Estimating computer performance

By J. A. Gosden

Methods are described and illustrated for measuring the performance of different computer systems and configurations on standardized problems, as they are being developed as part of *Standard EDP Reports*.

In order to provide reliable performance comparisons, the methods of measurement must be objective. Standardized problems must be precisely defined to minimize subjective error and be unbiased toward particular computer types, yet they should be flexible enough to exploit useful individual features.

One of the standard problems is a file processing run in which records from a detail file are used to update a master file, and a printed record of each change is produced. Record layouts are fixed for the detail and report files but are left flexible for the master file. Detailed block diagrams define the internal operations that must be performed.

Graphs of performance corresponding to different parameters are presented and discussed. The results provide reliable performance measures, are comparable among many different computers and tasks, and are economical to produce.

Introduction

In the early days of computers it was a significant achievement to debug and to run a job successfully, and the running time was not a vital consideration. But during the first phase of expansion of the use of computers the problems of underestimation were a common occurrence. Many lessons were learned and much experience gained during this time, until nowadays, there are competent people who can prepare accurate estimates of both program running times and program preparation times. The estimating procedures which estimators use are, in general, based not only upon wide experience of the problems being considered, but also upon the techniques, people and equipment to be used in the implementation. The techniques are usually dependent upon careful detailing of all the elements of a problem and some simple arithmetic. Encouraged by the growing accuracy of specific estimates, the editorial staff of *Standard EDP Reports* has tackled the problem of estimating the performance of computers in general. This paper introduces the contexts in which these estimates are required; the decisions we have made to produce reasonable generalization; the techniques we use for estimating; the types of results we obtain; and comments upon the value of the results.

General Background

Standard EDP Reports is a monthly subscription service which produces detailed reports on competitive computer systems. Each report is divided into many sections. Each section is devoted to an individual unit of a system; for example, a magnetic-tape unit or a COBOL translator. For each type of unit, all the important characteristics are recorded in a standardized format and the performance of the unit is measured. For example, for a magnetic-tape unit, we measure the basic elements such as the peak transfer rate, and also compound measures such as the product of number of blocks per second and the number of characters per block. We call this an "effective rate" because it allows for the interblock gaps and the start/stop times (see Fig. 1). We also note the time that must be devoted by the central processor in the operation of such a peripheral unit, which we call the *central processor penalty time*. Estimating the performance of individual hardware units is fairly straightforward provided the basic facts can be obtained. Estimating the performance of software units, however, is more difficult, particularly the translation times of compilers. As a part of a comparison service, we feel that it is also necessary to present some overall measures of performance of the whole system.

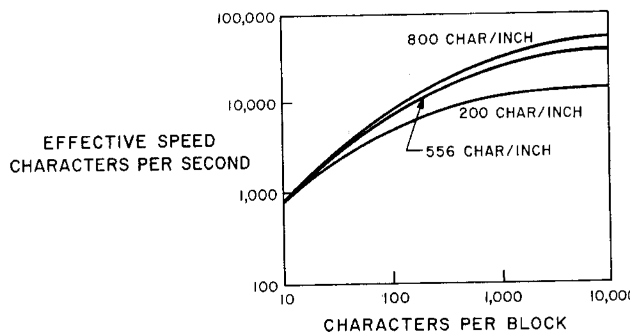

Fig. 1.—Magnetic-tape speeds

Standardized tasks

In order to measure the performance of the whole system, we established a number of standardized tasks, and estimate the times a computer would take to perform such tasks. The most elementary task we use is Matrix Inversion. A slightly more difficult task to estimate is Magnetic-tape Sorting. A still more difficult task is File Updating, considered as a typical commercial data-processing task. A more detailed discussion of each problem is given in Gosden and Sisson

