NAL ON-LINE COMPUTING FACILITIES

A. E. Brenner
Harvard University

A. R. Clark
Lawrence Radiation Laboratory

E. C. Fowler
Duke University

J. V. Jovanovich
University of Manitoba

L. Leipuner
Brookhaven National Laboratory

J. MacLachlan
National Accelerator Laboratory

and

H. A. Neal
Indiana University

ABSTRACT

The problem of on-line facilities for the experimental program at NAL is studied in some detail. A flexible approach utilizing a small dedicated computer at each experiment with a link to the Laboratory's large central facility is recommended.

I. INTRODUCTION

In order to do high-energy physics experiments at NAL effectively there is a great need for both on-line and batch-processing computing facilities. As part of the 1969 NAL Summer Study Program, we have examined the on-line computer problem and its relationship to batch processing, making reasonable extrapolations of future needs, and technology based upon our own experience and expectations. Many of the other participants of the summer study group have contributed to our study. The primary task to which we have addressed ourselves is that of defining the needs and suggesting a viable system for providing a general purpose on-line computer capability for experiments performed at NAL, by both "in-house" groups and external groups.

It is clear that there is no unique solution to this problem, and although there are many viable solutions, there are clearly more bad solutions. Most of the systems in existence at other laboratories leave much to be desired; in almost all cases there is major dissatisfaction at some level. A major reason for the less than optimum solutions have been constraints emanating from the structure of the computing organizations which were largely set up at a time when computing was primarily a closed-shop batch operation with little on-line utilization.

At this time NAL is in the particularly good position of being able to organize its computing facilities with no historical constraints. Furthermore the computer companies have now more fully recognized the utility of on-line operations in a general sense and have organized their product lines more appropriately to service simultaneously that type of computing, as well as the usual batch-processing requirement, both on a local and on a remote basis. In the body of this report we also deal with and comment on the large batch-processing facility we expect NAL to acquire, in so far as it is relevant to the solution of the on-line computing needs of the laboratory. Although there are several possible alternatives with respect to that acquisition we have mostly avoided any detailed study of those problems.

Finally, in the body of this report at various places we select various computer configurations so as to make our recommendations more concrete and to be able to make quantitative cost estimates. No weight should be given to the choice of any particular computer or computers used as examples. The details should be worked out by NAL personnel after studying in detail all the relevant factors and in consultation with representatives of the users.

Exclusions

We wish to emphasize at this point that we are addressing ourselves to the problems of the on-line experiments. We do not address our attention to such questions as the control of the accelerator, control of primary beam transport systems, or automatic film scanning devices. Those parts of the beam transport system which are particular to a given experiment might indeed be controlled by the on-line computer but the problem of general beam control is best solved by those involved in their design. We assume that while the values of accelerator variable parameters are available to experimenters and their on-line computers, the generation and control of these will be done elsewhere.

II. REQUIREMENTS FOR THE FACILITY

In general the requirements of an on-line computing facility are to provide the greatest flexibility and power at a minimum cost. While no one can predict what devices will be available for doing physics in the future, one can gain considerable insight by examing present devices in terms of their computer needs and hope to provide a flexible enough system for future ones.

A. Classes of Experiments to be Serviced

Perhaps the simplest on-line devices handled at the present are the counter-hodoscope systems. These consist of tens or hundreds of counters and gated latches which are triggered on an appropriate signal. They may or may not be coded to produce position information. These devices could produce hundreds of bits of information per μsec but nothing exists that could digest information at this rate. Generally, however, only a very small fraction of this data is meaningful and the experiment's trigger logic selects these events for further processing. The use of an on-line computer can often simplify the logic.

Spark-chamber systems are at present the greatest producers of on-line data. They produce trajectory information with a precision of up to 1/4 to 1/2 mm and can be built in a great variety of forms and systems. A "typical" system may consist of 20 planes with magnetostrictive readout. Each plane typically may be required to record up to 4 sparks per event and the whole system can be triggered at 10-millisecond intervals or less. Each spark in each plane requires a 12-14 bit scaler so the whole system might contain $20 \times 4 \times 14$ bits. This system would, at its full rate, be a 10^5 bit/sec system. Devices an order of magnitude larger can be made and are planned.

New detection devices are of course always being developed. At present the most promising is the proportional chamber. This device which is like a spark chamber in its geometry and size could be operated at speeds approaching that of counters. Timing and pulse height information from counters and proportional chambers may be useful in certain experiments. Special multi-output Cerenkov counters may be built. Most likely some combinations of the above devices will be used in the future. No one can predict which new or old devices will be most useful in the future, but most will require a computer in the system for optimum performance.

B. Classes of Jobs to be Done

The computer on-line to an experiment is required to do a great deal more than compute. The name reflects more its origin than its function. Surely new uses for this instrument will be developed, but a description of at least some of the general areas of its utility is possible.

