

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

A SPECIALIZED, MULTI-USER COMPUTER FACILITY FOR THE HIGH-SPEED, INTERACTIVE PROCESSING OF EXPERIMENTAL DATA**

Creve C. Maples*

Abstract

A proposal has been made at LBL to develop a specialized computer facility specifically designed to deal with the problems associated with the reduction and analysis of experimental data. Such a facility would provide, a highly interactive, graphics oriented, multi-user environment capable of handling relatively large data bases for each user. By conceptually separating the general problem of data analysis into two parts, cyclic batch calculations and real-time interaction, a multi-level, parallel processing framework may be used to achieve high-speed data processing. In principle such a system should be able to process a mag-tape equivalent of data through typical transformations and correlations, in under 30 sec. The throughput for such a facility, assuming five users simultaneously reducing data, is estimated to be 2 to 3 times greater than is possible, for example, on a CDC7600.

Background

Although advancing technology has succeeded in making modern computers faster and more efficient with respect to processing speeds, there is at least one class of data processing which has seldom been addressed directly in terms of both hardware architecture and system design. This is the general problem of highly interactive data processing. Such problems are typically characterized by reasonably large data bases, perhaps in the 60 to 500 mbyte range; by potentially requiring arithmetic transformations and many levels of conditional testing on each word of the data base; and finally by the necessity of having a person interactively involved with the process in order to evaluate the results of each pass over the data base and to determine the conditions of the next pass.

Interactive computing problems of this type are frequently encountered in processing the results of scientific experiments, particularly in the area of nuclear and high energy physics. At present there are two main approaches available to a user with such processing requirements: the use of a large, central computer facility, or a small dedicated computer. The large computer facility offers the advantage of large mass storage capability and fast processing speeds. Such systems, however, basically operate in a batch environment and are, in general, not easy to adapt to a high-speed, graphics-oriented, interactive environment. The necessity for optimizing system resources and CPU processing in such large facilities makes it difficult to guarantee an interactive user both memory and processing time, on demand and on a continuing basis.

A dedicated, smaller computer, on the other *Lawrence Berkeley Laboratory, University of California, Berkeley, CA 94720. **Supported by the Nuclear Physics Division of the U.S. Department of Energy.

hand, can be much more responsive to the dynamic demands of an interactive user. The smaller computers do not, however, have either the processing speed or bulk storage capability of the larger facility. Thus a single pass over a data base which requires several seconds to several minutes to process on a large computer (e.g., CDC7600) may require between 5 minutes and several hours to process on a smaller machine. A recent paper discussing the data reduction requirements at one national laboratory indicated that from 200 to 1000 hours of computer time were required to process the results of one experiment.¹ Recent advances in technology have significantly increased the processing and mass storage capability of modern, mid-range computers. Unfortunately the data processing requirements, particularly in the area of experimental nuclear and high energy physics, appear to be increasing much faster than technology. This is evidenced by increases in both the amount and complexity of the information.

Problem Definition

Experimental data normally consists of the detection of a physical occurrence (event) and the measurement of various various types of information (parameters) associated with the event. Currently in experimental physics the number of parameters per event tends to range from several to a few hundred, although new detection systems, currently planned or in development, extend this number into the thousands. Because of the enormous number of possible correlations and permutations present in such data, it is necessary for a scientist to be intimately involved in the reduction and analysis process. In this way the path of the analysis may be guided at each step by the scientist's intuitive and subjective evaluation of the results up to that point.

Currently the critical path in data analysis is getting the information back to the user for evaluation. In a small computer environment this is typically limited by CPU processing speeds. At a large computer it is determined by the time between job submittal and execution and by the time necessary to produce hardcopy of the results and get it to the user. Ultimately, however, the critical path is the user himself, and his ability to understand the relevant information contained in the complex, multi-dimensional data.