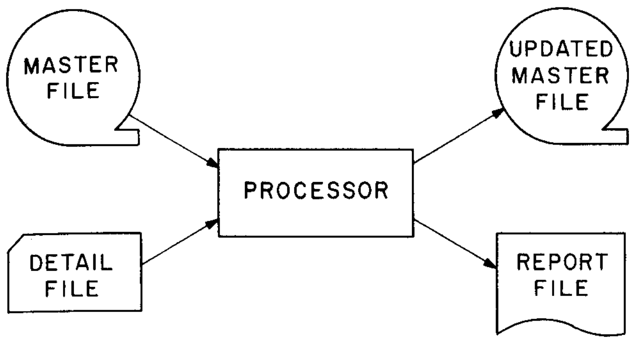

Fig. 2. Run definition for File Processing problem

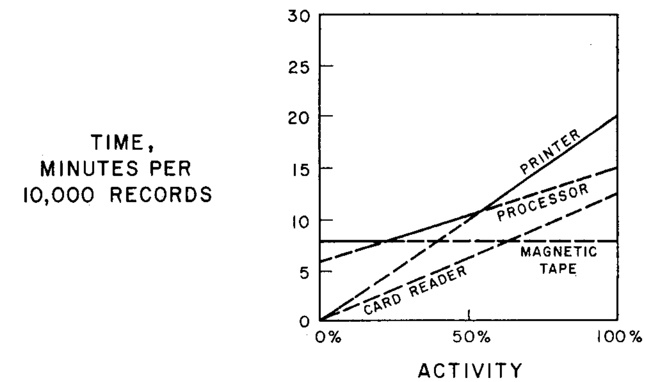

Fig. 3.—Typical results for a single configuration

(1962); in particular, the generalized File Updating problem is given in detail by Gosden (1962) and can be briefly summarized as follows.

We define two input and two output files (see Fig. 2). Each detail record causes a master record to be updated and a report to be printed. We define the detailed procedure we want the computer to perform by means of flow-charts. We estimate the central-processor time by skeleton coding, and the times for each of the input-output units from basic specifications. Then for a particular configuration of the units we calculate the times to execute one run of the job for 10,000 records of the master file, as they vary for different values of a parameter which we call the *activity ratio.* The activity ratio is the ratio of the number of records in the detail file to the number of records in the master file.

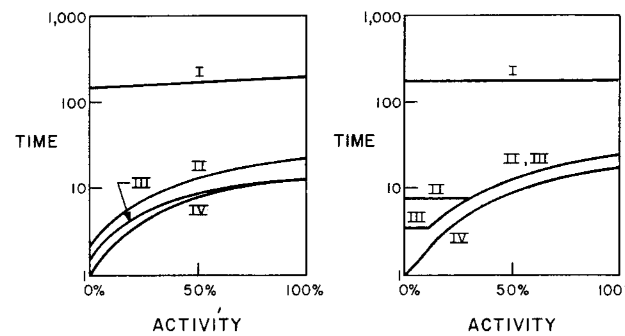

Fig. 4.—Side by side comparison of two systems

Presentation of the results

In the most elementary case, the times for the standard problem can be broken down into five separate parts: the times for the two input units; the two output units; and the central processor. If we plot the times for each of the units against the activity factor, we discover that the plots for the detail file input and the record file output are proportional to the activity ratio. The time for the master file is independent of the activity ratio, and the time for the central processor is a combination of a constant and a component which is dependent upon the activity ratio (see Fig. 3). In the simplest case, in which there is complete simultaneity of operation of all the units, the maximum of these individual plots at each ordinate would be the resultant plot, the solid line in Fig. 3. In the simple case of no simultaneity at all, the resultant plot would be the sum of all the other plots.

For the purposes of comparing times between one system and another we plot the time on a logarithmetic scale because ratios are more important than absolute differences (see Fig. 4). The graph on the left represents one computer system and the graph on the right represents another system. The four plots on each graph represent four different configurations of the two computer systems: a card configuration, and 4-, 6- and 8-tape configurations. The tape speeds and the degree of simultaneity increase with the large systems. When comparing several systems the amount of simultaneity and the speeds of the various units do produce significantly different patterns between different computer systems. Fig. 5 shows the patterns of times for four competitive systems on the same problem. Notice how the variations between high and low activity ratio values are quite different in the four cases.

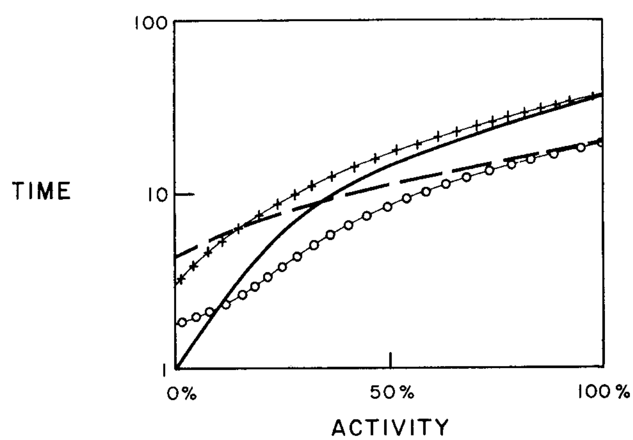

Fig. 5.—Relative times for four computer systems

Other Factors to consider

The value of the results, of course, is closely related to a proper appreciation of the limitations in the generalizing process. These fall into four groups.

Problem variations

Most file updating tasks will differ significantly from our standards. There may be intermediate record sizes, different detailing of the updating tasks to be executed, and widely varying format and editing requirements. Other input and output files may be required to fit one run into a scheme of many runs. A hidden difference is the non-uniformity of the detail file. It will usually be bunched and reduce the simultaneity obtained.

