General Information Manual

Improved Job Shop Management Through Data Processing

IBM

IBM®

General Information Manual

Improved Job Shop Management Through Data Processing

Contents

Introduction

Some of the most complex problems confronting modern business lie in the area of manufacturing control. At the same time it has become increasingly apparent that enormous savings may be gained by effective solutions to these problems. Efforts to effectively deal with them in the past have been hampered by two related difficulties. One of these is the great volume of information which must be taken into consideration. The other is the comparatively small amount of time which is available before decision is necessary.

The advent of the powerful data processing system in recent years has changed this picture considerably. Because of its large storage capacity and electronic speed it is capable of both processing voluminous information and arriving at results which are timely in relation to manufacturing processes.

This manual contains a description of the way in which an IBM 704 Data Processing System can be used to deal with difficult problems in the area of job shop manufacturing. It will be concerned with discussing the logic of the approach which has been taken by the 704 program rather than the technical details of the program. In this sense it is application-oriented. Along this line it will attempt, first of all, to clarify the nature of the problems involved in designing an effective job shop. Once attention has been brought to a focus on the problems, the manner of their solution will be discussed in considerable detail.

The specific program was developed by International Business Machines Corporation in cooperation with the General Electric Company. Its scope actually extends beyond the job shop context into the area of scheduling problems generally. The specific use to which it is put depends on how its variables are interpreted. In exploring its use in connection with job shop manufacturing only one application of the program is being considered. In this application the program design enables companies with job shop type operations to experiment with these operations on a computer simulation basis.

For the majority of practical situations an IBM 704 with at least 8,192 words of storage and the following system components is necessary: 5 magnetic tape units, 1 card reader, 1 printer (with standard SHARE control panels).

Readers of this manual who wish to avail themselves of the 704 program may obtain it by writing to: Applications Library, International Business Machines Corporation, 590 Madison Avenue, New York 22, New York. In addition to the program cards, a program listing, block diagram, assembly information, operating instructions, and information on how the program may be modified in certain respects, will be sent. The information and materials which are available have an immediate bearing on the 704 program as such and are necessary for its use.

Improved Job Shop Management Through Data Processing

PART I

The Nature of the Problem

The Goals of Management

A basic necessity for any manufacturing organization is to have sufficient machines and manpower. If sufficient machines and manpower are not available, orders will not be delivered to its customers on time. The responsibility for satisfying this goal lies with management.

If meeting customer delivery obligations were the only goal, it could be achieved very easily. The problem could be solved by having a very great number of machines and unlimited manpower. If this were done every customer order could be quickly completed. No order would ever have to wait until a machine or an operator became available. Since there would be no waiting, the amount of time required to complete an order would be no greater than the amount needed for actual processing. This would be ideal from a customer's standpoint.

Unfortunately, from the standpoint of management in a manufacturing organization, this solution is completely unrealistic. The cost of having so much manpower and so many machine tools would make the price of the product so high as to be noncompetitive. There is no point to having a lot of machines if you have no customers.

It can be seen from this that the problem of deciding how many machines to have is not simple. It is difficult because management must take many goals into account. Thus, the additional goal of successfully competing with other manufacturers affects the means which may be used to achieve the goal of meeting customer delivery obligations. The existence of both goals at the same time places management in the position of having to work with comparatively limited machine and manpower resources. The effort to achieve both goals dictates that these resources be used as efficiently as possible. As will be seen later, this is in itself no easy problem.

The Inventory Problem

Let us consider still another way of dealing with the customer delivery problem. Why not maintain a large inventory of finished products? If this were done, customer orders could be filled from stock. Customer deliveries could be immediate.

In many cases, however, the finished products in question are not known in advance. They must be manufactured according to specifications which are not known until the order is placed.

But even when this is not so, there are still other problems. It is necessary to consider the cost of maintaining an inventory. Capital tied up in inventory is not directly producing income. Instead, it is subject to such additional charges as insurance and storage. It also involves risks, such as the possibilities of obsolescence, spoilage and changing market conditions. These costs must be balanced against the advantages of meeting customer obligations quickly. The most efficient use of available capital funds is certainly another of management's goals.

Quite possibly management might decide to maintain a partial inventory. Such an inventory might consist of parts and assemblies of finished products for which there is repeated demand. If this were done, it would take less time to fill a customer order for the finished product.

The fact of the matter is that many means might be used to adequately meet customer delivery obligations. This does not, however, eliminate the goal of using manpower and machines as efficiently as possible.

The Machines and Manpower Problem

The efficient use of machines and manpower is actually not very difficult where manufacturing requires only one type of machine. In this situation experience will have shown that a certain quantity of this type of machine and operator is required. Predicted demand for the product involved will also make possible the provision of additional tools and operators for the future.

Actual manufacturing, however, is not usually as simple as this. Consider the situation where the product changes, and where its manufacture involves more than a one-tool, one-operator process. Suppose it involves many different types of tools and operators.

Job shop manufacturing tends to be of this type. Physically, the most common job shop consists of a collection of machine tools — lathes, drill presses, milling machines, and so on. Ordinarily these machines are of the general-purpose type, capable of being used in the production of a variety of finished products.

Job shop machines are usually arranged in physical groups on the basis of the general type of work they perform. The work lots handled by these machines tend to be variable in several respects — quantity, dollar value, and type. Such variations, along with the general-purpose character of the machines, make for a number of possible paths along which jobs can move. Different types of orders may use the same machines at different points along their work paths. This kind of manufacturing contrasts with what has been called the continuous type. In continuous manufacturing, work routes are fairly fixed. The machines are more specialized and are used in the production of comparatively fixed end products.

Let us take a closer look now at some of the problems facing management in the specific situation of a job shop. First of all, the general problem of machines and manpower also exists here. It is especially difficult, in the complex situation of an actual job shop, to decide how many of each type of machine and operation the shop should have. It was suggested at the outset that having too many machines would be too costly. The specific problem of how many is too many is related to the question of how efficiently available machines can be used.

An illustration of a practical job shop situation may help to clarify the complexities involved here. The various types of machines making up a job shop are functionally related to each other. This is true because each machine is a working unit in the paths along which orders may flow. Some of the machine operations necessary for an order are required earlier in time than others. In addition some of the operations take much longer than others. A short operation may follow a long operation but it cannot be begun until the longer one has been completed.

In this situation, as a given set of orders moves through a shop, waiting lines may develop and continue to grow behind certain machine groups. This would happen if orders are arriving at these machines faster than they can be processed by them. At the same time machines may be idle at other points of the shop. Under these conditions the machine and manpower efficiency level of the shop is suffering. The practical problem is to improve efficiency.

In an actual shop it may be at just this point that the possibility of installing an additional machine suggests itself. Perhaps this will solve the bottleneck problem and thereby increase efficiency. Before deciding, however, it is important to know just how much efficiency will increase. But this would be very difficult to estimate. Because of the complexity of the situation even intelligent estimates might easily be wide of the mark.

