The PC Based Control System of the NAC

H.F. Weehuizen, I.H. Kohler, M.E. Hogan, P.J.Theron
National Accelerator Center, P O Box 72, Faure, 7131, Republic Of South Africa

Abstract

The NAC has been operating a PC based control system since 1990. Currently it consists of 30 PCs grouped into functions of operator consoles, instrumentation nodes and a database node communicating over an Ethernet LAN. The PCs run the OS/2 operating system and use LANServer to communicate over the network. The operator consoles offer a mouse based graphical user interface while the instrumentation nodes interface to the electronics. The system grew to its present size from an original pilot system as it incrementally replaced a previous minicomputer base system which had come to the end of its useful life. This paper looks at the design of the system, discusses the integration process and the difficulties encountered in its development and integration.

1 Introduction

The National Accelerator Center operates a 200 MeV Separated Sector Cyclotron and two injector cyclotrons[1]. One of the latter is a light ion injector while the other has two external ion sources. One is an ECR source which is used to accelerate heavy ions while the other is a polarized ion source. The original control system for this facility was of classical design and consisted of a minicomputer controlling remote hardware which had little or no intelligence. However this system aged and the requirements of the facility began to outgrow the capabilities of the system.

There was thus a need to develop a new system that could cope with the anticipated future requirements of the NAC facility.

A distributed system was chosen as it offers superior reliability, flexibility in configuration and is readily expandable. It is described in the next section.

2 The distributed control system

In keeping with control strategy, this system is designed to be as homogeneous as possible. It has five separate functional parts namely, network, instrumentation nodes, console nodes, database nodes and graphics nodes. Its logical layout is illustrated in figure 1. The design choices for this system are discussed in more detail in reference [2].

2.1 The network

This is an Ethernet CSMA/CD 10BASE5 network which was chosen because the product was mature and readily available. Loading is kept light to avoid the probabilistic nature of Ethernet becoming a problem. Currently the average loading is about 2.5% with the peak loading at about 18%. The average rate of collisions is about one every 10 seconds, so the probabilistic

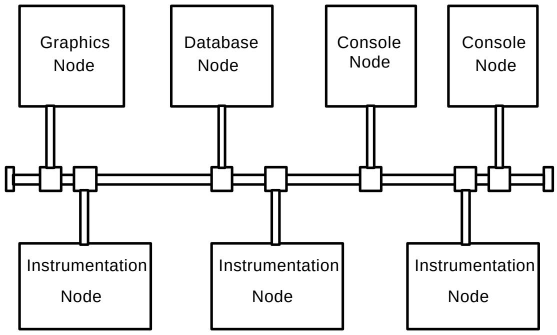

Figure 1. A block diagram of the distributed system

component is negligible. At this loading the transmission delay of the Ethernet system is small [3].

The communication strategy is designed to keep the traffic on the network low. In event driven systems, traffic is proportional to the number of events occurring in the system. Unfortunately when an abnormal situation develops, a large number of state changes can occur, generating a large amount of network traffic. Transmission delays could then become significant at times when good network response is needed. Thus a combination of event and state driven philosophies is used. All operator initiated traffic is event driven providing a quick response. The actual value readback is partly state driven, partly event driven. If events at a node occur at intervals between 1 and 15 seconds, then the communication is event driven. Outside these intervals, communication is state driven, limiting traffid at high rates or indicating that a node is still active at low rates.

2.2 The system nodes

The decision to use PC compatible computers for the nodes has proved to be good . The costs have remained low despite a dramatic increase in the performance of the hardware and they have proved to be reliable. A mix of 486 and Pentium based computers are used. The OS/2 operating system is used as it is a stable and reliable.

2.3 The instrumentation nodes

These interface to the electronics that control or measure the cyclotron parameters. Two interface types are used. The one is the CAMAC system while the other is based on the 8051 microcontroller. The latter is the slightly newer system and was introduced for economic reasons.

The function of the instrumentation controller is centered around a variable table which holds all the values of the parameters controlled by the node. There are several values for each parameter. The first group of values consists of a set point reference value and an associated reference status. The second group holds the actual value and its status, including an error status if relevant. Further fields are used for housekeeping and linkage requests to the variable. This forms the standard control system software for each instrumentation node. It communicates over the network as described above and provides a standardized interface to all the application software. The latter take the cyclotron related parameters and translate these into the device specific operations to effect the control of the hardware. Application tasks are activated when any change to the reference group occurs. Those tasks providing the actual value information are scheduled on a regular basis to update the actual value data in the table.

The communications task also receives commands from the console nodes to control the reference value of the variables. Before a reference value can be accessed, any contention for access from more than one console is resolved at the node that controls the variable and is done on a first come, first served basis.

Functions such as the acquisition of dynamic data would be handled by a separate process as these would be application specific.

2.4 The console nodes

The system controls some 1500 variables associated with the cyclotrons. For convenience these are logically divided into a hundred pages so that about 20 variables at a time can be presented on the console for access by the operators. Joysticks and touchpanels were used on the original system to provide operator input to the system. The operating personnel were very keen to maintain this form of interaction on the present control system but mouse control was provided as an alternative because it was easy to implement. With time mouse control became the preferred method of interaction, so the joysticks and touchpanel have now been removed. Page selection can still be done by keyboard as this provides a faster method of interaction than the mouse.

