[*Proceedings of the Institute of Electrical Engineers*, vol.103, part B: Supplement no.3 (1956)]

SESSION ON THE COMPUTER IN A NON-ARITHMETIC ROLE

INTRODUCTION

By A. D. BOOTH, D.Sc., Ph.D., F.Inst.P.

(Address delivered 12th April, 1956.)

I propose to mention briefly the types of activity which we can carry out on a computing machine and the operations that are suitable for a machine whose primary purpose is numerical.

One could treat this in a historical way, and it is perhaps not a bad thing to do. One of the first things which was not obviously a numerical operation and which was suggested as a computing machine application was the playing of games. There is a long history about games-playing machines in connections quite different from what we would expect with electronic digital computers. However, right at the beginning in, say, 1945 or 1946, various people suggested that it would be possible to play chess. So far as I know, no satisfactory programme has ever been constructed for playing a good game of chess. 1 remember one which always lost by 'fool's mate'. On the other hand, I believe there are better ones nowadays which can beat a bad human opponent by fool's mate, and which are almost as simple as the early programmes.

Later in this session we are going to be told of some work done on playing draughts. Draughts is a non-trivial game in the sense that we cannot immediately predict what the result is going to be. Some of you at least will have seen demonstrations of computing machines playing nim or noughts and crosses. Neither of these games is really a fair operation for a machine against a human being. A machine does not make mistakes, or it should not, and you can be assured it will take the best possible course in any given situation.

I shall not have time to say anything in detail about the way in which machines play games, but merely to state that games playing depends either upon working out a strategy which is based on a numerical discrimination or alternatively playing the game by means of some function which is maximized. The second approach is numerical representation of the game itself, and the first is the carrying out of a set of rules.

The second application, which I believe was made on a machine, was that of playing a simple tune. I have heard a number of these tunes at various times. They seem to vary in character with the country of origin. I remember hearing, oddly enough, 'The Bluebells of Scotland' played in America, and also lighter performances such as 'Frankie and Johnnie'.

Of course, those of you who know anything about the internal structure of a computer will realize that it is not at all difficult to make a pattern of digits which we can shift round in the register of a computing machine with a frequency which will generate tones, either pure or with a reasonable harmonic content.

In this way we can build up the different frequencies required to play a piece of music. The duration of the notes is simply obtained by setting up a counting index which arranges that the data are circulated in the register a number of times which in aggregate are equal to the required note duration. There is nothing very mysterious about this.

A third branch of machine activity which I think is by far the most interesting is that of character recognition; i.e. the use of machines to recognize shape in its more general context. We are going to hear later about efforts on machine character recognition. I have one small contribution to make, on some work we have been doing. But before I do so in detail, I want to mention the recognition of sounds.

Dr. Booth is the Director of Birkbeck College Computational Laboratory, University of London.

Actually the system we propose, which we have tried out on a limited scale, will work for either sounds or printed characters, because it represents the characters to be recognized in numerical form.

Let me try to make this clear with a simple example. Suppose we take a sound wave. I can draw this as a voltage/time picture (Fig. 1). You will see that we have the means of generating a

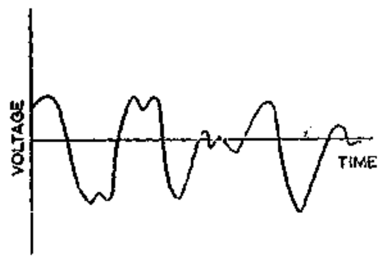

Fig. 1

number directly from this waveform by counting the number of times the waveform crosses the time axis. We sample the incoming waveform over discrete intervals of time. In our particular case they vary between l/15sec, which is convenient for our machine, up to times of 1 sec, which can also be generated. Shorter times do not seem to improve the characteristics of the system.

It turns out that merely taking a linear representation of the incoming sound considered as a voltage waveform does not give adequate recognition. What we do, therefore, is to take three channels of input, the first being the unmodified waveform, as 1 have drawn it; the second the first derivative of that waveform; and the third the integral of the waveform. From these three waveforms we can calculate three numbers in each sample interval. In some statistical experiments we have done, these three numbers have seemed to prove adequate to recognize with an accuracy of better than 1 part in 1 000 the digits 0 to 9.

That is as far as this particular work has gone at the moment and, indeed, probably as far as we are interested in carrying it. I do not see any reason why the technique should not be extended to recognize more complicated sounds, although when it comes to recognizing a word one has the difficulty that the principles of sampling are more complicated. If a word is considered as a relatively small set of samples, a number of words will give the same sampling numbers and will therefore be indistinguishable. We have not done a great amount of work on that, so I cannot say more about it.

With regard to the recognition of shape, particularly the recognition of printed characters, the technique used is roughly the same as that for sounds except for the addition of one feature which is unfortunately not available for sounds unless they are recorded. It is a feature I should like to call 'digital feedback,' and I will explain it in a moment.

