SLAC-PUB-1465
July 1974

INITIAL EXPERIENCE WITH A MULTI-PROCESSOR CONTROL SYSTEM*

K. B. Mallory
Stanford Linear Accelerator Center
Stanford, California

The SLAC accelerator has now been controlled through a system of two linked computers for a year. In the first four months of 1974 nine additional CPU's have been linked into the system. Figure 1 shows the present configuration. It is time to review the problems encountered in computer-control of the accelerator and to discuss their solutions--installed, pending and proposed.

BACKGROUND

When SLAC was built, an SDS-925 computer was installed in the switchyard's data assembly building, now the Main Control Center (MCC). Its functions were to monitor interlocks and to control a few of the switchyard magnets. There was no computer associated with the accelerator itself. Four years later a PDP-9 was installed in the Central Control Room (CCR) to log klystron performance and to help automatically select spare klystrons to replace failures. By making maximum use of existing relay multiplexers the PDP-9 was connected to nearly all accelerator control, analog and status signals within six months. As a result, however, it was limited to executing one control or analog readout at a time.

It was then proposed functionally to move CCR to MCC by linking the computers and using "touch-panel" displays (Refs. 1,2), without physically moving hardware. The pattern generator for defining multiple-beams was the last interface to be completed, early in 1973. Operations from MCC started shortly thereafter.

INITIAL PERFORMANCE

From the outset, nearly all of the accelerator control and monitoring signals could be made available at the touch panel display units. Operators designed their own panels using a convenient software "panel compiler". However, the system could only operate one control at a time, just as was the case at a single control position in CCR. It presented analog signals only when the operator explicitly asked for them by name. Its only alarm annunciator was a "scroll" of the last fifteen messages resulting from detected changes. Since a major fault is normally followed by a number of consequent faults, the most important message often was rolled off the "scroll" before the operator realized anything had happened.

In our operating system (Ref. 3), each status change and button-push creates one or more tasks. The PDP-9 can handle up to roughly 50 or more tasks at a time; the 925 can support twice as many. Nevertheless, there are many occasions when multiple changes create far too many tasks for the system to manage which, at first, caused many system crashes. The solution was to cause various kinds of tasks to be deleted before the system was swamped. To give an extreme example, not so rare as we would like, a fault can occasionally dump the entire RF system. This causes some 800 status changes to appear at the PDP-9 within two or three seconds. When too many status changes occur, the PDP-9 stops reporting individual changes and instead, with a single task, sends a total update of accelerator status to the 925. Similarly, most messages across the link can be aborted if there are too many tasks already in the system.

Since the computer system essentially replaces a pre-existing manual control system, there has been considerable pressure to "give us back what we had before." The annunciator problem was the first to be attacked. There are now several new programs that (1) display the status of all 245 klystrons on one panel, (2) display status of a system (e.g. personnel protection or vacuum) and soon, (3) display a list of accelerator faults in order of priority. This last list displays current status, rather than changes, and is thus independent of occasional missed tasks, (4) multiple scrolls so that messages of different types can be directed to different displays.

RECENT IMPROVEMENTS

These software changes have made the system more reliable and have improved its operation, but software alone could do nothing about the one-at-a-time control and very little about the slow analog acquisition. The original proposal for improving controls required 30 special-purpose processors in the gallery, with circuits to store a command and then drive local relays while the PDP-9 transmitted control signals to other sectors. These processors at first were very complicated; but even when all of the timing and much of the special logic was deleted, they cost as much as putting a computer into every fourth sector, and increased the amount of work to be done by the PDP-9! Since an analog multiplex system and a method to restructure the addressing of some of the control channels were also desired, the special-purpose processors never got off the drawing board. We bought nine PDP-8's instead.

The first of the PDP-8's arrived a year ago. An executive program similar to the one in the PDP-9 and 925 had been written and tested in a simulator in the IBM 360. A terminal-emulator program was written for the PDP-8; it asked the 360's text-editor for a binary "listing" of the executive program, loaded it into core, and we were in business. By August, the eight gallery processors had been linked to the one in CCR, in October the link into the PDP-9 was established, and we were waiting primarily for fabrication and installation of interface hardware in the gallery. (Fig. 2) Programs to control the new hardware and a program in the PDP-9 to store the PDP-8 programs on its disk were completed this winter. During this conference we have installed remote restart switches, as a pacifier to the operators who know that any system can go down occasionally.

By the end of this conference, when the accelerator is being started up, we expect eight of the nine computers to be fully operational, and to have 28 channels of control to the accelerator instead of the one we had before. The ninth computer is being used to test the analog multiplexing hardware as it comes out of the shop. We expect to have a form of analog multiplex system operating in September.

DEVELOPMENTS FOR THE FUTURE

I suggested above that we were restructuring our control addresses. The original control addressing scheme required the operator first to select a sector and then push a button for the desired control signal. When we installed pulsed beam guidance, we subdivided individual control addresses for adjusting up to six preset levels of a device. For manual operations, pushing one button at a time, it made no difference to an operator what the addressing hierarchy might be in the hardware.

