The use of a large computer on a bureau basis

By A. S. Douglas*

The problems arriving to be handled by C–E–I–R facilities are very varied and include almost every type of work. A large part of the load on a large computer consists of short runs for testing or simple production purposes, preceded and followed by considerable amounts of input and output. Moreover, much of the work done is subject to monetary or time constraints. The effects of these latter factors are discussed in this article in relation to operating procedures and program discipline for various machinery systems, including the IBM 1401/7090, of which C–E–I–R has practical experience. Some comments are offered on the implifications of this experience on the development and use of future systems.

The association of C–E–I–R with large computers is of long standing and largely springs from our interest in economic and industrial planning. In the U.K. we at present operate on an IBM 7094 model 2, and expect shortly to begin operating on the London University Atlas under the aegis of our majority shareholders, British Petroleum. In the U.S.A. our American parent operates two IBM 7090 and one IBM 704 system, and shares three IBM 7094 systems. Thus we have considerable accumulated experience in working on these large machine complexes.

Analysis of the load

In London we presently use about 120 hours per month. Approximately 80% of this time is concerned with economic studies, carried out mainly by the techniques of linear programming or simulation. Another 10% concerns general mathematical, statistical, physical and engineering problems for which some special programs (for example Time Series Analysis, Multiple Regression, Factor Analysis, Adaptive Forecasting, and Critical Path Planning) have been developed. The bulk of this work is written in FORTRAN. The remaining 10% concerns data processing and market research. For the first task COBOL and standard routines such as IBSORT are used. For market research a special program, OPAL, has been developed. However, a substantial part of the work—on very large surveys, or on parts explosion problems, for instance—is written in FAP, a symbolic machine language.

From the point of view of operating the machine the work done can be divided into four main groups, (1) program debugging, (2) model debugging, (3) short production runs and (4) substantial production runs. Generally speaking group (1) gives rise to runs on the computer of less than two minutes, (2) and (3) incorporate mostly runs of between 2 and 8 minutes, and only jobs of group (4) run for substantial periods.

It is interesting to look at the distribution of machine runs. A table of these was given in the *Journal of the Royal Statistical Society* by Dr. M. G. Kendall two years ago (Kendall, 1963). There has been no significant change in this pattern since then. This table shows that some 80% of all runs last for less than 2 minutes, 15% between 2 and 8 minutes and only 5% for more than 8 minutes. However, if one considers the total time spent on each group, approximately 80% of the time on the machine is spent in production runs of more than 8 minutes, 15% in runs of 2–8 minutes and 5% in runs of 2 minutes or less. It is, perhaps, fair to point out that the amount of development work on our machine has not been high, since the figures all refer to a period during which a standard LP package has been in use. During periods of intensive program development one might expect the number of short runs to increase and the proportion of time spent in debugging to be rather higher—though it could not normally go too high for economic reasons.

Requirements for operating

With this pattern of usage there are certain required operating conditions which we have found by experience must be met. Efficient utilization of skilled programming staff, coupled with the need to meet difficult deadlines, implies that there must be rapid turn-round of programs handed in for debugging. We attempt to provide facilities for four to six "shots" per day for the programmer. This is done by operating "monitor" runs three or four times during the normal working day, and allowing additional runs to be made at night on request. Since programs and data have to be stacked on tape on the 1401 initially, and the results printed out subsequently, actual turnround from the time the last cards are handed in to the time results are distributed is 1 to 1½ hours normally.

Of course, it would be convenient to have uninterrupted production runs for group (4) jobs. In practice this is difficult to achieve except at weekends. For one thing we share the machine with IBM and thus scheduling problems may make it essential to chop jobs up. Secondly the machine may fail (although it has rarely done so in fact) or (more frequently) faulty tapes may interrupt running. To cope with these situations production work is normally capable of being broken down

* *C–E–I–R Ltd., 30–31 Newman Street, London, W.1.*

into runs of 20 to 30 minutes duration, and all large programs are provided with get-off and restart facilities, the get-off being activated as a matter of routine at intervals of not more than one hour.

