DISTRIBUTED DATA ACQUISITION AND CONTROL FOR NUCLEAR PHYSICS EXPERIMENTS

Björn Dreher
Institut für Kernphysik, Universität Mainz, Mainz, F. R. Germany

Summary

A distributed system of one HP3000/68, three HP1000 systems, one HP9000-500 and one Perkin-Elmer PE3220 system is being used at the Institute for Nuclear Physics at the University of Mainz in West-Germany for the data acquisition during nuclear physics experiments at the 180 MeV Microtron. The microtron itself is controlled by two of the HP1000 systems. The backbone of the entire system is the distributed systems software, which has been developed in-house. It is using a packet switched message system to exchange data between processes on the various computers and can be used for local processes as well. The system is in operation since 1981. The HP3000/68 is the central node of the star topology system. Details about the data acquisition system with on-line analysis of the data on the HP3000/68 will be given as well as an general discussion of the control system of the microtron. Future developments of the communications software will include the usage of IEEE 802.3 LANs as hardware vehicle for the message transfer as well as the inclusion of other vendors MC68000 based VMEbus systems in the network.

Introduction

At the Institute for Nuclear Physics of the Johannes Gutenberg University at Mainz, West Germany, basic nuclear physics research is done in the field of medium energy physics with electromagnetic interaction using two different electron accelerators. The older linear accelerator produces a pulsed electron beam with a maximum energy of up to 400 MeV (million electron volts). The new racetrack microtron MAMI (MAinzer MIkrotron) yields in its current second stage a maximum energy of 180 MeV of c.w. electron current. The third stage, which is currently under construction, will deliver a continous electron beam of up to 800 MeV.
Already for use during the experiments using the old linear accelerator a hierarchical distributed computer system was developed to be used for data acquisition and control [1]. In addition, the calculations needed during the preparation of the experiments and the time-consuming analysis of the measured data should be done with this system. Fig. 1 shows the current configuration. The network consists of three HP1000s, one HP9000 series 550, one Perkin-Elmer 3220, and one HP3000/68.

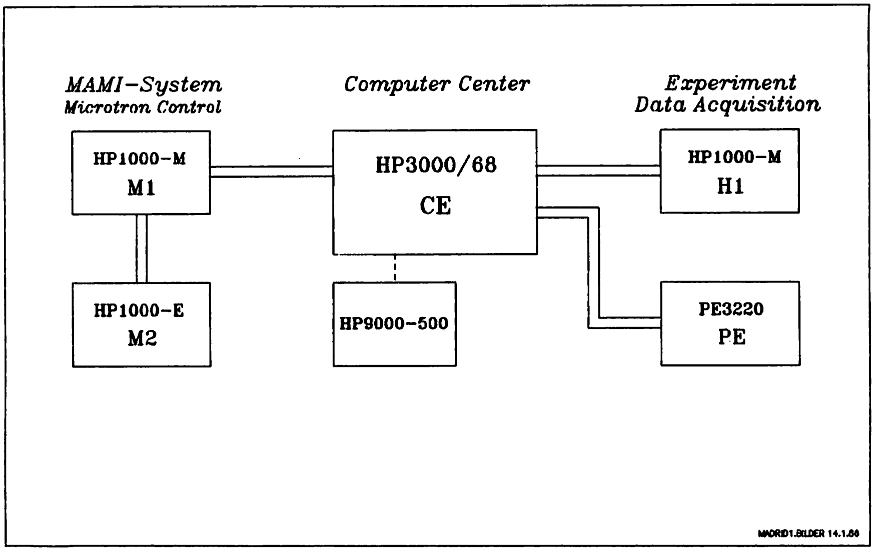

Fig. 1: Distributed Computer System

Overview

The HP3000/68 under MPE-V/E serves two major purposes. First, it is used as the main system in the institute's computer center for technical and scientific calculations. Second, it is the central node of the distributed computer system, which has essentially a star topology. Because the HP3000 is still a 16 bit computer system, the HP9000 system has been added to provide true 32 bit capabilities under the HP-UX operating system.

