A Simulation of Melting Shop Operations

by R. Neate and W. J. Dacey

Summary: This paper describes a model set up to study a typical industrial situation—a melting shop in an integrated steelworks. The operations of the shop are simulated with the aid of a Pegasus computer. A description is given of how the model operates and what information can be derived from it, and comment is made upon its flexibility and upon the desirability of such methods for solving complex operational problems.

INTRODUCTION

Operational research problems almost invariably require the setting up of some kind of model to study the operations of a system. Although occasionally the model is so trivial that it would perhaps be pedantic to insist on using the term, more often than not the description is quite justified. The subject of this paper is a model designed to study the operations of a melting shop, where the process involved is that of making steel in open hearth furnaces. The initial object was to determine the increase in production rate likely to be achieved by the use of oxygen to enrich the furnace combustion air, but this object has since broadened somewhat so that we now have a model capable of investigating many other features of the system. To understand the model, a description of the plant and process is necessary.

PROCESS AND PLANT

Steel is made in open hearth furnaces by a batch process from scrap and blast furnace iron. If, as in this case, the blast furnaces and melting shops are located on neighbouring sites, the iron can be transferred in the molten state, thereby saving the considerable amount of fuel and time required to re-melt it from cold. This molten iron is termed *hot metal.* Very briefly, the procedure is to charge the furnace first with scrap and then with hot metal, heat the charge until it is molten and determine its chemical composition; the molten charge is then refined to give the specified composition, and when the correct temperature has been reached, it is tapped from the furnace and cast into moulds. After a period of standing, the steel solidifies in these moulds to form ingots, and locomotives are on call to remove trains of ingots from the shop.

A plan of the melting shop is shown in Fig. 1. There are three separate bays, two of which are at ground level, while the centre one, which contains the six furnaces, is elevated about 20 ft. Each bay has a number of cranes operating from overhead rails. Within each bay, they are identical, and run on the same set of rails. Scrap comes to the shop in railway wagons, which are housed in the scrap bay. It is unloaded from the wagons by three cranes, which are equipped with electromagnets, into boxes on the edge of the furnace bay. Three charger cranes working in this bay charge the contents of these boxes to the furnaces. The addition of hot metal and the tapping of the furnaces require the use of cranes in the casting bay. Hot metal is kept in special ladles at the north end of the bay, and is taken by crane from there to the appropriate furnace. When a furnace is tapped, the crane has to carry the ladle containing molten steel from the furnace to one of four platforms, at which a train of empty moulds is berthed. The steel is run out of the ladle into these moulds, the operation being called *teeming.*

After the furnace has been tapped, the hearth is repaired by an operation known as *fettling.* When fettling is complete, the furnace cycle starts again with the charging of scrap for the next cast.

The weight of steel cast from each furnace is 100 tons, and to achieve this weight, about 55 tons of hot metal and 55 tons of scrap are charged. Most of the 10 tons lost in the process goes into a slag formed by adding limestone. The batch time, or cycle time, is about eleven hours on average, but it varies from about 8 to 16 hours, and is occasionally as high as 20 hours.

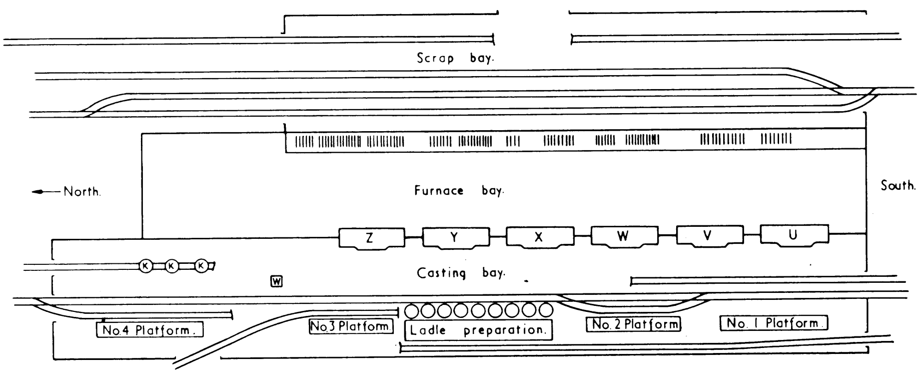

FIG. 1.—Plan of Melting Shop.

The main cycle time components, as recorded in routine shop records, are charging, melting, refining, and fettling. Typical times are:

Charging	4 hours
Melting	$4\frac{1}{2}$ hours
Refining	2 hours
Fettling	$\frac{1}{2}$ hour
	11 hours

but each one is subject to much variation. A more detailed breakdown would be as in Table 1.

