Management Control and Games – by Computer

By

F. M. ALVES

Data Processing
Accounting Machine Division

THE NATIONAL CASH REGISTER CO. S.A. (PTY.) LTD.

GENERAL INTRODUCTION

HAVE you ever been asked to speculate on the office of the future? And, if so, what sort of picture would you paint? Would you, for example, conjure a word picture of an airy, circular office, focussed around a data transmitting console? A few machine operators furiously punching out information in a common business language, that sped with the speed of light to a central data processing facility? There the data was digested, analysed and scrutinised, and back-flashed the printed results in televised form to be viewed by a radiant team of executives. Rapid, crisp decisions were made, ideas sprouted from a well-ploughed field of facts and the Golden Eagle of Prosperity flew over the land.

Or, would you envisage a picture more like a nightmare? There was the Sales Executive being crushed under four tons of sales analyses, and the Production Engineer being strangled by punched paper tape. Then there was the President of the Company, whose office was piled from floor to ceiling with reports, figures, indexes, computations and permutations. With his office stuffed and his exit plugged, there he sat emaciated and buck-eyed ploughing his way through the paper work. There was data everywhere—"and not a drop to drink!"

Neither sketch is very far from reality. They both exist within to-day's management practices. Whether industry can harness the computer technology and channel the massive information potential to useful purpose or whether the executive is about to be buried under facts and figures, lies to a great degree in the hands of the specialists engaged in electronic data processing.

When J. von Neumann and O. Morgenstern published *The Theory of Games and Economic Behaviour* as long ago as 1944, business economics took a long stride forward. For electronic computers provide simplified models of some of the functions of the brain while "games" are simplified models of business economic behaviour. The inter-relation between man and machine is strengthened by the fact that machines can be made to play games and that business economic life can be thought of as solving highly complicated computational problems.

Thus, whether we are discussing management control or management games by computer, it is at once clear that some form of decision making lies at the heart of the matter. And it is for this reason that I have chosen *Decision Making* as the central theme of my lecture to-day.

DECISION MAKING is an activity which has historically been performed by people. These people, including you and me, seem rather defensive about their exclusive prerogatives to perform this activity, that is, their prerogatives to "make decisions". To suggest that some inanimate object like a digital computer will soon be "able" (or is presently "able") to make decisions previously made by people is damaging to the ego and thus constitutes a major form of heresy. Nevertheless, this suggestion has been made.

The capability of the digital computer to make decisions has often been described as one of the fundamental characteristics of the set of equipment which constitutes a digital computer system. While some of this description has been factual and some admittedly fictional, the concept of automatic decision making has become inseparable from that of computers in the minds of most.

A definition of decision making

The process by which one arrives at conclusions —makes decisions—is interesting to examine. In the first place, while these conclusions are usually said to derive directly from a set of "facts", it almost always develops that these so-called "facts" are in reality an *estimate* of the true facts of the case, made from as many clues as are available to the estimator.

The question now arises as to whether the available clues have been dealt with systematically and logically, are of sufficient scope to provide the basis for a decision, and do not distort the facts (as clues in detective stories so often do). One may think of this situation as a competition between Life and the Decision Maker, where the role of Life is to choose the underlying state of things, and the role of the Decision Maker is to select an action dependent upon the strategy of Life and at the same time most favourable to himself.

In order to perform his function in a business environment, the Decision Maker often employs relatively costly data processing techniques and equipment. By utilising these facilities, he then hopes to obtain some information about the strategy which has been selected by Life. In theory this information can be improved (up to a limit) by increased processing of more and more data toward this end.

As a balancing factor, however, there are costs associated with this processing. The problem now becomes that of balancing the cost of data processing against the errors costed in decision making (which could potentially be lowered if more processing were carried out as a result of more moneys being spent for this purpose).

Finally, the Decision Maker must decide whether even all the data theoretically available and completely processed can guarantee accurate decisions. This is not possible, for example, if the decision involves such fundamental uncertainties as whether past history will prove to be representative of future events in general and, moreover, the particular future events about which he must decide. We have come to expect the sun unfailingly to rise each morning and to set each evening and it is highly probable that it will do so tomorrow; but it is *not certain* it will! It is substantially less certain that our boss will continue to operate our organisation as stubbornly and improperly as he has in the past, although the history of the situation certainly suggests he will do so.

Decision Making, then, is the job of determining the course of action which is most likely to produce the desired results according to some predetermined criteria. When the criteria and the rules for determining likeliness are explicitly provided to a device, such as a digital computer, which has the capability to manipulate this information logically, the computer assumes the ability to process data according to these criteria and rules, and hence to "make decisions".