The most common task to be performed by the on-line computer is collecting data. The data from the detectors will have to be read into the computer; scalers will have to be read; the components of the beam transport system may be monitored; etc. Usually data from the detectors will flag their own presence and one will want to empty the detector registers as soon as possible in order to prepare it for the next event. In present-day computers this can be done in something like 1 μsec per word of 12 to 24 bits once the process is started. Starting the process may take as little as 10 μsec. The beam, beam transport, and accelerator monitoring usually can be done between beam pulses and is often initiated by the computer. Data can be organized and edited by the computer before being written on tape or may be written exactly as received.

The experimenters' equipment may be controlled by the computer. Targets may be inserted or removed or changed by computer control if the experiment requires it. Counting periods may be determined by the computer and sequences of experimental conditions changed by it. In general it will be only those things which are changed often that will be implemented for computer control. The monitoring of experimental equipment is an important use of on-line systems. For some devices monitoring will provide control; that is, the computer will be part of a closed servo-loop.

One great advantage of on-line systems is that the physicist can have summaries of his data as it is being taken. If his equipment is misaligned he may be able to detect this in time to do something about it. It is important therefore to have a device capable of displaying some of the data stored in the computer. This device should be located at the data-gathering station. An oscilloscope and a teletype constitute a minimum for this purpose, although other devices are available. The physicist should be able to summon a display of some complexity in a short time, of the order of seconds.

For some fraction of experiments and for some fraction of the data quite complete analysis of the data will be required. Geometric reconstruction and kinematical fitting may be desired to sample the effectiveness of the apparatus. For this purpose the computational abilities of the facility, rather than its ability to manipulate bits, remember, or control, are important. Here, and perhaps only here, the speed of "computation" is important.

C. Desired Features of the Facility

We gather here some of the desirable features that the system is to have. Surely the expected large number of simultaneously active experiments at NAL, many in preparation and some in execution, demands economy per experiment. Each experiment will want a system which can be devoted entirely to that experiment. It will be necessary to attach devices and remove them at the convenience of the experimenter. The major components should be capable of being tested even before floor area is consumed with the setup.

One great virtue of computers is their flexibility. This flexibility should not be compromised any more than economics demands. The system should be programmable, at least in its computational aspects, in Fortran. This has become the most known and used compilation language for scientific purposes. Logical manipulation, however, is probably best done in machine or assembly language.

While NAL will have to supply some sort of system to its internal users, and to some extent to external groups, the system so evolved should be capable of accomoda- ting a variety of computers brought by outside groups. A reasonable policy on how much interfacing a group with a different computer must provide should be set.

The computer industry is a rapidly growing one and NAL's on-line system should be capable of growing with it. The systems should remain as flexible as the demands upon it require and it should be able to keep up with the state of the art.

III. BATCH PROCESSING AT NAL AND ITS RELATIONSHIP TO ON-LINE EXPERIMENTS

Although we are principally concerned with the needs of experimenters for on-line computing facilities there have been recent developments in the utilization of large batch processors that bear on this problem. Thus to be realistic we have considered the interaction of batch and on-line use. In fact, because we interpret "on-line" broadly to mean turn-around times sufficiently short to govern the course of an experiment, the remote job-entry facilities supported by the major manufacturers represent a useful partial solution which blurs the distinction between batch and on-line operations. We have, therefore, considered batch processing carefully in general terms with emphasis on matters affecting outside users and have formed some recommendations which we urge be included with internal requirements in choosing a large computer system which will serve the on-line needs of visiting experimental groups as well as the entire needs of the staff.

A. Some Organizational Recommendations

The successful coexistence of batch and on-line operations in the same processor seems to require special understanding from the computing staff. In order to attune the computing facility to the needs of the high-energy physics program, we strongly recommend that NAL establish its computer group as a service organization whose leader and upper echelon personnel are physicists or have substantial training in physics. It is of fundamental importance that the computer management and experimenters talk the same language. Also to the largest possible extent NAL should rely on the standard executive programs supplied by the manufacturer to maximize the amount of supported software, to reduce the maintenance effort and disruption caused by new system releases, and to reduce problems of compatibility with users' home installations. (Major research in computer science is not a desirable activity for the support group.) We believe the recommendations given here are fundamental to the viability of responsive high-energy computing at NAL both in the on-line and the batch processing modes.

B. The State of the Art

It seems clear that very recent developments have made remote entry of jobs into large computers a practical and attractive major mode of operation. The principal manufacturers of large systems have committed themselves to primary support of remote-entry software and they offer terminals with fast card readers and printers. Several of these terminals could be located around the laboratory including some in experimental areas where the experimenter could receive bulk output initiated by requests from his small computer or by a program deck read in at the terminal.