Management's decision here is complicated by the fact that other considerations are involved. It is also a management goal to keep the cost of machine tools down. If this particular goal were the whole basis for decision, no machine would be added.

The point is, however, that the problem does not lend itself to solution in terms of any one principle. In abstract it is certainly desirable to minimize capital investment in machine tools. But this is so only if order delivery obligations can still be adequately met.

The problem of what to do about the situation exists because these two goals are in conflict. If the comparatively low efficiency of the shop means that too many orders are delivered later than they were promised, this could be serious. The cost of customer dissatisfaction might be much greater than the cost of an additional machine.

An economic balance between these factors is necessary. But this can be properly achieved only when fairly precise information on the effect of installing an additional machine is available. Only such information could make clear whether or not it is worth it. Yet it is exactly such precise information which is so often missing.

The Work-in-Process Problem

The entire problem has still not been clarified. Let us consider another factor involved in a practical attack upon the bottleneck situation. This is the so-called "in-process" inventory investment.

Like finished stock inventory, work that is in the process of moving through the shop also represents an investment of money. Such funds are frozen so long as the work remains in process. If it remains in process longer than necessary, the money resource of the organization is being used inefficiently.

Investment in an additional machine, by speeding the flow of work in the shop, might result in lessening the amount of money tied up by in-process inventory. This saving could also easily overbalance the cost of the additional machine.

Unhappily, a usefully accurate prediction of the degree to which in-process inventory would be reduced is also difficult to make. Yet the question of degree is again central in importance. Without such knowledge management is operating in the dark.

The Changing Order Problem

Among other things, an accurate prediction depends also on the volume and type of customer orders which may be expected.

If the *volume* falls below a certain point, there may already be enough machines. Previous bottlenecks might vanish. In this case an additional machine might have no purpose. It might neither reduce in-process inventory nor help meet customer delivery schedules.

If the *type* of customer order changes, work routes will also change. This is a frequent occurrence in a job shop. As a result a bottleneck at one machine group might give way to one at quite another machine group. Since this is a possibility the question of *whether* to obtain an additional machine is not separable from the question of *which* machine it should be.

Only a few of the problems involved in so seemingly simple a management decision have been touched upon. It should be evident, however, that this question of additional machines cannot be separated from such matters as the volume and type of orders in the shop. Similarly, it cannot be properly considered in isolation from the goal of adequately meeting customer delivery obligations. Nor can it be considered apart from the goal of reducing in-process inventory and capital investment generally. All of these matters bear upon each other. It should also be pointed out that the same problems surround decisions regarding increases in operating personnel, or increases in the number of shifts of work, etc.

Although they are all related to each other, some of the problem areas we have discussed center on the work orders of the shop. Some are focused on the resources of the shop. A third basic problem area has to do with the way in which the resources of the shop are organized in relation to its work. This third problem area further complicates management decision at every point.

The Dispatching Problem

The question of how many machines are desirable depends partly on how efficiently their use is organized. A comparatively few machines can do more work than a greater number if they are used more effectively. In fact, not only machine tool efficiency but every one of the management goals discussed is affected by the "strategy" with which machines are put to work.

In a job shop this strategy expresses itself by decisions made at machine operating stations. These decisions must be made. They can, however, be made either on a haphazard basis or in terms of a guiding principle established by management decision. Such a principle is usually referred to as a "dispatching rule."

The strategy contained in a dispatching rule comes into play every time it is necessary to decide which available machine should be put to work on which eligible customer order. In effect, the rule decides on a priority rating for each order. Thus, given a machine and a collection of orders now requiring it, the priority ratings determine the sequence of assignment to the machine.

One very simple strategy is the so-called "first come, first served" rule. On this basis the highest priority is always given to the order which arrives first at a machine station. The application of this rule in the shop is a very simple matter. Orders arriving at a machine station are simply placed at the bottom of a list. Machine assignment is made from the top of the list. The priority of each order is fixed by its position on the list. It is not reviewed or changed as new orders arrive.

Since a great many different strategies are possible, management's problem is to decide which one offers the most advantages. There is no doubt that the "first come, first served" rule has all the advantages of simplicity. This, however, is not sufficient recommendation. What is important in evaluating it in relation to other possibilities is its consequences for the total job shop operation. What comparative effect does it have on the value of in-process inventory? Will it result in more efficient usage of machines and manpower? Would fewer customer orders be late if some other rule were used? These are the kinds of questions which must be answered. Only when they are adequately answered can a decision be made as to which strategy would best serve management's goals.

At first glance it would not seem difficult to predict the effects which some rules are likely to have. Let us consider a specific example. Suppose priorities were to be assigned to orders on the basis of their due dates. By such a rule each dispatching station might give highest priority to the order which has the earliest due date. An order arriving at a given station would be assigned a higher priority than that of any of the orders already waiting, provided that it had an earlier due date than any of them.

On the surface, the general effect of such a rule would seem no great mystery. Common sense tends to suggest that the extent of order lateness would be reduced. After all, the orders which have due dates close at hand would be processed first. It would seem that the shop is being organized specifically to achieve the goal of minimizing lateness.

This does not, of course, mean that if an "earliest due date" strategy were used there would be no late orders. That would be too much to expect. It has already been pointed out that the job shop situation is complicated. Other factors also have a bearing on whether or not the goal of meeting customer obligations is satisfactorily achieved. What one would tend to expect, however, is that those orders which would still be late would not be as late as they otherwise might have been.

Unfortunately, the actual complexity of a specific job shop situation could easily prove this expectation wrong. A simple illustration may make this clear. Let us assume that the relative priority of two different orders is under consideration at a dispatching station. Which order should be assigned to an available machine first? Order 1 has a due date of October 14 and order 2 has a due date of October 16. By an "earliest due date" dispatch rule, therefore, order 1 is assigned first. Order 2 is kept waiting.

But now consider the situation a bit more closely. The two orders may have altogether different machine tool routes to follow. In view of this it is possible that at this point order 1 has almost been completed. Order 2, on the other hand, may still have a great many machine operations to go through before it is completed.

Under these conditions the decision to keep order 2 waiting might have the effect of making this order very late by the time it is finished. This would be especially true if the remaining operations for order 2 were comparatively time-consuming. It might also be that the waiting lines for these remaining operations are very long. In addition, other factors might tend in the same direction. To the extent that some of these conditions exist in the shop, an "earliest due date" rule would not have the result suggested by common sense. The fact is that the actual situation is too involved for easy prediction.

In a similar way common sense might anticipate the general effect of other dispatching rules. It would be possible, for example, to have a rule which awards priority on the basis of the dollar value of the order. Such a rule might give precedence to those orders having the highest accumulated value. This strategy would insure that these orders moved through the shop as quickly as possible. Now it is tempting to assume from this that the capital tied up with work in process would thereby be minimized.