TABLE 1

COMPONENTS OF FURNACE CYCLE

	Approximate average duration		*Type*
	Hours	*Min.*	
CHARGING			
Charging solids		45	*i*
"Assimilating" the solids charged so far		45	*n*
Charging remainder of solids		45	*i*
Assimilating	1	30	*n*
Charging hot metal		15	*i*
MELTING	4	30	*n*
REFINING	2	0	*n*
FETTLING			
Tapping		10	*i*
Fettling		20	*n*

Types: *i* = interactive
n = non-interactive

MONTE CARLO METHOD

During some of the component times, the furnaces can interact with each other. For instance, during the solids charging periods the furnaces require the services of charger cranes, and these in turn require the services of scrap cranes. In the hot metal addition and tapping periods, the demands are made on cranes in the casting bay. As there are more furnaces than there are cranes of any one type, these demands sometimes conflict with each other and furnace delays ensue. The melting shop obviously has queueing problems, but while each one is itself fairly simple, the queueing situation in the shop as a whole is quite complex. Coupled with the fact that the frequency distributions of furnace component times are not easily expressible in mathematical form, this meant that we were unable to use mathematical queueing theory for this problem.

The method adopted was that of simulation, using the Monte Carlo random sampling technique. This involves simulating the shop operations for a period of time, and the normal method is to write down component times in tabular form, or represent them as distances on a chart. Some operational times have so little variation that they can be assumed to have constant values, but otherwise they can be derived by the Monte Carlo method. A more detailed description is given elsewhere (Neate and Dacey, 1958). Consider, for instance, the refining time of an open hearth furnace. Its frequency distribution may be as in Fig. 2, which shows a plot of 100 values. Let us assume firstly that tests have shown no relationship between refining time and any other time or factor occurring in the system, nor between the refining times themselves. It can then be said that the frequency of occurrence of any refining time value is governed only by the shape of the distribution. The relative frequency, with which the time values occurred in practice, is the only factor governing the likelihood of their recurrence in similar practice. The most likely time in this case is $1\frac{3}{4}$ hours (or more correctly between $1\frac{5}{8}$ and $1\frac{7}{8}$ hours), which occurs on 20% of occasions, while only 1% of the times are as high as 5 hours. By using random numbers in conjunction with this distribution, the times required in the simulation can be derived as and when required.

Occasionally, the time in question proves to be related in some way to another time or another factor in the system. For instance, there is a strong inverse relationship here between charging time and melting time. It has been assumed to be linear, within the charging time range 2–$6\frac{1}{2}$ hours, and the least squares method has given the equation:

Melting time $= -(0{\cdot}4966) \times$ Charging time $+ 88{\cdot}51$, where the times are in units of 5 minutes.

Now, when a melting time is required in the simulation, the charging time is already known. The corresponding melting time is found from the equation, but this cannot be used directly since the effect of variation in the melting times would then be lost. To overcome this a value is sampled from a distribution of deviations from the mean melting time and this value (which may be +ve or −ve) is added to the melting time value.

USE OF A COMPUTER

We have often used Monte Carlo methods to deal with problems which are not amenable to other mathematical treatments, but the employment of a computer to carry out the simulation constituted quite a step forward. It provided our first computer experience of any kind. The job had been in progress for a number of weeks, and we had reluctantly decided that a simulation was necessary; reluctantly, because large-scale simulations were rather unpopular with us at the time. One such job, just completed by manual methods, had occupied two experienced operators for six months, and even then there was some doubt whether the period of simulation had been long enough. During discussions with a computer manufacturer on data processing matters, we learned, incidentally, that computers had been employed on Monte Carlo problems, and this gave considerable stimulus to our interest in these machines. This job, however, was likely to be big, by our normal standards, and we were fully aware that if we took a bite we might spend a long time chewing.

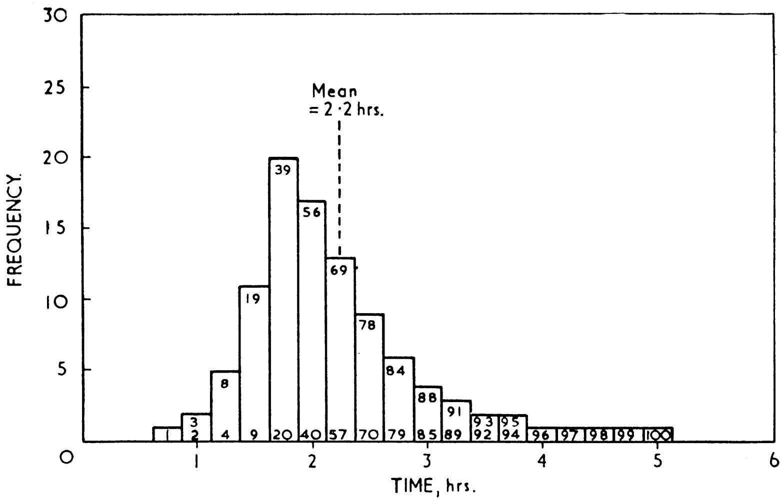

FIG. 2.—Frequency distribution of refining times.