Even if an experimenter has a small computer for data collection to allow for possible independence of the central facility when necessary, it is highly desirable that numerical work requiring floating-point arithmetic, like reconstruction and kinematics, be done in the large computer with its fast floating-point hardware. Small computers are excellent for data collection and organization, and in fact have basic cycle times comparable to the largest available computers, but they are slower by factors of 1000 or more in floating-point operation. Wide band-transfer, say to 50 K bits/sec or more, via peripheral processor, multiplexor, data channel, or some combination of these allows the experimenter very fast access to the central processor for suitable tasks. Many of the current operating systems on large third-generation computers have options permitting so-called "real-time" programs. This facility could be used to service these operations.

C. NAL Large Computer Usage Projections

To make explicit the scale of the central processor which is involved we have made some tentative estimates of NAL's internal computer needs for the next five fiscal years. The numbers for FY 69 and FY 70 have been derived from discussions with laboratory personnel and are reasonably consistent with independent estimates made by G. Duffy of the Accelerator Theory Section.[1] The projection for later years is more speculative and the bases for those estimates are given in Appendix A. The projections in CDC-6600 hours are given in Table I. Clearly the estimates contain large uncertainties and depend critically on the timing and existence of certain large pieces of apparatus, such as large spectrometers or automatic film-scanning and measuring devices. Furthermore if the facilities are available to any class of users at no cost to themselves, inefficient use can lead to inflated utilization and premature saturation. This is a well known problem at large laboratories, which should be recognized, and for which some mechanism should be introduced to maintain a sensitive level of observation and control of the facility utilization. However, it should also be recognized that as the laboratory grows, its sensible appetite for computing will grow at a rapid rate. Thus, whatever large computing system is chosen by NAL, as a fundamental part of the acquisition the capabilities for long-term expansion with a minimum of changeover chaos should be kept in mind.

It is clear that simply for NAL's internal use a rather large machine will be required. We believe that NAL should encourage outside users to do all of their post-experiment processing at their home installations not only to keep the size of the NAL central facility within reason, but also for reasons to be discussed in Sec. VI.

The major content of the above considerations on batch processing can be summarized as a recommendation for a single system dedicated to physics use with several remote batch-entry stations and with on-line users treated as highest priority batch users.

IV. THE PREFERRED SOLUTION TO THE ON-LINE PROBLEM

A. The Scheme

In our opinion the system offering the greatest advantages in flexibility, reliability, and economics is the following:

1. Each experiment with such a requirement has a small computer on-line to the experiment. This computer is considered part of the apparatus and is under the complete control of the experimenter.

2. Each computer is connected to the large NAL batch processing computer through a high-speed data link.

3. The large computer has connected to it standard remote-batch terminals, each with a line printer and a card reader and punch located in the experimental hall(s) (one per hall or area, not one per experiment).

4. All components of the small computer and of the interface to the experimental equipment should be under the jurisdiction of a general NAL pool similar to the electronics (PREP) pool, to enhance flexibility, interchangeability, ease of maintenance, and to insure compatibility between all the components to be acquired by NAL. Efforts should be made to acquire several identical, flexible, "standard" systems, and to avoid " one of a kind" systems.

5. Outside users should have access to pool equipment, including the small computers. Enough equipment should be available so that complete systems can be set up for testing and checkout purposes several months before the experiment is scheduled for floor space.

6. An independent teletype net should extend throughout the accelerator area for disseminating hard-copy information about accelerator operating and beam-transport conditions. (This net operates independently and is not directly connected to the equipment described above.) This teletype net is intended to provide a permanent record of various accelerator running conditions, rather than relying on intercoms as in the past. (Intercoms would still be necessary, however.) Message inputs to the net would come from teletypes in the several accelerator control rooms, the secondary beam-transport control computer and from the accelerator control computer(s) which may exist. While this system has not been priced, it should be relatively inexpensive since the connection to a remote unit only requires a single twisted-pair line, which could be run easily with the other communication lines required at each experimenter's area.

B. Configuration

A block diagram showing the various components of the preferred system is shown in Fig. 1; cost estimates based on published list prices[2] without any discounts are given in Table II. The PDP-15 and Sigma 2 were taken as representative of the small computer class for purposes of configuring and pricing the system. The functions of the various components are discussed below, along with our concept of the way in which the overall system should function.

The heart of the system is the small computer represented here by the PDP-15 or Sigma 2. This computer should have 1) a cycle time of approximately 1 μsec; 2) hardware multiply and divide (fixed point); 3) a priority interrupt system; 4) a word size of at least 16 bits; and 5) memory size of at least 16 K words (most of the presently available computers can be expanded by essentially plugging in additional core memory so this parameter could be modified relatively easily to fit a specific requirement).