Here again, however, the conclusion of common sense may easily be false. Assume, for example, that there are three orders awaiting assignment at a dispatch station. Suppose, in addition, that the dollar values are $1,000 for one of them and $700 each for the other two. Now let us assume that the $1,000 order requires 6 hours on the machine tool in question, whereas the other two require only 3 hours each.

In this situation, assigning the $1,000 order first would certainly move it through the shop faster. Surprisingly, however, it turns out that less capital would be tied up if the high-value order is *not* assigned first. The basic reason for this is that in calculating in-process inventory cost, time is as important as dollar value. Tying up $100 for 3 minutes is not as bad as tying up $1.00 for 3 days.

In the present illustration, for example, there are three possible dispatching priority patterns. These assignment sequences are: 1000-700-700; 700-700-1000; 700-1000-700. The calculation of an in-process inventory rating for each, in terms of both dollar value and time, is shown and explained below.

	Dispatching Priority Pattern			Dispatching Pattern Calculation in Time and Dollar Value	Rating of In-Process Inventory
	1st	2nd	3rd		
1	1,000	700	700	(1,400 × 6) + (700 × 3)	10,500
2	700	700	1,000	(1,700 × 3) + (1,000 × 3)	8,100
3	700	1,000	700	(1,700 × 3) + (700 × 6)	9,300

The first pattern above would be the result of a "high dollar value" dispatching rule. According to it, the $1,000 order is assigned first. During the 6 hours of its machining, the two $700 orders are kept waiting. Hence: (1400 × 6). When processing of the high-value order has been finished, one of the $700 orders is assigned. For the next 3 hours the second $700 order will be waiting. Hence (700 × 3). The rating or index of in-process inventory, considering both dollar value and time, turns out to be 10,500. Performing the same calculation for the two other dispatching patterns, the ratings turn out to be 8,100 and 9,300.

Thus, it is the second pattern which is really best, from the standpoint of minimizing capital frozen in unfinished orders. Interestingly enough, this is the pattern which gives *lowest* priority to the high-dollar-value order.

The example above has shown that it is necessary to consider the number of comparatively low-value orders and the amount of time they are tied up in waiting while a higher-value order is being processed. These factors vary, of course, at different times and at different dispatching stations in the shop. This makes it that much more difficult to predict accurately. The main point is that common-sense impressions regarding the likely result of a dispatching rule are very likely to be wrong. It is a little like assuming the earth is flat because it sounds pretty reasonable.

Many other simple dispatching rules could be used. It is also possible to combine more than one strategy in a single rule. Thus, to give a comparatively simple example, a rule might assign priorities on the basis of earliest due date and, where there are ties, give first place to the order having the higher value. The more complex the rule, the more difficult it is to predict its overall effect.

In addition to the unreliability of common sense in predicting, there is still another great difficulty. This problem stems from the fact that, even when it is possible to guess the *general* effect of a rule, its impact can still not be predicted in *detail*. Detail is, however, all-important.

To be sure, it is obvious enough that, all other things being equal, it is good to reduce the value of in-process inventory. But this merely states the goal in ideal terms. Such a statement pays for being obvious by being practically useless. In real job shop situations the advantage of reducing in-process inventory depends upon many considerations. Just *how much* would it be reduced, and just *how much* would the reduction cost in sacrificing other goals.

Suppose it should turn out that the price of lower in-process inventory is a lessened ability to meet some delivery obligations. This price could be too high. There is no way of knowing unless these consequences can be predicted in detail.

The basic question is always: Will the final effect of a rule increase or decrease the "health" of the shop's operation from the standpoint of management's goals?

The Interrelation of All Problem Areas

Clearly, the same sort of difficulty hangs over all the other considerations which have been discussed here. The question of whether an additional machine should be acquired is as much subject to this difficulty as is the question of what dispatching rule to use. It is certainly desirable to avoid investing in additional machines—but only so long as goals A, B, C, ------- Z are not affected too adversely.

This, then, is the heart of management's problem — *to achieve in its action a favorable balance among all of its goals.* The reason a decision regarding a particular action is difficult is that the action can affect more than one goal. The successful solution to the problem demands accurate, detailed information regarding the results of each proposed action. And the basic reason this is so difficult in the job shop situation is that this situation is so complex.

It should be evident now that this complexity comes from the fact that an operating job shop is really a system in which all the parts are related and dependent on each other. In some ways it is like a network of rubber bands which are linked together. Plucking one makes all the others vibrate. This in turn reacts on the band which was first touched, and so on. The main difference is that a system of rubber bands tends to return to its original state. By contrast, a job shop either gets better or worse and tends to stay that way. This is what makes for management's opportunity as well as management's risk. The amount and type of information available at the moment of decision is what makes it primarily the one or the other.

The Trial and Error Solution

Of course, where information is needed, ordinary experience can generally provide some. It is always possible to try changing something in the job shop and then simply wait and see what happens. The effect of a change will be known after it has happened.

In this actual trial approach information is gained from experience. Unfortunately experience also makes clear that where there is trial there is error. Error is the price of this method.

But apart from the fact that a change may not work out too well, there are also other circumstances which make trial and error difficult. In the first place a job shop is not a laboratory. It cannot be made a laboratory without upsetting the shop's work patterns.

In addition it cannot be made a laboratory because it does not provide the proper conditions for a controlled experiment. It would not be possible to be sure that all of the significant variables had been controlled. This is important because, if unplanned changes can enter the situation, the relation between a specific planned change and the final result cannot be established. The result may, to some degree, be caused by the influence of the uncontrolled variable. Has the experimental change improved the shop or not? The answer to this question would still not be clear.

Another difficulty is that it takes time to acquire experience. If every reasonable possibility for improving job shop operation were tried, it would take a great deal of time. Also, during this time the conditions under which the shop would operate might change. This alone could cancel much of the value of previous experience.

One change that is very common is a change in the type of order processed. It may not be possible to control such a change. If there are frequent and significant changes in the "order mix," learning from experience is rather hopeless from the outset.

Still another difficulty associated with trial and error merits brief mention. It might be called the "irreversibility" problem. The difficulty here stems from the fact that many of the decisions which may be made in a job shop are not reversible once they have been made. It may be desirable, for example, to find out what effect adding a machine would have. But the decision is an "either-or" decision: you either add one, or you don't. Suppose a machine is added and the results are disappointing. The die has nonetheless been cast because the machine has already been obtained.

A careful examination of the trial and error approach would reveal still other difficulties. Enough has been said, however, to suggest that while it may sometimes be a path to improvement, it is also a fairly rocky and uncertain path.

Unfortunately, however inadequate a method of trial and error may be, until recently management has had little choice but to rely on it. Management has therefore been in the position of having to make guesses. It has also been obliged to pay for the bad ones.