Proposal

A proposal has therefore been made at the University of California Lawrence Berkeley Laboratory to develop a specialized, highly interactive computer facility to handle general data analysis problems. The purpose of this facility is to provide a highly effective interface between the computer and the user and in this manner to create an environment designed to optimize the

Handwritten signature or initials.

effectiveness of the user and to hopefully minimize the real time necessary to complete a given analysis problem.

The proposed facility, although possessing all the attributes of a conventional mid-sized computer, is designed to optimize the problem of data analysis. To this end the following assumptions will be made:

- 1) That the users' data pre-exists on some physical media—normally disc or tape (a possible exception to this requirement will be discussed under the section of future possibilities).
- 2) That the data base is made up of packets of information (events) of either a fixed or a varying number of parameters.
- 3) That each event packet is essentially independent of other events, with the exception of a possible real-time correlation which could be implicit in the general ordering of the data (thus this does not exclude the possibility of events being organized into time-dependent regions, such as by a beam burst).

- 4) That for a given scan of the data base, the processing of each event packet or region is well defined and essentially repetitive (intraevent iteration and time or position interevent effects are permitted).
- 5) That the information resulting from one or more passes over the data base may be displayed to the user in a number of graphs of 1- or 2-dimensional plots.

In order to provide effective interaction from a user perspective, the design goal of the facility is to allow data processing at an average speed of approximately 4 μ sec/word for typical analysis problems. This processing includes data retrieval, transformation, correlation checking, and sorting into final displays. This average speed would correspond to processing the equivalent of 1 mag tape of information in approximately 30 sec.

Initial studies suggest that such throughput may be obtained by using current, state-of-the-art, microprocessor technology in a parallel processing environment. The overall layout of a proposed five user facility is shown in Figure 1. The general

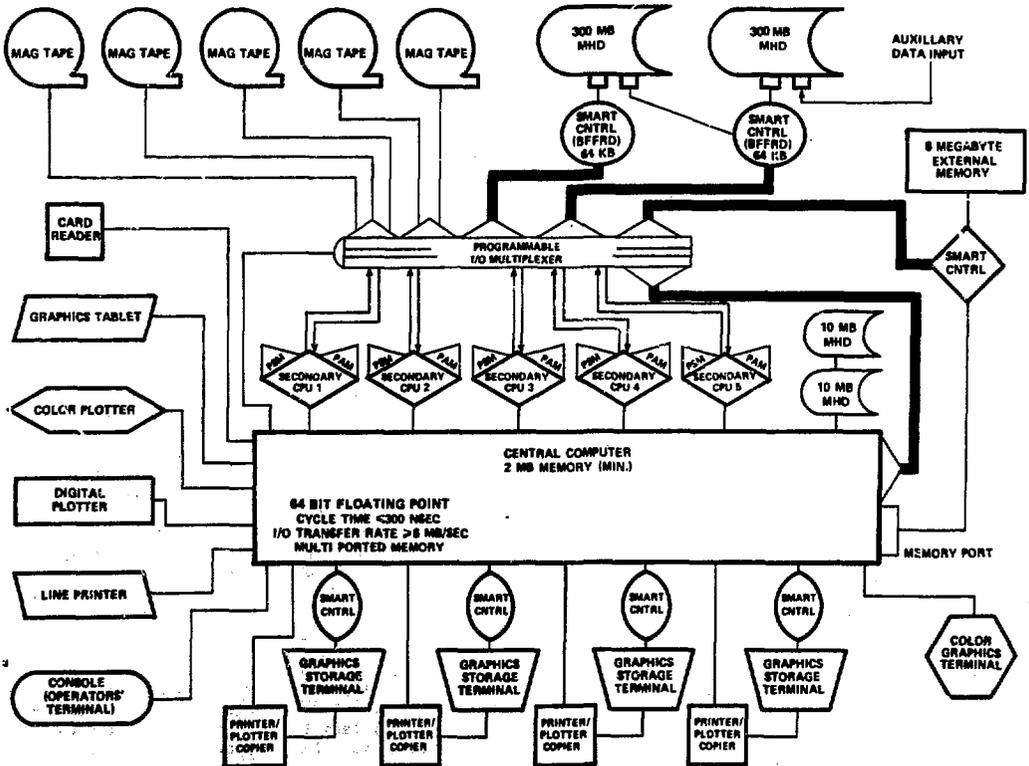


Fig. 1. Layout of proposed Interactive Data Analysis Center.