Configuration variations

The configurations that we use have not been deliberately selected to perform well on a particular task. Each case must be examined to see if a small increase in equipment can produce a large improvement in performance, or if a large decrease in equipment would result in only a slight degradation of performance.

Environment variations

We have not made any allowances for errors caused by equipment, operators, data, or the programmers. These may not affect all installations in a uniform way. Nor have we allowed for set-up or change-over times, because these are dependent upon the use of techniques such as supervisor routines and alternating tape units. Times for such operations are given in the reports on each unit. One factor which we will have to consider soon is how to allow for computers which multi-run several programs.

Programming variations

The efficiency of a running program depends upon both the skill of the programmer and the software he uses, assuming that he was adequately briefed about the task in advance. Frequent changes to a program can reduce efficiency. The availability of standard, efficient, input-output routines and operating conventions are also important. In general, we assume good clean programming with the best available software.

The results

The value of the results depends upon a good knowledge of the basis of the estimates and the particular use to be made of these results. It is obvious that they are better indicators of relative performance between computer systems than producers of absolute performance times. For any use, appropriate allowances must be made for the restrictions that we have placed on the tasks and the restrictions that we have placed on the configurations. Any person interested in a particular situation must adjust our figures both for price and

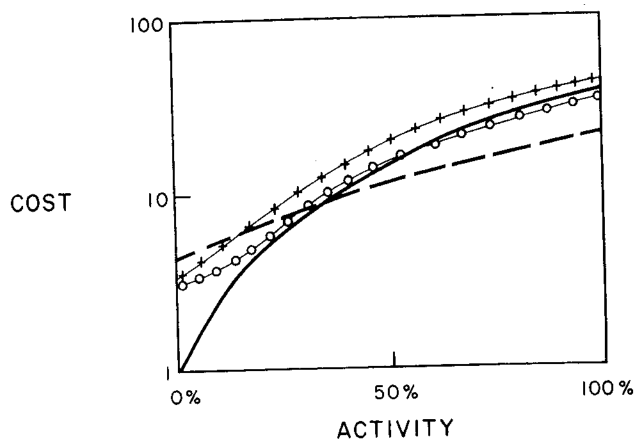

Fig. 6.—Relative costs for four computer systems

performance of each configuration. We specify the steps of the estimating procedure in considerable detail in order to simplify this problem for the analyst.

The results can be changed significantly by consideration of costs. Fig. 5 shows the relative times of four particular configurations performing a standard task, and Fig. 6 shows the relative costs. These plots assume that cost is directly proportional to prime shift rental, which is another simplification for which an adjustment has to be made. When relative costs are considered, we see that one computer system is strongly dominant at high activities and one at low activities.

Fig. 7 shows another interesting pattern which is beginning to emerge from our results. Such results must be interpreted rather carefully because these are based upon a small biased sample. The figure shows a pattern of the relationships between prices of configurations and their times to perform the standard task for ten per cent activity and one of the sets of values for block sizes and unit of computation. The broken lines are hyperbolas that represent lines of equal cost per unit task. The line nearer the origin represents double the value represented by the line further from the origin. If we imagine a line that envelopes the points of "best performance" of the systems we notice that the "optimum" rentals are in the middle of the range, and values decrease as extreme configurations are encountered. We feel that the figures which we produce

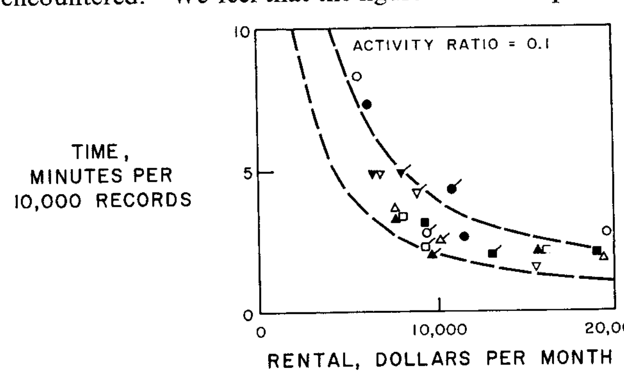

Fig. 7.—File Processing performance

are realistic. We have compared them with figures produced independently for similar types of problem, and there is close agreement. In addition, we have been able to use our procedures to estimate magnetic-tape sorting times on small business computers, and our estimates have been within ten per cent of the quoted times for the manufacturer's own standard routines.