The peripheral devices needed on the small computer are:

1. A magnetic tape unit capable of handling 800 bits/inch at 75 inches/sec. The choice between 7 or 9 channel tape drives should be made on the basis of compatibility with the drives available on the large computer. The tape unit is necessary for the stand-alone capability of this system; all data gathered from the experimental apparatus are written onto magnetic tape by this unit, whether or not they are sent via the link to the large computer. Experiments with high data rates may require a second tape drive to minimize time losses for changing tapes; accordingly, the tape controller should be capable of handling (not simultaneously) several drives. To handle cases with extremely high data rates, additional tape drives and controllers can be added so that data can be written onto two or more tapes simultaneously. (Data rates are discussed below.)

2. A bulk storage device such as a small disk or DEC tape is generally required for operation of the manufacturer-supplied software system, and would be used to store the system and program library. This device can also be used as scratch storage for histograms or for accumulation of data for transmission to the large computer.

3. A storage scope such as the Tektronix 611, to be used primarily to display various histograms accumulated from raw event data and to display results from the on-line program in the large computer. A small fast scope such as the Tektronix 601 equipped with a Polaroid camera can be slaved to the display scope so that hard copy of the displayed data may be obtained.

4. Teletype or typewriter for general communication between the experimenters and the system.

5. Paper-tape reader and punch, for system compatibility and general utility.

Additional equipment could be added to make the system more convenient to use (particularly for off-line data analysis), but which are relatively expensive, include a card reader and line printer.

The small computer is connected to the experimental apparatus through the Experiment Interface. This interface is composed of several modules which logically connect the small computer to the various registers holding the data for one event in the experimenter's equipment or to devices for controlling the experimental apparatus or for sensing various conditions in the apparatus, and to other registers containing data pertinent to the experiment such as magnet settings, beam conditions, etc. The detailed structure of the experimental apparatus will of course vary greatly from one experiment to another and consequently the interface may also vary to suit from one unit to another.

The small computer is connected to NAL's large batch-processing computer by means of a telecommunications link or a special-purpose coaxial-cable data link. Typical transfer rates vary from ~ 1 K words/sec up to ~ 1000 K words/sec depending upon the needs. The use of the data link is discussed in greater detail below.

C. Concept of Operation

The various tasks of the small computer and their relative priorities have been discussed above. The main program monitors the experimental conditions, controls the overall flow of experimental data from the apparatus to magnetic tape and to the large computer, keeps a log of the data as they come in, and sets up various configurations of the apparatus as requested by the experimenter. This program is also responsible for generating the various histograms and displays to monitor the progress and performance of the experiment.

As events (triggers) occur, the data are stored in (typically fast) register in the experimental equipment. Each trigger also generates a program interrupt, which transfers computer control to a program which:

1. reads the events registers into the computer

2. performs any necessary conversions (e.g. BCD scaler readings to binary)

3. blocks the data into the output buffer--(Magnetic tape usage is generally more efficient if long records are written, to minimize waste space from inter-record gaps; for an experiment with events of 200 words/event, one would block this data into units typically 1000-2000 words/record. The maximum tape record length may be limited by the operating system in the large computer.)

4. initiates writing of the buffered data onto magnetic tape when the buffer is full. All data are stored on magnetic tape at the experimental area. This allows an experiment to continue taking data when the large computer is down or unavailable for any reason.

5. checks the data for correct performance of the apparatus and adds the event to any histograms which are being accumulated

6. routes this event, if appropriate, via the link to the large computer

7. returns control to the main program.

The various tasks outlined above involve primarily data-moving, bit manipulations, or simple calculations which can be done quite easily in fixed-point arithmetic--all of which can be efficiently and rapidly handled by a small computer.

As data are sent over the link, a small resident program in the large computer accumulates the data on a large bulk-storage device, probably a disk. This program should be part of the central computer's system, as opposed to being a batch job. It is necessary that the large computer system be oriented around the fact that the on-line jobs have top priority for their alloted time; the program which accepts data from the link must be able to take it promptly so as not to delay or hang up the small computer and the experiment, which would in turn waste accelerator time.

Periodically, on a demand or on a time-interval basis, an analysis program is brought into the large machine to analyze the accumulated data; these data can then be discarded since they are already on tape at the small computer. Results of these more complex calculations are then returned to the experimenter. Suitable results can be returned to the small computer for display on the CRT, or short answers can be typed out; more extensive results would be returned to the line printer at the remote batch terminal located reasonably close by in the experimental hall. These batch terminals cost approximately \$60 K each for the smallest versions, including ~ 300 ℓpm printer. In some cases it may be appropriate and economically justified to install a line printer on the small computer and produce the hard copy at that point;

this decision requires detailed analysis in each individual case. The question of the time and memory required of the small computer to handle the printer should also be examined. A very attractive alternative scheme is to replace the standard teletype on each small computer with a device similar to the recently announced "ink-jet" teletype, which prints at a rate of about one full line per second.