discussion of this design will be broken into the following sections:

1. User Interaction
2. Mass Storage
3. High-Speed Processing
4. Data Transmission
5. System Communications

1. User Interaction

One of the most important aspects of this proposed facility is the ease with which users may communicate with the computer. To be effective this communication channel should place strong emphasis on graphics displays and minimize the necessity for typed input. The graphics terminal should be capable of reasonably high resolution (1024 x 1024 minimum) and should be capable of sustaining a substantial amount of information on the screen at one time, flicker-free. This requirement is necessitated by the typical multi-dimensional nature of the data. In order to simultaneously examine transformations and various projections, it is important that the screen be capable of displaying many 1- and/or 2-dimensional displays at one time. This could mean up to about 30 1-dimensional or 10 2-dimensional displays.

At the present time a storage tube appears to be the only practical unit capable of meeting this requirement. It is also felt that the screen size should be at least 19 inches (equivalent to a Tektronix 4014). A refresh terminal, with independent local refresh capability, does, however, offer some significant advantages over the storage terminal in terms of graphics capabilities. Modern refresh terminals, coupled with large local memories, are approaching the practical ability to maintain on the screen, flicker-free, the levels of information density expected.

Another difficulty associated with graphics display terminals is the transmission speed of information to the terminals, typically 9600 baud. Experience has indicated that this is approximately one order of magnitude too slow to serve as an effective interface in a highly interactive environment. For this reason it is planned that the CPU-to-terminal link, shown in Figure 1, be a high-speed DMA or communication link capable of graphics transmission at a minimum rate equivalent to at least 120K baud.

As indicated previously, the slowest link in the system will be the users' ability to perceive and understand complex interrelationships which may be contained within the data. One way to enhance the users' perceptive ability is to add another dimension to the interactive terminal--color. Color displays, particularly in the area of 2-dimensional information, can make the data far easier to understand. At least one such full color terminal, with a minimum of 512 x 512 resolution, is planned for the proposed facility. Since rapid advances are being made in the area of color displays, coupled with declining cost, the final design may, in fact, contain more color units.

graphics displays, but a fast and efficient means of communicating a user's requests to the computer. One method of dynamic communication, which has proven highly effective, is a series of switches and buttons which can be monitored by the computer but which have no intrinsic hardware function. Such switches, under software control, can supply an enormous spectrum of user oriented functions, can easily be adapted to highly interactive screen controls, and can change function as the needs and the programs change. Thirty-two such general purpose switches have been installed on many of the terminals on the data acquisition computers in LBL's Nuclear Science Division and have proven highly effective. Every interactive graphics terminal on the proposed facility would be equipped with these. Other highly interactive devices, such as touch panels over the screen, and a "mouse" (movable ball-bearing) are also being considered.

One final requirement is the ability of each user to generate personal hardcopy output. Capability should exist locally at each terminal for a user to produce printed output, screen copies, and simple plots. A matrix type printer/plotter device is therefore supplied at each user station (Figure 1). The ability to produce hardcopy of the terminal screen is considered important and therefore will also have an effect on the specification of the graphics terminal.

2. Mass Storage

The mass storage requirements may be broken down in several ways. First the system requirements for program storage, operating systems, libraries, etc., are filled by small 10-20 mbyte moving-head discs. Figure 1 indicates two of these units in the basic configuration with expansion capability of between 4 and 8. The reason the smaller discs are favored is two-fold: the ability to backup or archive discs on a single magnetic tape; and to provide an easy media for users to independently maintain their own libraries, programs, etc.