Conclusions

Our work to date shows that it is possible to produce economical and useful comparison figures of computer performance in a generalized format. This is achieved by specifying the tasks in considerable detail. It is essential to begin by making a thorough and detailed timing analysis of all the various units in the system individually. Then it is reasonably easy to combine these figures to suit different variations of the basic tasks and different configurations of equipment. In spite of the limited amount of work we have already done, we have obtained results, and confirmed the practicality and accuracy of our work. We have found no short cuts. The complexities are mainly due to the different varieties of simultaneous operations available, the intricacies of the demands on the central processor, and the wide range in power of different instruction repertoires. We hope later to extend this work, and prepare programs in languages such as COBOL and ALGOL and run them. This will then give us comparisons of computer performance in which the software plays a much larger part.

References

Gosden, J. A. (1962)ed. "Users' Guide: System Performance," *Standard EDP Reports*, p. 4

Gosden, J. A., and Sisson, R. L. (August 1962). "Standardized Comparisons of Computer Performance," paper delivered at IFIP Congress 1962, Munich. To be published in 1963 by North Holland Publishing Co.

Summary of discussion

Mr. W. S. Ryan (*General Post Office, London*): I should explain to you that the opening of this discussion was to be the privilege of the head of the Central Organization and Methods Branch of the Post Office, Mr. Smith, but he was unable to be present today and he has asked me to read his contribution. (*Mr. Ryan then read Mr. Smith's contribution.*)

Mr. C. R. Smith (*General Post Office*): Mr. Chairman, I am sorry that absence on the other side of the Atlantic has made it impossible for me to be present to hear Mr. Gosden's paper in person. I am therefore having to open this discussion as it were by remote control, and from no more than an abstract of what I feel sure must have been a very interesting paper.

I suppose that the only reason that this honour has fallen to me is that it is known that we in the General Post Office have been doing some work in this field; indeed I am to have the privilege of addressing the B.C.S. in London next month on this very subject. But now is not the time to talk about our work and, having stated my qualifications, I would like to open this discussion by welcoming Mr. Gosden into this very difficult field of computer research, where I am sure far too little work has so far been done.

I think that there is one point which we must recognize from the outset. The problem is to choose a computer, but in that choice there are probably a dozen or more different criteria which have to be taken into account. Many of these criteria—such, for example, as manufacturer's experience, quality of his backing service, even the likelihood of his remaining in the computer business—are not capable of a sound numerical evaluation. Final choice of a computer, therefore, is bound to be a matter of judgment, and this we ought clearly to recognize. There have been all kinds of attempts to produce comprehensive indices by all kinds of "points systems," but in my view these only obscure the area over which pure judgment has to be exercised and invest the final decision with an air of absolute authority to which it is not entitled. Much better, I think, to concentrate on separate delimitations of those criteria which are clearly capable of a worthwhile numerical evaluation, and then exercise a judgment over the whole field.

But one of the most, if not in fact the most, important criterion in this decision is that of value for money—what in our jargon we call the cost-weighted performance of a computer. It is to this factor that Mr. Gosden addresses himself, and in this we in the G.P.O. are in complete agreement with him. We think that it is possible to produce cost-weighted performance curves for any kind of computer apparatus or for any configurations of such equipment arranged in a complete computer installation. We know that such curves can be produced with far less expenditure of time and effort than involved in the more usual methods of computer evaluation, and for the purpose required they appear to be reliable.

I would like to emphasize at this juncture what I believe is another very important point, and one on which it is easily possible to become confused—what is the purpose for which they are required? It is a point which appears very clearly in the abstract which I read of Mr. Gosden's paper, in which he talks about the four levels at which computer performance can be measured, and says that measurement at the fourth level "overall operations of individual computer departments" are seriously complicated by a number of factors but that "measurements at the third level do provide useful preliminary estimates of the *relative* performance of alternative systems."

I think that once again Mr. Gosden and ourselves are in agreement here. What we are trying to do is to see which computer gives us the best performance for our money. We are choosing (on costed performance alone) between different computers. A *relative* measurement is therefore adequate for this purpose, provided that there is nothing in the technique of appraisement used which would make the *absolute* performance of different computers on the same practical job of work vary widely from their estimated *relative* performances. I think this will generally be true, but I think it must also be recognized that, having put a number of different computers into an order of relative value for money, the precise configuration of equipment which is to be ordered for a particular job still needs some careful thought. I do not think there is any method—except perhaps in the smallest and most simple cases—whereby the complete answer to this very difficult and complicated question will ever be had simply by "looking it up on a graph." What we can do is to save ourselves—and the manufacturers—a great deal of time and trouble, and get sound starting points for our detailed planning.

Our own work has tended to be concerned more with the medium and large computers, and we have tended to concentrate more on the central-processor performance on a tape-to-tape basis, leaving the amount of input and output equipment required to be assessed on a fairly simple balancing calculation once the central processing performance has been assessed. But our objective is slightly different from Mr. Gosden's. We were trying to find the "best buy" against an inexhaustible head of work for one computer, whereas on the curve drawn Mr. Gosden is showing the performance of a small computer configuration, including the card reader and printer, against a particular model job. Even so, I imagine Mr. Gosden would counsel some careful consideration in extrapolating from his model to an actual job on the ground. But what does this matter? There is still a very great gain if we can say with some certitude: "The order—value for moneywise—of these equipments on this kind of job is A, B, C."