D. Rates

The standard channels on the currently available small computers can run at maximum rates between 0.25 and 1.0 M words/sec. Using the slower rate, an event in a medium size experiment with 200 words/event can be read into the computer in 800 μsec; if the data are gathered from magnetostrictive spark chambers, the system must delay several hundred micro-seconds while the data are digitized before read-in can be started. These numbers give an electronic dead time of 1-2 milliseconds, which is still faster than most spark-chamber systems can be run. Using a 2 msec dead-time and a 1 sec spill time, the computer could conceivably acquire data as rapidly as 500 events/accelerator burst of 10^5 words/accelerator burst.

The limitation on data rates for this system is determined by the tape-writing speed, which determines how fast the computer can dispose of the acquired data. Sixteen bit words can be written, for example, on 9 track tape (800 BPI, 75 IPS) at a rate of 30 K words/sec; the 9 track controllers for 18-bit machines can also read or write full words at 3 bits/word, or 20 K words/sec. Words of 16 or 18 bits can be written on 7-track tape at 20 K words/sec.

The amount of data which can be disposed of during one accelerator beam burst is determined by the amount which can be written on tape during the spill plus the amount stored in memory at the end of the spill, which can be written before the next spill. Assuming a one-second spill, maximum tape rates and a reasonable buffer size of 8 K of memory, the maximum data output rate to one magnetic tape is ~38 K words/accelerator burst, or 190 events/acc burst using the above example of 200 words/event. Of the 190 events/burst, 100 are recorded during the 1-sec spill, limited by tape writing speed, and 90 additional events (assuming at least 0.9 sec interval between beam spills) limited by the capacity of the 8 K buffer. This rate is more than adequate for most experiments.

The output rate can be increased if necessary by adding more tape units and writing simultaneously on several tapes, by using a tape drive capable of 150 IPS (a factor of 2), by using a hypertape cartridge unit (a factor of 5), or by increasing the buffer size (with the addition of more memory). These options are all rather expensive. The use of a disk as a buffer device is not sensible because of the relatively slow average access time and the program complexity required to keep track of where the data is. The best alternative for handling extremely high rates would be to use the small computer to pre-process the data to eliminate events from spurious or bad triggers before writing the good events on tape.

E. Software Support

A programming support group should be set up which would be available to all small computer users on a consultation basis. The group should be under the ultimate direction of the major computing facility group leader, with immediate direction in the hands of a physicist, or at least someone sympathetic to the requirements of the experimenter. This group should provide useful general purpose routines, with good documentation, for the use of the experimenters. New subroutines should be written at the request of an experimenter only when they have general utility. An experimenter setting out to use the facility will have available a set of documented routines which he would put together with whatever additional control programming he needs to generate his working program. The purpose of the support group is to assist him in this task, but it should not be their function to do the job for the experimenter. It is assumed that direct software support is available through other channels to the "in-house" users.

V. OTHER POSSIBILITIES CONSIDERED FOR ON-LINE COMPUTER CAPABILITY AT NAL

There are other possible solutions to the problem of on-line computing at NAL. Clearly, one possibility is that NAL supply no on-line computing facilities for external groups. (A different philosophy would, of course, have to be applied to the in-house users.) Each outside group would then be expected to bring whatever computing power it needed. Examples of configurations that might satisfy the group's computing requirements are: (a) its own small computer of the PDP-15 class, or (b) if more computing capability is required, the small computer could be connected over common carrier lines to the group's large university computer, or (c) time could be rented on any of the various time-sharing services.

It is clear that there are many advantages in individual groups having their own small computers. This approach is certainly to be encouraged. The possibilities stated in (b) and (c) above are expected, however, to be impractical in most cases. For this reason, a link from the user's small computer to NAL's large computer appears to be the most efficient means of obtaining high computational power. This could be provided with only a modest increment in the cost of NAL's central computing complex. Also, due to the fact that all groups will not have their own small computer, the purchase of a few standard computers by NAL appears well justified.

Another possibility is for NAL to supply, from a laboratory pool, a small dedicated local computer (including a generalized interface to experimental equipment) to those groups requiring such a facility. The computer would have magnetic-tape writing capabilities, but would not be able efficiently to conduct extensive computations in real time. That is, it would have integer arithmetic but no floating-point hardware. Under this scheme, however, there would be no provisions for links to a large computer. This is not considered a reasonable alternative since, as stated above, a small increment in the cost of NAL's large computing facility could make available an extremely powerful arithmetic capability via a link to the small local computers. It is expected that this capability will be found useful in all experiments requiring any on-line analysis.