Long term storage requirements for data has traditionally been magnetic tape, so the proposed facility provides at least one unit per user. Multiple tape controllers are used so that more than one drive can be transferring data simultaneously. It is anticipated that the primary use of these units will be the transfer of information to disc for short term storage until processing and analysis is completed.

The mass storage requirements for active data is divided into two classes--that information requiring high-speed random access and that requiring high-speed sequential access. It is expected that sequential information would normally consist of either the raw data or a transformed data base. This information would be stored on disc drives of at least 300 mbyte capacity or greater. These drives would be specially configured for data storage only. The drives would be dual ported to permit disc-to-disc transfers without involving the main CPU. Figure 1 indicates at least two of these drives would be required in the basic configuration and each drive would have its own

Effective interaction not only means effective

specialized controller (discussed in the section on Transmission Rates). Each controller would be capable of handling 4 such drives and the basic design framework is capable of having up to 4 controllers.

It is frequently important to have a large storage area which is capable of high-speed random access. The obvious answer to this requirement is memory. Since this memory is to be used for data storage only (arrays, multi-dimensional histograms, etc.), it is far more cost-effective to utilize bulk memory storage external to the main-frame computer. This unit should be capable of randomly accessing a word of its memory in 1 μ sec or less. An advantage of this bulk memory is that its word size can be chosen for optimum effectiveness (i.e., 24-bit memory might be used). Figure 1 indicates the basic design contains at least 8 mbytes of such memory. The controller should be capable ultimately of handling at least 32 mbytes.

3. High-Speed Processing

The discussion of processing will be broken into 3 levels: the central computer; the secondary CPU's; and the third or array processor level. The central computer should be a modern mid-range computer, capable of handling at least 2 mbyte of memory. The computer should be as fast as possible, with basic instruction times (register-to-register) in the 200 ns range. It should be able to carry out basic arithmetic operations with 32-bit words and should have a multi-ported memory.

The function of the central computer is primarily as a supervisor and overseer of a distributed processing environment. It serves to allocate resources and determine priorities. Most importantly it handles all user communication. In this distributed environment, unlike many conventional networks, the central computer has the primary intelligence and decision making capability. The main computer is always aware of the status and current function of all other attached CPU's since their operation is dictated by the central unit. Thus this is a master-slave relationship as opposed to a distributed intelligence.

An important function of the central computer is to minimize the load on its own resources. Thus in handling communication with users in an interactive environment, e.g., by providing requested display buffers, monitoring switches, cursors, etc., the required response time is within human reaction times and presents no great demand on the main CPU. The main computer must, however, be able to distinguish the interactive, or non-time-critical requests, from the CPU-intensive, batch operations. This should not be difficult in the case of data reduction since, at some point, a user must request an analysis of the data base according to some previously described criteria. At this point the central computer hands this CPU-intensive batch "task" to an available, dedicated, secondary CPU unit.

These secondary CPU's are typical modern min-computers, capable of running at least 128K bytes of memory. They should have a flexible instruction set, preferably a multi-ported memory, and be easy

to interface to specialized hardware devices. These mini's are basically stripped CPU's with essentially no dedicated standard peripherals. (A possible exception to this in the future might be the installation of a small local mass storage device, such as a floppy disc, on each unit.) The function of these units is to take a single specific task (e.g., a set of analysis conditions) from the central computer and run it until it is finished or until it is instructed to stop. The layout in Figure 1 indicates there should be one of these units for every data analysis terminal.

Since the secondary units are unable to actually execute an analysis problem within the time frame required, their primary function is to further subdivide the problem into separate stages of processing and to load each stage into specialized array processor hardware. Having completed this task, the secondary CPU then starts the data transfer and processing and serves to monitor the overall throughput of the operation.