In spite of this there were several factors which favoured going ahead. The job, which concerned our oldest steelmaking unit, had fairly low priority, because the installation of plant which would make oxygen available on a large scale was not likely for nearly 2½ years. This was important, because other simulation jobs which arose a few months later, and which would have been more suitable for the first attempt with a computer, were extremely urgent ones, involving high capital expenditure.

Our Company has been expanding continuously since its formation about 12 years ago, and as a result, many of our simulation problems have involved the prediction of how plant will perform under new conditions. We have undertaken many comparisons and alternative schemes, and management has proved quite adept at conceiving new alternatives. Very often there have been "comebacks" when, after a job has been completed and reported upon, the management has requested the examination of new conditions.

Simulation methods have the strong disadvantage of requiring a complete "run" to examine one set of conditions, in contrast to mathematical methods, such as queueing theory techniques, where once an expression has been derived, the examination of new conditions involves only a re-evaluation with new parameters. The idea of having a computer program, which could be quickly amended to incorporate the new conditions and then run off in a few hours on the machine, was extremely attractive. This is perhaps the most important reason for using a computer for simulations, but like so many jobs in other fields, it is doubtful if this job in its present degree of detail would be practicable by manual methods. It would employ at least three people working together, and they would do well to simulate one melting-shop-hour per hour. As they would only work for about 30 hours per week, while the melting shop works for 168 hours, the effective ratio of simulated time to actual working time would be about 1 : 5½. The Pegasus Computer achieves a ratio of 350 : 1, which is nearly 2,000 times as fast.

When we decided to go ahead and employ a computer, we were in the rather peculiar position of having neither programmers nor a machine, nor yet any definite prospects of having a machine. B.I.S.R.A. were, however, due to take delivery of a Pegasus in about eight months' time, and as we, being a contributing steel company, would have access to it at reduced rates, the choice of machine was fairly straightforward. Two people were selected and sent for training as programmers, but as one of them left us in the early stages of programming, the brunt of the work fell upon the other one, who has been the only programmer employed on the job.

ANSWER REQUIRED

Before discussing the program, let us look at the overall requirements of this job. We wanted a realistic model, which would enable us to simulate the operations of the whole shop, in order to determine its productive capacity under different conditions. The model had to be flexible, so that such features as the number of furnaces, cranes, teeming platforms and scrap boxes, the types and quantities of fuels, the ratio of hot metal to scrap and the proportions of different qualities of scrap could be readily changed. The question of operating policy was also important. We wanted a program able to cater for different methods of furnace charging and different policies for dealing with the queueing situations, i.e. different queue disciplines. The main result to be printed was obviously the production rate in terms of the number of casts made in a fixed period, but, in order to find where production time was lost, we also required details of the various types of furnace delays caused by inadequacy of the ancillary equipment.

SIMULATION METHOD

Initially, the operations of one furnace cycle were written down in flow diagram form. The cycle time was resolved into component periods, during each of which the furnace was said to be in a particular "state."

Changes of state were signified by "events." The intention was to compute at each event the total time before the occurrence of the next event for that furnace. This system is quite satisfactory for the computation of periods when the furnaces are non-interactive; e.g. at the event signifying the beginning of a refining period, a value of refining time—derived by the random sampling method—can be confidently added to the cumulative furnace time, because, as the furnace is then in a non-interactive state, nothing can happen to alter this derived value.

The situation is not so simple, however, during an interactive period, such as charging. The duration of this period is very much dependent on conditions existing in the shop. If a number of furnaces are in the charging state, some of them may spend time waiting for the services of charger cranes. Also, the number of scrap boxes at the disposal of each furnace may be reduced. At the event signifying the start of the charging period, the conditions which will exist throughout the period are not known. In the case of simulations carried out manually, where the operator often has the complete situation displayed before him in graphical form, this "event-to-event" method is quite conventional and effective; but, when using a computer, it could be very wasteful of machine time. The computer, unable to view the operations of six furnaces at once, would have to scan in some regular or possibly random manner and determine after any event, not the time of the next event for the same furnace, but the time of the next event for the whole shop. To do this, it would first have to determine the next event for each furnace, and then find which of these was the earliest occurring one. This would be the only valid event, and all the others (or at least those occurring in interactive periods) would have to be discarded. Each furnace line would then have to be computed again up to the time of the valid event. Then the whole process would start again.

We therefore considered the alternative method of working the simulation to a fixed time base, so that all activities are reviewed at the end of each time base interval, and all necessary changes are carried out at that time. In view of our complete lack of computer experience, we sought the advice of B.I.S.R.A.'s Computer Applications Group. They strongly recommended the fixed time base method, which was therefore adopted. The time base unit (or scanning interval) chosen was 5·8 min, which is the average time taken to charge to the furnace a group of four scrap boxes. All times, with the exception of those involved in loading scrap from the scrap bay into these boxes, are expressed in multiples of 5·8 min.