The next level in the hierarchy consists of somewhat smaller minicomputer systems with true real-time capabilities. They serve for data acquisition and for various process control purposes. The PE and H1 systems are used for fast data acquisition during nuclear physics experiments (more than 10 kbytes/sec). The loosely coupled M1 and M2 computers are being used for the control of the MAMI microtron. The HP1000 systems run under RTE-IVB, the PE3220 under Unix.

The HP3000/68 system

The central HP3000/68 has 72 ATP ports for terminals and other RS232 peripherals, e.g. letter quality and graphics printer and plotter. Two HP7925 and one HP7933 disk drives serve as on-line mass storage, four magnetic tape drives (two HP7970, two HP7978) are used two archive the experimental data. The HP3000 system is used during the preparation of the experiments which are being performed at the institute, for the analysis of the experimental data, either on-line during the experiment or off-line after the experiment has been finished, and for various technical and scientific calculations, which stretch from complicated solutions of problems from theoretical physics to the calculation of the spatial distribution of the magnetic field strength of the main magnets of the microtron and its graphical representation in the form of contour lines. Secondly, the resources of the central computer are being used by the front-end minicomputers via the network. This can be part of the central computer's processing power, e.g. during the data acquisition and on-line analysis of the measured data, or the transparent access of peripherals, e.g. line printer, magtapes, file system, etc., by the smaller computers.

The Interprocess Communications System

Backbone of the entire system is the communications software, which has been developed in-house. When the first implementation was begun in 1977, Hewlett-Packard could not deliver a system with all the capabilities required for our purposes. Today, only the new NS/3000 software has similar features as our own software. In the future, we plan to adopt at least the low level protocol layers (according to the ISO OSI standard) so that we can use standard hardware components as defined in the IEEE 802.3 standard.

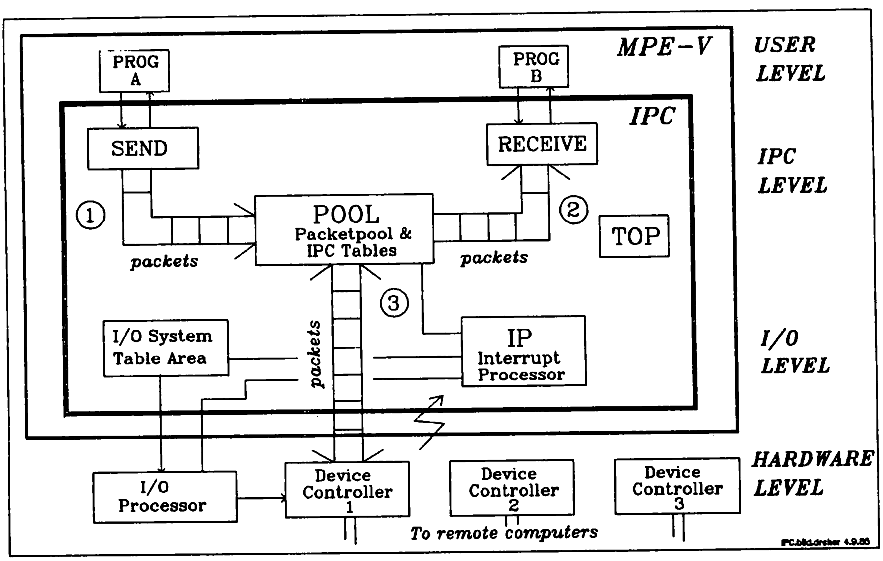

Fig. 2: Interprocess Communications System

Currently our Interprocess Communications System (IPC, not to be confused with HP's MPE IPC!) is implemented as a store-and-foreward packet switching network (Fig. 2). Information is being exchanged in the form of messages which are packed into one or more packets. They are sent from a sender to a receiver. It is relatively unimportant whether both reside on the same system or whether they are processes of different computers. If the data exceed the length of 2046 bytes, the message is split into two or more packets. Packets are the smallest entities that are transmitted between different computers. They consist of a header portion, which contains essentially address information in the form of symbolic names and the data area.

The main software interfaces for the programmer are two procedure calls: SEND and RECEIVE. Parameters are the symbolic names of the sender and of the receiver, a buffer with the data, and the amount of data to be transmitted. The communications software decides if this is a local communication (between processes at the same system), or if the packets have to be transmitted to another computer in the network via an appropriate I/O channel.