Elements of decision making

Decision Making can theoretically be carried out according to a systematic logical plan. It is in such cases called "scientific decision making". Scientific decision making by numerical means may involve techniques from the fields of mathematical statistics, information theory, logical nets, symbolic logic, game theory, linear programming, and dynamic programming. Decisions made by the proper application of these techniques are by definition "correct" since the fields themselves can be shown to be internally consistent and completely determinate.

For this reason, scientific decision making by the use of digital computer "simplifies" to the problem of the proper application of the techniques mentioned above. While this is not a simple problem in the sense that every man in the street is completely conversant with the application of these techniques, their application is well defined to the experts. Once the problem has been represented mathematically, its solution becomes a relatively simple problem for the trained technician.

In the past few decades there has been a considerable amount of study in the field of decision theory. A great many names have been associated with these studies. It is of interest to study the decisions with which this theory is concerned since they are logically quite similar to decisions in those business situations which are well defined.

Even though, as pointed out above, it is not the usual case that business situations are well defined, the role of the Decision Maker is to assume that the definition he is given is in fact correct. Concern about the gulf between the definition of the problem which he is given and the true nature of the problem is not properly that of the Decision Maker (at least not in his role as decision maker). By approaching the problem in this way, the Decision Maker is able to place his problem in one-to-one correspondence with that represented by this theory and so to employ techniques developed for solution to "gaming" problems in the solution of his problems of the "real world".

Most of the theoretical work which has been done in the field of decision theory has dealt with problems of a determinable nature. Thus the reasoning involved in playing nim, checkers, chess, contract bridge, and several other games has been explored. While nim and, to a lesser extent, checkers, are simple games, the other games mentioned are so complex that, even though they are determinable, i.e., a satisfactory solution can be drawn from a finite number of possibilities, this number is so large that it has so far been considered infinite.

Although the possible logical paths leading to a solution of these games are very many indeed, the rules involved are often quite simple. This is equivalent to saying that the mathematics (logic) defining the *rules* of these games is simple but that the mathematics describing a preferred *application* of the rules in order to obtain an optimal solution is very complex indeed.

A digital computer is a device with a considerable facility to manipulate (add, subtract, multiply, divide, compare, etc.) numbers and thus with some facility to apply these rules. Let us consider whether it is theoretically possible to build a digital computer which could:

1. Perform its activities within the rules of the game and identify any activities outside the legal limits of the rules.
2. Solve problems of limited scope within the rules of the game (as, for example, to indicate whether, in a given position in a chess game, white has forced mate in three or to indicate whether, for a particular deal of hands in contract bridge, a given sequence of play will or will not make the given contract).
3. Play a reasonably good game of chess (bridge), i.e., given any chess (bridge) situation which is not particularly unusual or difficult, indicate a reasonably good move (play) after several minutes of computer calculation toward that end.
4. Play chess or contract bridge, as defined above, and improve its play from game to game, profiting in each game by the experience gained in the previous games.
5. Answer questions put to it so that it is not possible for the questioner (if he cannot see the machine) to determine whether the answers given are those of a man or of the machine, i.e., given an appropriate number of wrong, partially wrong, or misleading answers.
6. Respond to any situation with conclusions based on feelings, biases, intuition, and other reactions familiar to the psychologist, just as man does.

The science-fiction writer is way ahead of this speaker; he has said for many years that all this can readily be done, although it may take a few years to implement. In most cases, he is probably closer to being correct than one might think he is.

You can readily see how (1) can be accomplished by a computer. After all, there are really very few rules to deal with. Likewise, (2) is almost as easy to accomplish since—in chess or bridge—even all the possibilities made available by this limited subset of the total game does not represent a very large number of possibilities.

To play a reasonably good game, (3) is also rather clearly possible, although somewhat more difficult. It is possible partly because we usually mean by "reasonably good" that the player (computer) must win only about half the games, against another player of average ability and, moreover, in order to do this it need not make the optimal play at every turn since it will have ample opportunity to recover from a reasonable number of errors in play and to take advantage of a similar number of errors made by its human opponent.