The Computer Simulation Solution

An ideal solution to the problem would eliminate the guesswork. It would have all of the advantages of an actual trial and still avoid its costs and its uncertainties. Let us be as clear as possible about this. An ideal solution would satisfy at least the following conditions:

1. Information would be obtainable quickly. It would therefore be appropriate to current shop conditions. Its usefulness would not be destroyed because conditions had changed while it was being gathered.
2. The process of obtaining information would not upset the normal operation of the shop.
3. The information would be reliable. There would be no problem of questioning results because of accidental or uncontrolled factors.
4. Information would be obtained without running the risk of bad guessing. A valid decision could therefore come before actual trial in the shop instead of following it.
5. Information would be in sufficient detail to enable decision.

On the surface, satisfying the above conditions seems a little like having your cake and eating it too. With the development of the high speed computer, however, this is no longer a dream.

In recent years there has been increasing interest in the use of computers for the solution of some of the more difficult problems of manufacturing control. Much work has been done in connection with the problems we have been discussing. It has become clear, for example, that the most important aspects of job shop operation can be expressed in numerical terms. This applies to such things as the due dates of orders, the number of machine groups in the shop and the number of machine tools in each group, the amount of processing time each machine operation takes, the number of shifts during which the shop operates, the amount of manpower available for each shift, the cost per hour of each operation, etc. In addition, the dispatching procedure used in a shop can also be expressed in numerical and logical terms.

What this means is that a description of a physical job shop can be made in terms of these variables. This description can be entered into a computer. In effect, the entire pattern of variables necessary for describing a shop – those having to do with *orders, shop resources* and *dispatching strategy* – can be read into a computer. The large storage capacity of a modern computer makes a fully workable description possible.

Now a computer is essentially a processor of information – a data processing system. Its processing abilities can therefore be used to turn a "static" description of a shop into a model of an operating shop. The application of a dispatching rule, which is what is involved here, uses the computer's ability to make logical decisions. At the same time, its processing ability enables it to calculate the changes which would occur in the value of certain factors because of the shop's operation. To sum it up, it is possible for a computer to *simulate* the actual operation of a job shop.

The important thing to note at this point is that simulating the actual operation of a shop is not the same thing as simulating an existing shop. In fact the usefulness of a computer here is that it can simulate a proposed shop. Thus, if a change in an existing shop is proposed, the computer can simulate the operation of the shop as it would be if the change were carried out. In this way it can test the effect of changing a dispatching rule, or the number of machines, or the number of shifts of work, etc.

Simulation, of course, has meaning only if the computer is able to present its results in a meaningful way. The results are all-important. They must enable management to decide on the advantages or disadvantages of possible changes in the existing shop. They should therefore be presented in terms of the ways in which they influence management's goals. In other words, the results of a proposed change for such things as the value of in-process inventory, the meeting of customer obligations, the percentage of machine tool and labor usage, etc., must be shown.

The introduction to this manual announced that a data processing application had been developed to do exactly this. It satisfies all of the conditions specified above. It is flexible in its ability to simulate the important characteristics of a shop. The results of simulation are shown in a set of printed reports. In this way, the effect of changing any of the characteristics of the shop can be conveniently displayed. The remaining sections of this manual will be devoted to a more detailed discussion of this application.

Before turning to this, however, the present chapter will be concluded with a few remarks on some of its more general uses.

Some organizations operate entirely as a job shop. In addition, almost every manufacturing organization contains some operations which are of the job shop type. Over and above this, however, the present application is pertinent to many work situations not ordinarily associated with an actual job shop.

To mention one example, a punched card data processing installation resembles a job shop in many ways. Like a job shop, it contains different types of general-purpose machines. It is capable of producing a variety of products. Such products may be sales reports, budget reports, inventory reports, payroll checks, etc. Each of these jobs is likely to follow a different work route and have different due dates.

Because of such similarities the operation of a punched card installation involves many of the same problems as a job shop. Thus, many of the principles for running a job shop efficiently also apply to the punched card system.

The general point in this is that the specific problems of a job shop sometimes find a close parallel in other business problems. This is so even when their immediate content is entirely outside the area of manufacturing control.

Perhaps the main reason for this is that the logical structure of the job shop as a whole is basically the same as that of a business as a whole. The situation always involves work, resources, and rules for applying resources to work. For this reason the task of operating a job shop is, at bottom, the same as that of running any business. From management's point of view it always consists of selecting and using resources as efficiently as possible to achieve the goals of the organization. The specific goals and resources involved may be different. Resources may be harbor or airport facilities, or traffic lights, etc. The important thing, from the standpoint of using the present application in other areas, is that the basic model is the same.

This is not the place to further explore such parallels. Their existence, however, gives the simulation approach to job shop problems additional significance. They are mentioned here only to illustrate the bearing which a computer, as a management tool, has upon the solution of business problems generally.

PART II

Describing the Job Shop for Simulation Purposes

The first part of this manual has endeavored to make clear that the basic question in managing a job shop is that of determining the best pattern for allocating and integrating resources in a shop. The job shop simulation application functions to provide a reliable basis for answering this question. Fundamentally, it does this by successively displaying the consequences which follow from alternative patterns.

The second part of this manual will be devoted to discussing the items involved in describing such a pattern for simulation purposes. The third part will describe and discuss the simulation procedure and its results. These results will contain information respecting such consequences as the number of orders which turn out to be late, capital investment in in-process inventory, extent of machine tool and labor utilization, etc.

As will be seen throughout Part II, a great many descriptive variables must be considered. Taken together, they constitute the information which a user of the application must provide. For the purposes of systematic exposition this information will be discussed under the following three major headings:

A. Information concerning the orders which activate the shop.

B. Information concerning the various resources of the shop.

C. Information concerning the principles which organize the shop's resources in relation to its orders.

All of this information is necessary for a complete description of an active or functioning job shop. The first general category of information should, however, be distinguished from the other two, in terms of both the shop situation and the structure of the application. In the sense that it is the entry of orders which forces the shop into motion, orders are external to the shop. Thus, a *static* description of a job shop would not include a description of its orders. It would describe the shop as it is prior to the entry of orders.

In addition, before orders enter a shop they have been scheduled. This is to say that at least work starting dates and completion dates have been calculated for each order. The completion dates correspond, in effect, to the delivery date promises made to customers. One of the features of the present application is that it is capable of performing schedule calculations. Thus, in the application's chronology the 704 program works first with order information for scheduling purposes. It is only when scheduling has been accomplished that order information is combined with other descriptive information (for the most part categories B and C above) for simulation purposes. This phase development of the application is roughly illustrated in the diagram below.

In view of the fact that a full description of the job for simulation purposes properly includes the factor of *scheduled* orders, the scheduling phase of the application will be discussed in Part II under category A.