But we now have three independent operating positions in the Main Control Room, and are talking about

*Work supported by the U.S. Atomic Energy Commission.
(To be published in Proceedings of the IXth Inter. Conf. on High Energy Accelerators, Stanford Linear Accelerator Center, Stanford, California, May 2-7, 1974.)

adding automatic controls that might be initiated by the computer. The operator now selects which beam he wishes to tune on his touch panel (subdevice address). Then he selects what function he wishes to control (control address, e.g. beam loading or vertical steering). Finally he selects the particular location (sector) where he will make his adjustment. Thus our addressing logic has been turned completely upside-down. In particular, for instance, the controls for beam current and spectrum sharpening of all beams share three control addresses at the injector, which are accessed only by one of the new PDP-8 channels. Thus operators could still find themselves interfering with each other.

We plan to install special interfaces to allow parallel control of selected multilevel devices. The first was scheduled for installation in July, but it now appears that parallel control for "phase closure" (spectrum sharpening) and for fine energy control may become operational before the end of May.

When the analog multiplexing system comes into operation, a new problem will arise: the amount of link traffic is expected to be at best doubled, perhaps trebled from what it is now. Two potential solutions are being studied -- a more-efficient message-switching system for the PDP-8 in CCR (which will not do much to reduce link traffic from the PDP-9 to the 925) or a new data link, for analog signals only, direct from the PDP-8's to an auxiliary processor at the 925.

Eventually, we will probably have to replace the 925, reliable though it be now. We are proposing to install new minicomputers to buffer several of the I/O devices now connected to the 925, and later to connect them to a new major processor which can first share and later, if necessary, take over the tasks of the veteran - the original computer installed for control of the BSY at SLAC.

REFERENCES

1. D. Fryberger and R. Johnson, "An Innovation in Control Panels for Large Computer Systems," Particle Accelerator Conference, Chicago, Illinois, March 1971.

2. S. Howry et al., "SLAC Control Room Consolidation – Software Aspects," Particle Accelerator Conference, Chicago, Illinois, March 1971.

3. S. Howry, "An Operating System for Process Control," 3rd ACM Symposium on Operating Systems Principles, Stanford, October 1971.

4. K. B. Mallory, "Some Effects of (Now Having) Computer Control for the Stanford Linear Accelerator Center," Particle Accelerator Conference, San Francisco, California, March 1973.