After all, let's face it, if one wants to *know* just how long a specific job will take on a particular equipment the only safe way is to program the job and run it on the machine. Surely there are enough grisly examples both here and in the States to make this point. *Anything* short of this, and I include here specimen sections of program based on system studies of part of a job, are bound to be estimates. The whole point of Mr. Gosden's approach—and our own—is that the time and effort involved in programming a job for several different machines just is not on. What is wanted is a quick and reliable lead to one or two machines, and a more intensive study of them.

Having recognized then that anything short of full programming is estimating, there will always be room for argument about the estimating technique used. The technique has to take account of the characteristics of the equipment and, in some fashion, of the characteristics of the work, but we do not think that these two things are connected quite so closely as some people think. The real connection is that some computers do certain things better than other computers do, and the effect of the nature of the job is to provide varying amounts of opportunity for particular features to come into play. Mr. Gosden recognizes this by postulating standard jobs of work, and we do something of the same kind by something we call a Post Office Work Unit.

The abstract does not contain enough information for me to comment on the detail of Mr. Gosden's method. Our final presentations are quite different, but it appears that the underlying work runs on very similar lines. We—because we are embarking on a large-scale extension of computer working on service-centre lines, and therefore are not buying equipment for fixed and specific jobs—have concentrated on producing cost-weighted maximum potential performance curves over the whole range of equipment configuration, for a wide range of data transfer rates and a wide variation of program length per record. We find that we can quite easily construct Mr. Gosden's uncosted performance curves from our own, and if one had a sufficient number of Mr. Gosden's curves for a particular equipment I think our uncosted curves could be drawn from his. We are not quite so happy about Mr. Gosden's method of introducing the cost factor, but we have not the detail for useful comment.

One more comment before I close, not directly related to the subject-matter of Mr. Gosden's paper but to a comment I found in the abstract. He visualizes coding his standard problems in plain-language programs and comparing the running times with those of programs written at lower levels. This is very much to be desired, and if my observation on a recent trip to the States is any guide, I would forecast that the result would be to read the funeral rights over generalized plain-language programming. The sooner the body is decently interred the sooner we can get on with our work.

Well, Mr. Chairman, this has been a difficult operation. I have necessarily had to comment on pretty fundamental lines. I hope that I have done more than merely re-present Mr. Gosden's paper, but if this is all I have done it merely indicates how closely we agree. I can assure him from my own experience in this field that there will probably be an awful lot of people who will not agree with either of us, but I am personally quite convinced that more and more people interested in buying computers will come to using a preliminary appraisal process of the basic kind described here this morning.

In conclusion, Mr. Chairman, my apologies once again for my absence, and my thanks to my colleagues, Jack Rayner and Wilfred Ryan, for putting this paper over and being available to share in the discussion. As they have done most of the work, they are probably better qualified to do this than I am anyway!

Mr. Ryan: Mr. Smith ends up with a comment about his colleagues, saying that as they have done most of the work they are probably "better qualified to do this than I am." Those who know Mr. Smith will understand that this is very much of an understatement.

Mr. J. A. Gosden: I should like to comment on three specific points while they are still fresh in my mind. First of all, Mr. Smith did not understand how I included relative prices; the answer is that I did not include them in my summary.

Secondly, I noticed, I think, that they were concentrating rather more on the central processor and appeared to be disregarding the rest of the equipment. This, if it is true, is a serious omission. We have discovered that sometimes the central processor is the one thing that we can ignore in certain configurations, whereas there is a great deal to be looked at in the linkage of the input-output system, the amount of blocking and unblocking to be done, and the available amount of simultaneous operation. The way that these are provided may, in some cases, have an extremely important bearing on the timing.

Thirdly, I should like to dissociate myself from Mr. Smith's remarks about machine-language programming. I will not make any excessive claims for COBOL, ALGOL, JOVIAL, or any other language; neither will I support much of the evidence ranged against them which is based upon a number of *ad hoc* experiments, with uncontrolled variables all over the place, and uncontrolled interpretations. It is more prudent to withhold judgment based on such evidence.

Mr. E. N. Duke (*A. E. Reed and Co. Ltd.*): What weighting, if any, is given for the provision of software?