THE PROGRAM

In the program, one computer register is assigned to each of the furnaces, box groups, teeming platforms and casting bay cranes. The shop houses 96 scrap boxes on stands along the edge of the furnace bay, and the arrangement of these stands is such that the boxes can be conveniently considered in groups of four. There are thus 24 groups. (In fact, there are 26 groups, but two of them are used for limestone and other additives which do not contribute to the required charge weight.) Twenty-four registers are therefore required to represent the box groups, and these are in 3 blocks of the Pegasus main store. The first four registers in another block contain the data relating to the four teeming platforms, and the next three registers in this block are used for data on the three casting-bay cranes. The furnace registers are in the computing store, where they occupy six positions of block 5, viz. 5·1 to 5·6. Location 5·0 is used to record time, and its contents are increased by 1 at the end of each scanning interval. Location 5·7 is used to store the number of charger cranes working at any time. As the program does not, at the moment, cater for cases of crane breakdown, there is always a scrap crane to work with a charger crane. The scrap cranes need not therefore be brought into the picture.

TABLE 2

CONTENTS OF PEGASUS REGISTERS

Sign digit	Modifier	Counter
1	13 bits	25 bits

FURNACE REGISTERS

	Furnace state	Time or weight

BOX GROUP REGISTERS

State	Furnace modifier Weight	Loading time

CASTING BAY CRANE REGISTERS

(*a*) NORTH AND SOUTH CRANES

		Block position number

(*b*) CENTRE CRANE

	Block position number	Block position number

PLATFORM REGISTER

State	Teeming or Standing	Furnace modifier or time

The main program consists of scanning the furnace, box group, and teeming platform registers once per interval, in that order. To reduce any bias, the furnace registers are scanned alternately in reverse order. The scan of a furnace register determines the furnace state, and the program is then directed into one of a number of subroutines associated with the various states. The number (between 0 and 14), which signifies the furnace state, is in the modifier position of the register as seen in Table 2. The counter position contains a time or a weight. When the duration of a period, such as the melting period, has been determined, the time value is put into this counter position and is reduced by 1 during each interval until finally it is zero. The state then changes. Since all such times are integral multiples of the interval time, the number can never fall below zero. During the solids charging periods, the required solids charge-weight is put into this position, and as each box group is charged, the weight of the group is subtracted. The state changes as soon as the remainder is $\leqslant 0$. By this method, the charge-weight will be on average half a box-group-weight too high, and to correct for this the required solids charge-weight is reduced to the nominal 55 tons minus this excess amount.

During the time when a box group is allocated to a furnace, the first 3 digits of the modifier position contain the appropriate furnace modifier; the remaining 10 digits show the weight of scrap in the group. The loading time, in the counter position, is the time required to fill the group with scrap. When a box group is to be loaded, the type of scrap is found by sampling at random from a frequency distribution of the various scrap types. The scrap cranes are assumed to load this into the boxes at a constant rate, to give a box-group-weight which is constant for each type of scrap. Thus, for each scrap type, there is a corresponding time taken to load a box group, and this is the one case of a time not being an integral multiple of the interval time. When an empty group is to be loaded, this derived loading time is written into the counter position, and an interval time is subtracted from it. For most types of scrap, the loading time is between 1 and 2 intervals, and so some loading time will remain, meaning that the group is not yet full. During the next interval, the subtraction will leave a negative number. The group state is then changed to "full," the group weight is written into the appropriate position in the register, and the negative number is written into the counter position of the next box group register, where it becomes a sort of credit in hand. When a furnace starts charging, it receives an allocation of box groups, and the procedure in the program is to work through these groups in turn. Each of the three scrap cranes deals with one set of allocated groups during an interval. When there are less than three such sets, meaning that there are less than three furnaces charging, the surplus crane capacity is used to deal with any non-allocated groups which may be empty.

In the teeming platform register, the sign digit shows whether or not the platform is reserved for a furnace. A mould train arrives in the shop some hours before it is next required for a cast, and it has to undergo preparation during the waiting period. We have assumed that it arrives at the point where the refining period of the furnace begins. The selection of a teeming platform for the mould train is made by using a preference list stored in the machine. As soon as a platform has been so chosen, its state is changed. The appropriate furnace modifier is written into the register's counter position. This remains there until the cast is tapped, at which time it is replaced by the teeming time. Also, at this point, a 1 is written into the modifier position to signify the teeming operation. The teeming time is reduced each interval until zero, and then a time known as the "standing time" replaces it. This is signified by writing 2 into the modifier position. Again, the standing time is reduced each interval by the interval time until it is zero. The mould train can then leave the shop. The standing time is due to a metallurgical requirement that the steel be allowed a minimum time to solidify before it is moved. It depends only on the grade of steel. After this treatment of the teeming platforms there only remains the job of adding one to the cumulative time register, and the whole procedure begins again.

