method of establishing coefficients of very low value without using inconveniently high values of resistance is shown in Fig. 2. If R_1 and R_3 are 10 megohms or more and R_2 is taken to the signal earth of the amplifier with which it is associated, quite good results can be obtained for a relatively wide range of coefficients.

Conclusion

In conclusion, it should again be mentioned that no true reactor data have been used in the examples which are only included as illustrations of the methods frequently used. Where a large and complex simulator is concerned it is impossible in a short paper to outline more than a few aspects of its design. It is hoped that those matters discussed will be of some general interest.

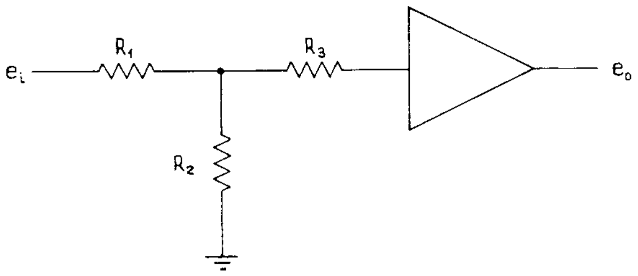

Fig. 2.—Method of avoiding high resistor values

Acknowledgement

In a project such as the training simulator for C.E.G.B. there is need for considerable liaison between customer and contractor, and the author wishes to record the close and friendly co-operation which existed at all stages of this machine and to express his thanks to the Central Electricity Generating Board for permission to use the simulator to illustrate this paper.

An Analysis of a Hydro-electric System

By P. F. King and D. A. Peel

A mathematical model of one of the hydro-electric schemes operated by The British Aluminium Co. Ltd. in the Scottish Highlands has been programmed for an Elliott computer. This has facilitated the study of various ways of operating the system and also enabled the evaluation of the gain in overall power production which might be obtained if some possible modifications to the system were effected. This paper was presented at the Harrogate Conference of The British Computer Society on 5 July 1960.

Introduction

This is an illustration of a class of problem whose solution was virtually impossible prior to the introduction of automatic computers.

The British Aluminium Company operates several hydro-electric schemes which produce power used in the smelting and refining of aluminium. This particular system, the Lochaber Scheme, is situated in the Scottish Highlands. Due to the highly seasonal rate of rainfall, a fair proportion of the water which enters the catchment area during the winter months is spilt over the tops of the dams, while to add insult to injury, it is quite normal to have to reduce load in summer due to water shortage. Several proposals have been made for altering the system so as to increase water usage, and the problem is to evaluate the benefits which would result from each of these modifications.

After a fair amount of preliminary study, it became clear that it would only be possible to discriminate between the various modified systems by simulating them in considerable detail over an extended period. Day-by-day reservoir levels and water consumptions were available for a period of fourteen years, and some trial simulations carried out by hand indicated that it was necessary to work in time increments of one day to achieve the desired accuracy. This obviously precludes the use of desk calculators.

The Lochaber Scheme

Before dealing with the actual computing involved it is necessary to give a thumbnail sketch of the system itself. The catchment area is of rather more than 300 square miles and includes much of the Ben Nevis range. The rainfall drains into a series of reservoirs, but for the purpose of this study, only two need be considered. These are connected by a tunnel through which water flows, not unnaturally, from the upper one to the lower one, at a rate dependent upon the water-levels at each end. The functional dependence is quite interestingly complex, since there are one or two discontinuities caused by the fact that there is a sluice gate at one end which has to be closed whenever it appears likely that the rush of water would be so great as to wash away the Fort William railway line! Water from the lower reservoir travels through a 15-mile-long tunnel to the powerhouse near Fort William. There it passes through the turbines which drive the generators. These produce the electricity used to convert alumina into aluminium. The rate at which power is generated depends upon the rate at which water is used, and also the level in the lower reservoir. The more water there is in the lower reservoir the less of it has to be used to produce a given amount of electricity.

Ninety-five per cent of the loss of water which occurs is spilt from the upper reservoir. This is caused by the relatively small volume of storage within this reservoir. After a rainstorm the water-level rises rapidly (as much as 50% of the total range in a day) and, although the flow in the connecting tunnel also increases, it is frequently insufficient to prevent the water-level topping the dam. In the absence of further rain, however, the level falls rapidly. The more promising modifications that have been suggested are those which tend to increase the effective working storage in the upper reservoir. Other modifications which have been investigated are aimed at increasing the rate of flow in the connecting tunnel, either by means of structural changes or by the installation of pumps.

Choice of Computer

It was not clear at first whether this type of problem was more easily solved by analogue or digital means. The decision was eventually made in favour of the digital approach for a number of reasons, the chief ones being the greater flexibility and accuracy of a digital computer, together with the difficulty, in this case, of preparing the input data for an analogue machine. So far, all the work has been done on the Elliott 402E Service Bureau computer, but the programs are now being re-written for an Elliott 802. This is being done because the smaller machine (having only a 1,024-word core store) is not only adequate for the problem but, by virtue of its more powerful instruction code, will do the job in much the same time, so reducing the cost per run. In addition, the simpler instruction code of the 802, and the fact that the storage is all of the immediate-access variety, means that program alterations may be incorporated without upsetting the timing.

Computer program

The first piece of work was to prepare a paper tape giving the *run-offs*, that is the volume of water entering each reservoir for each day of the 14-year period. This involved reading in the observed water-levels and daily water-consumption figures, converting the levels to volumes of water in storage, and then calculating the run-offs by forming the difference between the volume at the beginning and end of each day, allowing for water-consumption at the powerhouse, flow in the connecting tunnel, and spill over the dams, if any.