The third level of specialized array processors permit the system to achieve the high level of processing speeds desired. Since each event packet is essentially independent, and each will be processed through the same analysis sequence, high effective processing speeds can be obtained by the parallel processing of a number of event packets at one time. Rather than attempt to design an array processor which performs all possible functions it was decided that a modular approach, with specialized array processors, would be more flexible and cost effective (see section on Future Possibilities). The basic design (Figure 1) divides the batch problem into two phases, transformation and correlation, and requires two specialized array processor units, a PAM and a PSM, for each secondary CPU.

The problem of data transformation is frequently encountered in analysis problems. A number of parameters within an event are combined, through an appropriate algorithm, to create new parameters which may be added to or replace the original information. For this purpose a Programmable Arithmetic Module will be used. This processor has a specialized, arithmetic instruction set, and the necessary transformation operations are loaded into this unit by the secondary CPU. The transformation of the basic data is then carried out, in parallel, on different event packets according to the fixed procedure, and transformed event packets are produced.

Because most physical data acquired in physics are integers of limited accuracy (usually less than 13 bits), it was felt that the PAM unit could be an integer, as opposed to a floating-point, processor. As long as internal registers of at least 48 or 64 bits are used, results within 13 bits of accuracy can easily be obtained. In addition many operations which traditionally require floating-point operation, trigonometric functions, roots, logarithms and floating-point exponentiation, can, due to the limited accuracy of the data, be carried out by means of table lookup techniques. To this end the PAM unit contains special memory to hold such tables. These tables would then be filled once with the necessary information by the

secondary CPU at the beginning of a processing sequence. (For operations requiring floating-point arithmetic, see section on Future Possibilities.)

The Programmable Sorting Module (PSM) performs a completely different function. It scans an event, (either raw or transformed) looking for correlations specified by the user. Each correlation might typically consist of a variety of logical .AND., .NOT., and .OR. type of operations on conditions within or between parameters in an event. Normally a user might request many such correlations be examined within each event. The result of such tests would be whether or not each set of correlation conditions were met within a particular event. A special class of correlation which must be handled by this unit is a functional or 2-dimensional correlation. This would occur, for example, by constructing a map of the intensity of parameter x vs y (with intensity in the z-direction) and dynamically selecting a region (non-rectangular) of this map to correlate with another parameter u. In this case the correlation condition must be expressed as a function of both x and y.

The PSM will obviously have specialized masking, bit testing and branching capabilities. The correlation conditions would be loaded initially by the secondary CPU. During analysis the PSM would operate in parallel on several events at a time. Also, of course, it would be operating in parallel with the PAM unit. Since all fast, batch-type processing is handled on a secondary CPU together with its third level of array processors, it is obvious from Figure 1 that, from a CPU standpoint, all users may process data simultaneously and in parallel with minimum impact on each other.

4. Data Transmission Rates

Although CPU processing speed is usually the limiting factor in data analysis, transmission rates may also impose limitations to high-speed processing. One technique to enhance transmission speeds is to overlap data transfer with data processing. To this end it is proposed that specialized disc controllers be developed to handle the large data discs. Such controllers would be fast, specialized computers with at least 64K bytes of local buffer. These controllers would have the task of handling all data transfers to and from the disc and would be completely responsible for determining the file structure and layout of the data discs. Thus the disc (or disc.) on which a given set of data is placed, its actual location and geometry, is determined by these controllers to optimize retrieval.

In order to use these discs to read data, a setup call would be required. This call would inform the controllers of the data base desired, the total number of words to be read and of the size of the internal buffer into which the data would be placed. A special read command would then be used to retrieve successive buffers of data. In this fashion the disc controller can anticipate future requirements and can continue to read information from the disc into its own local buffer for high-speed DMA transmission to the next available memory buffer. Each controller will probably have

more than one bus available to it. Additional advantages of the independent, intelligent disc controllers will be discussed in the following section on system communication.