We recognize that in some experiments the combination of a small dedicated local computer ($\lesssim$ \$100 K) connected to NAL's large computer might be less efficient than a single dedicated medium-sized computer ($\lesssim$ \$200 K - \$800 K). Such situations are expected to arise so infrequently, however, that they should be handled on an experiment-to-experiment basis. A group proposing an experiment for which it felt a medium-size computer (e.g. PDP-10 or SIGMA 5) was required should formally request such a facility in the proposal. We are recommending that NAL not purchase computers of this class until specific proposals justify their procurement.

We further recommend against the multiprogramming of one of the medium-sized computers to service several experimental groups simultaneously, in the manner that BNL is now operating. Recent advances in the technology have made this approach considerably less efficient, because the large machines are already configured to be multiprogrammed. To configure a medium-sized machine for multiprogramming introduces additional hardware expenses and additional software complications.

VI. SOME NAL COMPUTER POLICY QUESTIONS

It was practically unanimous among the users at the 1969 NAL summer study that NAL needs a large computing facility for both bulk data analysis and on-line computing.

It is scarcely feasible and almost certainly not desirable for NAL to provide the major part of the computer needs of the external users. It is recognized that the computing problem of high-energy physics groups at universities is a major one, which has a large variety of complications. However, it is practically certain that an attempt to alleviate these problems by NAL "picking up" the responsibility is out of the question. There must be limitations on the use of the NAL computer facilities by university and other outside users of the NAL accelerator.

A. Limitations on the Use of the On-Line Facility

Even the direct use of the on-line facility must be limited according to the pressing needs of the user as related to his specific experiment. Consequently, it is recommended that a quantitative estimate of the on-line computer use for a given experiment be made an integral part of every proposal at NAL. This should be an important consideration in the consultation and decision of the scientific programming advisory committee for NAL. When a proposal is approved, it is assumed that the on-line computing needs are as stated and that NAL will do everything possible to meet those needs. If major changes in computing needs become necessary, it should be assumed that these changes are just as critical as changes in accelerator use, and the new requirements should be reviewed by the programming advisory committee. Within these limitations it is assumed that, wherever possible, NAL will attempt to provide the requested on-line computing needs of both the user groups and the NAL groups.

B. Bulk Data Analysis

On the other hand, NAL shall not undertake to satisfy the overall needs of the user groups for bulk data analysis. Clearly, a reasonable fraction (typically 10%) of the experimenter's data should be carried through in an on-line mode, to a level of analysis such that the experimenter may be certain of ultimate success in meeting the requirements of the experiment. With its goal of meeting the needs of many varied experiments carried out by ~25% "in-house" and ~75% outside users, it would be nearly absurd for NAL to undertake the problem of supplying most of the bulk data analysis for all. Since individuals, or quite small university groups, often collaborate with "in-house" groups, insistence on a rigorous interpretation of separating the computing load according to its source would be equally absurd.

There will be problems of dividing the bulk data analysis load when a large university group undertakes collaboration with an in-house group. It seems that this is a special case, in which again the details of the division of the computing load between NAL and the university should be decided by the programming advisory committee in its consideration of the original proposal for the experiment.

C. The Administration and Operation of NAL's Computer Facility

This section deals with a most important question of NAL policy with regard to administration of computing. The authors believe that the computing needs of the various experimental groups and the means of solution are sufficiently well established that the emphasis at NAL should be on supplying computing services and not on computer-oriented research. In other words, we feel that is would be a serious mistake for NAL to plan to supply computer service within the framework of a

Computer Sciences, or Applied Mathematics Department. Instead, it is urged that the computer facility be strictly for physics and, to the greatest extent possible, that it be planned and manned by high-energy physicists.

D. Administrative Data Processing

Although the authors do not include experts on administrative data processing, following the same line of priorities indicated in C above, we feel it is extremely unlikely that the computing needs for handling payrolls and other similar administrative data-processing jobs are compatible with the needs for processing bulk physics data and the high priority on-line service envisioned here. The administrative data processing work has special needs for:

1. File security
2. Large permanent file storage with frequent random accesses
3. Hard-and-fast calendar deadlines.

These features have proved to be exceedingly difficult to integrate with the needs of computer use for scientific purposes, even where the on-line use from active experiments was not present. The file-security problem, in particular, is foreign to the needs of the scientific community and its introduction is likely to cause major complications. Thus it is felt that adding the administrative data processing (which does not in general need a large central processor with major arithmetic capability) would add a tremendous responsibility without producing appreciable advantages or savings. Indeed, it is probably more economical to have a smaller, separate, appropriately configured computer for these purposes.

VII. ADDITIONAL COMMENTS AND SUMMARY

We briefly summarize our major conclusions here. With respect to on-line facilities we recommend that NAL take a flexible position receptive to the needs of any given experiment, proposed either by an "in-house" group or an external group. However, we do propose that some small computer system be adapted as the "standard" NAL system available from a laboratory pool. This system, which should be supported by a small group of systems programmers, should also have an easily set up connection to the large NAL computing facility for major computing power. This capability for on-line usage of the large computer comes with the highest priority, but also is limited to a reasonable fraction of the available computer time averaged over some reasonable period.

