

INSTRUMENTATION TECHNIQUES IN NUCLEAR PULSE ANALYSIS
National Academy of Sciences—National Research Council
Publication 1184

IV-6 A TIME-SHARED COMPUTER FOR REAL-TIME
INFORMATION PROCESSING

J. Leng,¹ J. A. Quarrington,¹ P. K. Patwardhan²
and G. Bell³

Presented by J. Leng

1. Introduction

The use of small high-speed computers for multi-parameter analysis is now an accomplished fact in several North American physics laboratories. A report by W. F. Miller and H. W. Fulbright (1) lists many more planning such projects.

Many of the systems, including the large fixed-program machines, tend to be most suitable for use by one person at a time, although they are capable of concurrent operation on more than one problem.

In order that two or more experimenters be able to make use of such a facility simultaneously, steps have to be taken to ensure that each does not interfere with the operations of the other. In fact, to keep each user completely happy, complete unawareness of the others operations should be achieved. The latter is difficult to realize in practice but nevertheless one can go a long way toward meeting these requirements with present day computing machinery.

One way of coping with this problem is to use a high-speed computer on a time-sharing basis. This technique is normally applied when a very large and expensive machine, in order to make use of all available time and thus to reduce unit costs, shares its time on the concurrent operation of several programs. However, when time-sharing between experiments, the computer time is secondary and the need is for the computer to be able to process as quickly as possible, on demand, the requirement of an external device. It is this latter method of operation that is most suited to the problems of nuclear pulse analysis. Here it is not a question of the computer being required to perform several long and complicated computations simultaneously, but rather to interrupt one computation in order to perform a set of simple operations on request, from one or more external experiments. When increasing use of the computer is to be made for the accumulation of events such as for pulse-height analysis, rather than interrupt a computation to add one count into storage for every event, it is preferable to by-pass the processor and provide direct access to the computer memory (2). In this way it is possible to use part of the memory concurrently with, and independently of, the processor for the accumulation of events. At the same time the processor has the ability to perform further analysis on them whenever required. The original concept of EDAC, the Electronic Data Analysing Center, for nuclear physics research at Chalk River,

¹Chalk River Nuclear Laboratories, Chalk River, Ontario, Canada.

²Atomic Energy Establishment, Trombay, Bombay, India.
At present a Colombo Plan Fellow at Chalk River.

³Digital Equipment Corporation, Maynard, Massachusetts.

was planned so that this approach could ultimately be used. The first step, however, was to provide the basic computing facility. Now that this has been put into operation and demand for its use is growing, the ultimate system is fast being implemented.

2. Existing System and Method of Use

EDAC is located in the control room of the Tandem Van de Graaff accelerator and has been used up to now exclusively on work associated with that machine. At present all of the major analysing equipment associated with the accelerator is electrically connected or is in the process of being connected to the computer. This includes two 100-channel kicksorters, a 900-channel coincidence kicksorter with its associated magnetic tape recording system, a variety of scalars and switches, and eight pulse-height encoders. A simplified diagram of the system is shown in Fig. 1.