A specialized, intelligent controller, similar to those for the data discs, but without the large memory buffer, would be utilized for the external memory module. This controller should also support multiple paths to the memory and would be responsible for dynamically mapping and allocating memory according to the demands of the users. Normally such operations might be handled by the operating system of the central computer. By relegating all such functions to an intelligent controller, the impact on the operating system is minimized and the flexibility of the bulk memory, both in terms of ease of usage by other processors and/or in terms of any subsequent hardware or software design changes which might be desired, is enhanced.

The intelligent controllers for the graphics terminals (Figure 1) are designed to take the routine problems of screen updating and refreshing (and possibly simple display functions) away from the central computer. It may well be possible to purchase such controllers with the terminals themselves. In the case of a fully-buffered, refresh screen, or the color terminal, this is certainly true.

5. System Communication

The basic philosophy behind the design shown in Figure 1 was to isolate every expected problem area to a separate, almost stand-alone, module. It was hoped that in this manner the standard operating system for the central computer could be utilized with minimal modification. The central computer should be the only CPU in the system which must really deal with multiple processing. Every other processor in the system has essentially only one job to do, either fixed, as in the case of the controllers, or variable, as in the case of the secondary CPU's. In an attempt to further simplify the interaction between various CPU's, a switching unit called a Programmable I/O Multiplexer is proposed.

This I/O multiplexer is designed as a large, multiple-bus switching unit, under the control of the central computer. In a typical operating situation a user would indicate to the central computer the location of his data, the number of words to analyze, the details of the analysis conditions, the type of results and possibly where they are to be stored. The central computer would verify that all requests were valid and would select an appropriate secondary processor unit to handle the problem. Using the I/O Multiplexer, the central computer would connect the appropriate transmission lines between the designated mass storage device and the input channel of the selected fast processing unit (secondary CPU plus third level processors). In a similar manner, the output lines would be connected. At this point the central computer would signal (interrupt) the selected secondary CPU to stop all operations and prepare to receive new analysis conditions or code. This operation is not particularly time critical.

and can be carried out at medium speeds. After completing transmission the central computer will instruct the secondary unit to begin processing.

At this point the main computer no longer directly concerns itself with the execution of the problem. The secondary CPU, as indicated before, loads the conditions into the appropriate array processors, begins the data transfer and analysis and monitors the progress. Should some problem occur, the secondary will report the condition to the main computer and await instructions. The main computer will periodically scan the status buffer associated with this task to determine whether it is executing, finished, or an error condition has occurred. In the latter case the central computer will take whatever action appears appropriate-- attempting to rerun the problem, instructing the secondary to run diagnostics, switching to another secondary unit, notifying the user, etc. At no time do the two units get involved in a time-critical or asynchronous dialogue, nor does the central computer ever attempt direct communication with the array processor units.

Assuming that all appears to be proceeding normally, the main computer will periodically wish to verify the progress of the analysis, both for itself and for the user. To this end the computer will perhaps wish access to the results being obtained in order to display these partial results to the user. In this or any situation where the main computer wishes access, for whatever reason, to devices currently in use, the computer takes that access directly without the necessity of communication with the processor currently using the device. This is accomplished by disconnecting the processor from the device in the I/O Multiplexer. A fast processing system has only one task to complete. If the units are disconnected from a data channel they will simply wait until the channel has been restored to continue processing. In this manner the main computer can obtain access to any device as it decides it is necessary and with a minimum of overhead. The user can also see continuing updates of the status of his analysis problem to verify its progress.

Another area where system communication and involvement is kept to a minimum is in the area of the smart disc and bulk memory controllers. With respect to the discs, the main computer need not concern itself with the file structure or partitioning of the data discs. It needs only to be able to obtain the status and mapping information when needed, primarily for users or diagnostic purposes, and to be able to obtain data from the disc if desired. Since the high-speed disc operations proceed independently of the central computer it becomes much easier to try new techniques for improving access and transmission. A controller can handle simultaneous read and write requests from two processors by, for example, storing the write in its local buffer while letting the read proceed.

Another area which may be investigated independently by such controllers is the problem of gather-reads and scatter-writes to the discs. Considerable work has been carried out in studying ways to create an efficient transfer algorithm.