Since all packets are always buffered in a global data area before they are delivered to the receiver (packet pool), it was very easy to implement a store-and-foreward function. Therefore, if a packet comes in at the CE computer from e.g. the H1 system and the final destination is the M1 computer, it is simply put on the outgoing queue to the M1 computer and then eventually transmitted via the corresponding I/O channel.

The communications hardware is currently still a 16 bit parallel interface, which originates from the "Programmable Controller" product that was offered many years ago by HP for the HP3000-II. When we upgraded our Series III to the HP3000/64 we had to build a converter from the HP-IB (the internal I/O bus of the HP3000/64) to the old parallel interface. Therefore we had not to change anything at the side of the front-end computers. The maximum hardware data rate can be up to 1 MByte/sec. As an effective transfer rate we achieve 90 kBytes/sec if we transfer large blocks of 16 kBytes.

A general server program on the HP3000 and on the HP1000 systems allows the use of their file system and their peripherals by the other members of the network. As an example, the HP3000 serves very well as a remote spooling system for printer output of the smaller systems.

Data Acquisition and On-line Analysis in a Distributed Computer System

The systems H1 and PE are both being used for fast data acquisition during experiments. As the interface to the experimental set-up we use CAMAC, which is a very popular interface standard in nuclear and high-energy physics. It allows to transmit 16 to 24 bit words in parallel with a rate of up to 1 M words/sec.

During a typical experiment high energy electrons hit a piece of material that is to be investigated (target). A certain physical process happens at the point of interaction between the electron and the atomic nucleus. During that process the electron looses some amount of energy and changes its direction of flight (scattering). It may happen that the electron is absorbed by the nucleus and/or other elementary particles are produced and emitted out of the volume of interaction. All particles coming out of the target have to be detected in appropriate detector systems and their characteristic data, such as energy, angle of flight, and exact time of detection, are measured. With the help of Analog-to-Digital and Time-to-Digital converters (ADCs, TDCs) the data are transformed into digital computer readable data. These data are supposed to describe the physical process under investigation as completely as possible. They are read out via the CAMAC system and written in compressed form on disk or - for archival purposes - on magnetic tape. Like a transaction logging file they contain an exact history of the experiment being performed. Later, after the experiment has been finished, the experiment can be replayed on the computer simply by reading the magnetic tape.

This method of using a distributed computer system for data acquisition uses the small computer system with real-time capabilities for the time-dependent tasks that need fast real-time responses. The data is then sent on-line to the central computer, where it is archived on mass storage and where there is enough computing power to do a first on-line analysis of the data.

This is accomplished by the receiving process on the HP3000 by not only writing the data to magnetic tape or disk but by writing the data