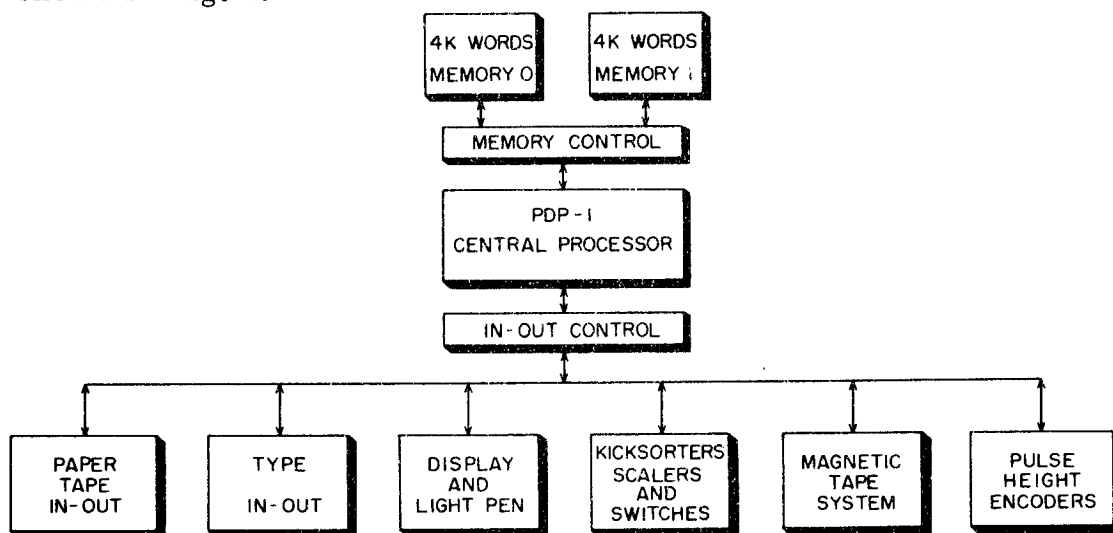


Fig. 1. Present electronic data analyzing system.

The computer has two blocks of 4K, 18-bit words of memory, (1K= 1024 words), attached to its central processor. The various input-output means together with the external analysing equipment are connected to its IN-OUT register. The display consists of a 16-inch cathode-ray tube capable of operating at 20,000 points/s and a "light pen" is used to allow points of interest on the CRT to be indicated to the computer.

At present EDAC can be used in three basic ways, namely:

- a) For examining and processing information gathered in the associated kicksorters, etc.
- b) For sorting information directly from pulse-height encoders.
- c) For performing more complex computations on the data gathered by methods (a) and (b).

The most common mode of operation to date has been that outlined in (a) above. This in fact was the most logical approach to the use of

the system, as the physicists were experienced in the operation and capabilities of their existing analysing equipment and first wanted to extend its use with the aid of the computer. For this method of operation, memory "0" of the computer and approximately the first 200 words of memory "1" are set aside for a variety of program subroutines designed to handle the data transfers and perform certain operations on it. The remainder of memory "1" is available for storing spectra and appropriate words to identify them. Sixteen such words are stored with each spectrum; several of these indicate such things as which kicksorter the data was derived from, which of a series of spectra it is, who the experimenter was, and what day the information was gathered. Other words indicate the number of channels stored, various energy constants, the duration of the kick-sorter run, the location of certain peaks of interest, and between which limits areas are subsequently to be calculated. Some of these words are stored at the time of spectrum entry, others are entered subsequently, and in the case of area limits are indicated using the "light pen."

The executive programs for handling and further analysing kicksorter spectra have been written so that the majority of communications between the operator and the computer take place using the input-output typewriter. The programs so far produced are able to perform all the normal requirements for data transfers, including such things as reading from kicksorter memories to the computer memory, reading out whole blocks of computer memory or individual spectra onto punched paper tape or via the typewriter, and displaying spectra on the 16-inch CRT. Other programs involve the use of the "light pen" in indicating channels of interest on the display to the computer, and also provide for such operations as subtracting spectra, calculating areas and the like. Most of these operations are initiated by typing into the computer identifying codes specifying the type of operation, followed by labels indicating which kicksorter the data must be derived from or which spectrum already in memory is required. The following list of instructions are included in the present spectrum handling program. Each of the codes is normally preceded by the computer typing out nc=, meaning that it is requesting the next command.

<u>Code</u>	<u>Mnemonic</u>	<u>Decimal Input</u>
dy	day	day number
rn	run number	run number
zl	zero labels	
ss	standard set-up	Kicksorter No. (mode), sequence of labels
su	set-up	Label, No. of ch., 1st ch. No., Kicksorter No.
kr	kicksorter read	set of labels
dp	display positive	set of labels
dn	display negative	set of labels
tp	type out	set of labels
pu	punch	set of labels
rp	read paper tape	No. of spectra
ad	add/subtract	3 labels, factor (a/b)
ar	areas	set of labels
ll	list labels	

(instructions continued)

<u>Code</u>	<u>Mnemonic</u>	<u>Decimal Input</u>
ec	energy calibration	2 sets of relative channel No.
sp	start program	and energy (KeV)
nd	end program	
go	enter program	
gp	general punch	
gr	general read	
ej	eject page	

After typing one of the codes the computer immediately types out the mnemonic in full and one usually then types in certain decimal numbers to indicate kicksorters, to label spectra, or to specify channel numbers. A carriage return operation will then cause the computer to request the next command.