BASIC APPLICATION PHASES

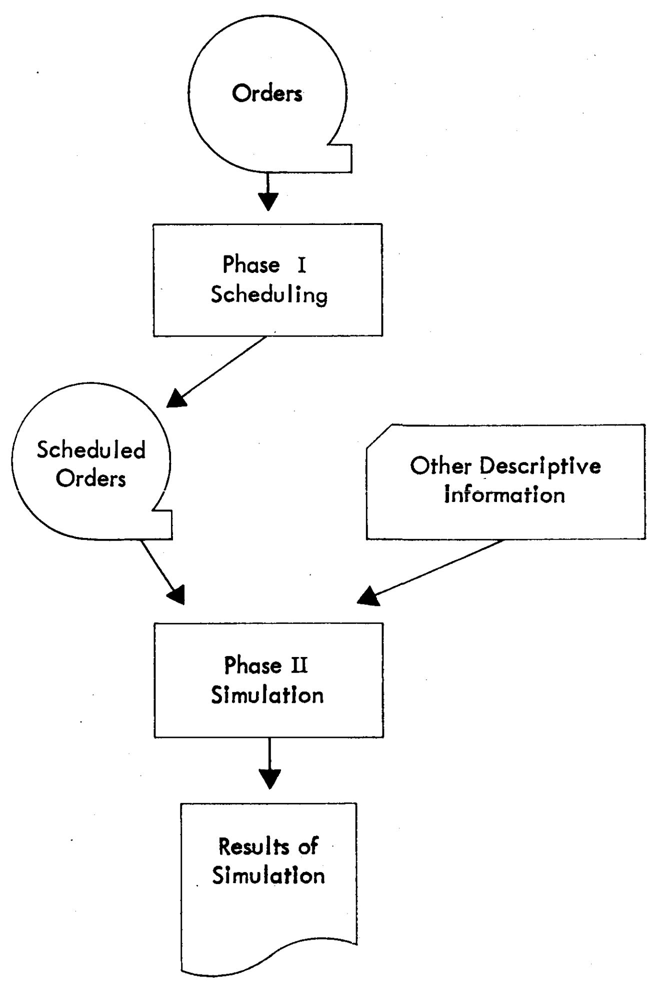

The first problem of systematic description, whatever its purpose, is always that of selecting attributes which are *relevant*. Exhaustive descriptions are neither necessary nor possible. From the standpoint of management the significant characteristics of a job shop are those which have direct or indirect bearing on the achievement of management's goals. The present section of this manual will discuss the relevant descriptive variables which must be specified if we are to have what is, in this sense, an adequate description of a job shop. From the standpoint of the 704 program these items represent its input.

A: Orders — The Job of the Job Shop

One of the most important of these variables consists of *orders* defining the work which the shop must accomplish. Without orders the shop is a dormant system. It does not move. It should be kept in mind that we are interested in a model of a functioning job shop. Thus, before the job shop simulator can begin its simulation of a shop's activity, a set of customer orders must be made available for it to work upon. The simulation moves, so to speak, on a stream of customer orders.

What is an order? We may say that in principle an order is a demand upon a working shop to produce a given quantity of a given item. Fulfillment entails working upon a given raw material with given machines in a given sequence. The demand which an order makes on a shop is ambiguous unless it contains, or is associated with, information regarding all of these variables: item, quantity, raw material, machine tools, machine tool sequence.

As this information is necessary to the operation of an actual shop, it must also be provided in the simulation of a shop's operation. For the purposes of simulation it is convenient to incorporate all of it in order records maintained on magnetic tape. Let us briefly consider the relevance of some of these variables in the total job shop context.

Order Items and Quantities

So far as the particular product item and the quantity of it to be produced are concerned, these matters are, within the limits of the shop's capacity, determined by the customers of the job shop organization or by considerations outside its immediate province. They are therefore basically external decisions. Only one point needs to be made with respect to the recording of these decisions in the order record.

In one sense the job shop simulator simulates a pure job shop. Once an order begins to be processed at a particular machine station, it will be assumed that it must be completely finished at that station before it is moved to the next one. The order cannot be "bumped" and each unit item of the order must wait until all have been processed before further machining can take place. In view of this, the quantity of an item specified in an order is significant to the simulator only in so far as it affects the total amount of machine time required by the order. Thus, so long as the quantity is taken into account later in the specification of machining times, it is not absolutely necessary that quantities as such be stipulated in the order records.

Raw Material Costs

It has been indicated that an order implies work upon some raw material. Different orders will require different types and amounts of raw material. In job shops where the cost of raw material is significant, the cost of this raw material should be included in the order records submitted to the simulator.

One of the governing objectives in the operation of a job shop is to minimize the dollar value of in-process inventory. It will therefore always be important to know how much money is tied up in this way. For this reason it will be necessary for the simulator to track the increase in each order's dollar value as it moves through the successive machine operations to which it will be subject in the shop. The costs of these operations must be accumulated and associated with orders. In this undertaking the raw materials cost is significant by virtue of its establishing a base upon which further costs will accumulate.

Order Routings and Processing Times

In addition, if the simulator is ultimately to provide information on how long it will take to process orders through a proposed shop, the order record must not only contain a specification of the successive machine operations required (i.e., "routing") but must also include the amount of time each is expected to consume.

The amount of time during which any machine in any given machine group is tied up in the processing of a particular order may be divided into two component elements: (1) the direct machining time required for the order and (2) the amount of time it takes to set up or adjust the machine for its job of direct machining. Setup time is especially significant in the job shop where general-purpose machines are involved in processing a range of different orders. The simulator assumes that each order requires a separate setup period.

Machining time and setup time must be specified separately. The specific times allowed for machine operations, both setup and machining time, may be based originally on engineering estimates. A way of adjusting these, if they are not realistic, will be described later.

Lower Limits for Simulation

The three variables just discussed – setup time, raw material cost and machining time – all have specific values for each order in the order file. The simulator will use these values or any adjustments to them which may be made. It is necessary, however, to advise it directly regarding the lowest value which each of these variables may be expected to assume.

For internal programming reasons, having to do primarily with storage economy, the values of these variables are recorded in 704 storage in relation to a lower limit. Lower limits for setup time, raw material cost and machining time represent items of information which must be submitted directly to the simulation program on punched cards.

The first three cards shown in Appendix I establish lower limits for each of the items in question. Appendix I displays the proper format for these cards. Throughout the remainder of this manual all descriptive information which is discussed and which must be submitted to the simulation program on cards will be listed in this appendix. The appendix contains all the card input information required. As each separate item is discussed in the text it will be accompanied in the text by numerals in parentheses. These numerals are intended to identify the position of the cards in question in the appendix listing. Thus in the present instance the identification is: (1-3).

The specific input information contained in the listed cards has been used in an actual simulation run. The output shown in the final section of this manual is a result of simulating the shop described by this information.