Mr. Gosden: I must admit that some of my answers are going to be flavoured by—what shall we say—a practical commercial approach. One ground rule is that we try to give an equal coverage to all systems; therefore, we must have equivalent information. Normally, having information about somebody is to their detriment; and to publish a detailed report on one machine is to the advantage of the machine on which you do not publish a detailed report. (*Laughter.*)

Therefore, we feel that it is only fair to publish only to the level at which we can be uniform. At the moment, we assume the use of a good assembly system for most of the medium to large machines, and are considering process oriented languages. We use the normal facilities such as standard conversion routines and standard IOCS (input–output control systems).

For the small desk size machines where there is a large number of interpreters and special languages, we time the problems in both an assembly system, if it is available, and the popular interpreter system. On a mathematical problem where floating point is not built in, we publish the floating point time using subroutines and also a time for fixed point operation. We do not publish any measure on how difficult it is to program one way rather than another. Unfortunately we do not have a convenient control group of neutral people whom we can time on programming these problems. Nevertheless, the software is taken into account in a consistent way in some detail for the small machines, and at a reasonable level for the medium to large machines.

Mr. B. R. Taylor (*Ministry of Aviation*): You have obviously made a considerable effort to make your figures objective here. I wonder how objective you feel this has been? How much can your own staff influence the times quoted, and the fact that one program is better than another to give a completely different answer?

Mr. Gosden: Of course they can influence it considerably. This is an interesting question, as they say. (*Laughter.*) The answer is best given by describing the way in which we organize our work. We have a group of about seven people, all of whom work close to each other. The time to cover thoroughly a machine is about six calendar months, although the effort spent is less than that in man-months of work. An analyst begins by being biased for or against a machine, depending upon his past experience. Then he begins to be irritated because of the difficulty of getting information. That puts him in the right frame of mind. (*Laughter.*) When he begins to examine the performance, he is somewhat challenged, and he really starts to look for good ways of doing each problem. Sometimes it is necessary for us to say, "Well, you may have thought of that clever way, but we must assume that the *average standard programmer* would not." Such conversations on most of the controversial figures are normally at least three-way discussions. I do not think that we have any two people with compatible backgrounds; therefore we often get three different answers, which we resolve. Although we cannot claim that we are absolutely objective, we do claim that we are more objective than one single person, and that we are more objective than a group with a common background or limited experience. Remember that each of our staff may be studying three machines at a time.

Mr. K. D. W. Janes (*H.M. Treasury, London*): We in the Treasury are particularly interested in the problem that Mr. Gosden has been dealing with, because in the Government service we naturally have to do our best to arrive at a completely impartial decision on the equipment that we buy. We cannot decide on the basis that the Chairman of the company has had lunch with the Chairman of the other company and likes him, and the sort of work that Mr. Gosden has described is very similar to the work which we have been doing in a variety of ways. I was particularly interested and glad that Mr. Gosden stressed the importance, after having prepared these generalized curves and generalized information, of then looking at the job in considerable detail, and that the particular problems of the job to be dealt with are really the factors which decide what configuration of equipment and indeed what equipment one finally buys. The meeting may sense perhaps here a difference of opinion with the statement read from Mr. Smith, which implied that it was possible to arrive at value for money without considering the job parameters.

One of the problems which we find particularly difficult in any question of deciding what equipment configuration to buy is in discovering what size of program would have to be contained within the machine, and hence very largely the amount of memory storage space, and also the amount of processing time required, because in many of our jobs we find that we are computer limited and not tape limited, and that processing time within the central processer is particularly important. This is complicated by the fact that in the job there may have to be a wide variety of items dealt with, items of varying types; the frequency of the varying types of items and the number of instructions, the spread of the number of instructions which have to be obeyed to deal with this particular item, or items, can be important factors in determining the running time. The spread can be very wide. One case which has been analysed, a job which is actually running, showed the number of instructions obeyed ranging from 7,000 to 25,000 per item; the distribution of such items is quite important. This is something which can only be decided on the basis of experience, analysis of the job and experience of similar work. Another factor which Mr. Gosden referred to as being a difficult one to assess, and which he said, I think, depended on the judgment and intuition of the particular analyst, is the comparison of machines with varying types of address: the single-address machine, the multiple-address machine, variable-length instructions, and so forth. One system which we have tried to use here is, recognizing that this must be an exercise of judgment and approximation, is to do our best to find a short program which looks reasonably typical; then to get our program for this part of the system by each manufacturer and from the number of instructions required and the running time get some quantitative comparative factors. But this is still only to help estimation and cannot replace the ultimately only satisfactory solution of writing programs for the whole job, which is clearly impossible.

Mr. Gosden: I would like to make two remarks to show how we try to be objective and make it easy to adapt our results to other problems.

First, we actually do code out in detail all the central part of the program in skeleton form, and this is the basis of our estimates for the central processor. The difficult aspect is laying out storage and deciding on variable or fixed lengths. Often we do the layout both ways to see which way is better. We sometimes have to try the layout of the master file two ways, in binary and in decimal, to see which way is the better.