Typically this problem has been handled by special programs. In this framework the procedures could be utilized in the controllers themselves to make such transfers more efficient. This would have the effect of making the techniques immediately utilisable by both the operating system and existing programs, and essentially transparent to both.

Similarly the intelligent bulk memory controller minimizes interaction and communication with the operating system. Basically the operating system needs only to be able to read the memory when necessary, and to obtain status information on demand. In a manner similar to the discs, the memory controller only needs to be informed which units are currently permitted access, and the amount of memory requested by each unit. The problems of actual memory management, diagnostics, dynamic mapping, allocating and deallocating memory, are then functions of the controller and need concern neither the operating system nor the individual processors.

Future Considerations

The general design proposed for this data analysis facility, as shown in Figure 1, attempts to be as modular and flexible in nature as is practical within the indicated goals of the facility. Rather than attempting to anticipate all future analysis requirements, the design emphasis was placed on the identification and isolation of those potential areas likely to lead to future restrictions in the performance of the facility. It was felt that by isolating potential problem areas, in both hardware and software, the system could easily continue to develop along the lines dictated by the demands placed on it.

An example of this implicit growth capacity is the third level of specialized, high speed array processors. There is no intrinsic reason a secondary processor can not have more than two of these units; there is no restriction as to the function of a specialized array processor; and there is no reason why secondary CPU's need to have either the same number or class of processors. If, as will undoubtedly occur, the necessity arises for floating-point processing, these units could be obtained, either by construction or commercial purchase, and also be attached to the secondary CPU's. The software modifications would then essentially be confined to the secondary CPU's, which need only to be instructed in the loading and operation of this device. (Some user interface subroutines would also be required in the central computer libraries.) If the demand for floating-point array processing only comprises, for example, 20% of the facility's work, one could consider installing only one of these units on a single secondary CPU. In this situation, the central computer would simply be made aware that a particular processing station had special capabilities and would allocate it accordingly.

There is also no reason why these array processors must have a completely general function. If future usage indicated specialized requirements in a particular area--multi-parameter ray tracing

through a magnetic field, for example--there is no reason that a specialized processor could not easily be constructed to perform this complex operation. The general interfacing criteria of such units would already be indicated by the existing processors. The cost of these units should also be relatively inexpensive. The current anticipated equipment cost of the planned PAM and PSM units should be well below \$10K per unit.

An overall percentage breakdown of the equipment costs for the proposed basic facility is shown in Figure 2. It is obvious from this figure that the third level of array processors comprises only a small percentage of the total cost of the facility. The function of such processors is relatively simple and is well defined. If, for example, technological advances within a subsequent 3-4 year period indicated that substantial improvement in throughput could be achieved by replacing the existing array processors, this replacement would be rather straightforward and relatively inexpensive. Additions or replacements could also be phased since every secondary system is independent.

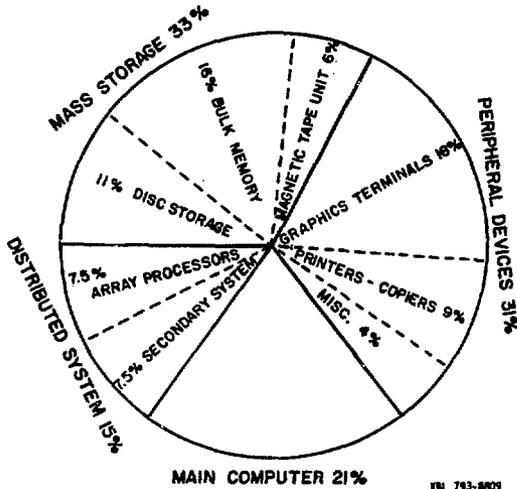


Fig. 2. Percentage Breakdown of Equipment Costs

Advances in chip technology rapidly give way to advances in computer technology. On a somewhat longer time scale, say 5-6 years, replacement of the secondary CPU's or special controllers might be considered. Again, the function of these units are well defined and the percentage investment, relative to the total system cost, is small. In a similar fashion advances in mass storage technology should also be easy to accommodate. New devices, such as video-discs or bubble memory, for example, could be added to the system by constructing or programming another special controller, basically similar to the existing controllers. Since these controllers already operate independently and are viewed similarly by the operating system, integration of a new controller should present no special software problems either.