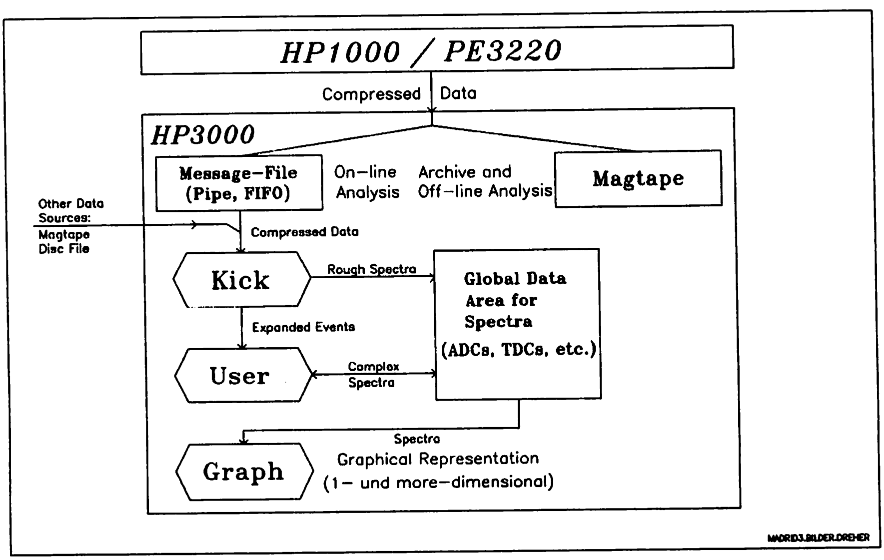

Fig. 3: On-line Data Analysis

to an MPE message file as well. At the other end of the message file the data analysis process reads the data for on-line evaluation (see fig. 3). This process acquires several large Extra Data Segments and sorts the data from the various sources (ADCs, TDCs) into one- or two-dimensional histograms. These spectra can then be displayed graphically by another process which has access to the same Extra Data Segments. Therefore already during the running experiment the experimentalist is able to analyze the incoming data. This is absolutely necessary for an efficient usage of the sparse and expensive beamtime. In fact, the same analysis program is used for the thorough off-line analysis as well. Therefore all the capabilities of the off-line evaluation are available on-line already.

The MAMI Control System

The M1 and M2 systems are used to control the MAMI accelerator. Figure 4 shows an overview of the MAMI control system. The two computers are coupled through the IPC message system and in addition by accessing a common file system on a shared HP7925 disk drive. In fact, our inter-process communications system as it exists today was originally designed for use in the MAMI control system and only later it was also implemented on the other systems. Particularly for the MAMI control system it turned out to be indispensable.

Here again the connections between the computers and the external process peripherals (operator desk and individual components of the

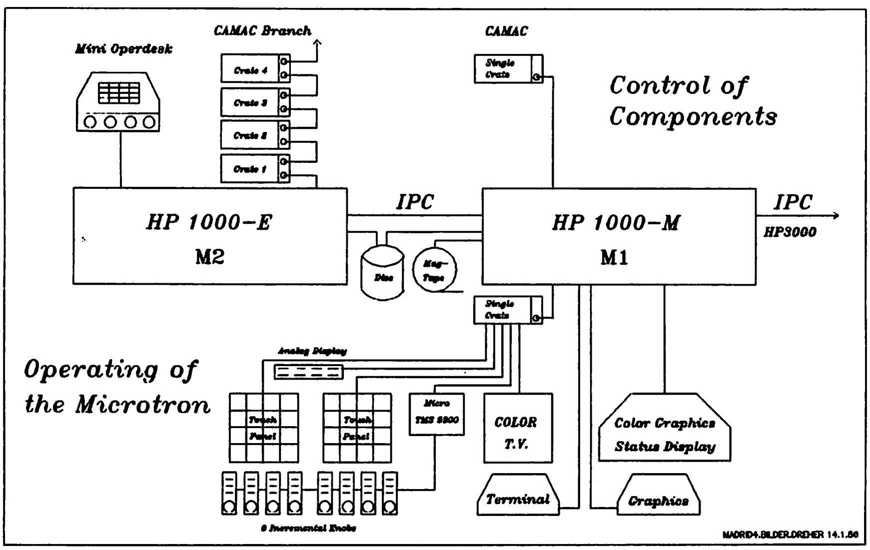

Fig. 4: The MAMI Control System

microtron) is through CAMAC. The communications tasks between the operator and the microtron run mostly on the M1 system while the M2 system is mainly used for the actual control and automatic optimization of the accelerator. The entire control software is, according to the necessary individual task, split into many small processes that are distributed over the two computers. The processes communicate through the message system.

Each component of the microtron that has to be controlled, or each class of similar components, has a special "service routine" assigned to it (Fig. 5). This service routine has the detailed knowledge of the hardware that is to be controlled (similar to an I/O driver in an operating system, but this is a regular user program). Characteristic data of the individual components, like hardware addresses, nominal values, limits, etc., are kept in a data base. They can be read or, if necessary, modified by the processes of the control system. In particular, the data base contains at any time a complete image of the actual status of the microtron (as far as the values are known to the computer). On a touch of a button this image can be stored in a file. Later, the data from that file can be loaded again into the actual data base. Thus on the basis of these values the microtron is brought into the same operating state as before.

