

Program Interrupt on the Univac Scientific Computer

JULES MERSEL

IN the use of high-speed digital computers the general situation has been to have the program or the programmer at the console in sole control of what the machine is doing. In recent years, however, a situation has often occurred where a program, so to speak, has only been lent the machine until a higher priority problem is ready to go on the machine. At these times the machine was stopped, the low priority problem was taken off the machine, and the high priority problem was allowed to run. Though such a situation is preferable to having a computer idle while it awaits the high priority problem, the time lost in getting the old problem off and the new problem on has been an unfortunate loss.

At the National Advisory Committee for Aeronautics' (NACA) Univac Scientific installation at Cleveland, Richard Turner foresaw that the foregoing situation would be the usual thing rather than merely an occasional nuisance. The way Mr. Turner intended to run his machine, and in fact does run it, the wind tunnel has first claim on the services of the Univac Scientific. It is important to NACA's operation that information from the wind tunnel reach the computer as soon as possible, that it be processed as soon as it reaches the computer, and that the results are returned to the wind tunnel immediately after processing.

The fast handling of wind tunnel information was achieved by running input-output lines directly from the wind tunnel to the Univac Scientific. Inasmuch as the Univac Scientific has the capability of having any kind of input-output, this was not a difficult task. The problem of quick access to the machine was solved by Mr. Turner's suggestion of the "interrupt" feature.

The interrupt feature has the following characteristics. When an interrupt signal is sent to the Univac Scientific, the computer, immediately after finishing its present instruction, notes the location of which instruction it would normally do next, jumps to a fixed address which will allow the value of this location to be stored

and from which the computer will find the information as to what it is to do because of the computer having been interrupted.

As an example, in the wind tunnel operation, when the wind tunnel is ready to use the computer, an interrupt signal is sent from the wind tunnel. At the completion of the current instruction, the Univac Scientific notes where it is in the current problem, stores the current problem away, reads in the program for the wind tunnel problem, and starts that program. The wind tunnel problem then reads in its own input from the wind tunnel, executes the needed computation, and returns the answers to the wind tunnel. As soon as the answers are returned it reads the previous problem back into the rapid access storage and continues it from the point where it had been interrupted. A delay of microseconds has been incurred in getting ready for the switchover rather than the much lengthier delay which might be expected without the use of a program interrupt.

This application of a program interrupt inspired much thought as to other applications on which the program interrupt might profitably be used. The wind tunnel application is an input-output application with a device that is not normally considered part of a computer. Thought was naturally given to whether the program interrupt might be profitably used with ordinary input-output equipment. Such pieces of peripheral gear as punched card input and output equipment turned out almost to demand the aid of the program interrupt.

In handling punched cards the normal procedure is that, for 12 very short intervals during the whole card cycle, information must be either read from the rows of the card or punched onto these rows or both. Since on the Univac Scientific less than 1 per cent of the time of the card cycle is required for the reading or writing of rows, the programmer usually attempts to perform computation during the remaining time. This, unfortunately, requires that he carefully segment his computation into 12 groups none of which is time-wise long enough to interfere with the processing of row information. This

is an onerous task and usually results in either safe but inefficiently short segments of computation or segments that cause machine errors due to being too long.

The card input-output device on the Univac Scientific was modified so that it would send an interrupt signal when it was ready either to release or receive row information. The need for the careful segmentation of the computation disappears. When the interrupt signal is received, the computation is stopped, the row information is either written, read, or both, and then the computation is continued from the point of interruption. The elimination of the need for careful timing of computation usually allows much more computation to be performed during the card cycle.

Probably the most interesting uses of the "program interrupt" center around the task of using two computers on the same problem.

Some Univac Scientific users intend to use their machine in conjunction with analogue computers. In this case the analogue computer interrupts the digital computer whenever it is ready to transfer information. The Univac Scientific then processes the analogue data, returns it if needed to the analogue computer, and then awaits new data from the analogue computer. If much delay is expected in receiving information from the analogue computer, computation on other problems may take place during the wait. This, of course, greatly increases the efficiency of use of the computer.

