

The Manchester University Atlas Operating System

Part I: Internal Organization

By T. Kilburn, D. J. Howarth, R. B. Payne and F. H. Sumner

The organization of input and output on the Atlas is described. This is controlled by the supervisor program, and uses varying amounts of core and drum store, supplemented where necessary by magnetic tapes, to obtain, when possible, the maximum overlap between input, computing and output. The system will be used on the Manchester University Atlas, but is suitable for any Ferranti Atlas Installation.

Introduction

Atlas* is the name given to a comprehensive computer system designed by a joint team of Ferranti Ltd. and Manchester University engineers. The computer system comprises the central computer, fixed store, core store, magnetic drum store, magnetic tapes, and a large quantity and variety of peripheral equipments for input and output. The Manchester University Atlas has 32 blocks of core store each of 512 forty-eight bit words. There is also a magnetic drum store, and transfers between core and drum stores are performed automatically, giving an effective one-level store* of over two hundred blocks. The average time for an instruction is between 1 and 2 microseconds. The peripheral equipments available on the Manchester University Atlas include

8 magnetic tape decks	90,000 characters per second
4 paper tape readers	300 characters per second
4 paper tape punches	110 characters per second
1 line printer	600 lines per minute
1 card reader	600 cards per minute
1 card punch	100 cards per minute

The operating system described is applicable to any Ferranti Atlas computer and will be used on the Manchester University Atlas. The system is designed to operate with any configuration of one-level store, magnetic tapes, and peripheral input and output equipments. If no magnetic tapes are available to the system the efficiency may be reduced.

The design of Atlas is such that the transfer of information to and from the core store is controlled by programs contained within the fixed store. In the case of drums and magnetic tapes, information is transferred in "blocks" of 512 words. Word by word transfers proceed automatically; fixed-store programs are only required at the beginning and end of a block. In the case of all other peripheral equipments (e.g. readers, printers, punches) information is copied to and from one-character buffer registers, and is transferred between these registers and the core store by fixed-store programs.

* Papers have been written on the Atlas and on the one-level store, and will be submitted to the Institution of Electrical Engineers and the Institute of Radio Engineers respectively.

These equipments are collectively referred to as "slow peripherals." They may include line printers, paper tape and card punches and readers, special magnetic tapes, graphical output devices, etc.

During the operation of most slow peripherals, the central computer is required for about 1% of the time. Thus even with many slow peripherals operating, the central computer is available for direct execution of a problem for a large proportion of the time. The design of the Atlas is such that these problems cannot interfere with the operation of the fixed-store programs controlling the slow peripherals, which themselves are prevented from interfering by the *Supervisor Program*. The time taken to switch control between the main program and the fixed-store programs is about 10 μ sec, and is included in the above estimate of 1%. This time-sharing of the central computer is generally referred to as the "overlapping" of input, output, and computing.

If the slow peripheral equipments could always transfer information at the rate required by the central computer for any problem, then the maximum overlap of input, computing and output could be easily achieved. However, the central computer requires and produces information at widely varying rates, with an upper limit in excess of three hundred 512-word blocks per second. On the Manchester University Atlas, with four paper tape readers and one card reader, the maximum input rate with all equipments operating simultaneously is half a block per second. The maximum output rate with four paper tape punches, one card punch, and one line printer is also half a block per second. Magnetic tape input and output can increase these rates to sixteen blocks per second per channel. The use of magnetic tape for input and output requires a large amount of off-line editing equipment, and unless a large number of channels are available, it is still possible for the central computer to be input or output limited for short periods of time.

One method by which the amount of overlap could be increased is to hold several problems within the central computer at any one time, and, by means of the supervisor program, to arrange that control is switched between these problems in such a way that the rates at which information is required or produced by the central computer are reduced to those which can be dealt with

by the available slow peripherals. In order to achieve this, it is necessary that the available problems are such that, whilst some of them are producing results, others are computing and others require the input of further information. This needs either a large number of programs available within the central computer, or off-line organization of the problems presented to the computer. Furthermore, this switching between problems can considerably reduce the efficiency of the computer, as it is necessary to store and reset all the common working registers, requiring up to 1.5 msec on the Manchester University Atlas.

An alternative solution is to designate two areas of the one-level store as input and output wells, and to arrange that the central computer draws information from one well and supplies it to the other. These wells are filled and emptied by the slow peripherals operating in parallel. The use of these wells "smooths out" the variations in the rates at which information is required or produced by the central computer. The larger the wells, the greater the "smoothing out" and the greater the overlap obtainable for any sequence of problems. This method overlaps the computing of one problem with the input and output of others; only one problem need be in the central computer at any one time, and no switching of control between problems is necessary. The division of the one-level store between the central computer and the two wells is arranged by the supervisor program, and is capable of continuous variation. The size of the wells is increased by the use of magnetic tapes, and another magnetic tape is provided to enable the suspension and dumping of problems whose continued operation would disrupt the efficient operation of the system. Suspended problems are re-entered when the supervisor program considers such a re-entry to be worth while.