The bulk of the program, of course, is in the various subroutines initiated by the examination of furnace registers. These straightforward routines need not be described here, but two elements which are the basis of Monte Carlo simulations, viz. the generation of random numbers and the treatment of queueing situations, merit some discussion.

Random Number Generation

There are many instances where the random sampling procedure is used. Fettling, assimilating, and refining times, and melting-time deviates, occurring in the furnace routines, are all derived in this way. For every box group loaded, a scrap type is sampled, and the time for which the moulds stand at the teeming platform after teeming is also obtained by sampling.

Numbers which are genuinely random could be produced only by some kind of special generator, such as, for instance, the "ERNIE" machine. So many are required in the simulation, that it is not practicable to store random numbers taken from published tables, and the only solution is to use pseudo-random numbers which are derived in the machine when required. The fact that the numbers are derived by a mathematical routine means that they obviously cannot be random, but they prove quite satisfactory. One method (Mayer, Ed., 1956) which has been used is that of selecting a four-figure number, squaring it, and extracting the centre four digits; squaring again and extracting the centre four digits and so on. These four-figure numbers prove suitable for use as random numbers. The method fails as soon as any four-figure number repeats, because the process then cycles. A suitable choice of initial number may, however, give a cycle large enough to afford an ample supply of numbers for the particular job being undertaken. Another method, due to Lehmer, uses a congruence:

$$X_{n+1} = kX_n \, (\text{mod. } M)$$

by choosing $k = 23$ and $M = 10^8 + 1$ this method gives 8-decimal digit numbers with a period of nearly 6 million. Messrs. Courtauld Ltd., in a simulation of textile machine operations, have employed this congruential method.* We have used a B.I.S.R.A. subroutine, which works as follows. An eleven decimal digit number is squared in the double length accumulator 6 and 7. The contents of 6 and 7 are logically added (i.e. without carry) and the required pseudo-random number is selected from the resulting number. This resultant is then squared, and so on.

In some Monte Carlo models, it can be advantageous to use pseudo-random numbers. Suppose that a particular model requires only one random number per cycle, or, otherwise, always demands the numbers in a fixed order. Then, if the model is run under a number of different conditions, and the same sequence of numbers is used for each condition, the variance of the result, and consequently the running time of the simulation, can be reduced. In this melting shop model, the pattern of demand for random numbers is quite unpredictable. If the model, starting from the same initial state, is run under different sets of conditions, the use to which sequential random numbers are put will be quite different for each condition. The n'th random number may be used to derive a fettling time in one case, a type of scrap in another case, and so on. Hence, while this model cannot make use of the feature of using the same random number sequence to reduce the variance of the result, it is itself such a good randomizer that there is little chance of it being invalidated by the lack of true randomness.

There is the additional point that truly random numbers may add considerably to the difficulties of program checking and debugging, since any run would not be repeatable unless special provision were made.

QUEUEING SITUATIONS

This problem exhibits several occurrences which can be formally described as queueing situations. They occur whenever there is a demand for some service, e.g. cranes, which is in limited supply, and this conflict of demands is what is meant by interaction between furnaces. The order in which a queue is served, or the "queue discipline" as it is called, is the point of most interest to the programmer. Normally, the melting shop works on a first-come-first-served basis. Now, consider what happens in this type of simulation, which employs the fixed time-base system, with the furnaces treated sequentially. The last operation in the program cycle is to increase the content of the cumulative time register by one. At any point midway through the furnace routine, the furnaces which have already been dealt with are one interval ahead of those still to be treated. Only at the beginning or end of the routine is the time situation the same for all furnaces. Thus, when a furnace, other than the first or the last, happens to require any form of service, we should really wait until the routine is complete, before dealing with it. While no such service may be available when the routine deals with the furnace, perhaps another furnace coming later in the sequence will relinquish this service and make it available.

The logical procedure, then, is to place any furnace requiring service in an appropriate queue, and then deal with these queues when the furnace routine is complete. This form of treatment, however, requires a lot of program, and has only been used for the case of casting-bay cranes. Furnaces do, of course, have to wait for the services of charger cranes. At the end of the fettling period, when the subtraction of an interval time has finally reduced the counter content of the furnace register to zero, the furnace state is changed to "ready to charge." The next step is to try to find an available charger crane, and this is done by subtracting 3 from the contents of register 5·7, which contains the number of chargers working. If the result is zero, meaning that all three are already working, the furnace suffers a delay and remains in this "ready to charge state" for at least one interval. The interval time is added into a register containing charger delays, and the program proceeds to the next furnace. Suppose that during the next cycle, this furnace still can't acquire a charger and it suffers another delay. If the next furnace happens to relinquish a charger, and the next after that requires a charger, the first furnace will be cheated of its turn and receive unfair treatment. However, in the simulation, this should affect neither the total shop production nor the total delays, while it saves a lot of program. The queue discipline, which approximates to random service, is quite unrealistic in itself, but it doesn't affect the final result. If, when the test is made on register 5·7, the result is <0, at least one crane is available, and in the case of charger cranes, it does not matter which one, since they can be readily switched from one furnace to another. The furnace state is changed to "charging" and the contents of register 5·7 are increased by one, since another charger is now working. There is then a jump to a subroutine dealing with the allocation of box groups. When this is complete, and the required total solids charge-weight has been written into the counter position of the furnace register, the program proceeds to the next furnace.

