THE ATLAS 2 SUPERVISOR

by D.W.Barron (Cambridge University) and C.R.Spooner (Ferranti Ltd)

1. INTRODUCTION

This paper describes some aspects of the supervisor system being developed for Atlas 2. Its primary concern is to expound a philosophy as to the best way to utilise a large, fast computer, rather than to describe a particular system in detail, and it treats the general structure and concepts of the supervisor rather than the details of implementation. We first give a brief account of Atlas 2 so as to indicate the sort of machine we are interested in.

Atlas 2 is a medium-large to very large computing system, developed from the Ferranti Atlas 1 to give much of the power and facilities of that machine at appreciably less cost. It has been designed jointly by Ferranti Ltd and Cambridge University, and the prototype is currently being commissioned at Cambridge. It can have a core store of 32K, 64K or 128K words capacity (1K = 1024), with a cycle time of 2.5 or 5 microseconds. Overlapping of operations can produce an average rate of operation (with the faster store) of about 2½ microseconds per instruction. The core store is supplemented by magnetic tapes, up to 8 decks with two independent transfer channels, or up to 16 decks with 4 independent channels. Extensive program interrupt facilities are built into the hardware; in particular all peripheral devices are controlled on an interrupt basis. The machine runs under the control of a supervisor program which controls the peripheral devices, schedules jobs with time-sharing where advantageous, and provides all the facilities of an elaborate operating system to the user.

2. TIME-SHARING

Atlas 2 is a time-sharing machine. Unfortunately, time-sharing is now a suspect term, and indeed is almost meaningless without some qualification. Sometimes, time-sharing features are incorporated in a machine apparently with little thought as to how useful they will be, and some time-sharing machines have been severely criticised, with justification, by programmers who found the time-sharing more of a hindrance than a help.

- continued -

We therefore start by defining a useful and practicable time-sharing system as one which, by allowing the parallel processing of several activities, maximises the ratio of efficiency to cost of the computing system without prejudicing either safety or user convenience.

To be safe a time-sharing system must incorporate full protection, that is to say it must be impossible under any circumstances for a program to affect any other program with which it is being time-shared. To be convenient the system must be designed so that programs are written and compiled separately, without the programmer or the compiler having to take account of other programs. Also the time-sharing must not dictate a rigid pattern of day-to-day operation of the installation.

Protection can only be provided by hardware. Convenience is provided primarily by software, though suitable hardware facilities can help. A time-sharing system such as we have defined opens up the possibilities of new modes of operation, which we develop in the remainder of the paper.

3. INCREASING EFFICIENCY

To obtain maximum efficiency it is necessary that the central processing unit shall be fully occupied for the greatest possible proportion of the time. The things that can prevent this are delays due to

(a) slow peripheral devices

(b) program loading

(c) magnetic tape transfers

It is only when these delays become large compared with the average execution time for a program that they assume any significance. Thus it is only fast computers that are affected, and to that extent our problem is a new one. However, as computers increase in speed the situation will get worse; this means that what is marginally worth while today will be essential tomorrow. The fact that there exist techniques which alleviate the difficulty on present machines should not blind us to the need for a more elaborate approach to deal with the faster machines which are now emerging.

Computers have for some time been substantially faster than the slow peripheral devices, that is paper-tape readers, card readers etc. Once these devices became too slow to connect on-line, the most common development was the use of satellite equipment to buffer information to and from magnetic

- continued -

tape. At first special purpose converters were used, but nowadays the use of a small satellite computer is almost standard practice for a large machine. The satellite system also helps to remove the second cause of delay, program loading, since a stream of programs can be assembled on a tape and executed in sequence. However, this gain in efficiency is at the price of some convenience, since it is not easy to execute programs other than in the order they appear on the tape. This means for example, that special expedients are needed to run an urgent program of "crash" priority, and also that any scheduling is the responsibility of the operator who loads programs on the satellite computer. As machines get faster the pressure on operators will increase and it is unlikely that any complicated scheduling to balance the load on the various parts of the system will be possible. This will mean a consequent reduction in the overall efficiency of operation.

The alternative to the satellite system is the "self-contained time-sharing" system, in which peripheral devices are controlled by time-sharing on an interrupt basis, and the operation of the computer is controlled by a supervisor program. Such a system allows reasonable scheduling to be done internally. It also reduces the load on the operators by removing the need for programs to be loaded as units; the system can accomplish the necessary sorting of multiple input streams into separate programs.

The third type of delay, magnetic tape transfers, only becomes significant for a machine of the "microsecond" class such as Atlas 1 or Atlas 2. To remove this source of inefficiency there is no alternative to time-sharing between programs in order to absorb the waiting time, and this implies more or less elaborate internal scheduling since the jobs which are to be time-shared must be carefully selected. Effective scheduling can be done by the self-contained system defined above, but not by a conventional satellite system.