The question becomes slightly more sticky when we come to consider (4), which really amounts to learning, at least as far as chess or contract bridge is concerned. Although it is intuitively very difficult to accept the fact that a computer is capable of learning in a sense which is quite real, it must be remembered that the scope of learning suggested here is extremely narrow (being limited to the single subject of, say, chess) as compared with the normal scope of human learning. The digital computer is therefore certainly capable of recording in its memory events which have occurred in past games, searching through its memory to find already proved solutions to situations similar to those now confronting it, and even of erasing the record of a problem solved ("forgetting") when it has been unused for some time and the space in which the record is stored is now required for the record of some recent event. In the course of such an endeavour, the computer now begins to face limitations imposed by storage access time or storage size which tend to bind its effectiveness in the field of learning. These limits are advancing at a rapid rate along with the general state of the art of digital computers. Even today, computers are able to play the game of chess on a par with players of considerable skill, beating them a reasonable number of times and playing very rapidly indeed. This "ability", to be sure, is a clever definition of the problem by omnipotent man to the computer doing the job.

The problem of building a computer and programming it so that it will *appear* to think like a man (5) is considerably broader in scope than that of playing a proper game of chess. But it is probably logically identical. A great many facts must be available to the computer to accomplish this, and it must be able to retrieve them very rapidly and manipulate them logically according to modes of reasoning (logic) selected from a large set of logical alternatives. While such a computer is surely out of the question even at the level of today's research, it would seem overly pessimistic to suppose that it will never be possible to build one.

One question (6) then remains: Is it theoretically possible to build a computer which will respond to any situation with conclusions based on feelings, biases, and intuition, even as you and I. This is, as Bowden points out, something I can never know, just as I do not know whether *you* respond to situations with conclusions based on feelings, biases, and intuition, as I do. It seems possible to take a more optimistic point of view on this question than Bowden suggests; namely, studies currently under way in the theory of complex information processing may reveal that the fundamental nature of intellectual response in people is such as to permit us in the far distant future to duplicate these processes on a very large, very fast digital computer.

Complex information processing as a tool in decision making

Most problems which deserve to be called "complex" are characterised by the dual facts that the processing required to solve them cannot be known until some processing is completed and that a solution to each such problem is required within a limited amount of time. In addition, they appear to be characterised in varying degrees by a nearly inexhaustible amount of information relevant to their solution and by the existence of an extremely large set of potential solutions.

These characteristics of complex problems immediately suggest that a computer capable of solving them must be able to solve any set of logical equations and be able to solve them so fast that the degree of their complexity is usually of minor consequence. But this is useful only when some entity (human or machine) understands the particular problem sufficiently well to describe it in detail, i.e., write the logical equations theoretically sufficient for its solution.

Decisions are easily made when the alternatives and all their ramifications are unambiguously presented and when all the necessary data are available. The fact that different people decide different things in the same situation (and even that the same people decide different things faced with the same situation at different times) seems certainly to indicate that not all the data and possible alternative decisions and decision criteria are available to each person.

The human approach to this difficulty is to approximate the optimal solution as described above by the simple device of *continuous human cognizance* as the solution proceeds, i.e., continuous redefinition of the problem and the consequent identification of alternate possible solutions, and introduction of new data as they become available. Such an approach recognises the problem as one having different forms at different stages, and the path to the optimal solution as one found by testing and trying many different solutions at each point of decision.

Developments are in process today which will enable a digital computer to make step-by-step decisions in these heteromorphic problems by such probative means. This is done by eliminating from further consideration at each step those alternatives which, on the basis of the information available at that point, have the lowest probability of being the most favourable alternative. This is done until the number of remaining alternatives is sufficiently small to enable the computer to deal logically with them. While the exact number of alternatives which can be dealt with on a logical basis is of course a function of the size of the information storage of the computer and of its speed, the number often proves to be sufficient to make "good" decisions. Thus the computer can approach the solution of a problem exactly the same as problem solvers do; i.e., neither computers nor people always have available a complete set of data and alternative decisions at every given decision point. The difference is a matter of degree.

This point can be illustrated by an example in the field of programming a digital computer to play chess. On the basis of the rules and objectives of the game and values assigned for each piece, the computer assumes the role of a player.

First it is necessary to define the basic rules of the game in a form which can be interpreted by the computer. Such a form is the computer program, a sequence of coded instructions to the computer. In our example, this portion of the program enables the computer to solve the set of mathematical equations which reflect these rules.

Next we assign a quantity to each chess piece. Such quantities enable the computer to "know", for example, that a queen is worth more than a bishop and so should be protected more zealously. It is often desirable to assign a whole set of quantities to each piece so that their changing worth as the game progresses can be accounted for.