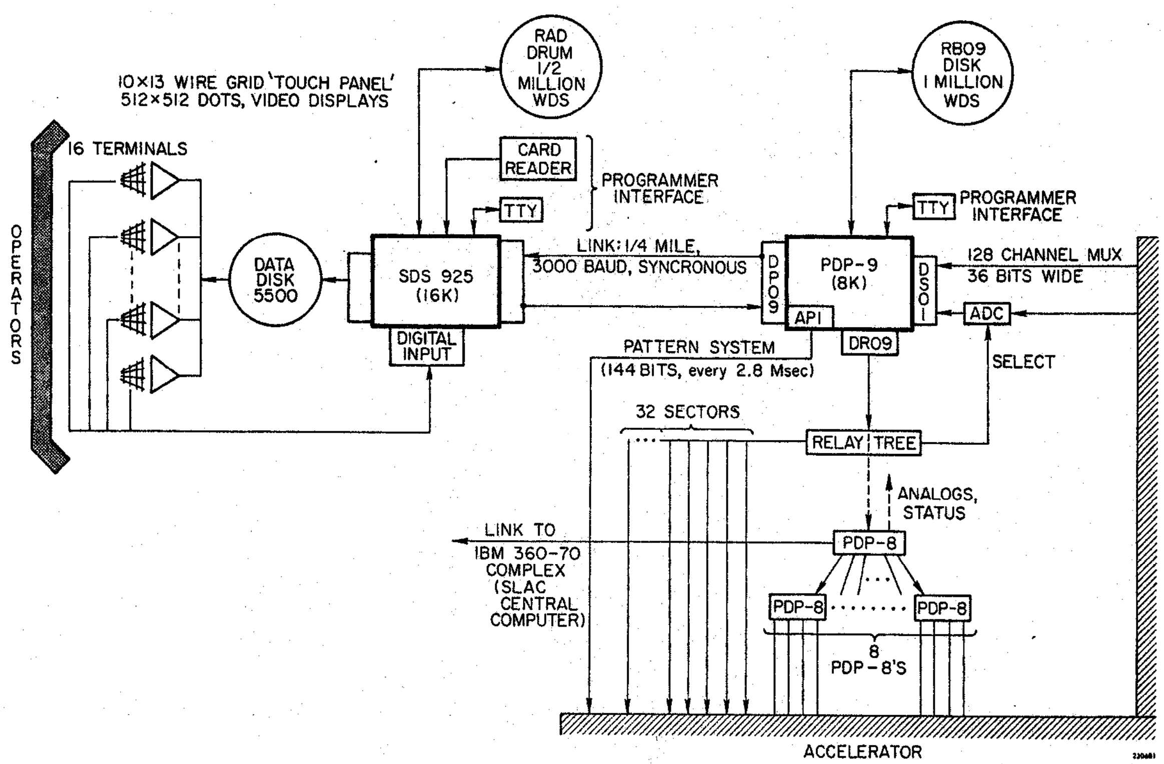

Fig. 1

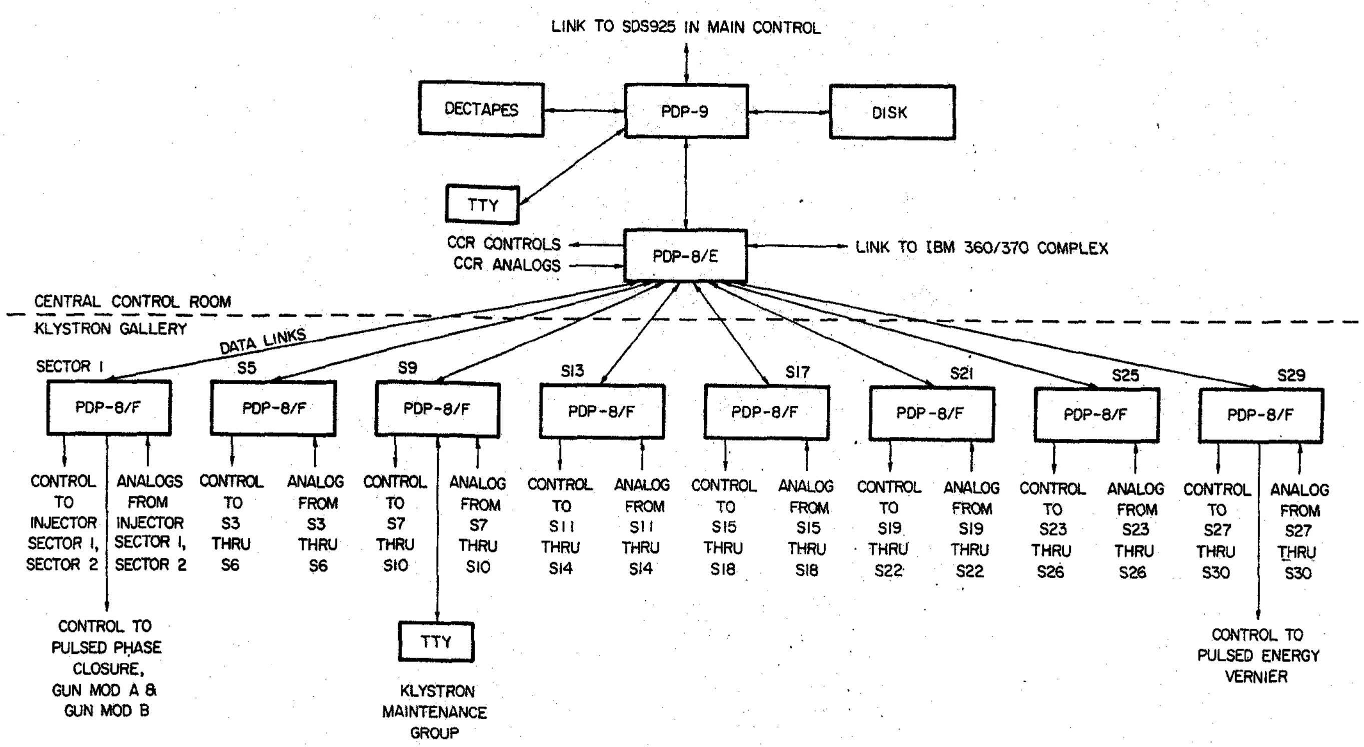
LINK TO SDS925 IN MAIN CONTROL
DECTAPES
PDP-9
DISK
TTY
CCR CONTROLS
CCR ANALOGS
PDP-8/E
LINK TO IBM 360/370 COMPLEX
CENTRAL CONTROL ROOM
KLYSTRON GALLERY
DATA LINKS
SECTOR 1
S5
S9
S13
S17
S21
S25
S29
PDP-8/F
CONTROL TO INJECTOR SECTOR 1, SECTOR 2
ANALOGS FROM INJECTOR SECTOR 1, SECTOR 2
CONTROL TO S3 THRU S6
ANALOG FROM S3 THRU S6
CONTROL TO S7 THRU S10
ANALOG FROM S7 THRU S10
CONTROL TO S11 THRU S14
ANALOG FROM S11 THRU S14
CONTROL TO S15 THRU S18
ANALOG FROM S15 THRU S18
CONTROL TO S19 THRU S22
ANALOG FROM S19 THRU S22
CONTROL TO S23 THRU S26
ANALOG FROM S23 THRU S26
CONTROL TO S27 THRU S30
ANALOG FROM S27 THRU S30
CONTROL TO PULSED PHASE CLOSURE, GUN MOD A & GUN MOD B
TTY
KLYSTRON MAINTENANCE GROUP
CONTROL TO PULSED ENERGY VERNIER

Fig. 2