The Order Record

Returning now to the order file itself, order records containing all the information discussed above might look somewhat like the following:

(ORDER IDENTIFICATION)
(NUMBER OF UNITS PER ORDER)
(PREDETERMINED PRIORITY CODE)
DUE DATE (Discussed below)

Initial raw materials cost: $310.00

(Operation 1)	Machine group	3
	Setup time	2.0 hours
	Machining time	5.0 hours
(Operation 2)	Machine group	1
	Setup time	1.5 hours
	Machining time	7.0 hours

(ORDER IDENTIFICATION)
(NUMBER OF UNITS PER ORDER)
(PREDETERMINED PRIORITY CODE)
DUE DATE (Discussed below)

Initial raw materials cost: $400.00

(Operation 1)	Machine group	5
	Setup time	0.5 hours
	Machining time	3.7 hours
(Operation 2)	Machine group	1
	Setup time	0.2 hours
	Machining time	6.2 hours
(Operation 3)	Machine group	5
	Setup time	1.5 hours
	Machining time	3.2 hours

At the present time the simulation program is assembled to utilize a format such as the above. It may, however, be reassembled to use some other format. It should also be mentioned that currently the simulator does not utilize the first three items of information shown in the above format. If it is desired to use such information or any other information placed in the order record, additional routines must be added to the simulation program.

As has already been mentioned, in the present 704 application, magnetic tape is the input vehicle for entering order records.

The Problem of Scheduling

The above records also contain information regarding due dates, which represent in effect the delivery commitments which the shop has made. The various due dates which have been established for the orders current in a shop may be thought of as defining the final limits of the shop's work schedule.

Inasmuch as meeting this schedule is one of management's goals, it serves as one of the major reference points for evaluating a proposed shop's operation. To the simulator the due date of an order will be the measuring point for determining whether it is completed early, late or on time. Before any actual simulation can occur, therefore, each order must be associated with a definite due date. How is this date arrived at?

The procedure actually employed by a shop for quoting delivery dates to customers can be fairly primitive. It may consist of no more than uncontrolled commitments made by sales personnel of the organization. Such a procedure is capricious if it does not take into account how long it will actually take for the shop to fill an order. The natural effort of a salesman to please his customer in this respect may inadvertently betray the shop which must back him up.

Sometimes the quoted delivery date does consider the number and kinds of machining operations involved and how long they will take. But even in this instance the calculation is often no more than a rough approximation. It tends to presume for the shop a fictitiously infinite capacity in the sense that it is based on standard operating times and does not allow for such matters as (1) how long it takes to move work from one machine to the next, (2) how long it takes to set up a machine for a new operation, or (3) how long and how often an order may have to wait until a machine is available. Before the advent of fast, large-storage-capacity devices which could perform these calculations, their complexity relative to practical time limitations tended to enforce a more or less intuitive approach.

One of the accomplishments of the present application is that it will deal with this scheduling problem more systematically. Before this is discussed, let us again clearly locate the point at which this problem comes up for solution in the flow of job shop activities. The activities of a shop can be naturally divided into two parts — those which take place before commencing the actual production of a customer's order, and those making up the manufacturing process itself. The relevant activities which take place before are broadly covered by the term "scheduling." At this point attention is being focused on the way in which the application handles these activities. (This is Phase I of the application. See page 14.) The reason for this is that a complete description of a functioning job shop for strict simulation purposes includes a description of orders for which a schedule has already been prepared.

The Nature of a Schedule

In current usage a certain degree of ambiguity attaches to the term "scheduling." For example, it is sometimes used interchangeably with the word "dispatching." In the context of the job shop simulation application, "scheduling" denotes the process of developing a manufacturing time table. Like a train schedule, a manufacturing schedule represents a time *plan* for using resources. It is therefore always developed before the events which it schedules. From the standpoint of such a plan, a dispatching procedure is an integral part of the events referred to. It is true that a dispatching procedure may be constructed in such a way as to shape the timing of events in the shop as closely as possible to the predetermined schedule plan. It is also quite possible, however, that the dispatching procedure will not contain a reference to the schedule plan. (See pages 28-32 for a clarification of dispatching.)

A fully developed schedule includes more than the stipulation of due dates. It contains answers to many timing questions: When should work on an order start? When may we expect it to engage the facilities of machine group 3 within the shop? At what point in time may we expect the work on the order to be completed?, etc. All these questions are related to each other. The answer to the last one appears in the order record as the due date.

Some of the questions are significant primarily in their bearing on matters external to the shop's operation. Others have some of their significance in connection with internal matters. The question of when an order will be completed is clearly of importance to the outside recipients of finished orders. These may be customers of the job shop or, if the shop is producing components for inventory, it may be those whose responsibility it is to store these components or perhaps to integrate their assembly with other components. On the other hand, the question of when work should begin is manifestly of immediate concern within the shop. Without a specific decision on when to begin, no beginning would presumably ever be made.

As was intimated above, the basic elements which enter into the development of a schedule are the expenditures of time and the shop's resources. It is the capacity of these resources which ties together both of the time poles of a schedule — starting date and completion date. In the same way, the speed of a train determines the spread between its departure time and its arrival time.

Forward and Backward Scheduling

It follows from this that, given a definite set of orders, and information regarding processing times, it is possible to proceed from a known start date to the planning of a completion date. The converse is true when a desired completion date is known. In this case, an appropriate start date may be determined.

The scheduling portion of the present application accommodates itself to either approach. It is capable of calculating the theoretical due dates of orders when the start dates of the orders are already known. It is also able to calculate the proper start dates of orders when desired due dates have been given. In this sense it can schedule either forward or backward through time. In either case the theoretical start dates of each intermediate operation step in the processing of an order are also produced.

Whether the forward or backward alternative is used in a given instance depends upon the practical situation of the shop. It may depend on the kinds of problems which its environment presents.

Possibly the job shop is one department in a larger manufacturing organization. If this is so, an order on the shop may be the result of a "requirements explosion" which establishes component requirements for the manufacturing cycle of the total organization. In this situation the date when a job shop product is needed for the larger cycle is likely to be known. The general expectation may be that the job shop will adjust the timing of its activities to the rhythm of manufacturing in the organization as a whole. In this situation the organizational objective might be to steer a path between (1) having to keep the product too long in stock and (2) being sure of having it available when needed. For the shop itself the basic problem is when to start. Backward scheduling from a determined due date would probably be employed in this context.

To consider another possible situation, the products of the job shop may be built for a variety of outside customers. It is evident that if due dates or delivery dates are quoted to customers without reference to what the shop is capable of or what its current commitments are, some proportion of them is likely to be seriously unrealistic. In the light of the shop's capacity they may even imply start dates in past time. This approach would entail continual problems of expediting and/or customer dissatisfactions with all the direct and indirect grief which these things entail. The shop will probably use some scheduling procedure for the purpose of determining a reasonable due date — one to which it may safely commit itself. Where the problem situation is structured in this way, forward scheduling is called for. A customer may then be told, with a certain amount of confidence: "You can have what you've ordered by such and such a date."

The scheduling technique as such, however, is essentially the same in either case. Forward scheduling begins with a start date and works toward a due date, whereas backward scheduling simply reverses the process.

