R. T. Braden
CGTM 14
May 1967

CONCERNING THE REAL-TIME COMPUTING ISSUE: CENTRALIZATION vs. DECENTRALIZATION

Real-time computation will play a vital and increasing role in the activities at SLAC. Computers currently at SLAC are being connected and programmed to gather data on line and to exert control over spectrometer and wire chamber experiments, measuring machines, CRT, film digitizers, and the BSY. Trends at SLAC and at other laboratories indicate that such real-time computing will rapidly proliferate at SLAC.

As each new experiment is proposed with a real-time computing need, the experimenters will have to take a stand on the issue of centralization: should they obtain their own computer and dedicate it to their experiment, or should they hook up to the central machine and use time-shared cycles as required? This note sets forth my own philosophy on this subject, as it relates particularly to the existing situation at SLAC.

When this issue is discussed (vigorously) between proponents of centralization and those of decentralization, the discussion rages around two extreme models. The centralist model would have just one extremely large central computing engine, gathering data direct from apparatus in the end stations and other sources, manipulating this data in a time-shared fashion, logging the data onto its tapes and disks, and driving displays and other output devices to summarize the data for the experimenters. The extreme decentralist model would have a large number of computers of small to moderate size, each dedicated to a particular task. This issue in real-time computing mirrors a much larger and more general competition in the computer world between large centralized time-shared machines and smaller local computers. On the one hand, recent studies have shown clearly the large economies-of-scale which exist, for example, in the IBM 360 line of computers.[1] Empirical data[2] have been found to support the computer industry folklore that the cost of a computing system is proportional to the square root of its computing power. Thus, economics, plus the very real advantages of centralized files of programs and data as well as sophisticated programming systems, favor the centralist view. On the other hand, new technologies have dramatically reduced the cost of electronic circuitry, so one can now buy for $50,000 computing power which would have cost over $500,000 ten years ago. As the hardware cost continues to decrease, decentralized computing becomes more and more attractive.

To return to the real-time issue at SLAC, let us list the apparent disadvantages of each of the two approaches.

Disadvantages of centralized real-time computing.

1. The greatest disadvantage arises from the potential unavailability of computing cycles from the central machine. At the very least, distributing cycles will entail scheduling computer time as well as beam time, and will, therefore, lead to bureaucratic frustrations for the physicists. Furthermore, any such scheduling of resources will inevitably imply situations where the physicist is unable to operate because of schedule conflicts. In general, the sharing of a central machine implies managerial and bureaucracy difficulties which the physicist would not experience if he had his own dedicated computer.

There is another availability problem, that of downtime on the central machine. It is unfortunately true that the failure of almost any component of the SLAC central computer will take the whole system out of operation. The duplexing of computer hardware for high reliability and continuous availability is expensive and, furthermore, it tends to negate the economy-of-scale factor which is one of the arguments for using a central computer. In summary, the physicist who chooses to use the central machine for his real-time computing puts himself at the mercy of a big, complex mass of electronics which is outside his ken and control. Furthermore, there will be times -- hopefully seldom -- when the central machine is out of commission for significant periods of time, so that the experimenter is not able to use the beam time available to him.

2. Another disadvantage of centralized computing is the introduction of communications difficulties. The computer will be located at some distance from the experiment, and, therefore, the experimenter will have to rely

upon teletypes, displays and telephones for communication with the central machine and its operators. The result is to put additional demands upon the software.

3. Finally, there is a psychological problem for the physicist who does not have his own private computer, but instead, lives in (paranoid) fear that his program is being subtly and mischieviously affected by other users of the central system, and that he ultimately does not have control of his own experiment.

Disadvantages of decentralized real-time computing.

1. The most significant disadvantage is the cost. It is my belief that in almost every circumstance it will cost more to distribute the real-time computing load among a number of smaller dedicated machines than to handle the same load on a single centralized machine. The savings realized with the centralized machine comes both in the hardware cost and in the personnel cost. It requires more system programmers to handle the decentralized machines than will be necessary for the one central machine. Furthermore, the very real economies-of-scale in computer hardware (speed $\propto$ cost^2) mean that the cost per cycle for a single large machine will generally be significantly less than for smaller machines. Some, but not all, of this apparent economy will be lost because of overhead in the more complex software necessary for the centralized machine.

2. There is also a loss of flexibility when decentralized machines are used. Once such a dedicated computer has been selected and purchased, the physicist is forever limited to the capacity of that machine in doing his experiment. On the other hand, there is potentially considerable flexibility available in the fraction of the big machine which he uses, and therefore, as the experiment progresses, the magnitude of his real-time load can grow more (subject to the demands of other physicists) on the big machine.

In almost every case, physicists have ordered (and probably will continue to order) real-time computers with half of the core memory which they will ultimately need for their job. Here again there is more flexibility with the central machine whose very large core memory gives the physicist's program room to grow as his needs(and aspirations) increase.

I can summarize my argument as follows:

In a tight money situation, the most economical way to handle many real-time applications is through a shared central machine. Individual decentralized machines are frequently more convenient from the point of view of the physicist, but they cost more money.

There will always be some applications which simply can't be done by the central computer because of excessively short response time requirements or high data transmission rates. At the other end of the scale, there may be certain applications for which a decentralized machine is simply not sensible because of cost or acquisition timing. As the fiscal situation eases, a wider class of real-time problems should be handled by individual machines.

The decision on centralized vs. decentralized computing should be made on an individual basis for each experiment proposal at SLAC, in the light of the current fiscal situation, the current software situation, and previous experience with other experiments, both off-line and on-line.

References

1. "Economies of Scale and the IBM System/360," Solomon, M. B., Comm.ACM 9, 6 (June 1966).

2. "Changes in Computer Performance," Knight, K. E., Datamation 12, 9, (September 1966).