A block diagram of the system is shown in Fig. 1.

The Input System

In order to accelerate the filling of the input wells by the slow peripherals, all slow peripherals are permitted to operate simultaneously. Furthermore, in order that the input system should be as flexible as possible, it is not necessary for the input of one problem to be on one "tape" (in this context, "tape" implies a paper tape or a stack of cards) or to be through one equipment, or even for the component parts to be supplied in any particular order. All sections of input are headed by a title for the purpose of identification (see the paper by Howarth, Payne and Sumner on p. 226).

A complete problem is referred to as a "job" and one of the component sections contains a job title together with a list of the titles of the other sections. The input information is organized within the system by the supervisor program, which maintains lists of the titles and lists of incomplete and complete jobs. Only jobs whose slow peripheral input is completed are made available to the central computer. The system is best described by considering what happens when "tapes" having the various permissible headings are read into the core and drum store. A tape bearing a "job" title causes this title to be added to the list of incomplete jobs, and a system number is allocated to this job. The component titles of the job description are added to the list of titles. Subsequently, any "tape" bearing one of these titles causes the relevant entry in the title list to be labelled as "present," and the information is stored, together with the title and the job number. The first of these "tapes" to arrive is stored after the job description and, whenever any piece of information overflows from one 512-word block into another, then the last line of the first block contains the number of the next block, and the first line of the second block contains the job number. When all the component titles have been labelled as "present," then the job title is transferred to

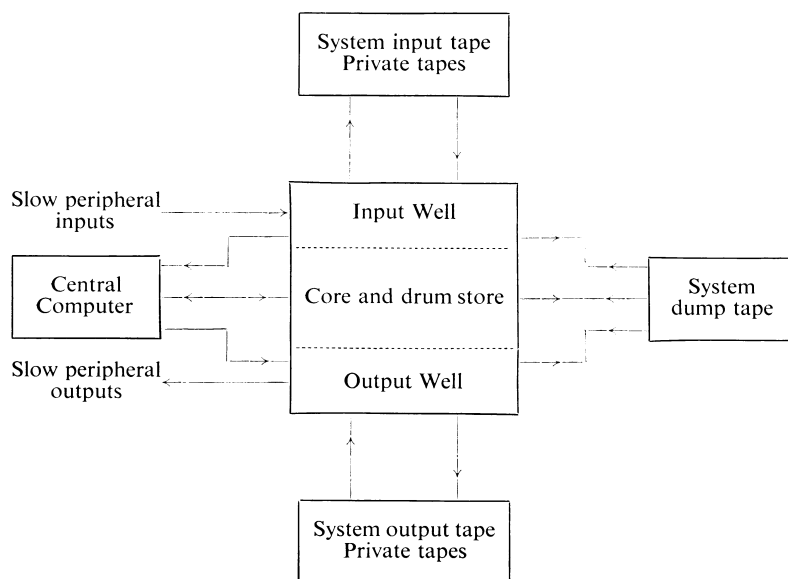


Fig. 1.—Block diagram of system

the list of complete jobs. Whenever the central computer requires a new job it extracts a title from this list, and all the appropriate blocks from the input well, which are then declared to be empty. If a component title is read before its job description, then the title is added to the title list, and the first action on reading a job description is to scan the title list for any component titles which may already be present.

So far no mention has been made of the system input magnetic tape; the use of this tape is not essential to the system provided that the input well in the core and drum store can be made sufficiently large to assemble complete problems read in from slow peripherals, and to retain them until they are required by the central computer for execution. However, without the system input tape, it is possible for the organization to become inefficient under certain circumstances. For instance, if computation on one problem continues for a long time, the amount of input supplied by the slow peripherals will exceed the size of any reasonable well; again, if a problem requires a large amount of store for its execution, then the size of the input well has to be reduced.

When the system input tape is in use, the input well within the core and drum store is divided into two parts. One part is filled by slow peripherals, and complete blocks of information are transferred from it to the system input tape. The other part is filled with complete problems by scanning the system input tape, and these problems are then transferred to the central computer as required. If information is being written to the system tape at point A, then the length of the system tape which may be scanned without effecting the flow of information from the slow peripherals is directly related to the size of the first part of the input well. For example, suppose this part of the input well consisted of five 512-word blocks (out of a total of over two hundred blocks in the core and drum store on the Manchester University Atlas). These can be filled by four paper tape readers and one card reader in 10 sec. The length of the scan is that amount of magnetic tape which can be scanned in 5 sec, that is 75 blocks. Thus the effective size of the input well is 80 blocks, only five of which need be in the core and drum store. These 80 blocks can be filled by the slow peripherals in 2.7 min. If this represented 80 programs in source language, then this well itself is capable of smoothing out an irregularity of 80 : 1 in computing to input time for an isolated problem. If at any time problems not yet dealt with by the central computer extend further along the system input tape than can be efficiently scanned, then either the input well can be increased, which increases the length of the permitted scan, or, if this is not possible, part of the system input tape is transferred to the system dump tape in order to reduce the length to be scanned. Problems so transferred will be treated later at a suitable time chosen by the supervisor program. The size of the second part of the input well will be allowed to change between the widest possible limits, since the rate at

which blocks of information can be processed by the central computer will vary extensively.