This treatment of the queueing situation takes place instantaneously at the event signifying the end of the fettling period. The current interval time has already been used up in completing the fettling period, and no actual charging can take place until the next interval.

The queueing situation in the casting bay is more complicated, and is treated on a more realistic basis by forming queues which are served when the furnace routine is complete. Unlike the situation in the charging bay, where all the jobs are identical, viz. feeding scrap to the furnaces, the jobs in the casting bay occur at different points in the furnace cycle. The cranes here have two jobs, viz. tapping (and teeming) and the addition of hot metal. The first, by a rigid melting shop

* See p. 90.

rule, always takes precedence over the second. Hence, there are really two queues; a tapping queue which is served first, and a hot-metal queue. For each queue, the discipline is first-come-first-served. Unlike the charger cranes, the casting-bay cranes cannot be instantaneously switched from one job to another. Both their jobs involve the handling of molten metal, and once started, they must be completed. Hence we have a feature which we call "blocking," i.e. a crane working on a job which cannot be interrupted. Let us term the three cranes, north, centre and south. Suppose that hot metal is being added to the *Y* furnace by the north crane, which is the only one working; and suppose that the *X* furnace becomes ready to tap and that the appropriate mould train has been berthed at platform No. 3, a quite feasible situation. The *X* furnace now suffers a delay, not because there is no available crane—the south and centre cranes are available—but because the north crane, working at the *Y* furnace, creates a "block" there.

When the furnace program reaches the end of the refining period, the furnace state is changed to "ready to tap," and the furnace is placed in the highest available position in a tapping queue, which utilizes six registers. When the time comes to deal with this queue, the first test is to determine if any cranes are available by subtracting three from the contents of a register showing the number of cranes working. If all three are working, all the furnaces in the queue suffer a "delay to tapping." If one or more cranes are available, the next step is to consider the first furnace in the queue, and find which teeming platform has been reserved for it. The number in the queue register, i.e. the furnace modifier, is subtracted from the number in the counter position of each of the platform registers having a sign digit of 0 (indicating that they have been reserved for furnaces). As soon as the result of this subtraction is zero, the correct platform has been located.

In the operation of tapping and teeming, the crane spends twice as long at the platform as it does at the furnace, and so the platform position is the factor that decides which crane should be used. Each platform has its own preference list for cranes, which is stored with the data. The first crane in the preference list is tested for availability, this being indicated by the absence of block position numbers in the crane's register. If none of the cranes in the list is available, there is a delay to tapping. Assuming that one or more are available, the next test is to ensure that no blocks exist between furnace and platform. This is done by assigning numbers to relevant positions in the casting bay, viz. the furnace centres, the ends of each platform and the hot metal storage position marked K on Fig. 1. The numbers indicate the order of these positions from south to north. When a casting-bay crane is working and thereby constituting a block, the relevant position numbers are put into its register. The centre crane register has two numbers to signify the northern and southern extremities of a block. This is because the block extends over a distance rather than being located at a single point. For instance, when the furnace is tapping, the distance from furnace centre to teeming platform is blocked. During hot-metal addition the block extends from furnace centre to hot-metal position K, and during teeming, from one end of the platform to the other. The north and south crane registers need only contain respectively the southern and northern extremities of any block, since they are the end cranes. Hence, when a crane is selected for the tapping and teeming operation, the procedure is to determine, by a series of subtractions, if any more northerly crane which is working will come south, and if any more southerly crane will come north, of the positions which indicate the extremities of the contemplated range of working. If blocks exist, there is a delay to tapping. Otherwise, the state of the furnace is changed to tapping, and a time of two intervals, which is the duration of the tapping period, is written into the counter position of the furnace register. Tapping is one of the times which have little variation and are, therefore, assumed constant. The position numbers of the furnace and of the remote end of the teeming platform are written into the register of the appropriate casting-bay crane, since a block now exists between these positions. At the end of the time of two intervals, the furnace state is changed to fettling, while the block is changed to extend over only the teeming platform and remains there for the constant teeming time of five intervals. The crane register contents are changed accordingly. The program for dealing with furnaces queueing for cranes to add hot metal is, of course, the same.