Thus although the satellite system is acceptable for most present machines, it is inflexible and lacks development potential. The self-contained time-sharing system, as used for Atlas 2, offers all that the satellite system can, with at least the same efficiency, and it can also cope with the special problems presented by a very fast machine.

4. REDUCING THE COST

So far we have considered efficiency as an end in itself. Equally important is the cost of obtaining it, since absolute efficiency is of little interest if it is too expensive to achieve. The true purpose of time-sharing is not solely to increase efficiency, but to increase the ratio of efficiency to cost, in other words to increase the "value for money" of the installation.

- continued -

Part of the cost of a satellite system is explicit - the cost of the satellites. The counterpart of this in a self-contained time-sharing system is a proportion of the central computer time used by the supervisor. However, this simple comparison is misleading, for we must take into account the cost of the central computer hardware, i.e. core store and tape decks, used by the system. Some of this is obvious - the tape decks used by satellite systems and the core store occupied by supervisor systems. But there may also be hidden costs. In a system that relies on having, say, four programs simultaneously in core-store to time-share, there is the cost of the core store occupied. On a satellite system, extra tape decks are required whenever multiple input or output streams are desirable (e.g. for monitoring), whenever mixed external media are used, or whenever an urgent job must jump the queue.

In choosing a system for Atlas 2 we considered both efficiency and cost in its widest sense. We rejected satellites as the normal method of input and output both because their potential efficiency is limited, and because the efficiency they do provide is unnecessarily expensive. The self-contained time-sharing system which we have developed provides the required efficiency and benefits at less cost, but it can work with satellites on those occasions when satellites have an advantage, e.g. remote stations. (Though it is arguable that a time-sharing machine allows a better way of treating remote stations - namely to provide remote consoles connected on-line to the computer by a data link).

The potential of the new system is far greater than that of a satellite system, for with comparatively little extra equipment it can make the computer easy to use, and it can perform the intelligent internal scheduling that makes time-sharing between programs economically advantageous. However, the system does not depend on time-sharing between programs, nor require it, when it is not the most economic solution.

5. STATIC FLEXIBILITY

Most existing supervisor or monitor systems use a fixed amount of core store and a fixed number of tape decks. This means that many installations do not obtain the best value from this equipment because their requirements from the system are such that some of the tape decks, or some of the core store, is hardly ever used. It is impossible to meet all requirements in the most efficient manner with a fixed allocation of equipment. There are essentially four variables to take account of:

- continued -

(a) type of jobs to be run

(b) facilities to be provided

(c) amount of core store to be used

(d) number of tape decks to be used.

If efficiency is ignored these are independent variables: as soon as an efficiency constraint is imposed the variables interact in a complicated manner, and a compromise "best" solution must be sought. The Atlas 2 supervisor allows a wide range of choice for items (c) and (d); thus an installation with very large programs might restrict the supervisor to 6K of store and 2 tapes, whilst another installation with a very large throughput of small programs might obtain best efficiency from a supervisor using 6 tapes and 12K of core store. The system can be run with any amount of core store from about 5K upwards, and any number of tapes from 0 to 6; in general fewer tapes means more core store used, but no hard and fast rule can be laid down.

6. DYNAMIC FLEXIBILITY

The hardware configuration is chosen to meet the "normal" needs of the installation. If the installation processes only one type of job this provides a very efficient system, but if there is a wide variety of jobs then the hardware requirements of the object programs will vary. Many will not use all the nominally available core store, but there may be the occasional very large job. To prevent decks and core store being under-utilised, the supervisor varies its requirements for tapes and core store dynamically so as to absorb any surplus capacity; thus whilst small jobs are being run it will expand to use the available core store; but when a big job comes along it will automatically contract. Spare tape decks are absorbed in a similar way: the changeover is smooth and rapid, often taking only a few seconds, and during the changeover the system continues to operate efficiently. This full utilisation of equipment maximises the "value for money" of the installation.

Time does not permit us to detail the techniques that are used to obtain this combination of flexibility and efficiency. Since in the interests of reducing cost the basic Atlas 2 does not have drums or discs, it has been necessary to devise a novel method of amalgamating a variable number of tapes into a semi-random access backing store. This is combined with an elaborate system of dynamic store allocation, which is assisted by the fact that all programs are relocatable, the relocation being done by hardware at run time.

- continued -

7. CONCLUSION

In conclusion we must confess to a stroke of luck. We set out to design an efficient system within the constraints of the machine's hardware. There was no reason to suppose that the most efficient system would also be a flexible system: however, when we had designed an efficient system we found that the flexibility followed with very little extra work.

Acknowledgements

The work described in this paper forms part of the Atlas 2 project by a joint team of Cambridge University and Ferranti Ltd., whose permission to publish it is acknowledged. The authors also wish to record the influence of the corresponding Atlas 1 project by a joint team of Manchester University & Ferranti Ltd.

D.W.B.
C.R.S.

9th April, 1963.

[redacted]