In addition, in the vicinity of the experimental halls, remote batch terminals connected to the large computer should be set up for conventional batch processing from the experimental hall. This facility will also provide a major hard-copy capability for those experiments which have such a need.

With respect to the major computer, NAL should recognize the continually expanding role large computers play in high-energy physics and should plan accordingly. Negotiations should be started at the earliest possible time, with a programmed compatible expansion of the facility projected several years into the future. NAL is in a very good position to induce temporary loans of computers, short-term rentals possibly with option-to-buy features, purchase with buy-back clauses, and also major software support for special features which may be required by NAL.

Finally, it should be noted that there is no fundamental reason why any one of the larger computers is right for high-energy physics. The choice should be made taking into account only the technical matters of system viability, long-term potential growth capability with minimal expansion pains, and on the best-buy principle.

APPENDIX A. ESTIMATED BATCH COMPUTER USAGE AT NAL

Basis for Estimates Given in Table I

1. From a survey of bubble-chamber users[3] excluding Berkeley and Brookhaven, in 1968 the average BC group used 80 hours of 6600 equivalent computer time. (This number has a typical spread of approximately ±40 hours.) On the assumption that start-up and development at NAL and higher energies will require more computer time, we have chosen 100 hours of 6600 time as an average to cover bubble-chamber experiments.

2. A similar survey of counter experiments[4] shows a very much wider variation. Taking numbers from the survey 30 ± 10 hours per year per experiment covers many experiments. However some experiments take as much as ten times that number. At Harvard, where the detailed numbers are available to us, that number appears to be about twice as high, i.e., 65 hours/year. It is also important to note that the rate of increase of computer usage by counter groups is much higher now than for bubble-chamber groups because of the change over to wire spark-chamber systems. It is reasonable to believe that in approximately two years the average counter-spark chamber group will be spending about 100 hours of 6600 computer time per year. For this projection we choose 30 hours per year for FY 70, 60 hours for FY 71, 100 hours for FY 72, 130 hours for FY 73, and 160/hours year for 1974.

3. The introduction of an automatic film scanning and measuring machine such as HPD, PEPR, or POLLY will increase the load enormously. We have assumed NAL has such a system at the beginning of FY 1973 and that the large 25-ft bubble chamber will start up in FY 1974. We have chosen for FY 74 the current operating level for bubble-chamber physics at BNL.

It should be noted that these are very difficult projections to make, and can be greatly modified by schedule alterations such as changes in:

1. machine turn-on date
2. automatic scanning and measuring turn-on date
3. the 25-ft bubble-chamber construction schedule.

If the schedule is about right then the estimates of computing need are probably good to about ± 30%. Also it should be noted that Parkinson's Law with respect to computing at "no-cost facilities" quickly saturates any facility.

REFERENCES

[1] G. Duffy, private communication.

[2] The PDP-15 prices are from the (Preliminary) Price List dated 3/13/69, and the SIGMA-2 prices are from the GSA Federal Supply Schedule Price List Contract No. GS-VOS-76148.

[3] Informal Poll of Bubble-Chamber Users by E. Fowler, R. Plano, and A. Rosenfeld.

[4] Informal Poll of Counter Groups by H. Taft.

Table I. Projections of Large Computer[a] Use at NAL.

(See Appendix A for basis)

FY	No. of Experi-Ments	Hrs/Year	No. of Counter Exp.	Hrs/Year	Machine design Hrs/year	On-line Expts. Hrs/year	Totals
1969	-	-	-	-	435	-	435
1970	3	300	5	150	550	-	1000
1971	3	450	5	300	600	-	1350
1972	4	800	6	600	600	200	2200
1973	5+auto	2000	7	900	600	400	3900
1974	+25-ft	4000	8 (2 large +6 small)	2000 (1000 + 1000)	600	1000	7600

[a] All units are in CDC-6600 equivalent times.

Table II. Cost Estimates for On-Line Computer System.

Costs based on current GSA list prices; no discounts have been applied. Two representative systems are shown.