My second remark concerns both the last two speakers. We try to leave enough of our basic working data in the reports, so that anybody who disagrees with the way we have synthesized the data can resynthesize it for themselves another way. This is necessary when you want to estimate the time for one particular job of your own. You should go back to the basic data, which we present in a convenient standard way. Then you can easily make an estimate of a job with different values of the parameters. What you do not have to do is any digging for the basic data.

The Chairman: I should like to ask Mr. Gosden whether his work indicates that—taking his particular job which shows that at a low activity some machines are best and for a high activity that others are best—does this indicate that the characteristics which make a machine the best at either end are mutually exclusive? Therefore, is it impossible at a reasonable price to design a machine which will be best all along the line?

Mr. Gosden: The easy answer to this is that I do not have enough information to answer it. However, we have now looked at 15 systems out of a total of something between 50 or 80 systems available in the United States, so that we do have some idea. Unfortunately, no clear patterns are emerging. I think Fig. 7 of my paper illustrates this rather well. This covers about eight or nine small business systems. Now, there is little pattern in this. Just look at one particular system—the square boxes. They are almost arranged in a horizontal line. The circles are another system. Different patterns occur because some of the systems have been extended from the small size by adding extra items, and some have been contracted by cutting items off. Now, if you imagine that there is an envelope drawn around all the symbols, then you may deduce that there is an optimum size for this range of computer systems. It occurs at about the 9,000 dollars a month mark. If we drew a figure for expensive systems, as opposed to this one of smaller business systems, you might find an optimum somewhere else. Then we note that this figure covers only one problem and one value of the activity factor. If we take another activity factor, we should get a different pattern. Fig. 4 shows that sometimes there can be a wide divergence of performance between similar configurations.

Mr. T. P. Goldingham (*International Computers and Tabulators Ltd.*): I believe that the United Kingdom journal *Which* would claim to have had some influence on manufacturers. Could Mr. Gosden tell us what the reaction of manufacturers has been to his investigations, and whether he envisages that in future they might publish their specifications in a standard format, possibly quoting performance on a standard test program?

Mr. Gosden: The interesting question is: Will somebody design their machine so that it rates well? Perhaps we should publish a report on a hypothetical machine to see if anybody builds it. I do know at least one manufacturer who is rating a new machine against our performance measures. Another manufacturer noted an answer "none" against a feature which his software package lacked, and immediately decided to incorporate it.

Mr. D. H. Kelley (*B.I.S.R.A.*): One of the major decisions to be faced at present is whether to use large-scale random-access storage or magnetic tapes. I imagine from Mr. Gosden's remarks that he would calculate his critical parameters for both types of configuration and leave his client to judge the issue. That is fair enough. I suggest, however, that the use of random access will simplify systems design and programming. Has Mr. Gosden any comments to make on this topic please?

Mr. Gosden: At present, we are dealing with problems that are susceptible to our type of approach and we do not deal with others that are not. This does not mean that they are always to be excluded; it means that we do not think they are susceptible to this particular type of approach. You can time a system both for an on-line or a batch processing procedure; but you cannot really compare one against the other. You can compare different batch processing times against each other. You can compare different on-line processing times against each other. But a comparison of apples to pears is not straightforward. The conveniences of one approach against the other concern the organization around the machine, the conveniences of the fast response of on-line type of system, the different recovery procedures, and whether there are critical peak loads. These are application-oriented answers and I do not think we can tackle them in this way. This paper describes one particular area which is susceptible to this type of approach, and how we have tackled it.

Mr. E. M. Blair (*A. E. Reed and Co. Ltd.*): Would it be unfair to conclude that unless you have a very clear idea of what work you are going to put on the computer you might almost do just as well in choosing one by putting the names in a hat?

Mr. Gosden: No.

Mr. J. E. Sachs (*International Computers and Tabulators Ltd.*): Could Mr. Gosden say a little more about the examples he has taken, the applications which he is using in measuring the performance of computers? Has he built his own artificial examples of, say, a stock up-dating, or a payroll, or has he examined a large number of actual jobs and compared those with the assumptions that he has made, and does he think that the examples he has shown are sufficiently typical and cover a sufficiently wide range of file up-dating problems?

Mr. J. A. Gosden: The answer to the last question is: in theory, no; in practice, probably yes. The way that we did this is as follows: a group of about four people sat down and designed typical problems based on our experience, wide experience, both European and American. We are limited by having to produce our report for a budget. This has influenced us to avoid producing hundreds of different problems in different variations. Therefore we publish in considerable detail in the reports all the basic detail we have used to make our estimates, in such a way that anybody who wants to look at our estimates can see exactly what we have done. If he wants to use our figures for a different case, he must make his own adjustments. We cannot do more than this within a reasonable budget.