With this programming system it is possible to set up a whole series of commands which can be initiated one by one, by the operation of a console switch. No definite word locations have to be specified for holding spectra since the program automatically finds vacant space and stores accordingly.

An example of the increase in capability this simple set of operations gives the nuclear physicist is emphasized by experiments involving Doppler-shift measurements of gamma rays (3).

These experiments usually involve taking gamma-ray spectra both at a forward angle and at a backward angle to the incident particle beam. In order to derive the Doppler shift one requirement is to subtract one spectra from the other and measure the area under the difference peaks. Previous to using the computer system it was necessary to type out the results of the two spectra gathered in the kicksorter, and then embark on a long and tedious process with a desk calculator. This had to be repeated many times with various experimental parameters modified, until optimum results were achieved. With the present system it is possible to transfer the kicksorter information to the computer and have it display either the spectrum taken at one of the angles, or the difference spectrum, as shown in Figs. 2 and 3 respectively, in no longer time than it takes to type in the appropriate instructions via the typewriter. It will be noted that the spectrum identification code and also a scale factor are displayed at the same time. Following this, one can have the areas under the difference peaks calculated just as quickly and the appropriate answers typed out. The operator is thus able to assess the results rapidly and modify appropriate experimental conditions accordingly.

Now that some experience has been gained with the computer in conjunction with the existing analyser equipment and techniques, method (b) outlined above is being used. In this case the only equipment used external to the computer are pulse-height encoders and appropriate coincidence circuits. The pulse trains from the encoders are counted in one-word buffer stores and subsequently transferred from here under




Fig. 2. Cu^{63} spectrum taken at 45° .




Fig. 3. Cu^{63} difference spectrum 45° - 135° .

program control to memory, where one digit is added to the appropriate word. The accumulated spectra can be gathered from up to eight inputs simultaneously which can be used for single or multi-parameter analyses. Since the count rate is low for most experiments (100 c/s to 1 kc/s) compared to the computer clock rate of 200 kc/s, the major part of the machine's time is unused in this mode of operation. In order to utilize this time therefore, a sequence break system is being used. This allows normal use of the computer, as already described for spectrum handling problems, and simultaneously with this, events can be fed to memory "1" from the encoders. Priority is given to the encoder inputs in this case so that whenever an event arrives the computer interrupts its current operation and processes the appropriate input. The time taken to interrupt and process the event can vary from 25 to 100 microseconds or more depending on the type of analysis being performed. At the completion of this time the machine returns to the point in the program where it previously left off and resumes operation. Thus an experimenter can perform all the usual functions such as punch, read, type, display and arithmetic, etc. on previously accumulated spectra, while further information is being gathered.

The computer has not as yet been used for the type of operation outlined in method (c). However, this is a programming problem and a library of routines is in the process of being prepared which will allow more complex operations to be performed on various spectra. One such program to be used frequently will allow least squares fits to be carried out (4). This will identify the various components in unresolved spectra and measure minute shifts in peaks in a matter of minutes. A much more complete analysis of the data will thus be possible as an experiment proceeds.

3. Time-Sharing Separate Experiments

Now that the Tandem accelerator laboratory is fully integrated with

EDAC, it is intended to make the system available to other nearby laboratories. Initially, the experimenter will be required to spend a certain amount of time at the console in order to perform subsequent spectrum analyses; eventually they will be able to operate completely from their own laboratories and effectively have independent systems.

Two remote experiments are in the process of being integrated with EDAC, one of these making time-of-flight measurements with a rotating crystal spectrometer and the other taking pulse-height spectra from an electron-gamma angular correlation table. The first of these will be recording individual time-of-flight events on magnetic tape for durations of a day or more and subsequently feeding them back into the computer via the tape deck associated with the 900-channel coincidence kicksorter. This input will operate in the sequence-break mode, as described for the separate pulse-height encoders, thus allowing available computer time to be used for examining the accumulating spectra.

