SOME EFFECTS OF (NOT HAVING) COMPUTER CONTROL FOR THE STANFORD LINEAR ACCELERATOR*

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

Abstract

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. Although computer control has not materialized, this design influence has produced unusual features in the system. Video information was to be digested at the source and spoon-fed to the computer: the present system for bringing video signals to the central control room is very sketchy. Large special-purpose data recording systems were considered non-competitive with a small computer: there is now no automatic data-recording system. Sub-systems which are basically nothing but local switch-gear were to be fully implemented outside of central control: Central Control is only incidentally concerned with automatic phasing of the accelerator, personnel protection and beam shutoff interlock systems, and the automatic replacement of failed klystrons.

Introduction

The control system of the Stanford Linear Accelerator was designed with convertability to computer control as one of the main objectives. It was intended that such a computer should monitor status changes of the accelerator and its components, provide an operational log and allow analysis of operation. It also was to assist in programming of the beam, energy control, and beam guidance. Ten years ago computers were not what they are now, and the concept of computer control appeared to impose more stringent restrictions on the system than it does now.

One of the restrictions was that it appeared unreasonable to ask a computer to monitor directly the twenty-five thousand signals which would be required for complete monitoring of all equipment associated with the accelerator. A second restriction was that video signals could not be handled directly by a computer. A third was that it was inefficient to use a computer as a piece of switchgear if a job could simply be performed by hardware switching.

Data Handling

If the accelerator were instrumented like Stanford's Mark III accelerator, twenty-five thousand signals would be required in CCR. Two design decisions allowed this number to be reduced by an order of magnitude to a quantity which could reasonably be fed into a computer. The first decision was that reflex action should not involve equipment or people in Central Control. Thus, any signal which represents a modulator fault is used locally to turn off the modulator directly. The operator is informed about the action after the fact. The second decision was to summarize local indications; Central Control receives just enough information for the operator to determine if abnormal operation is due to equipment malfunction and for the maintenance personnel to identify the piece of equipment causing trouble.

Details are sent only as required for the operator to determine the appropriate corrective action. For the majority of equipment troubles, the only action he can take is to dispatch a technician to work on the equipment and to try to re-program the accelerator to work without the missing device. In general, then, one go/no-go signal for transmission to Central Control is derived for each piece of equipment. Since equipment can only be serviced locally, the details of the trouble are identified when the technician arrives at the equipment.

The equipment is normally operating unattended. All equipment is therefore designed to be self-protecting, to be stable, and if frequent operational adjustments are required, to be remotely programmable. Predictable or sequential adjustments, such as protective action in case of rf breakdown in the accelerator waveguide and the steps of the automatic phasing cycle, are programmed locally in order to reduce the number of individual commands required from Central Control.

To simplify eventual interface problems with the computer, a limited number of signal formats were specified. There are (1) binary status signals, to be used for logging equipment troubles and, if necessary, for initiating corrective action; (2) slowly-varying analog signals, to be scanned as often as prove necessary; (3) pulse-amplitude-modulated signals for monitoring the beam; and (4) control signals, all of which are momentary closures.

The original concept of computer control has thus led to a limited display of equipment status in Central Control, with a rigid format, and to extensive use of local control devices.

Data Logging

One "obvious" requirement in the control system was a means of automatically logging changes of status of the machine and its components. In fact, however, there is no automatic system for recording status or analog values for the accelerator. During the initial conceptual design, automatic fault-recording equipment was investigated, and the cost was found to run about \$50 per point; we had some 3000 points to monitor; for this cost, we could have a rather large computer, with control as well as logging capability. All consideration of special-purpose data-logging equipment was dropped.

Small computers were investigated (installed cost under \$50K) and found wanting. One process-control computer, for example, would have been saturated just monitoring and reporting changes of 3000 binary points. Input, detection of changes, editing and output would have occupied at least 80% of its time. A computer in the \$100K class would be capable of keeping up with the data-logging and other programs such as steering, energy control, and consideration of equipment status in the programming of the accelerator itself (in particular, reprogramming when a portion of that equipment failed during a run). A tape library of these additional programs might also include off-line

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

(Presented at 1966 Linear Accelerator Conference, Los Alamos, October 1966)

work associated with scheduling, maintenance and analysis of operational efficiency of various modes of operation. This size of computer, however, lay far outside our budget.

The net result is that we still have no automatic data-logging. Nevertheless, it has been found that the equipment is so stable that operator-scanning of the status signals is adequate for routine operation and trouble-shooting. The major exception is keeping track of the number of klystrons on-line at any moment. The klystrons cycle automatically after faults, but sometimes they do not come back on line. Analog values have also proved so stable that automatic warning of changes is of marginal utility. Typical analog signals used at CCR are modulator supply voltage (at present set to the same value for the whole machine and seldom changed from one day to the next), dipole steering currents, radiation monitors and readings of selector switch settings.