The operating of the accelerator is done through three "touch panels", which function very similar to the HP150 touch terminal. Our touch panels have 16 fixed touch-sensitive areas which can be written through CAMAC (which is much faster than via RS232). By touching only few fields the operator is able to find in a tree-structured way the

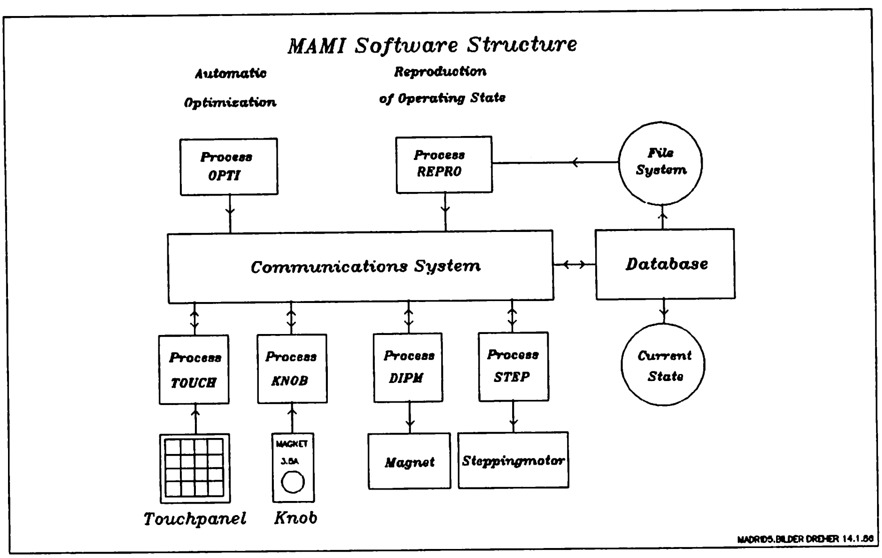

Fig. 5: MAMI Software Structure

final position, where he can for instance switch on or off an individual component of the microtron. To change analog values, e.g. the current of a magnet, manually, an incremental knob is logically connected (upon a touch) to that magnet. This knob has an alphanumeric display field where the computer displays the MAMI name of the magnet and the digital value of the actual current. A turn of the knob is converted to a digital signal which is sent to the corresponding service routine. If the requested value is legal, the service routine will adjust the magnet current accordingly. All this happens so fast that the operator has the feeling of using an analog potentiometer.

The global status of the entire microtron is displayed in graphical form on a high-resolution graphics display (HP2700). Its real zoom function allows the display of finest details.

Conclusion

The design and implementation of a distributed computer system for data acquisition, process control, and scientific calculations, was a great success for our institute. Many applications were only possible by features such as distributing different functions to the most adequate computers and then being able to communicate easily between the individual processes. Tranparent access from the front-end computers to the powerful peripherals of the central computer is a very economical solution for providing to all members of the network the access to those peripherals. Given the fact that the new DS products from HP, such as NS/3000, provide very similar functions as our system, many ideas of our design can probably very be easily adapted to all HP systems with the new software.

We ourselves will use in the future the hardware provided by these new products (IEEE 802.3 LAN) as a replacement for our old parallel link to connect the HP3000/68 or new upcoming products to other computer systems with a compatible LAN interface. Hopefully Hewlett-Packard will provide us with an appropriate software interface to the LAN at a level that is low enough to enable us to adapt our existing system to the new hardware. The LLA (Link Level Access) that is available in HP-UX for the HP9000 Series 500 is a good example for that.

Biography

Björn Dreher
is head of computing at the Institute for Nuclear Physics of the Johannes Gutenberg University at Mainz, F.R. Germany. He got his first computer experience with a Control Data 1700 minicomputer in the late 60s. The first Hewlett-Packard computer he worked with was an HP21MX in 1974, today known as the HP1000-M. In 1976 it was decided, that a distributed

computer system should be set-up at the institute using HP computers. Late 1976 a HP3000-II was installed together with two more HP1000 systems. Today the HP3000-II has grown to a HP3000/68 and the institute is anxiously waiting for the availability of the first member of the Spectrum computer family.

[1] Proceedings of the 1981 Berlin International Meeting of the HP3000 International Users Group, paper I.4