One user is planning to use two Univac Scientifics in tandem. His is a real time problem. One computer smoothes the raw data that comes in from the outside; the other computer processes the smoothed data. In such a case whenever the second computer is ready to receive more data it interrupts the first computer and through the Univac Scientific's versatile input-output system the smoothed data are transferred from the first machine to the second machine. This allows the maximum amount of smoothing to take place within the time limits of the problem. Many other applications can be visualized where in real time problems it is necessary to use two computers in tandem since for these applications no one computer is fast enough to do the job by itself. Since real time problems at such installations would occupy only a small part of the computing day, the two computers are freed to be used separately during the rest of the day.

The last application occurs when, because of the form of graphs displayed on

JULES MERSEL is with The Remington Rand Corporation, St. Paul, Minn.

the Univac Scientific's scope display system, the programmer wishes to change the flow of his problem. If he has previously decided on the alternate course or courses he wishes to pursue due to scope display information, he need only press the interrupt button on the console to

have the computer proceed on its new course. This avoids the stopping and resetting of the computer and the many errors that are inherent in trying to do this quickly.

Like many other new ideas in computer design it is expected that the stimulus of

the needs of the many computing installations around the United States will father many other and exciting uses for the program interrupt. The program interrupt is confidently expected to be a feature of most digital computers in the future.

A Pulse-Duration-Modulated Data-Processing System

J. R. LOWE J. P. MIDDLEKAUFF

A COMMON engineering problem is that of measuring physical quantities which cannot be directly observed. In the case discussed here, these quantities exist in a missile during flight. Since it is not practical to mount recording instruments in the missile and recover them after the flight, the data must be recorded somewhere outside the missile so that they can be reduced to graphical or tabular form.

The usual solution to this type of problem is some form of telemetry. This paper considers one such form, namely, pulse duration modulation (PDM) telemetry, and describes a system for reducing data so obtained. The system represents a great improvement in speed, flexibility, and accuracy over previous reduction methods. Considering the time factor alone, the system described in this paper is a striking improvement. From 3 to 6 months were formerly required to obtain analyses of the data from a missile firing. Now, these can be provided in 2 or 3 days.

In the system described here, the data are transmitted by radio link from the missile and recorded in analogue form on magnetic tape at the ground station. This analogue tape is converted to digital tape which is processed by an IBM type 701 Electronic Data-Processing Machine. The 701 punches cards which are used to produce graphic presentations on an IBM type 407 Accounting Machine. The system provides for a "quick look" at the

data, handles nonlinear calibration factors, plots only selected portions of the data, allows for many kinds of errors in the telemetry, permits subcommutation and supercommutation and numerous other exceptional circumstances. The reduced data are also left on magnetic tape ready to be used as input for other 701 programs.

Brief Description of PDM Telemetry

This is a standard Applied Science Corporation of Princeton (ASCOP) system and will be described very briefly.

The missile carries transducers which measure such quantities as deflections, temperatures, pressures, and accelerations. All these quantities are developed as percentages of a battery voltage.

The transducers are sampled by a commutator (see Fig. 1). In the example shown here, the commutator has 30 contacts (can sample 30 transducers) and rotates 30 times per second, thus producing 900 samples per second. However, the system is entirely flexible in this

regard, and other numbers of segments or commutation rates can be readily handled.

Square wave pulses whose durations are proportional to the sampled voltages are formed, and are differentiated twice. The resulting pulses frequency-modulate a radio wave which is transmitted to a ground station, where the pulses are demodulated and recorded on magnetic tape in analogue form. Simultaneously, a series of 0.01 second time pulses are recorded on a separate track of the magnetic tape. The data track is in the form of a series of groups of 30 "channels," each channel corresponding to a point of the commutator. Each group of 30 channels, corresponding to one revolution of the commutator, is called a "frame."

In practice, two adjacent segments of the commutator are not used. This produces two blank channels in every frame which are useful in identifying the frames. Also, two segments are used to measure the two sides of the battery called "ground" and "reference" voltages. If two batteries are used, as is frequently the case, there are two reference voltages and a common ground.

Fig. 2 is a sample telemetry record called a "lines record." This is a 35-millimeter film containing the same information, and produced by the same type of signal, as the analogue tape. Here the pulse widths are represented by line heights and the time marks by dots along the bottom of the film.




Fig. 1. Commutator

J. R. LOWE and J. P. MIDDLEKAUFF are with the Douglas Aircraft Company, Inc., Santa Monica, Calif.