The other experiment also runs for a day or more on occasions. In this case the information is gathered in a 100-channel kicksorter. The requirement here is to read out the kicksorter at regular intervals of 10 minutes and to store the information for subsequent analysis. The data will be stored on magnetic tape, which will be one of several systems shortly to be added as an integral part of the computer. When the 100-channel analyser is ready to dump its information the computer will first be interrupted on sequence break, thus allowing the spectrum to be transferred to memory "1". From here it will be transferred onto one of the magnetic tapes, in the correct format and along with editing information. At the completion of the data gathering, some several hundred spectra may be held on the tape and will require examination and further analysis to be performed on them. It will be necessary at this stage for the experimenter to occupy the console position on the computer in order to carry out this work.

Thus in a few months time it is likely that two or more people will want to use the computer console and display simultaneously for considerable periods of time even though the functions they are performing are quite simple in concept. Many of their operations in fact could be performed with little or no requirement of the central processor time if direct access were available to the computer memory. Such a scheme is in the process of being implemented and will provide a considerable increase in the flexibility of EDAC with a potential for further development as the need arises.

The basis of this improved system is a time-shared and directly accessible memory as shown in Fig. 4. A memory switch will allow up to four processors or similar devices to address up to four 4K blocks of memory (expandable to eight blocks of memory) in order of a pre-selected priority. Thus if two processors require simultaneous access to one memory then the input with the highest priority acquires immediate access and the other waits until the memory is free. If, however, the two processors require use of separate memories then they each have immediate access. In fact, four processors can work completely in parallel and with immediate access if they are calling up separate memories.

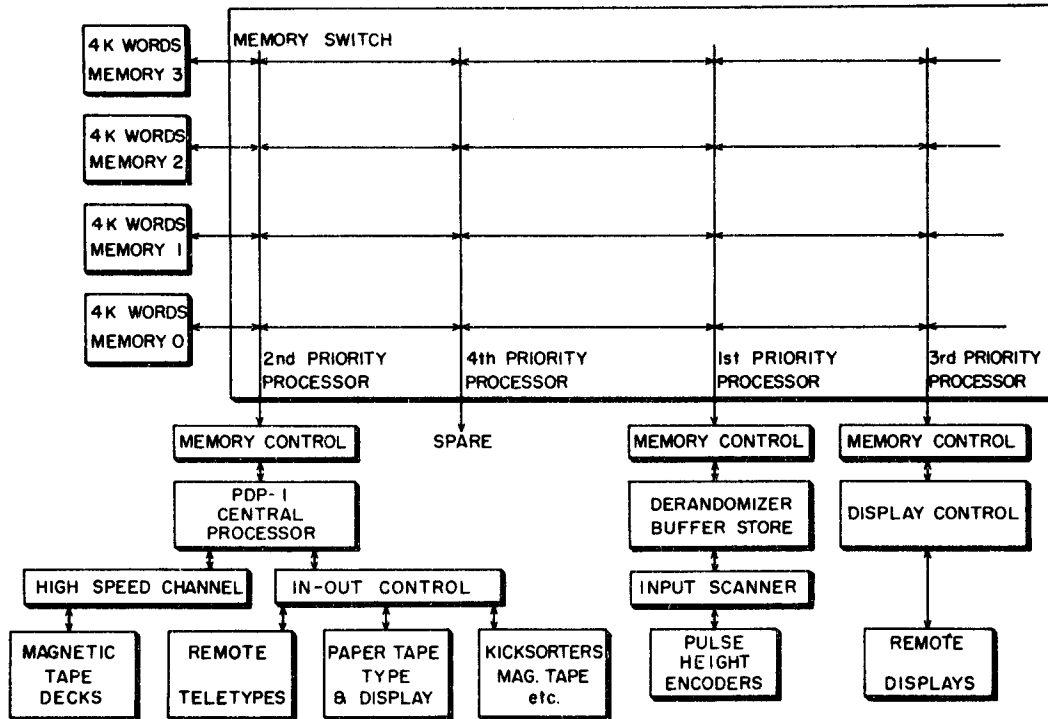


Fig. 4. Time-shared memory.