Simulation as an Aid to Improved Scheduling

Quite apart from the question of forward or backward scheduling, actual practice in a shop may embrace any one of a variety of possible scheduling techniques. These may be more or less well adapted to the situation of a particular shop. They may be simple or they may be complex. In designing the present application, special attention has been paid to the possibility of incorporating new scheduling procedures. The task of programming these and incorporating them in the application has been made a comparatively easy one. It is assumed that users of the application will take advantage of this to create more valuable scheduling procedures than those currently in use. Technical information on how to do this is available and will be distributed to users with the program package mentioned in the introduction.

In connection with the development of more refined scheduling methods, the nature of a schedule as basically no more than a plan might well be kept in mind. Unless occurrences in the shop are intentionally organized to fulfill the plan, it exercises no governing influence and is merely a prediction. It is certainly not a strait jacket. Once a schedule is developed, therefore, the question of the extent to which the actual course of events in a shop will conform to its timing is an open one. In making its prediction, the scheduling procedure really operates in terms of a simplified model of actual events. Thus it may turn out that there will be a considerable and costly discrepancy between the specifications of the plan and the complex reality which it presumes to predict. A more realistic schedule would be one which predicts better.

The gap which may exist between a schedule plan and actual shop results may also be in evidence when these results are produced by simulation. For this reason shop simulation can be used as an instrument for developing realistic scheduling methods. It is one of the important functions of the simulator to provide a testing ground for the validity of scheduling assumptions and to provide clues as to how they might be improved.

Apart from the possibility of harnessing feedback from simulation to improve scheduling, there is no necessary relation between these separate portions of the present application. Logically they represent two entirely distinct parts of the program. The scheduling section does not perform any shop simulation; it simply calculates theoretical operation start dates and order due dates which may or may not be met during the simulation. From the standpoint of the simulator it would be quite possible to bypass the scheduling section of the program completely. This would simply require that both start dates *and* due dates be given in the order record.

The Built-In Scheduling Procedure

As the program now stands, it contains a single scheduling procedure which has been given the name SR01. If a user elects to use this procedure, an input card must be prepared to indicate this (Appendix I, card 4).

Since the procedure is capable of both forward and backward scheduling, a choice must be made between the two. The reference point for identifying this choice is the due date. The program may be told to use a due date (i.e., schedule backward to determine a start date) or not to use a due date (i.e., schedule forward to determine a due date). Card 5 in Appendix I requests that the scheduling be forward and implies that there is no due date in the order record. Contrariwise, if a 1 (instead of a 0) is punched in column 12 of this card, the implication is that a theoretical due date is specified in the order record and that the scheduling should be backward.

When, as in the present case, the due date is omitted and must be calculated, the SR01 procedure uses as the order start date the day on which the order "arrives" into the shop from the order tape. The arrival rate of orders is controlled by card input information to the simulator (see page 22) and this information is also available for scheduling purposes at this point.

Using a shop arrival date as scheduling touchstone, a start date must be determined for each of the successive machine tool stations in the order's work route, as well as a due date for the order as a whole. Both types of information will derive from the same computation inasmuch as an order's due date is simply the calculated completion date of its last operation.

Given the starting time of the first operation its completion time may be calculated by adding the processing time of this operation.

A very naive approach might assume that the start date of each succeeding operation is the same as the completion date of the operation which precedes it. By this way of scheduling, the problem reduces itself to a matter of adding together the theoretical setup and machining times, as given in the order record, for all the operations of an order. If this were done with order 1 (shown on page 16), the theoretical due date would be 15.5 working hours after that order's arrival in the shop.

The tacit assumption in this is that no order will ever be obliged to wait for access to a machine. As we have pointed out previously, however, this is a profoundly unrealistic assumption. If the schedule plan is to be more than an exercise in fantasy, it must take into account the fact that many orders will actually be waiting in line in front of machine tool groups while other orders are being processed on these machines.

The SR01 procedure acknowledges this. It was designed to operate in terms of waiting time estimates which must be specified to it on input cards. These estimates will first of all stipulate the *average* amount of time orders may be expected to wait for service at each of the machine groups in the shop. In effect, the program is being supplied with a table of "lead times." In determining operation start dates and a final due date, the program can work directly with an average. There is no need to determine a probable specific lead time for each operation.

The greatest uncertainty in scheduling is likely to center about the estimation of lead times. At the same time, scheduling can be only as good as lead time estimates. It is therefore at this point that previous simulation experience is especially valuable. Simulation output affords direct information on waiting times.

A glance at such output (e.g., page 37) reveals that average waiting times may differ not only from one machine group to another, but also from one dollar-value class to another. In other words, the amount of time that an order must wait may be partly dependent upon whether it is a high-, medium- or low-dollar-value order. This will certainly be true if the shop uses a dispatch procedure which differentiates between these categories. If such a rule gives machine tool assignment priority to high-value orders, then, on the average, they will spend less time waiting in lines than will the low-priced orders.

The lead times designated for scheduling purposes should represent as closely as possible the waiting times which will actually occur in the functioning shop. Dispatching is related to scheduling at this point only because the nature of a dispatch rule may affect what will actually happen in the way of waiting times. If such a dispatch rule is employed, this influence should be reflected in the lead time estimates given to the scheduling procedure.

Cards 6-10 in Appendix I show lead time variations by both machine group and dollar-value class. The times are expressed as fractions of a 24-hour day, regardless of the number of hours actually worked. Card 6, for example, indicates that at machine group 1 the lead time for high-value orders is .054, for medium-value orders .076, and for low-value orders .147 of a 24-hour day.

The scheduling rule SR01 requires this variable information. If some alternative scheduling rule is programmed, it might require that other variable information be submitted as input. The number of variables which are pertinent and necessary depends, of course, on the particular scheduling rule employed. Card 11 specifies for this rule that 15 items of variable information have been supplied. As indicated above, this includes the lead times for each value class, for each machine group.

Given this information, the scheduling procedure can take into consideration the actual capacity of the shop. It will add together the *lead time,* setup time and machining time for each operation and come up with theoretical start dates for each operation and a theoretical due date for each shop order. It will simultaneously be obliged to keep track of the increasing dollar value of the order as it proceeds from operation to operation. This must be done if it is to assign the proper lead time at every point of an order's routing.

If the lead time estimates and the processing time estimates were reasonable the scheduled dates should prove fairly accurate. If not, the output of the simulation run should enable better estimates for later runs.

Adjustment of Recorded Processing Times for Scheduling

If the *processing* times given in the order record turn out to be unrealistic, the scheduling procedure provides a chance to adjust them on later runs. A routine has been built into the program which will convert the theoretical processing times given on an order tape to more realistic figures.

Suppose, for example, that it had become apparent that the processing times given on the order tape were uniformly about ten percent too small. It would be possible to advise the scheduling program to adjust the given processing times upward ten percent. This could be done by means of the following mathematical expression: $1.1\,X + 0$.

This polynomial expression represents the generic type of input information which may be used to adjust order processing times. It affords a good deal of flexibility in making such adjustments. The X in this expression stands for the original processing time as given on the order tape. Multiplying X by 1.1 increases the processing time by ten percent. The polynomial admits of further adjustment through the use of exponents of X (i.e., $X^2, X^3, \ldots . X^n$) and by the addition of a constant value. In the present case the constant has been made zero but it can be made any value.