GENERAL POINTS

We now consider a few details of more general interest. The complete program occupies 150 blocks of Pegasus, and another 102 blocks are used to hold the data. There is virtually no mathematical calculation in the accepted sense. The print-out, which is undertaken once per equivalent melting-shop week, i.e. once every 29 minutes in machine time, is accomplished in one line. More detailed print-outs have been made in order to assess the length of run required to achieve a given accuracy of result, to help in detecting program errors, and to ensure that the furnace component times were reasonable. They may, however, prove useful in "selling" computer simulations to management. Steel-makers are very practical people, who may not be content with a brief print-out saying that in a particular week so many casts were made and so many hours of furnace time were lost by delays. A detailed print-out may help to combat scepticism.

Several program faults were rather difficult to locate, but the debugging stage proved somewhat less difficult than anticipated. The greatest difficulty, of course, has been due to not having a computer on the spot—travelling 200 miles to use the machine is a considerable handicap. The time spent on programming is not easy to estimate, since the programmer was involved on other matters for some of the period, but from the start of programming to the first successful run on the machine was about 10 months. The program development stage occupied four of these months.

So far, one planned run has been carried out to determine the difference in shop output with 4, 5 and 6 furnaces working, 2 and 3 casting-bay cranes in operation, and with pitch creosote and fuel oil as alternative furnace fuels. This involved a plan with twelve treatments, and the simulation was run for one melting shop week per treatment. The whole plan was replicated once.

CONCLUDING REMARKS

Finally, let us consider the desirability of computer simulation methods as a means of solving operational research problems. This job, while taking considerable time, has provided us with a great deal of experience which is directly applicable to most simulation problems. The program is flexible and is enabling us to study the melting shop's operations under quite a variety of conditions. Unfortunately, this particular shop is not quite the typical melting shop, since the conventional practice is to add hot metal to the furnace from the furnace bay and not from the casting bay. Hence, to transfer the program to one of our other melting shops would involve something of a major change. Other people, viz. B.I.S.R.A. and United Steel Companies, are working on general simulation methods. B.I.S.R.A. are compiling subroutines which can be used on a building-block principle, while United Steel are developing a generalized program which can be readily adapted to any industrial situation. It remains to be seen which of these two methods will prove the more useful, but it is clear that both will lead to considerable reduction in the time and effort required to set up a model for simulation work.

One of the great difficulties, in this type of work, is trying to formalize operational procedures which tend to defy formalization. Each open hearth furnace has a "first-hand melter" who is responsible for the operation of his furnace, and each melter may have his own ideas of how to make steel in the shortest possible time. To work a three-shift system, there are four teams of operators per furnace, and so there can be quite a variety of operational practice in the shop. While great efforts are being made to standardize practice, in an industry which is somewhat steeped in tradition it will take time. This means that any formal rules, deduced from studies of furnace data, give only a rough representation of how the shop operates, and it accounts for some of the large variation found in operational times. The object of these studies, however, is usually to compare two situations rather than to make an absolute forecast about any one of them.

To conclude on a rather more optimistic note; we have many problems which have been appreciated for some time, but which have always been considered too big to be tackled, even though their solution would probably lead to very large financial savings for the Company. This project has clearly indicated that they are now practicable, if a computer is employed.

ACKNOWLEDGEMENTS

The authors wish to thank the Directors of The Steel Company of Wales Ltd., for permission to publish this paper. Figs. 1 and 2 are reproduced from the authors' earlier paper by permission of B.I.S.R.A.

REFERENCES

NEATE, R., and DACEY, W. J. (1958). "A Simulation of Melting Shop Operations by Means of a Computer," *Conference on Computers in the Iron and Steel Industry: 15th April* 1958, The British Iron and Steel Research Association (*Proceedings* obtainable from BISRA, 11 Park Lane, London, W.1, at £1 per single copy, or 15s. per copy for orders of six or more copies).

MEYER, H. E. (Ed.) (1956). *Symposium on Monte Carlo Methods* (March 1954, at University of Florida, sponsored by U.S. Air Research and Development Command). New York, Wiley & Sons; London, Chapman & Hall Ltd.

SUMMARY OF DISCUSSION

The following points were made during the discussion which followed the presentation of the above paper to The British Computer Society in London on 17th February 1959.

Dr. F. Yates (*Rothamsted Experimental Station*): In proposing the vote of thanks, Dr. Yates (Chairman) stated that he considered that the simulation method provided a very important tool for tackling problems of this kind, when they were not amenable to mathematical treatment. The formal mathematical approach often led to difficulties. In the first place, in order that a mathematical approach should be possible, mathematicians tended to over-simplify the specification; in his experience such over-simplification frequently invalidated the solution. Secondly, even after simplification, the mathematics often proved intractable. It was, of course, important when using simulation methods to pay great attention to the detailed specification. The authors had clearly done this in the present problem.

Dr. Yates considered that there was no real danger in using pseudo-random numbers rather than truly random numbers. Indeed, many processes of generating so-called random numbers did, in fact, lead to sequences of numbers which were less random than those provided by pseudo-random processes.