The existing PDP-1 processor will be connected in as a second priority device in this arrangement. The first priority input will be available for direct entry of data, via a derandomizer, from pulse-height encoders, etc. Third priority will be available for oscilloscope displays and the fourth input will be spare for the time being. The central processor is to be increased in usefulness, as previously mentioned, by the addition of its own magnetic tape decks. These will be available for storing all the program sub-routines and large quantities of experimental data not able or required to be held in core storage. Transfer of data to and from the tapes will be via a high-speed channel which after being initiated by the processor allows for automatic operation. The in-out control will provide for teletype inputs which can be located at the remote laboratories. In this way an experimenter can communicate with the processor without being present at the console.

It is expected that most pulse-height analyses will be performed by feeding encoded data via the derandomizer input. The derandomizer, the logic of which is shown in Fig. 5, will have a five-word buffer store capable of holding 13-bit word addresses and thus able to specify and add one count to any one of 8K words. Addresses will be fed to the buffer store from ten inputs, these being scanned at a 5 Mc/s frequency. It is to these inputs that remote encoders must present their word addresses. Normally several of the bits on each input will be set-up by toggle switches in order that a spectrum of 256 channels, for example, can be routed to an appropriate space in memory. By appropriate settings of these switches

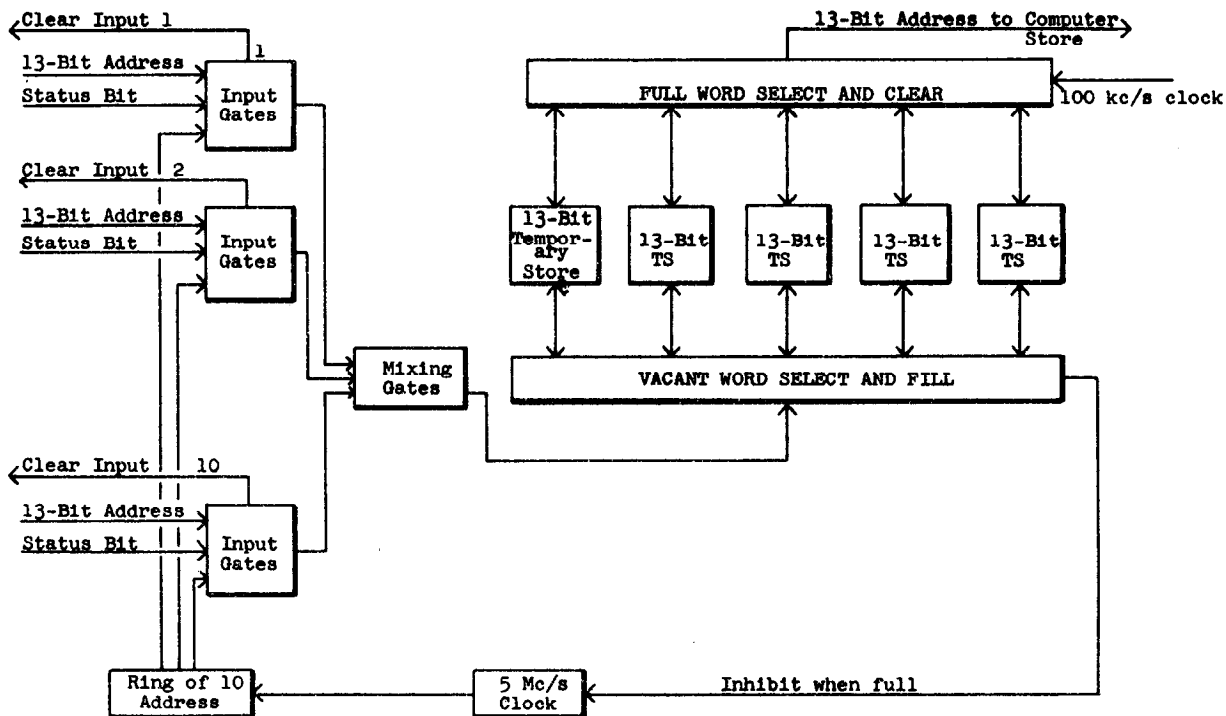


Fig. 5. Derandomizer logic.

ten separate experiments can be routed to individual areas in storage, provided, of course, their total word requirement does not exceed that which is available.