Finally, the objectives of the game are defined in terms of the computer program, again as a sequence of coded instructions which represent a set of mathematical equations. Thus the computer "knows" that an objective is to capture its opponent's bishop, but not at the loss of a queen, and ultimately to capture the opponent's king. This portion of the program, when executed, provides the computer with a set of values against which it can measure its situation at any point in the game, or by means of which it can project its situation into the future. It is clear that no statement of these values can be exhaustive in the game of chess because there are just too many possibilities, but the incomplete statement that can be provided has proved sufficient to do a remarkable job of playing the game.

The computer is now ready to begin play. Its moves and those of its opponent (which may, incidentally, be another computer) occur alternately, the computer "deciding" on its own moves by solving the logical equations of which its program consists and making each solution known to the operator, and the moves of its opponent being indicated, in turn, to the computer. At each turn, the computer computes a value for a group of alternative moves according to the value scale which it has been given and selects the move with the highest value. This sequence of events proceeds until the game has been completed.

At the same time that it is deciding on its moves for the present game, the computer is recording in its memory the effects which it "experiences" from the moves it makes, again according to the same scale of values. Thus it is accumulating experience. As the amount of experience accumulated increases, as each move must be made the program directs the computer to search its memory for situations which it has already experienced which are similar to the present one, rather than to compute a new value for several of the possible moves available to it. In this way, the number of alternative moves available to the computer increases substantially as it plays more games, much in the same way as people gain experience by repeated play.

Eventually, however, either the computer storage is filled or the amount of time required to search the storage to review all the experience stored there becomes excessively great. At this time, instead of simply ceasing to collect experience, the computer program can direct the computer to discard the experience which it uses most infrequently whenever it can substitute more pertinent experience in its stead. This corresponds to forgetting in human beings and, because the computer "forgets" according to plan, may even prove to be a more organised form of forgetting than that of which people are capable.

It is noteworthy that a computer has actually been programmed to play checkers in an analogous manner and has demonstrated a rather good game. This, then, is truly an example of complex information processing by means of a digital computer. Over a group of games the processing is heteromorphic and the method employed probative. The computer wins and loses, much the same as a man would do (although it is not likely to lose twice in exactly the same way). It seems already to have achieved a "human" solution to a problem in complex information processing.

When digital computers were first introduced commercially, after World War II, many of the designers were very vocal in proclaiming vast capabilities for the machines with respect to imitation of human thought (thereby suggesting the name "brain" for the machines) by the machine as a whole and even by its individual components. But before the middle of the next decade the pendulum had swung the other way and the more vocal of the so-called "experts" seemed to be those who were denouncing the pre-eminence of the computer as a thinker and relegating it to the accomplishment of only routine tasks. When, in 1949, E. C. Berkeley published *Giant Brains*, most computer experts of the day cringed at the choice of words.

Now we can detect a growing tendency to regard digital computers as devices capable of carrying out deductive thought processes. Considerable evidence exists to support this new attitude. This approach to the problem departs from the idea that computers are fast and dumb devices, and considers that they, in combination with their programmes, may be more "logical" than people and, in addition, that they can go to places (such as the moon, planets, or radioactive areas on earth) where people cannot go and hence offer a means of locating logic in such places.

It remains a philosophical or semantic question whether problem solving of this kind performed on computers actually constitutes "thinking" by the machine. Proponents of this point of view contend that, since a knowledge of the program which was designed by human beings only theoretically enables these human beings to predict the process by which the problem will be solved by the computer, it is the machine and not the human which is capable of executing the theory and hence it might be said to be "thinking". The evidence offers some support for this view.

The human element in decision making

The manager in this newly automated computer-infested environment must face still another major problem. He must agree in advance on the decisions he would make in a great variety of circumstances as a function of the decisions made by his counter-part in every affected department in the company. It develops that this sort of advance commitments of himself to a decision in a situation wherein he will not be given the opportunity to review his decision in the light of the decisions of other departments of his firm is very difficult for a manager to make. This is true even though he is informed of their decisions at the time they are made and whenever they are changed. Our habits are such that it is disturbing not to have the opportunity to reconsider each decision personally at the time it is made, even though we do not do so and, in fact, did not really have the *capability* to do so in our earlier unmechanised way of doing things.

This strong reluctance to change may indirectly provide us with the explanation for an unexpected development which has taken place as new computer systems are installed. These new systems were billed initially as great savers of personnel in terms of the jobs they were "taking over". But as the machines were installed, it developed that a significant percentage of each of these jobs involve the performance of eminently human functions such as letter writing, telephone answering, and so forth (at least until the entire system was redesigned to eliminate the need for some of these functions). Nevertheless, although the computers failed to "take over" the jobs represented by the functions, they performed as advertised. Postinstallation experience, in those cases where expansions in work load would have required doubling or tripling of the work force before computers were installed, did not actually involve any personnel increase and represented a negligible increase in the work load of the computer.

