

SOME EFFECTS OF (NOW HAVING) COMPUTER CONTROL FOR THE STANFORD LINEAR ACCELERATOR*

K. B. Mallory

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

Introduction

Amongst the original guidelines for the design of the control system for the Stanford Linear Accelerator was suitability for ultimate computer control of the accelerator. A major degree of computer control has materialized only after several years of operation. The delay has produced curious features in the system. Elaborate hardware has been built to set up interlock and control circuits; modern minicomputers could make this hardware obsolete. From the beginning, the beam switchyard had a computer for control of the major magnets; the switchyard is now the most difficult area in which to interface additional equipment to the computers. The system provided sequential, one-button-at-a-time control by the operator; present operational procedures demand parallel control, even of two subdevices on the same control channel. The computer system is a newcomer; it is required that the accelerator remain operable even without the computers.

Several years ago I presented a review of some anomalies in the SLAC control system resulting from an attempt to make the controls adaptable to future computer control.¹ Now that we have had a control computer for the accelerator for four years, I shall review the development of the computer control system with some emphasis on problems resulting from the delay of installation of the computer.

Development of the Present System

In November 1968, a PDP-9 was installed in the central control building; by February 1969, all the binary status and control signals from the klystron gallery had been connected to the computer, and by May, we had interfaced most of the binary status information existing in central control. A program had been written to scan the allocation of klystrons to various beams and to select automatically the appropriate reserve klystrons for replacement of ones which failed. A stable executive program with the beginnings of an on-line file-management system was completed in August and by December 1969 we had on-line program editing and assembly, had interfaced nearly all the dc analog signals, and had a program which would record quadrupole settings along the accelerator, and could restore the values according to a previously saved table.

It was at this point that we began to find that we were always a bit behind in satisfying operations requirements. The quad-setting program used the same control channels as the operators' manual controls, and therefore could not accomplish its task any faster than an operator (a point to which I shall return later). Although various methods for speeding up the action were studied, their implementation was deferred because we were in the middle of an operational crisis.

SLAC was built with a central control room for the accelerator itself and a separate console area for the beam switchyard. The production of an electron beam and its delivery to an experimenter had remained two separate jobs, with two separate operating staffs. While we were estimating the costs of picking up one control room and moving it to the other, I stuck my neck out and suggested that by linking the PDP-9 to the control computer already in the switchyard control area, one might be able to control everything from one location without moving any equipment at all. The idea

took, and so were born the main control center, in the switchyard's data assembly building, and the "control room consolidation" project which has just been finished.

The control computer for the switchyard is an SDS-925, which had been used to scan interlocks and to control the beam transport through the switchyard between the accelerator and end station A. It was already operational when the accelerator was first turned on. Although its core had been doubled from 8 to 16k words and a drum memory had been added, its functions had remained essentially unchanged for almost five years. Under the new proposal, another disk was connected to the 925, and used to drive a set of TV display channels. Over each monitor, a transparent overlay control surface ("touch-panel") was mounted, so that the sensitive points of the control surface could be relabeled each time the display was changed. The project also required linking the 925 to the PDP-9, rewriting its executive program to take advantage of the drum, and writing the required programs to create displays and interpret commands from the touch-panels.

In the original schedule for consolidation, I stated that it was "hoped that all (accelerator control) functions might be handled from the data assembly building by July 1971." Except for the program generator for beam triggers, all hardware had been developed and interfaced, and the configuration was as shown in Fig. 1. The executive program had been rewritten, on-line program and file management utilities had been provided, and an operable systems program created displays and executed touch-panel commands. However the original ultrasonic touch-panel turned out to be difficult to produce, so we were using the first copies of a second model with a crossed-wire grid for our control overlay; while the programmers, not to be outdone, were just completing their third draft of the panel management program. Then started the process of developing suitable display groupings for controls on the panels and a fourth version of the panel management program to minimize visible delays in response. This was a trial-and-error process which, indeed, continues to the present. The program for compiling the individual panel displays was expressly written to make it convenient to rearrange the panel layouts.

Meanwhile the scope and requirements of the job were changing. By January 1972 we were getting ready to turn on SPEAR. The operators had figured out how to run six beams at once (the maximum number available on the beam programmer) and up to six experiments were regularly scheduled. SPEAR was to require two more beams, at low duty cycle, and was not to impact the rest of the experimental schedule. The solution was to turn the entire trigger programming function over to the PDP-9. A program was written to represent each of the more than 400 toggle switches in the original beam programmer by a bit in an array in core. A data channel was installed to load a 144-bit register which was then connected to the same output drivers used by the original hardware programmer. Finally, panels were written to allow the operator to set each bit on or off. The result was a system which could handle any number of beams. The core array was set up for twelve beams. (Only eight beams were given space on the control displays.) The extra two beams for SPEAR were therefore provided at no extra cost as part of the "consolidation" effort. But eight experimental beams then became the standard for experimental scheduling, and the operators began to fall behind.

*Work supported by the U. S. Atomic Energy Commission.

At that point, better control programs in the computer could have relieved the burden on the operators, but they were too accustomed to managing every detail for themselves. A demand arose instead for a third operating position at the console in addition to the two originally planned. Much rearrangement and some rebuilding of the existing control panels was required to make space for the new touch-panels and displays, and it is only this month that we have finished installing a full set of display units, with a suit of panels to cover all operational controls.

I may have suggested that all of our effort has been devoted to interfacing devices to the computer system and providing a one-to-one replacement of hardware buttons by touch-panel buttons. This is not entirely true; a number of specialized programs have been written to help the operators and maintenance personnel.