The regular readout rate of the derandomizer to the main memory will be limited to once every 10 microseconds, i. e., half of the 200 kc/s possible. According to Alexander *et al.* (5), with this output rate from a five-word buffer store, counting losses experienced will be approximately 1% with a mean input rate of random events of 70 kc/s. This will thus allow an average of 7 kc/s per input, which is more than adequate for most work. Even a total counting rate of 90 kc/s will only give 5% counting losses. In fact, the largest contribution to counting losses will probably, for some time to come, be due to the encoding systems which precede the derandomizer.

Important measurements to be made, concurrent with the pulse-height distributions, are those of "live time" and "real time," for it is only with these that a true measure of counting losses and counting rate are possible. The most desirable way of recording "live time" and "real time" is to store each along with the associated spectrum in memory. To do so in this case requires that the derandomizer must have the property of providing zero counting losses once it has accepted information at the input, and also that it allows each input the chance to submit information at least an order of magnitude more often than the rate of "live time" and "real time" counting pulses. These two conditions are in fact met by the derandomizer shown. This is because once

the five-word buffer store is full, the 5 Mc/s scanning address is inhibited, thus temporarily halting the input flow and consequent loss of data. However, as soon as one word is cleared in the buffer store, by virtue of the 100 kc/s regular readout rate, the input scanning resumes. Thus the slowest rate to which scanning on the front end can temporarily fall is 100 kc/s, i. e., 10 kc/s per input, which is more than adequate to clear "live time" and "real time" codes which will probably be appearing at 10 c/s each. These codes can be easily generated in the individual encoders, and the probable way will be for "real time" pulses to be routed to channel one and "live time" pulses to channel two, the pulse-height spectra occurring from channel three onwards.

The problem now is to provide adequate displays which can be located along with the encoders at the individual experiments. Referring again to Fig. 4, it is seen that a third priority access line is available for this purpose. Since the derandomizer only uses half of the memory clock rate, a further 100 kc/s rate is available for other work on the same store. The display control unit will provide digital outputs for any group of words requested by the individual experimenters setting start and finish addresses. A new word will be provided every 10 microseconds and it will be necessary to convert this at the remote location to an analogue signal suitable for driving an oscilloscope. Since display readout is performed at 100 kc/s it will only be possible to provide flicker-free traces for a maximum of approximately 4K words (i. e., 25 frames/s). However, since display will be on demand only, it is unlikely that more than this number of words will be requested at the same time. An important feature of the display, heretofore not usually provided on pulse-height analysers, is that it will always be present regardless of whether the input counting rate is high, low or even zero. The only exception to this case is when the central processor, whose priority is one above that of the display control unit, requests access to the same memory for purposes of further analysis of the data. However, this will be momentary and at infrequent intervals so that it will usually pass unnoticed and at worst only cause short-term flicker.

4. Conclusion

The completed EDAC will provide a flexible source of nuclear pulse analysing equipment. It will be possible for one or more experimenters to quickly assemble encoding systems and have at short notice all the fast-access and magnetic-tape storage required for their particular problems. Displays and input-output controls will be available and, in addition to the accumulation of data, it will be possible to perform on-the-spot assessments of the information as the experiments proceed.

5. Acknowledgements

The authors acknowledge the considerable contribution made by F. S. Goulding in the conception of this system and by L. B. Robinson, A. Pearson, A. J. Ferguson and B. Miles in its subsequent development. Also R. E. Archinuk, R. Borbas, J. C. Irvine, J. S. Milton and G. B. Tucker have provided valuable work in getting the system operational.

References

1. W. F. Miller and H. W. Fulbright, Multi-Parameter Analyser Conference, Liberty, New York, November 12-15, 1962.
2. J. Leng, CRL-74, October 1962.
3. M. A. Eswaran, H. E. Gove, A. E. Litherland and C. Broude, BAPS 8, 375 (1963).
4. A. J. Ferguson, CRP-1055, November 1961.
5. T. K. Alexander, H. G. Reddering and J. M. Kennedy, CREL-779, November 1959.

Discussion

TOBACK: Are you planning on standardizing your remote locations with special purpose computers to aid in reducing the amount of time charged to large computers?

LENG: Well, this, in fact, is not a large computer; it's slightly larger than the ones talked about previously for analyzer functions, but no where near approaches the Atlas size and cost. We are talking of the complete cost of this system being \$300,000. Individual experiments will just have basic pulse-height encoders and displays. Their equipment will be quite simple and can be quickly assembled for any particular requirement.