A further resistance to change is exemplified by the middle manager who wants to make his decisions based on the same reports in the computer system as he has always received before the computer was installed. He expresses his requirements by grandly stating that he believes this expensive computer should provide him with "at least" the same information he had in the precomputer era, and at least as rapidly. Such an attitude, inadequately taking account of its resultant effect, gain or loss, on the new system for data processing which the computer has made possible, represents one of the major factors acting to delay the wide-spread achievement of benefits from the use of digital computer systems in business organisations.

The responsibility of top management, whose interests and authority extend over all the functions and objectives of the firm, to select a suitable set of constraints within which the system redesign must take place and then to support the indicated revision of procedures within these constraints must be exercised if we are ultimately to realise the potential of digital computers.

It seems clear that confidence—derived from considerable amounts of successful experience—is the only way in which the practical businessman, can, will, and indeed should be convinced that digital computers can be used for decision making in business. With this experience will come an attitude of trust in one's decision-making counterpart in another department in the firm so that operation without continuous human surveillance will become possible. With this experience too will come confidence in automatic decision making.

The broadening scope of mathematical applications

The recent advent of the digital computer has made it practical to extend the application of relatively complex mathematical techniques into fields in which they were previously unknown. These methods have made it possible to introduce systematic, rational, and completely unemotional problem-solving techniques into areas where they seemed to have been completely lacking. The scope of application has been very wide and we mention only a few of them here for illustrative purposes.

One area of business in which decisions have seemingly been based on judgement, intuition, and background of knowledge gained only from experience has been that of stock-market investments. Although these aspects of human decisions have not been entirely eliminated, a digital computer has now been applied to a major portion of the problem. By means of the mathematical techniques of "linear programming", a digital computer has been programmed to compare all the possible investment opportunities of which it is informed according to criteria perviously established. On the basis of these criteria—which include the objective of the investment program as well as facts about the stocks such as price, past earnings, dividend record, and the estimated future performance—the computer can then produce its recommendation as to the most suitable portfolio for the purpose indicated. Because of its very high speed and capacity for storing facts, this computer is able to select the "best" portfolio from among a group of possibilities which is very large indeed.

Medical research is another field which is beginning to get some rather important direct benefits from the digital computer. Equipment and techniques now under development will, when perfected, assist medical diagnosis by correlating data on disease occurrences and computing their probabilities accordingly. It is further expected that, by these methods, the analysis and correlation of symptoms by computer techniques will permit the detection of trends in a particular case much earlier than presently. Such a system might perform its analysis on data entered directly from the patient himself by means of devices which would read instruments attached to the patient so as to reflect his condition. Another possibility, of course, would be to record the desired data in a more or less usual manual fashion and then transcribe it into machine language and enter it into the machine system. Probably the main advantage of any of these systems is that diagnosis of the available information is by well-defined systematic means and is far more likely to be accurate and sufficiently broad of computational scope than diagnosis by present techniques.

By means of the capability to compute by digital computers, the techniques of mathematical probability statistics have been employed on an ever-widening scale to problems of our everyday lives. One such application, for example, deals with the problem of purchasing an automobile for private use. According to the alternatives given for a particular case which was presented to the computer—the costs of the car specified, the operating expenses and their rate of increase, the decreasing probability that the present car would last another three months, and the trade-in deals available—analysis revealed that the most favourable deal from the buyer's point of view, financially, is to buy an automobile three years old and trade it in on another three-year-old car when the first one is $6\frac{1}{2}$ years old.

Another problem of similar structure deals with a question of interest to baseball players and fans alike: For each situation which may be encountered, what strategy of base-stealing, bunting, and "swinging away" should be employed? The computer solution to the problem was alarmingly simple: Always swing away. The analysis carried out by the computer of the statistics involved indicated that, no matter the skills of the particular players involved as regards base stealing or hitting, it is "better" to try for a hit than to resort to any bunting or stealing strategy. Having defined "best" to cover games won over an entire season rather than any individual game, the computer analysis found further that each game should be played for itself without regard to considerations such as fatigue and even strategy which seemed to pertain to future games.

CONCLUSION

Machines can't think for us, but they do force us to think and re-think in solving our business problems. Perhaps one of the tragedies of computer use today is the automation of current confusion. It just isn't sufficient to restate the old rules. New rules must be designed to take advantage of the new technology available. Computerised decision making is knocking at management's door—the big challenge, however, is to harness this massive information potential to useful purposes. The wisdom for this also must come from management.