Mr. D. V. Blake (*National Physical Laboratory*): It has been mentioned that, in a simulation of this type, the performance of the model is very sensitive to the shape of the input distribution function. One hopes that this means that the actual system is also as sensitive. This sets an upper limit on the accuracy of results obtainable, particularly where the distribution cannot be measured with sufficient accuracy or where it changes with time. I should like to ask the authors whether they have been troubled by this in their very interesting steel-melting simulation.

It would be very useful if the simulation could be extended to help in detailed planning. This might be possible when simulating a port; here the most important input variable is probably the time of arrival of ships. This may happen sufficiently infrequently to make it worth while doing simulations to aid planning, using the actual input data, i.e. the position at a given point in time, instead of working with artificial positions generated by pseudo-random numbers from expected distribution curves.

Dr. A. S. Douglas (*University of Leeds*): The final result after a process of simulation may not always be the best answer. Interference is difficult to record mathematically, but the simulation presumably selects the path of least interference: were many trials necessary to get the best answer to given conditions of interference?

Mr. J. A. Gosden (*Leo Computers Ltd.*): It is probably not realistic to try to work out the best possible procedure for each set of situations that may arise, for this would involve having a computer on-line. I would suggest, however, that a model of the form described tonight would be suitable for testing out various simple sets of rules that foremen could use, and to improve shop efficiency by suggesting these rules to foremen.

Concerning the sensitivity of models to the distributions used, I think it is true that this effect is noticed mainly in "arrival distributions" and I would not have expected this particular model to be sensitive to small changes.

Mr. L. A. J. Verra (*National Physical Laboratory*): A difficulty appeared to arise in establishing the initial conditions at a point of time from which the simulation would run; would it be convenient to run the simulation for the equivalent of a few days, and then to use the conditions then obtaining as the starting point?

The authors in reply:—

Mr. Blake: In this model we have used the actual distributions of times occurring in practice in the shop, which are stored in the machine in their original form and not as mathematical functions approximating to them. There is no reason to believe, from previous work, that this system is particularly sensitive to distribution shape, although no actual tests have yet been made.

Distributions of the type used here, e.g. furnace component times, seem always to retain the same shape, but the mean values do change with time and this is somewhat perplexing from the point of view of simulations. In this shop, tonnage output has increased gradually over the years, partly because of improved raw material quality and partly because of improvements in operational practice. The best solution seems to be to use data from the most recent periods.

Regarding the second point, the state of affairs at any time in the simulation has no particular significance since it is not related to any point in real time, and so the model normally could not assist on-line planning. The idea is interesting, however, because it is conceivable that in making an on-line decision, e.g. deciding which of several wharf unloaders to lay-off for repairs, one may be faced with a probability calculation which is beyond normal mathematical methods, but which could be solved by the Monte Carlo technique.

Dr. Douglas: The simulation does not necessarily select the path of least interference. These interference conditions are treated by applying sets of rules similar to the rules used in the melting shop. For instance, given that a particular furnace is ready to tap, and that the moulds are berthed at a particular platform, the selection of a crane is made according to a predetermined preference list.

The number of trials necessary to arrive at an answer of given accuracy is dependent on the variability of the system. This stems from two sources, viz. the variability of component times as seen from the frequency distributions, and the additional variation imposed by interference. The former is by far the biggest contributor to the variation of the result.

Mr. Gosden: No attempt has yet been made to vary the sets of rules used to deal with any situation. In order to make the best decision in a particular case, it would be necessary to forecast events, e.g. when deciding which crane to use to teem a cast; it is really necessary to estimate if other furnaces are likely to tap within the following 40 minutes, since a wrong decision may result in the establishment of blocks preventing such furnaces from tapping. This forecasting is not incorporated in the present model and it would involve some lengthy programming. The relatively small amount of time lost in tapping delays (as found from the results) seems to justify this omission.

It is agreed, however, that this type of model could be used to test the relative merits of various decision procedures.

Regarding the point on "arrival distributions," various workers have encountered trouble, mainly on simulations of port operations. The models have proved very sensitive to the shape of the distribution of ship arrivals and the difficulty has been due to using approximations to the actual distributions.

In this melting shop model there is no arrival distribution in the normal sense, although it can be said that the furnaces "arrive" at, say, the point of "ready to charge," and so there is an implied distribution of "arrival per unit time" and of "time between arrivals." These are, of course, products of the simulation rather than inputs to it.

It is hoped to carry out tests with various forms of distribution, probably extreme forms (e.g. rectangular, exponential, single valued), and to see to what extent output and delays are affected.

Mr. Verra: There has been no difficulty with initial conditions. Previous work on melting shops has shown that any patterns due to furnace bunching (i.e. a number of furnaces in the same state at the same time) do not persist for more than about 48 hours. The initial conditions were therefore chosen for convenience, since there is no trace of the initial pattern when the result is printed out at the end of the simulated week.