We have found that both during the debugging of large program systems and during the debugging of models it is essential to have the program writer or model builder in the machine room from time to time. However, we do not regard it as efficient for such people actually to operate the machinery, as they are liable to cause "finger" trouble, inefficient tape loading and so forth. In some installations this might not be of importance, but we are contracted to pay for all time used (except for an allowance for setting up and taking down tapes at the start and finish of a period in the computer room) and thus cannot afford unnecessary mistakes or delays. We have thus evolved a scheme known as "command and control". Under this scheme operation is done by the operating staff under the verbal direction of the programmer or model builder, who assumes responsibility for any time wasted in carrying out his instructions, unless clearly due to incompetent operation. This scheme enables the advantages of "closed shop" operation to be combined with those of "open shop" use of the machine by professional staff, and we have found that it meets adequately a very real need which can otherwise be a source of friction between operators and professionals.

Scheduling problems

One of the most difficult problems for bureau operation is how to schedule the mixture of loads efficiently. With an adequately large complex of machinery this would cause little concern. However, for economic purposes it is usual to work with the minimum complex capable of handling the total load. In order to provide efficient debugging turn-round we normally arrange 30 to 45-minute monitor runs at fixed times during the day. Long production runs usually take place in the evening or at week-ends, since our share of the day shift on weekdays is limited. The major scheduling difficulties do not arise on the large machine itself, but in the 1401. This is because the 1401 is also used for data processing, production runs for which tend to clash with its use for tasks peripheral to the larger machine. Nor is our problem here unique. Users of 7070 and 1410 systems associated with 1401's have, I know, had similar difficulties, and these have occasionally been exacerbated by the introduction of tape switching. In one such instance known to me, the use of tape switching has constricted the use of each machine to half its nominal capacity due to mutual interference of the systems.

Although for economic reasons it is desirable to combine the data processing and peripheral activities on a single 1401, the problem of mutual interference implies that this combination of work cannot be extended to load the machine fully. We are presently using over 350 hours per month on the 1401, and have found it necessary to buy extra time to ease the clashes of interest. In our particular type of work I suspect that we shall continually be forced to extend our capacity long before the nominal capacity of our machinery is reached, if we wish to maintain a high standard of service to our clients. Nor do I believe this to be a bad thing operationally, since I consider that accounting considerations, pressing us to squeeze the maximum out of our machinery before change, tend to hold back the progress necessary to healthy survival in our rapidly changing technology.

Some remarks on future progress

The difficulties we are experiencing in effective scheduling of work lead me to question whether the manufacturers have not gravely underestimated the complexity of mechanizing this task on computers such as Atlas, the CDC 3600, and the large IBM 360 systems. To give a complete solution either requires, as I have said, a very large (and somewhat inefficient) configuration, or a solution of the job shop scheduling problem. The latter problem has so far proved notoriously intractable. The present networking and resource allocation schemes are rudimentary, and depend heavily on human intervention by experienced engineers or schedulers. We are still far from mechanizing even the comparatively straightforward problems of train signalling or airport control. I would commend ICT-Ferranti for a courageous attempt to overcome this problem on the Atlas. But it still has to be proved to me that their scheduling system can work efficiently, or reasonably so, under all mixes of programs. I also look forward to the implementation of schemes similar to Project MAC at M.I.T., since this should greatly improve turn-round time for program debugging whilst ensuring a reasonably prompt execution of production tasks. However, I do not feel this will be entirely satisfactory until more rapid data transmission becomes economically practicable. I feel that the GPO must be pressed on this, since it is essential to the future development of the art.

We in C–E–I–R are looking forward to extending our knowledge and understanding of the problems I have outlined and, perhaps, to contributing to their solution. It is fitting, I feel, that I should conclude this article by a reference to the help and support which the Atomic Energy Authority—and especially Aldermaston—have given us and others in the field by training (unintentionally) for us several of our operating staff. Unintentionally, no doubt, in the sense that they did not train them specifically to come to us, but for their own purposes. But it has been fortunate for all of us using large computers that their training has been so thorough and effective, and it is a pleasure to acknowledge publicly how helpful this has been to us.

Reference

KENDALL, M. G. (1963). "Address on the installation of the new President", *J.R.S.S.*, (Series A) Vol. 126, part 1, p. 42.