Mr. B. V. Piggott (*Eastern Electricity Board*): Considering the computers of a small/medium size used for commercial purposes, there are relatively few suitable for any particular. job and one is working near to the minimum size of each type. It is suggested that in such cases the theoretical methods described by the speaker are of limited value and that it is essential to assess the facilities of each machine against the job; a series of feasibility studies should be undertaken exploring the matter as far as practicable.

These studies would indicate the small number of suitable machines from which the final selection would be made, based on the cost, the manufacturers' experience, and the backing expected from the manufacturer.

Mr. Gosden: I think I agree with the conclusion but disagree with the initial premise. I do not really think that you can say there is not much variation possible from the manufacturer. Nor that there are very few to choose from. I agree that people often buy small or minimum machines, and some of our timings look extremely high because they reflect some of these limitations. We consider both a card configuration and a minimum tape configuration.

Mr. Piggott: Perhaps I might just explain my fundamental point, that the user that I have in mind would not consider the large tape system, he is on the border line as to whether he can use magnetic tape at all, in which case he would use a small tape system anyway.

Mr. Gosden: The times differ widely among the different small computers, particularly because of the number of simultaneous operations. This is often more significant than the speeds of the units. These times affect the small user considerably. The performance is not always obvious from the usual specifications.

Mr. E. M. Woodhams (*Whitworth Gloster Aircraft*): In his analyses, has Mr. Gosden considered the effect of differing support equipment and staff on the operational cost of the computers concerned?

Mr. Gosden: The answer to: "Have we done such an analysis?" is "Yes." The answer to: "Is it included in this service?" is "No." We feel that it is not possible at present to report manufacturers' support in a uniform, reliable, fair, and useful way. However, we do prepare individual particular reports for individuals.

We give the basic figures from which one can produce one's own summation of costs; but we cannot particularize this in any way for individuals; all we can give is the basic details, and individuals must add up their own totals.

The Chairman: I am sure that Mr. Gosden realizes that our appreciation of his paper is shown by the number of questions and the fact that people have stayed here and not gone for coffee. We have all enjoyed this session.

Future Papers—Notice

The following papers are among those which will be published in Volume 6, April 1963- .

Business Applications

D. J. Dace (*Commercial Union Group*)
Experience in the practical use of data transmission
J. Drummond (*M.P.N.I.*)
Some aspects of recording graduated insurance contributions
A. F. George (*S.A.S. Airlines System*)
S.A.S. aids for the jet age: data transmission for electronic reservations
J. R. Hopkinson (*U.K.A.E.A.*)
Integrated accounting using a variety of equipment
W. S. Ryan (*G.P.O.*)
LEAPS—the first three years

Mathematical and Scientific

C. W. Clenshaw and H. J. Norton (*N.P.L.*)
The solution of non-linear ordinary differential equations in Chebyshev series
A. J. T. Colin (*University of London Computer Unit*)
Coding of reverse Polish expressions for single-address computers with one accumulator
D. M. Collison (*Elliott Bros. (London) Ltd.*)
A method of forming a sorting key for a partly ordered list
David Elliott (*Basser Computing Dept., University of Sydney*)
A Chebyshev series method for the numerical solution of Fredholm integral equations
L. Fox (*Director, Oxford University Computing Laboratory*)
Partial differential equations
M. J. R. Healy (*Rothamsted Experimental Station*)
Programming multiple regression
B. Higman (*General Electric Company Ltd.*)
What EVERYBODY should know about ALGOL
J. W. Lewis (*Leo Computers Ltd.*)
Time sharing on LEO III
M. R. Mills (*Honeywell Controls Ltd.*)
Operational experience of time sharing and parallel processing
D. Mustard, J. N. Lyness, and J. M. Blatt (*University of New South Wales*)
Numerical Quadrature in n dimensions
M. R. Osborne (*Imperial College of Science and Technology*)
On iterative procedures for solving finite-difference aproximations to separable partial differential equations
R. Palmer (*U.K.A.E.A.*)
Computer calculations on the initiation of high explosive detonation
A. H. Stroud (*University of Wisconsin*) and D. Secrest (*University of Illinois*)
A multiple-precision floating-point interpretive program for the CDC 1604
L. H. Underhill (*U.K.A.E.A.*)
The growth of complexity in a general-purpose program

Conference Proceedings—Notice

The Proceedings of the Congress of the International Federation for Information Processing, held in Munich in September 1962, will be published in April 1963 by North Holland Publishing Company, Amsterdam. The Proceedings will contain full texts of the invited papers, contributed papers, symposia, and discussions. Further details will be circulated with the March 1963 issue of *The Computer Bulletin*.