If the operator were concerned only with keeping track of current status of the machine, there would be no problem. But it is also necessary to keep a record of equipment failures, operating hours, and of settings to which it may some day be desirable to return. All such logging must be performed manually by the operator.

We find that much of their time is occupied with trying to record essential status changes, some useful details are still forgotten, and, particularly during set-up of a new beam, they cannot keep an adequate log and devote enough of their time to beam management simultaneously.

Beam Programming

Early in the design of the accelerator, it was suggested that half the klystrons of the accelerator might be operated at 360 pps and the other half of the klystrons at 60 pps. The accelerator would then deliver 300 pps at half energy and 60 pps pulses at full energy. A dc magnet would bend them in different directions depending on the exact energy difference. It was then suggested that there would be no trouble building a pulsed magnet which will divert the half-energy beam to one target area and the full-energy beam to another target area.

Thus the concept of multiple beams was born. The most attractive use of such interlaced beams is to allow one experiment to take data at a high rate while another experiment is being aligned, tested or calibrated. Also the beam for one experiment could be set up, at a low repetition rate, during the last half hour of the previous experiment without interfering very much with the rate of data collection.

The next "improvement" to the multiple-beam concept was to allow a bubble chamber to use one pulse per second at maximum energy, while allowing the klystrons to operate at no less than 60 pulses per second. The half-energy experiment might now demand every pulse left over from the bubble chamber experiment. This would be solved by providing a "normal" and a "delayed" trigger for the klystrons. The beams are obtained by again operating half the klystrons at 360 pps and half the klystrons at 60 pps - BUT: now all but one of the 60 pps pulses occur just-too-late to accelerate a beam. Thus, three different kinds of machine pulses would be programmed: most with only half the klystrons running (specifically 300 times per second), many with half the klystrons accelerating the beam and the other half of the klystrons running late (specifically 59 times per second), and one (per second) with all klystrons accelerating the beam. It is clear that the first two groups produce a half-energy beam at 359 pps; the third creates the full-energy bubble chamber pulse.

This led to a concept of controlled triggering of klystrons so they could be operated at a constant rate, and a "Pattern" which determined exactly which pulse would be delayed and which would be prompt.

Further improvements led to the installation of three grid modulators for the gun, so that successive beam pulses may be of different amplitude or even of different pulse length.

Thus one might consider steering and tuning up a very-high-current-beam by having two beams set up: One, the low current beam of the desired energy which is left unchanged and one, the beam which is to be tuned up with heavy beam loading. At any time nevertheless, if the beam be lost, one can instantly switch back to the low current beam which has not been disestablished, and check if the beam loss is because of the properties of the beam or because of some trouble in peripheral equipment.

As soon as one grants the possibility that there may be many beams (e.g., null beams, alternative intensities of the same beam, alternative rates of the same beam), it becomes clear that one can readily switch back and forth between two beams identical but for some selected parameter, so long as the parameter is provided by a piece of equipment whose characteristics can be changed from pulse to pulse.

Equipment now exists for programming as many as six interlaced beams. The programming is accomplished with over 600 toggle switches. The settings of the toggles must be changed if additional sectors must be used to maintain an existing beam, to change its repetition rate, to change the mode of interlace of various beams, or to change the basic repetition rate of various groups of klystrons.

Computer assistance would be desirable, both for initial programming of a beam of specified energy, current and pulse rate, and for modifying the program when, for example, an entire sector of 8 klystrons suddenly becomes unavailable. A good-sized computer could handle the entire programming job with ease; the toggles represent directly the registers the computer would have to set. But each beam is programmed by an independent row of toggles. The operators can easily set them to produce a desired beam, so long as they need worry only about one beam at a time. No one has yet been rash enough to try to program more than three simultaneous beams, probably because no one has yet performed an experiment which requires more.

Beam Guidance

The beam position monitoring system uses three rf cavities in the drift section at the end of each sector. One is a normal pillbox-mode cavity; its rf output is proportional to the instantaneous beam current. The other cavities operate in the separator mode: they have no output when the beam is centered in the cavity, but have an output proportional to the product of beam current and average beam displacement when the beam is misaligned. One measures horizontal displacement, the other, vertical. None of the cavities is particularly sensitive to the diameter of the beam. The three rf pulses are combined in rf bridge circuits, detected, integrated, and normalized, using a logarithmic amplifier, and finally transmitted to Central Control as a sequence of three pulses, about 500 μs each, representing horizontal displacement of the center of the beam, vertical displacement, and the logarithm of the total charge in the last beam pulse. The signals from the drift sections, from the injector and from the Beam Switchyard are delivered on 36 wire pairs to a multiplexer which converts them to a pulse train of 36 10 μs pulses representing horizontal position, followed, after suitable intervals, by 36 vertical positions and 36 log Q signals. This output is fed to three scope displays, each of which is intensified during the appropriate (H, V or Q) interval.