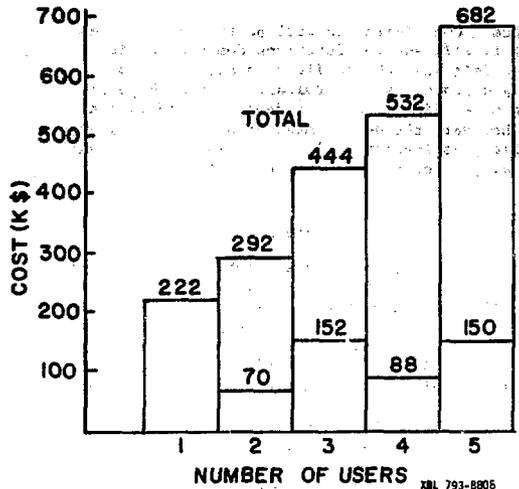


Fig. 3. Equipment cost of facility vs. number of simultaneous high-speed, data analysis users supported.

An interesting future potential is the possibility of monitoring live data acquisition almost directly. Although there is no intention of equipping such a facility with on-line acquisition hardware, there is no reason why a real-time experiment cannot be monitored by using a disc as a buffer. Figure 1 indicates the possibility of a separate acquisition computer using one part of a large, dual-ported disc to record live data. This information could then be accessed by a user at one terminal and analyzed in a completely standard fashion. The coupling of the analyzing power of the data center in such a quasi-on-line manner, could provide an effective means of monitoring future complex experiments.

Summary

Figure 3 indicates the current estimates of the hardware equipment costs on a per user basis to develop the facility as shown in Figure 1. For a system to support five high-speed, interactive users, as indicated, the current estimate of the total equipment cost is approximately \$700K. The estimated manpower requirement is for 3 software and 3 hardware personnel for a total of 2½ to 3 years. This time and manpower estimate is based on the utilization of some currently existing hardware and software. The 2½ to 3 year real-time estimate is the time span from funding to opening the facility to users. This does not require that the facility be operating at maximum throughput or capability. Indeed, as was indicated previously, it is anticipated that the developments and improvements can, and will, continue to be made as requirements dictate.

The ultimate number of high-speed terminals which could be supported by such a facility is not

obvious. One limitation will be the central computer itself, and the functions demanded of it. It was felt that the facility as proposed could safely support 5 such terminals, with perhaps 10 as an upper limit. A basic design consideration was, however, the development of a "local" computer facility, as inexpensively as possible, which could basically be operated by the users themselves.

Conclusion

By dividing the general problem of interactive analysis into separate parts and by utilizing a parallel processing approach, it appears possible to develop a facility which can support the interactive analysis of data by a user at speeds up to half those possible on a CDC7600. Further, by utilizing a distributed environment (as indicated in Figure 1) it is possible to support multiple users on the facility, each capable of independent data analysis at these speeds. By adopting a relatively straightforward master-slave hierarchy within the distributed environment, it should be possible to minimize the general problem of communication and to greatly simplify the overall system software. This software architecture, together with the separation of many functions into intelligent hardware, offers the possibility that the facility could develop and adapt to future demands and technological advances in a relatively easy fashion.

Reference

1. "Data Analysis Facility at LAMPF," D.G. Perry, J.F. Amann, H.S. Butler, C.J. Hoffman, R.E. Mischke, E.B. Shera, H.A. Thiessen, Los Alamos Scientific Laboratory Report LA-7034-MS, Nov. 1977, unpublished.