When the system input tape is being used, all information from slow peripherals is stored on this tape, and a permanent record of all input is preserved. Thus it is possible to correct programs within the computer by reading in a short correction; the amount of off-line editing is therefore correspondingly reduced. The second half of the input well may also be loaded from previous system tapes and from private magnetic tapes, this operation being under the control of the supervisor program.

The Output System

The output well is divided between the slow peripheral output equipments in the approximate ratio of the rates at which they can transmit information. This division can be varied by the supervisor whenever the output for any slow peripheral exceeds its well size, provided such a change does not affect the operation of the other equipments.

The maximum rate at which the output well can be emptied when all the slow peripheral output equipments are operating simultaneously is 1,760 characters per second on the Manchester University Atlas, where the output equipment comprises one line printer, one card punch, and four paper tape punches. For a problem which produces a large amount of output for one or more peripheral equipments, the output well could be filled very quickly. If the wells cannot be extended, and if the system output tape is not being used, then the program producing the results must be suspended and possibly replaced by another.

The organization of the system output tape is similar to that of the system input tape. The output well is divided into two parts, the first part of which is filled by the central computer, and the second of which is filled by the system output tape. The system output tape is partitioned into sections; each section contains one block for each paper tape punch and card punch, and ten blocks for each line printer, i.e. fifteen blocks on the Manchester University Atlas when all slow peripherals are required. A complete section will keep all the slow peripherals operating for approximately 40 sec. The second part of the output well consists of one or more of these sections. If the slow peripherals are to operate at maximum efficiency, it is necessary that only complete sections should exist in the second part of the output well. These complete sections could be assembled by the supervisor in the first part of the output well, but this would require a very large output well if any problem produced several blocks of output for any particular slow peripheral. This could be overcome by permitting a program to produce only one block of output, and then suspending it and obeying other programs until a complete output section is obtained. The organization of such a system would decrease the efficiency of the central computer due to the necessity of program switching. The

solution adopted is to organize the compilation of complete sections along the system output tape. In this way a program can produce several blocks of information, partially filling several sections, before it need be suspended. The supervisor organizes the allocation of output channels in such a way that only complete sections are transferred to the second part of the output well.

Suppose the system output tape is transferring complete sections to the second output well from the point A, and a program has produced so much output for a particular slow peripheral that this output is being written into a section of the system output tape at the point B. Then if the system output tape can be scanned from A to B and back whilst the slow peripherals are dealing with the contents of the second well, the program producing this data need not be suspended. On the Manchester University Atlas, if the second output well is one section of fifteen blocks, this will occupy the slow peripherals for 40 sec. During this time, 630 blocks of magnetic tape may be scanned. Thus the maximum separation of the two points A and B is 315 blocks, or 21 sections. Therefore, in such a system, a program can produce 21 blocks of information, even for the slowest output, before it need be suspended. Ten times this amount of information could be produced for the line printer. This technique of writing to the system output tape is possible on Atlas because pre-addressed, fixed block-length, magnetic tapes are used.

The first part of the output well is not divided into sections, and if necessary the whole of this well can be filled with information for one particular slow peripheral output equipment. The size of the first part of the output well is varied by the supervisor and if at any time more output is produced than can be contained in this well, then this output is written to the dump tape and transferred into the well as soon as possible. Alternatively, the program supplying this large amount of output can be suspended.

The slow peripheral output equipments cannot operate continuously for more than about 10 to 15 min, because they have to be disengaged for loading a new reel of tape or more paper. The reloading operation may take 2 min. It is possible for the system to take

account of this situation. At any time the supervisor program knows how many blocks of information have been supplied by the central computer to any slow peripheral, and can know when the tape or paper requires changing. It can then inform the operator of this, and suspend sending information to this output for a reasonable time during which the tape or paper can be renewed. This is easily done by writing incomplete sections on the system output tape. Furthermore, if a section arrives in the second output well and requires a slow peripheral which is disengaged, then the appropriate part of the section is transferred to the first output well, and thence to the system output tape again. If necessary the supply of information for this slow peripheral output from the central computer is suspended. In a similar way, if a slow peripheral output is permanently disengaged and information for it still exists on the system output tape, then this information is recirculated from the second to the first output well and reassigned to one of the slow peripherals which is operating.

The System Dump Tape

The use of this magnetic tape to store excess input and output information has already been described. Also, in the above description, reference has been made to the suspension of programs. When this occurs, the program is transferred to the system dump tape to make room in the core and drum store for a program which succeeds it. Another occasion when a program may be dumped is when it computes for so long that either the input well overflows or the output well is emptied. In either case, it is more efficient to proceed to a different problem in an attempt to maintain continuous overlap of input, computing, and output.

Acknowledgements

This work forms part of the Atlas project. It has benefited from many helpful discussions with the authors' colleagues at Manchester University and Ferranti Ltd., whose permission to publish is acknowledged.