ISTITUTO NAZIONALE DI FISICA NUCLEARE
Laboratori Nazionali di Frascati

LNF-84/48(P)
10 Luglio 1984

O. Ciaffoni, M. Coli, M. L. Ferrer and L. Trasatti:

CAMAC AND LAN's

Presented at the "Topical Seminar on
Perspectives for Experimental Apparatus
at Future High Energy Machines",
San Miniato, 21-25 May, 1984

Servizio Documentazione
dei Laboratori Nazionali di Frascati
Cas. Postale 13 - Frascati (Roma)

INFN - Laboratori Nazionali di Frascati
Servizio Documentazione

LNF-84/48(P)
10 Luglio 1948

CAMAC and LAN's

O.Ciaffoni (*), M.Coli (**), M.L.Ferrer (*), L.Trasatti (*)

(*) INFN-LNF, (**) ENEA-LNF

Abstract

We are implementing an Ethernet interface for our intelligent CAMAC crate controller, CANDI 2, to connect it to the Ethernet Network which will tie together all of the Frascati National Laboratories computing facilities. We feel that Local Area Networks can add a lot of power and flexibility to a multicrate CAMAC installation.

CAMAC is a very well known standard for data acquisition, at least in the field of High Energy Physics.

Unfortunately, it suffers from its age. In particular CAMAC is based on the concept of a single computer driving lots of modules directly (through the Branch Highway). This is understandable, since when it was first proposed a computer was roughly the size of an average CPU board and its cost was very high.

Nowadays, the situation is somewhat reversed: the host computer has shrunk in size and price, while the CAMAC crate has become bigger and bigger, with an increasing number of gadgets being packed in a single module.

The trend is now toward more intelligence being put inside

a) single CAMAC modules

b) CAMAC crate controllers.

The microprocessor group at the Frascati National Laboratories (O.Ciaffoni, M.Coli, M.L.Ferrer, L.Trasatti) has been following both directions.

As far as intelligent modules are concerned, we have built a Streamer Tube readout system using a microprocessor.

A Streamer Tube system readout (e.g. the NUSEX Mont Blanc Nucleon Decay Experiment) consists of several long shift register chains which are read serially by a CAMAC controller.

By putting a microprocessor inside the controller we have gained something like a factor of 20 in speed and/or price (the compromise depends on the single experiment).

This is because the microprocessor can read in parallel several shift register chains, and store the data directly into memory without any prelaboration (not even the suppression of zeroes).

This is possible since memories are very cheap nowadays, and you can put a lot on a single CAMAC board.

After the transfer the apparatus can go back to data acquisition mode, and the microprocessor can do some preanalysis of data without generating any dead time. Moreover, the type and quantity of preelaboration is software controlled, so it is easy to adapt it to various experiments.

As far as intelligent CAMAC controllers are concerned, we have built our own, which is now commercialized with the name of CANDI 2. It is interesting to make a short list of what we packed

inside a 3-unit CAMAC module since it proves the point that a CAMAC Crate is big

- 320 KBytes RAM
- 32 KBytes ROM
- A 512x512 pixel color graphics generator
- All of the standard NIM-ESONE subroutines, a subset of PLOT-10, a Tektronix, 4006 emulator, and a complete set of interaction routines to a PDP or VAX host computer
- Finally, the Crate Controller, but with a difference: we did not use the Branch Highway.

The reason for this choice is that once you have built intelligent Crates by packing computers inside the Crate Controllers and you probably still have some kind of host computer at least to hand mass storage devices, printers, and so on, the Branch Highway seems a little out of its depht.

What you would like to have is a more flexible and more powerful structure to exchange data, programs etc. between various computers. We believe that such a structure can be provided by a Local Area Network (LAN). What is a LAN? It is a collection of different objects, both masters and slaves, talking one to the other by a single wire. The main features are high speed (typically 10 Mbit/sec), short distances to minimize the error rate and therefore the software checking overhead, and communication at the peer-to-peer level.

Since LAN's are becoming an industry standard, it is reasonable to assume that also the software, at least at the lower levels, will become standard and will be easily available.

Two major schemes are being proposed: one sequential (token passing) and one based on a statistical approach.

The token structure consists of a sequence of stations with a packet of information (token) going around all the time carrying a flag to indicate if the line is busy or free.

A station wanting to send a message on the line waits for a free token to arrive, changes it to busy and sends the message.

This scheme has a disavantage in that every station, in order to be able to change the status of the busy flag, must introduce a delay of at least one bit time even if it has nothing to say.

The statistical approach (CSMA/CD which means Carrier Sense for Multiple Access with Collision Detection) is completely different. Here every station monitors the line continuously, and is free to try to send a message if the line is momentarily free. If two stations try to send at

the same time (i.e., start sending within the propagation time on the line) they must both sense collision, send a jam message and wait to retransmit for a time depending on how many collisions have already happened. This scheme has a disadvantage in a data acquisition-process control environment in that, due to the statistical approach, it is impossible to guarantee a maximum access time to a station. This can be bad if, for example, the station we are talking about is a fire alarm. On the other hand, this problem only becomes important at high traffic levels. If you compare the two schemes, the token has a disadvantage at low traffic, because the sequential mechanism is always operative and wastes time, while the statistical approach is worse at high traffic because the number of retransmissions increases.

One more thing to be noticed is that all of these schemes have been designed for a commercial environment, and that it is not so obvious to extrapolate their behaviour to data acquisition.

A definite advantage of the statistical scheme, at this time, is that it exists, in the sense that DEC, INTEL and XEROX have published specifications for a LAN called Ethernet based on the CSMA/CD scheme. Several other firms are now producing Ethernet equipment and it is expected that VLSI interface chips for Ethernet will be available soon thus making the interfacing to Ethernet much simpler and cheaper. In the meantime commercial boards implementing the Ethernet contention scheme are already available, thus allowing to start working on the higher levels of software, waiting for the hardware to become more reasonable.

Ethernet has by now become on accepted standard: it will be used in CERN at LEP, and tests have been made to connect it to CAMAC at CERN.

The Frascati National Laboratories of INFN have adopted it as the scheme to interconnect the various computing facilities of the Laboratories (see Fig.1). At the beginning, 2 Ethernet segments connected by an optical remote link will tie together a VAX 11/780 and two PDP/11's, while later on same office automation equipment for the Administration, some CAD/CAM for the Technical Division and an unspecified number of additional PDP/11's will be added. Also two CAMAC crates driven by our CANDI 2 controllers will be inserted into the Network, using commercial interface boards tied directly to the bus of the microcomputer.

As a conclusion, we feel that our choice of an Ethernet LAN to connect our intelligent crate controller to the rest of the facilities of the Laboratories will give us a powerful and flexible situation using standard hardware and software and going toward the goal of distributed intelligence (or, as somebody said, "what's come to be known as intelligence").