The console software is designed around a large variable table which holds a copy of the information in the instrumentation nodes together with the nodes from which they originate. The information in this table is updated by the transmission of the variable data from the instrumentation nodes. When a page is selected for display on the console, information is selected from the local table process from where the commands are directed via virtual circuits to the correct instrumentation nodes.

At startup the page layouts and variable groupings are fetched from the database node and held in memory thus forming a local memory resident database. The latter is used in preference to a networked or disc based database in the interests of speed.

This design has resulted in a great deal of flexibility in the control of the cyclotrons as up to 8 consoles are currently used in the control of the cyclotrons. This contrasts sharply with maximum of two that could be used in the original system where simultaneous use of the consoles consumed all the available resources of the minicomputer system.

2.5 The database node

The NAC has had a long period during which there has been construction, development and modification of the facilities on the site which still continues. Therefore it is an important criterion that the control system be flexible and readily expandable. This is most easily met by using a database to hold the configuration information for the control system. Thus it is possible to reconfigure the parameters of the system at a single point to allow for expansion and modification.

When the system is started, subsets of the main database are transferred to each node of the system as required and remain memory resident for access by the local node. The console nodes each have a complete set of the page configurations and display formats while the instrumen-tation nodes have a list of variables pertinent to their respective nodes.

Another function of the database node which is still to be developed, is to log the status of the system at regular intervals for collecting trending information and for system status recovery after interruptions in operation or return to operation at a previously used energy level. The use of the database node is such that it is not crucial to the operation of the system. If it did fail, operation of the system would continue but several important facilities would be lost. Thus a second node is available on standby if the primary node should fail.

Currently SQL Server is used as the database server running with LAN Manager to provide the network interface. As this is no longer supported under OS/2 it is in the process of being replaced with DB2/2. This is fully network aware and so requires no supporting network software.

2.6 The graphics node

This is a specialist operator interface that was designed primarily to provide a dynamic display of the beam profiles as derived from profile grids (harps) and profile scanners. The latter devices are situated on the various beamlines throughout the facility. Local instrumentation nodes control these devices and acquire the raw data from them. Each scan derived from the profile devices is sent as a packet to the graphics node where it is processed and presented for display. Virtual circuit communication is used to provide flow control over the network as the graphics display task is computation intensive and is thus slower than the acquisition task.

This node provides a graphical representation of the
Faraday cups, neutron shutters and the profile monitors in the beam lines. Their positional status is displayed on the monitor, any error conditions being indicated by the color of the background. Further status information can be obtained by means of a pop up menu. The position of any device can be controlled by the mouse. When the profile monitors are inserted the data acquisition system is started and the profile is displayed in the top two thirds of the screen. Up to six devices can be displayed simultaneously. Hard copy of the dynamic displays can be made on a laser printer.

This system has been very well received by the operators because it gives a clear indication of the status of the devices at a glance and the control of the devices is intuitive. The performance of the system both in terms of response to operator commands, indicated response and refresh rate of the dynamic display is far superior to that of the original system. The refresh rate when two or less devices are displayed is such that a delay has been inserted between scans to slow down the display to make it more readable for the operators.

3 Integration

Installation of the system had to be done during scheduled service periods. The only suitable times available are a three week shutdown period in January and a one week shutdown period in July. Thus only a section of the system could be integrated during each period. The process was divided into six stages which took three years to complete. Control remained available at all times on the original system while the present system was integrated. Where CAMAC is used, it was easy to arbitrate between the old and the new systems at the CAMAC level as the standard supports more than one source of computer control. The newer interface does not have this facility, but they require a local processor to interface to the original system via CAMAC. Thus the original local processor was replaced by the instrumentation node of the present system. The node was then also interfaced to CAMAC, retaining the link to the original system.

Except for the instance of dynamic data acquisition, no attempt was made to avoid control conflicts between the original and the present systems. Few problems arose as a result of this. Those that did were secondary effects and were eliminated by disabling control of those functions on the original system. The only instance where the complexity of control was such that the facility had to be disabled on the original system before commissioning on the present, was that for the control and display of the beam profile information. The process of changing over was very simple and took only five minutes to effect.

The operators experienced little inconvenience in the process of installing the new system. However considerable planning had to be done by the control division and many hours of overtime had to be spent commissioning the new hardware and software in the limited time available during each shutdown.

4 Conclusions

The new distributed control system has proved to be successful. The use of personal computers throughout the system has been justified as they have proved to be reliable and extremely cost effective. The chosen operating system has also proved to be reliable, most problems arising from lack of familiarity and lack of adequate time for fully testing the application software owing to very severe time constraints during installation. The system has proved to be flexible and easily expandable with no deterioration in system performance as the system has expanded.

Machine parameters are readily available for system analysis and diagnostic purposes. Specialist programs for machine optimization can now be readily interfaced, saving significantly on development time.

References

[1] G F Burdzik et al, Proc. 11th Int. Conf. on Cyclotrons and Their Applications (Ionics, Tokyo, 1986) pp. 422-425.

[2] H F Weehuizen et al, Proc. 13th Int. Conf. on Cyclotrons and Their Applications (World Scientific, Singapore, 1993) pp.656-659.

[3] H F Weehuizen et al, Proc. 7th IFAC Workshop on Distributed Computer Control Systems (Pergamon, London, 1987) pp. 73-77.