 30 twisted pairs, only extend to distances of the order of 100m, and the noise immunity is not as good as that of a coaxial cable (or fiber optic) LAN.

Nevertheless we feel that this is a viable alternative to conventional networks, at least for a small machine in a distributed processing environment.

4. ACKNOWLEDGEMENTS

We wish to thank the LISA group for continuing and stimulating discussions and suggestions, and in particular the other members of the control group for their invaluable support. We are grateful to A. Stecchi for his very important contribution to the data presentation.

REFERENCES

[1] A. Aragona et al., "The linear superconducting accelerator project LISA", Presented at the European Particle Accelerator Conference, Roma, Italy, 1988.

[2] M. Castellano et al., "Diagnostic and control system for the LISA project", Presented at the European Particle Accelerator Conference, Roma, Italy, 1988.

[3] B.G. Taylor, "The MacVEE Hardware User Manual"; Rev.4.1, CERN EP Division, 1987.

[4] C.E.S. SA, Geneva, Switzerland

[5] J. Bonthond et al., "Specifications for a Control Protocol for power converters", CERN PS/CO/WP 88-23, 1988.