The quantity of information is unreasonably large to digitize and feed into a computer for analysis; nor was it ever intended to do so. The pulse-by-pulse monitoring of beam information is dictated by the multiple-beam capability of the accelerator. Each successive beam pulse may be a different beam, of different energy, current or even charge. The assignment of successive machine pulses to different beams may be entirely arbitrary. Averages of beam information are therefore quite useless unless one knows that only information pertaining to one beam is being averaged. Selective filtering, by ignoring all pulses but those for a particular beam, can be performed in Central Control. This is already being done to allow displacing the baseline or suppressing the display of any beam on the monitor scopes. The filtered information could then be averaged, digitized, and used for computation quite readily.

The problem of teaching a computer to steer the beam is an interesting one, but is economically unfeasible. It has been found, for example, that an operator can make minor adjustments required to steer one beam to within 1 mm of the accelerator axis in less than one minute, two beams of different energies to within 2 mm of each other in five minutes, and three beams to within 4 mm of each in about 15 minutes. During less than 5% of this time is any beam lost due to oversteering.

Video Signals

It was presumed that a computer would not analyze video signals directly. Some sort of predigestion must be performed. Typically, a video signal will be integrated or analyzed by a sample-and-hold circuit; the resulting analog output might then be used as an input to the computer. An early proposal for an extensive video system was turned down because of its high cost. It was then decided that all video signals were to be conditioned locally. Except for video signals originating in Sector 27 (adjacent to the Control Building) and in the Beam Switchyard, only conditioned signals would be transmitted to Central Control. Locally conditioned signals include klystron modulator voltage, current and reflected rf power, which are used for local protective circuits; rf phase, which is used locally in the automatic phasing system; rf output power, which is transmitted to Central Control as a status signal indicating that the klystrons and modulator are working satisfactorily; beam charge per pulse and horizontal and vertical position signals, which are sent from each sector to Central Control for beam guidance, and the signals from secondary emission foils adjacent to the energy-defining slits for presentation of the beam spectrum.

Only recently has any video system been installed. Its function was to allow presentation of certain injector signals to the operator for comparison with the observed beam at the end of the accelerator and to permit transmission of television signals from profile monitors at various points along the accelerator.

Because a great deal of the initial testing has been performed using a bending magnet at Sector 20, the video current pulse originating in Sector 19 has been used in place of the signals from the end of the machine. The intensive testing associated with beam break-up has also resulted in a greater demand for video signals than was anticipated. Thus the original concept of conditioning all signals into forms suitable for direct insertion into a computer has resulted in a video signal system which has appeared, at least for initial testing, to be somewhat under-designed.

Existing Computers

I must not suggest that there is no computer control at SLAC. The beam switchyard area has an SDS 925 computer which scans status of all interlock signals in the switchyard, reports changes, monitors selected analog signals on demand, set the magnets of the switchyard to the correct field strength for a specified energy, and communicates with the operator via typewriter. Its input language reads somewhat like English and its compiler is remarkably tolerant of typing errors and readily accepts corrections. Not over 30% of the computer's time is occupied with these matters, but unfortunately, at last report, only eight addresses of its 4K memory remain for the programmer to utilize. Another 4K memory unit, on order, will allow addition of a data link to a computer in the counting house associated with end-station A.

The counting house computer is an SDS 9300 with 32K of core memory, a disk unit and 4 tape units. It is to be used to log experimental data, cont ol magnets, test electronics and will still have capacity for some on-line data analysis, such as displaying histograms. It might be noted, however, that only one particle identification is to be recorded per beam pulse and that much of the input data from counters is to be analyzed in analog circuitry before being digitized and fed to the computer. The computer is already in use compiling wire-lists for connections of experimental equipment.

Summary

Turn-on of the machine is reasonably quick: under favorable circumstances it is accomplished in under half an hour once the beam areas have been cleared. An half-hour is perhaps required for establishing a new beam, once the first is running. But these times would be reduced significantly if the operator were relieved of his clerical duties.

The form of the Central Control data displays, the extensive use of automatic control circuits outside Central Control, and the sketchy video system have their origin, in part, in the requirement for easy conversion to computer control.

Of the four jobs initially foreseen for a computer, but one remains as a problem. Steering has proved easier than anticipated (primarily, I believe, due to the remarkably good alignment of the machine) and can be readily handled by the operator. Programming of the accelerator and energy control, while somewhat complicated, are also techniques which have been learned by the operators. Computer assistance would be a luxury which would be appreciated but which could be justified only as a byproduct of some more pressing need. Off-line analysis of machine and equipment performance can be accomplished on any computer. Performing such functions in the control room could only be rated as a pleasant luxury. But the original problem – the logging of changes of machine status – still remains to be solved.