First of all, let me describe the basic technique. Take a letter, say A, and apply to it a scanning raster by means of a cathode-ray projection tube. The first intersection point of raster and character is recognized by a suitable photocell, and the output from a standard oscillator is gated into a counter, so that on each scanning line numbers proportional to the positions of the intersections are generated. This set of numbers may be called a_{ti}, where i ranges from 1 to n, n is the number of lines in the raster and t is the number of the intersection for the given scan line. For the moment, we are trying to recognize a set of characters that I will call B_i and it is assumed that we have calculated the sampling numbers a_{tij} for these lines. The machine calculates the value of:

$$\delta_j = \sum_{ti} |a_{ti} - a_{tij}|$$

and it is asserted that δ_j is a minimum when the character to be recognized corresponds to the standard character B_i.

It may be correctly objected that the recognition may not be unique. There are certain letter pairs, the most elementary of which is perhaps O and Q, which cannot be searched out by a pattern of that sort at all readily. To get over that we propose to use our principle of digital feedback. If there is more than one of the quantities δ_i near the minimum considered adequate for recognition, then the machine is told to apply a second scanning pattern to the character being recognized. In this case we shall get another set of quantities which can be compared with another set of standard quantities. This second scanning appears to remove most of the confusion between one letter and number and another.

We have not completed this work yet, but the idea that the computer can perform one particular sort of operation and as a result of it, and only if necessary, perform another operation which might clarify a doubt is a powerful one and can be applied in many other situations.

The fourth category about which I want to say a few words is machine translation. At Birkbeck College we pioneered the application of machines to translation, and recently we have extended our machine activities to cover in considerable detail the translation of French into English. Not only do we produce a translation in the word-for-word sense, which was the original object of machine translation experiments and is not difficult; but we also take account of the construction of sentences and the re-ordering of words, and so on.

The essential point about the computing machine applied to machine translation is that it has too small a store to make possible the storage of an adequate dictionary. I think this is still the case on almost all the machines in this country. I have not had an opportunity to read the paper which is to be presented later, so I cannot speak of the Russian machine.

In our particular machine, as constituted, we can operate with a dictionary of 250 stems and about the same number of endings. The essential point is that in taking a foreign-language word one breaks it up into the stem and the ending and proceeds to recognize the stem simply by comparing it with numbers held in the store. The input data, of course, are on the standard teletype record, which represents each of the alphabetical characters by a numerical symbol.

When we have detached and identified the stem several things are possible. We may output the stem translation or suspend any operation until more data are available. This is necessary in interpreting such things as homonyms and prefixes in expressions like 'to have' in English, which may not occur in the French word.

Having dealt with this part, we proceed to the ending, and at this point we either output an English stem with an appropriate ending or an English word accompanied by some preformed word. Alternatively, it may be necessary to suspend all operations until it has been possible to decide whether more information is required.

You will realize that this is not a practical proposition at the present time. The speed of translation on our own machine is of the order of 1 000 words an hour with an average word length of five letters. I believe a good human translator can do about 3 000 words an hour. On the other hand, there are conflicting reports on this matter; and one authority whom I consulted told me that, although it was true he could do 3 000 words an hour, at the end of that hour he had to rest for the remainder of the day. I am not sufficiently adept at translation to do more than 100 an hour from a simple language, and at this rate the machine wins hands down. The point is that if these figures bear any resemblance to the truth, the machine is not likely to be economic at the moment, because it costs between £10 and £60 an hour and human labour costs 7s. to £2, so you can see there is no great profit to be made in doing translations by machine.

I want to finish with two other branches of the art. One of them is automatic control in factories, which is a strategic planning operation of an essentially mathematical type. The other is automatic control of machine tools, which does not at first look like a numerical operation but can be turned into one rather simply. It has a very extensive literature at the present time, so I need not go into detail.

Here again the principle of feedback is used. The machine tool is set to do something and a signal from it tells the calculating part of the machine to what extent the operation has been satisfactorily performed. The whole cycle of operations proceeds in this way, governed by feedback.

Finally, some of you may be impelled to ask: Since we have gone so far with machines, what about replacing the human being altogether? This is a very difficult question. It is unfortunate for our self-esteem that as soon as an attempt is made to analyse any particular phase of human activity in detail one sees fairly clearly how it could be mechanized—I except at the moment the mysterious process known as 'original creative thought'.

All other activities seem capable of rationalization and thus of execution on calculating machines. To this extent it is not unreasonable to hope that most of the more or less routine operations, let us say, of management—maybe of higher administration—would be done considerably better by machines than by human beings. We can hope that eventually the machine will replace the human being. Undoubtedly it will give just results according to its own programme but perhaps not very popular ones in some cases. Personally, I would prefer the justice to the popularity.

What about original creative thought? It is very difficult to specify what one means by this. I always used to think that Sir Isaac Newton, when he propounded his theories of gravitation and of the calculus, was an absolutely original thinker. In fact, it has been pointed out by many mathematical historians that if he had not discovered these things a number of other people would have been bound to do so. This being the case, it may perhaps appear that all of the data on all of the world's problems at a given time may be supplied to a computer which would produce all of the possible answers based on strict reason.

I shall not attempt to suggest how this could be done. But what we mean by an original element in thought may be a purely random effect and not something which can be deduced from the existing corpus of knowledge. I do not know, but perhaps the discussion will throw some light on the point.