TOBACK: Have you given any consideration to the use of small remote computers for initial predigesting of data before you present this data to your main computer for data reduction? For example, in hodoscope experiments where one is interested in measuring angles between two particles, could you perhaps instead of recording basic address information, calculate trajectory angles and store these data thereby reducing the main computer memory load?

LENG: We are not faced with that problem as such. However, our approach of only using half of memory for data storage and leaving the other half available for computation, which can go on concurrently and independently, would probably lend itself well to the type of operation you described.

TOBACK: I was concerned with the problems where the available memory capacity is inadequate, when one is doing multiparameter measurements, where you need an unusually large memory or when one is doing real time totalization, where you might want to make use of the central processor in order to first reduce the data before doing real time totalization.

LENG: Well, the first thing that one must do with every event that comes in is to store it in memory, so it doesn't matter whether you go through the processor, or directly in. Once you have got that event into memory, then you can invoke an interrupt to instruct the processor to take a look at the number and see if it is of interest or not, and perhaps store it away on magnetic tape.

GOULDING: I think I would like to elaborate on an example which, in a rather different field, illustrates Toback's point.

In looking at tracks in spark chambers, for example, one might like to store the end points of the track, rather than all the intermediate points, just to save storage, and I think you can do this in your system. In fact, you can process the tracks as they occur and store a much reduced form of the original data.

LENG: One can do this either under program control, or alternatively it might be best to do it by hardware externally. One is going to use a combination of the two eventually, I think.

STRAUSS: I have two questions. In the event of high counting rate where you use the memory to accumulate the incoming events from the 5-word buffer, is the machine going to be busy most of the time just storing these events, or is there also sufficient time for other operations?

LENG: The derandomizer output will go to a maximum rate of 100 kc/s. That allows still another 100 kc/s of access speed to the memory for other data transfers. The derandomizer using the full 100 kc/s rate will accept random events up to 70 kc/s with only 1 percent counting losses.

STRAUSS: My second question is: How did you solve the problem of transmitting the information from your encoder or scaler into the computer, which is off line, off the experimental site? Or isn't that the case?

LENG: The computer is located in the Van de Graaff control room, so we don't have a problem with data from these experiments. For other equipment such as the beta ray spectrometer the distance is only 500 ft. and we don't expect to have a problem there. We have done some experiments on transmission of events over about 500 feet of line with complete success. What we do at each station is to encode an event, store it in a temporary register and present dc signals to the input of the derandomizer gates.

STRAUSS: But that you can do that only at a relatively slow rate, I wonder, if somebody else can comment on the transmission problem when you have more than 500 feet; maybe a mile or so, like from one building to another.

LENG: I think it's considered the reasonable thing to transmit these sort of signals up to half a mile over normal twisted-pair lines at a rate

of 50 kc's. We don't expect an encoder to produce events at that rate.

BELL: I want to clear up this DC business. The DC that you are talking about is 1 or 2 microseconds. If you have a long line you have to use fairly nice cable, but coaxial cable, which isn't terribly expensive, can go fairly long distances. And the worst thing that can happen is you may need some sort of repeater. I think anything more elaborate than that can be handled quite nicely with microwave.

LENG: We are prepared to wait for the signals to settle down before we feed them into the computer. We really have plenty of time; tens of microseconds if we want, because our event rate from our pulse-height encoders is not going to be too high.

CHASE: At Brookhaven we are giving a little bit of consideration to the problem of data links from large accelerators to computer centers which may be some distance away. We have considered laying coaxial cables. We have a similar problem to what they have at Harwell, and that is the ditch digging equipment is continually cutting coaxial cables, which would be rather embarrassing in the middle of an expensive accelerator run. We have been thinking in terms of microwave links to connect several buildings. We are concerned with whether or not there may be some problem associated with the radar interference of low-flying planes.

KANDIAH: I am glad you mentioned the problem of radar interference because we asked someone to carry out some field tests just as a probe to see what the problem was. And the first answer was that if all the radar beams were cut down the system would work perfectly up to quite a long distance, but we never knew when someone would come in and completely make rubbish of the signals.