The next stage was to use these run-offs as input data for the simulation proper. The basic simulation program was very like the program used for calculating the run-offs, but running in reverse as it were. That is, the run-offs were read in and the predicted water-levels printed out. It was first of all necessary to check that these water-levels matched those which had occurred in the past, assuming the same water usage, and this stage was completed satisfactorily. It was originally intended to drive the simulation by feeding in the observed power figures and calculating the volumes of water required to produce them, given the current water-level in the lower reservoir. But this led to grave trouble. When using water as input the model is quite stable, but the introduction of the power-conversion subroutine destroys this stability. If at some stage the level in the lower reservoir is incorrect then the quantity of water used to generate a given power will also be incorrect. This in turn causes a further error in the volume of water remaining in the reservoir, and this error is in the same sense as the original, thus aggravating the situation. As a consequence a lot of time and trouble was spent in trying to reduce day-to-day errors within the model. Even so, when the program was tested the results were almost unbelievably bad. The trouble eventually turned out to be due to the power conversion equation. This had been formed by fitting a polynomial to a scatter diagram which showed the relationship between water consumption and the ratio between power production and water-level. The assumption underlying such a procedure is that the errors about the curve are random and not autocorrelated. In this case, however, this assumption proved to be incorrect: the errors about the curve tended to be of the same sign for long periods at a time. The most likely reason for this is gradual fluctuations in the overall efficiency of the powerhouse. The effect was to produce errors large enough for the instability just mentioned to dominate the proceedings and produce absurd results. This was overcome in production runs simply by turning the power conversion around, that is by scheduling water usage and deriving the power production which resulted.

Results so far obtained

The final stage was to develop a subroutine to simulate the present operating procedure. That is, a set of rules to determine at each stage throughout the simulation the appropriate water usage according to the current situation. A start was made by determining just how much power could have been derived from the existing system had it been known in advance what rainfall to expect. This was done by means of a short series of runs in which the total amount of spill which occurred was reduced gradually by a trial and error process. In carrying this out quite a lot was learned about the fundamental effect of the bottleneck caused by the connecting tunnel, and this has helped in the design of the required set of rules. The runs using foreknowledge were used to determine the seasonal pattern of water input to the lower reservoir, when operated in a near-optimal manner. It was at the same time possible to determine the best level at which to aim to keep the lower reservoir. The problem here is that one wants simultaneously to keep the water-level both high and low. One wants to keep it high so as to get the maximum power from a given quantity of water, and low so as to increase the flow in the connecting tunnel and thus reduce spill. Obviously, to maintain this desired level the normal rate of water usage should equal the derived rate of input. The scheduling rule eventually used was thus to determine the current volume of water in storage in the lower reservoir, to compare it with the desired volume, and to modify the normal rate of usage for the time of the year by an amount depending upon the difference between these two volumes.

The effect of varying the desired volume throughout the year, instead of keeping it fixed at the optimal level, was also tried, with the purpose of providing additional temporary storage in the autumn months, that is, prior to the spill season. Another short series of runs was sufficient to select the best values of the parameters involved in this simple form of strategy. The best of these runs was quite pleasing since the power produced was only 1% lower than the best which had been achieved assuming foreknowledge of run-off, while the degree of load-changing involved was not excessive. Moreover, the power was several per cent higher than had been achieved in the past. This last result is misleading since, in the early years of the period considered, there were periods of major furnace-room reconstruction and tunnel cleaning during which power output had to be drastically curtailed quite independently of the water position. The best of the strategy runs, on the other hand, represents the power which could have been attained had these factors not been operative. In fact, the last three or four years of this run correspond closely to actual practice. It was decided, therefore, to use this run as a yardstick against which to measure other runs representing the modified systems.

In all the runs which are controlled by an automatic built-in strategy, a quantity has been accumulated which gives a rough measure of the amount of load-changing involved. Quite a large proportion of the cost of producing aluminium is due to the cost of rebuilding furnaces. Each furnace must be run at a steady power throughout its life, and so overall load-changes are made by changing the number of furnaces in operation. If a furnace has to be temporarily frozen-up, either due to water shortage or for some other reason, then its life is adversely affected. In addition, a furnace which has just been started or restarted does not reach its maximum efficiency for a period of several weeks. Both these factors cause a significant increase in operating costs. A fundamental conflict which has, therefore, to be resolved is between the desirability of increasing power from the hydro-system, and the danger of increasing costs due to greater load-changing. Unfortunately the increased operating costs are very difficult to assess accurately, but nevertheless the data which have been derived from these runs give a truer perspective of the matter.

This then is the present stage of the problem. Most of the questions which were originally put have been answered, but in the meantime the problem itself has expanded. Several fresh aspects have been brought to light, and it is these that will be attacked with the new 802 program.

Annual Prizes

As announced in *The Computer Journal*, Volume 2, page 150. The Editorial Board have considered the papers which were published between April 1959 and April 1960 in this *Journal* and *The Computer Bulletin*.

Awards of twenty guineas each were made in respect of two papers, and these were presented to the authors at the Annual General Meeting of The British Computer Society in London on 29 September 1960.

The winning papers were:—

Mr. T. C. Hickman (Unilever Ltd.)
"Early Experiences with an E.P.D. System"
published in *The Computer Journal*, Vol. 2, p. 152.

Mr. R. Neate and Mr. W. J. Dacey (Steel Company of Wales Ltd.)
"A Simulation of a Melting Shop Operation"
published in *The Computer Journal*, Vol. 2, p. 59.

A further competition will be held on the papers published between June 1960 and April 1961.