A. PDP-15

1. Basic configuration

PDP-15/20 (CPU, 8K memory, high speed reader/punch, extended arithmetic element, DEC tape control with two DEC tape transports KSR-35 teletype	$36.0 K	
8K additional memory (MM15A, MK15A)	14.0 K	
Automatic priority interrupt (KA15)	3.0 K	
Mag. tape control and bus converter (DW15A + TC59)	10.0 K	
Mag. Tape, 7 track, 800 BPI, 75 IPS (TU30) (For 9 track, add $1.0 K	21.0 K	
Storage tube display (modified Tektronix 611) (VP15A)	5.8 K	
Interface and link to large computer	10.0 K	est.
Basic system cost, PDP-15	$99.8 K	

2. Possible additional equipment:

Additional tape drives (TU30)	$21.0 K	
Additional tape controller (DW15A)	8.0 K	
Line printer and control, 300 ℓpm (6470)	25.0 K	
Card reader and control, 200 cpm (CRO3B)	5.2 K	
Disk and controller 262 K word (RF15, RS09)	15.0 K	
Storage scope slave (Tektronix 601 or equivalent, camera, custom interface)	1.5 K	est.
Graph Plotter (e.g., Calcomp)	10.0 K	est.

B. Sigma 2[a]

1. Basic configuration

Sigma - 2 Control Processor (8001)	$12.0 K
16 K Memory (8051 + 3X8053)	42.0 K
Multiply/divide (8020)	6.0 K
Keyboard printer (8091)	5.0 K

Paper tape reader/punch, control (7060)	12.0 K	
Disk and Controller (minimum) 0.75 Mbytes (7201, 7202)	26.0 K	
Priority interrupt control (8021+8022+8070)	4.0 K	
Magnetic tape control (7320)	20.0 K	
Mag. tape, 9 track, 800 BPI, 75 IPS (7322) (for 7 track, add $1.0 K)	27.0 K	
Display scope Tektronix 611 or equivalent with custom-built interface	6.0 K	est.
Interface and link to large computer	10.0 K	est.
Basic system cost, Sigma 2	$170.0 K	
2. Possible additional equipment		
Additional tape drives (7322)	27.0 K	
Additional tape controllers (7320)	20.0 K	
Line Printer, 600 ℓpm (7440)	35.0 K	
Card reader, 400 cpm (7122)	16.0 K	
Storage Scope slave	1.5 K	est.
Graph Plotter (e.g., Calcomp)	10.0 K	est.

[a]Note added in proof: Compilation of this table preceded the announcement of the Sigma 3, a successor to the Sigma 2, whose price schedule is more competitive than the Sigma 2 with that of the PDP-15.

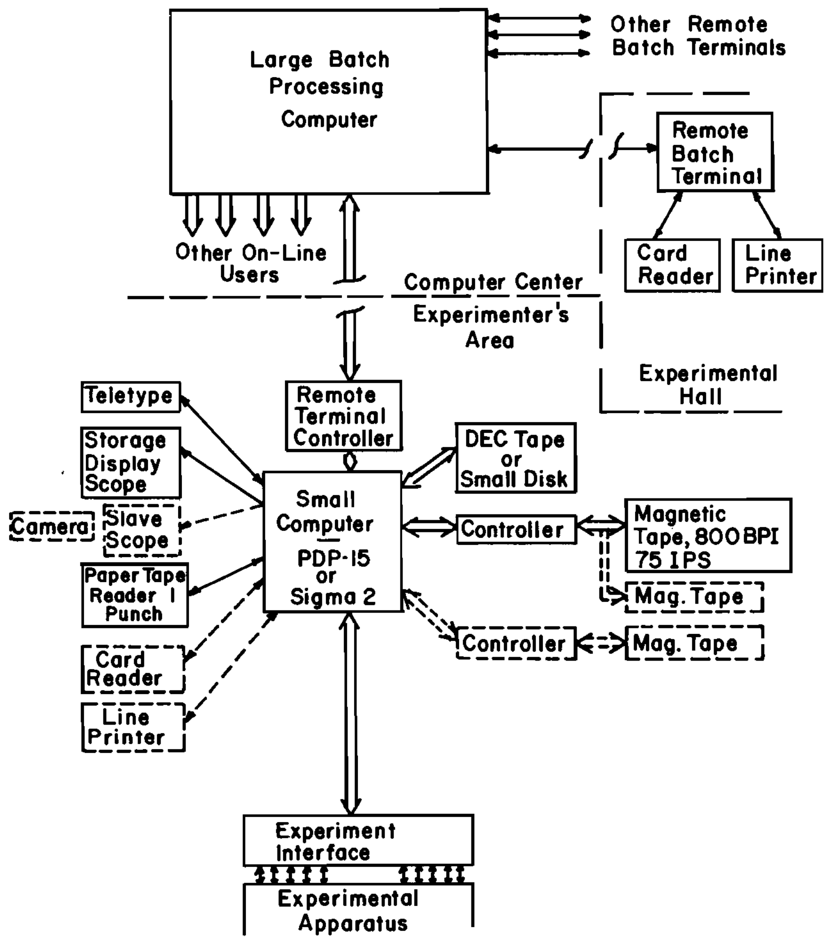

Fig. 1. Block diagram of recommended on-line computing facility. The peripheral units drawn in dashed lines are possible additions to the basic system.