In the PDP-9:

- Klystron replacement
- Klystron fault logging
- Klystron utilization log
- Subbooster fault log
- Beam time logging
- Substation voltage monitoring
- Personnel protection tone loop checks
- Beam program logging
- Quadrupole setting
- Beam position monitor zeroing

In the SDS-925:

- Interlock monitor and recording
- Magnet logging (beam transport to end station A)
- Magnet degaussing
- Automatic energy change (transport to end station A)
- SPEAR magnet tolerance scan

Multilevel devices for pulsed beam guidance have generally been built with six levels. When eight beams are operated, the assignment of levels to a beam varies from device to device. The operators found it difficult to remember which level had been assigned to each beam. The computer system now provides an automatic correlation of beam identification and device level. The operator selects which beam he wants to tune and makes his adjustment. The computers direct his control signal to the correct level of the device.

Work in Progress

Now is the time to increase the number of control programs, but we will still be up against the slow speed of the original multiplexing system which has been used to connect the PDP-9 to the accelerator. In the original control console, every button was connected via an encoder in the control room and a decoder at the destination to the relay or dc motor to be controlled. The computer was connected so that it replaces the buttons and encoder, but its output register is

required to transmit the control signal continuously as long as it takes to produce the desired result. This is typically 250 ms per signal, but is occasionally several seconds when a motor-driver device is being adjusted. Only one control can be actuated at a time. A much faster system of distributing control signals is needed. A PDP-8 computer will be installed every few sectors to take over the chore of holding output relays as long as required. For each command, the PDP-9 will send messages via another linking PDP-8 to the appropriate peripheral computer. Its control I/O channel will be tied up for a few milliseconds instead of 1/4 second or more for each command. The expanded configuration is shown in Fig. 2. The operating system is written and has been tested in a simulated network of four computers. The new processors have been ordered. At time of writing none has yet arrived, but the linking processor is expected momentarily; I hope to have it operational by the end of the next operating cycle.

There are a number of new operational programs we wish to write, especially now that the trigger programming is accessible to the computers. One would allow frequent semi-automatic rephasing of the entire accelerator in a minute or two with considerably less interruption of all beams than the present procedure. Another would reset all parameters of a beam to the values used by an experiment the last time it ran. The resetting of pulsed beam guidance values is curious, however. Since the operator originally could only make the proper adjustments by watching beam performance, the analog representation of the parameter was chosen to be the actual device output, not the set-point. As a result, beam pulses must be allocated during presetting in order to produce the monitor signal, whether or not a beam is delivered on those pulses.

The most controversial new programs are those which are written to replace existing hardware when new requirements arise. An example in point is the "master beam control," which the operator uses to suppress or to reduce the repetition rate of any beam. The hardware is organized to turn off the beam going to any specified experimental area. The new requirement is the result of the computer system's effectiveness in correlating beam guidance levels to beam identification numbers. The operators now require that the master beam control be reorganized to control beams according to the same beam ID that is used for the pulsed beam guidance. This feature is easily implemented through the computer. In fact a panel to do it has already been written. But this time I may have gone too far. The operators have not yet accepted the software replacement. At least part of the panel, I suspect, will have to be rebuilt as a piece of hardware.

Reference

1. K. B. Mallory, "Some effects of (not having) computer control for the Stanford Linear Accelerator," presented at the 1966 Linear Accelerator Conference, Los Alamos, October 1966 (Report No. SLAC-PUB-221).

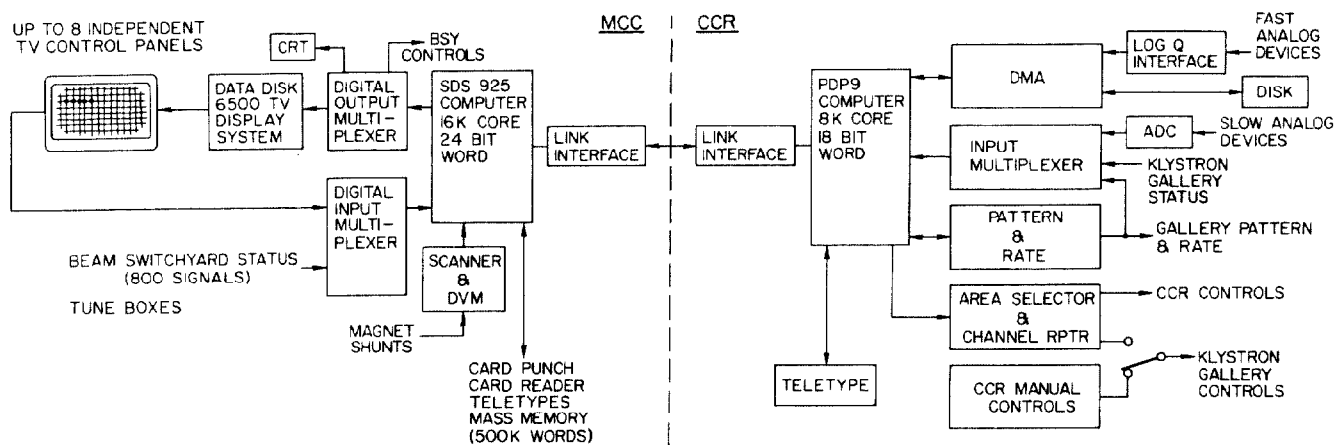


FIG. 1--Computer configuration for "Consolidation".




FIG. 2--A map of computer links for increased I/O interface capacity.