In our example, no change in the order processing times will be made. This is indicated to the program in the same format as though a change were being made. The following "no change" polynomial expression is used: $1.0\,X + 0$. The multiplication and addition which this expression calls for would not affect the value of X, and processing times would remain as they were.

Two of the significant elements in the polynomial are the 1.0 and the 0. These "coefficients" are given to the program on input card 12. In addition, information on exponents is significant and must be given. Since no exponents higher than 1 are to be used in this case, the "degree" of the polynomial is 1. This is indicated in card 13. If X^2 were the highest degree term used, column 12 of card 13 would contain a 2; if X^3, it would contain a 3, etc.

Cards 12 and 13 completely define a single polynomial expression. It is possible, however, for several polynomials to be used. This would indeed be the case if it were desired to change the processing time in different ways for each machine group. Suppose, for example, that these times, as given in the order record, have turned out to be accurate for some machine groups and inaccurate for others. Adjustments need therefore be made only in the latter cases. In addition, the type of adjustment required might be different in each of these instances.

A coefficient card, similar to card 12, must be prepared for each polynomial which is to be used. The program exacts one restriction: if different polynomials are used, they must all be of the same degree. For this reason additional degree cards are not required when more than one polynominal is used.

The total number of polynomials employed is indicated on an input card. In our case only the single "no change" polynomial will be used for all machine tool operations. This is established for the program by input card 14.

The only other information still to be stipulated is the machine group to which each polynomial applies. An input card must therefore be punched to assign the various polynomials to the various machine tool groups. Let us assume for the moment that the shop with which we are working contains five machine tool groups. Since each of the five machine tool groups in our example would use the same polynomial, this is indicated by card 15. In this card each of the entries beginning in column 12 corresponds to a machine group in the sequence specified in a Machine Group Code Table card discussed on page 26. If more than one polynomial were being used, the second (by order of input entry) would be identified by a 2, the third by a 3, and so on.

The use of polynomials to adjust processing times implies that experience has shown discrepancies between the standard processing times given in the order records and actual processing times in the shop. It will be convenient for purposes of exposition to think of plotting such discrepancies on a graph. Such a graph would make visually clear, for example, that where the standard time was 2 hours (expressed, say, along the horizontal axis of the graph) the amount of time actually consumed was 2.3 hours, or 1.9 hours, or the like (expressed along the vertical axis).

The graph would consist of a collection of such points. Each point would locate the relation between a standard time and an actual time. The polynomial expression which is used to adjust processing times actually represents, in terms of the graph, a line calculated to "fit" this collection of points.

Now it could turn out in a given case that the line which best fits the plotted points in the graph cannot be identified by the particular polynomial expression we have employed. If there is to be a close fit, it may be obliged to move between the points in a more complex way. In this case higher-degree polynomials should be used.

The program provides an additional corrective routine which may be used. This might be called the log-antilog routine, and its use is at the option of the user. It operates in basically the same way as the polynomial transformation routine. However, instead of using the standard processing times as given on the order tape, it first converts these to logarithms. It then applies the polynomial to the logarithm, and finally produces the antilogarithm of the result.

In our example logarithms will *not* be used. This is indicated to the program by input card 16. If it is desired to use the routine, a 1 should be punched in column 12 instead of the 0.

Adjustment of Recorded Processing Times for Simulation

Up to this point we have been discussing the transformation of processing times for the sake of schedule improvement. The new processing times are, however, recorded in the order records and are therefore also used during simulation. If for some reason it is desired to employ different processing times for the simulation, further changes must be made to the times given on the order tape. Changes may of course be made for simulation purposes even when originally given times are used for scheduling. It may indeed be desirable to learn just what the effect of using incorrect times for scheduling is. The correction would, in this case, be made only for the simulation.

Processing time changes for simulation are accomplished in essentially the same fashion as changes for scheduling. The necessary cards are numbers 17 through 21. They are identified as applying to simulation in columns 1-6 but in other respects are the same as the cards used for scheduling. The transformations called for will be effected by the program in the same fashion in either or both instances.

However accurate original processing times or adjusted processing times may be, it would be carrying innocence too far to assume that actual machine times in a shop will in every instance exactly bear out expectations. Real events are always subject to variations introduced by adventitious factors not specifically captured in a model. Thus, it is realistic to assume that an operation which may be said to take two hours will sometimes take more than two hours and sometimes less.

Such natural variability is to be anticipated. It is also relevant to simulation in that it will affect such matters of primary concern as the length of waiting times or the value of in-process inventory. For this reason the option of simulating this kind of processing time variability has been built into the program. The routine involved assumes that the magnitude and frequency of the increases will be about the same as that of decreases. In more precise statistical terms it is presumed that variability has a normal distribution with a mean of zero.

In principle, the information which must be supplied to the program is information for assigning specific time values to positions on the curve. This can be done by using a "variability polynomial" such as $1.1\,X + 0$. The form of input information is thus the same as in the cases previously discussed. In this context, however, the X represents the adjusted processing time. The polynomial represents a measure of variability around the adjusted processing times. Given this information, and assuming the normal curve, the program can compute the amount of variation to be expected for any given proportion of the cases. In such and such a proportion of the cases the variations will be within such and such a number of minutes, etc. Once this is done the simulator can use a random procedure to select a variability time value from the distribution.* It will do this in each instance of a machine operation and then add the selected variability time to the adjusted processing time for that operation.

In our example no variability factor will be super-

* The random number procedures actually employed by the program are complex. They are described in the IBM manual, "Random Number Generation and Testing" (form C20-8011). The techniques discussed in this manual are used wherever random numbers are required. In the present context we shall simply note that in generating random numbers a base or starting number (which must be odd) is required. Such a number must be supplied as input to the simulator. In actuality, since such a random number seed is required in two separate portions of the application, two input cards are required. Cards 22 and 23 fill this need.

imposed. This implies that the polynomial: $0X + 0$ will be used as input. As in the case of the polynomial used to adjust standard processing times, the coefficient and degree information must be specified on two cards 24 and 25.

Since there may be a separate polynomial for each machine group (i.e., variability differing with different types of machines) the total number of polynomials and the machine groups with which they are to be used must be indicated. In our case only the one "no variability" expression is to be applied to all machines. Cards 26 and 27 establish this. Card 28 indicates that the log-antilog routine, which may also be used here, should not be used in this case. As before, a 1 in column 12 would signify that its use is desired.

In reality, once orders have been scheduled they are ready to be processed by a job shop. For the purposes of simulating this processing, however, certain conditions associated with the introduction of orders into a shop are still to be described. Let us turn to a consideration of these at this point.

Order Arrival Rate

Before any of these orders can be considered for assignment to the machine tools specified in the order record, they must first arrive in the shop. The actual rate at which orders arrive from customers is a relevant attribute of a real shop and it must therefore be specified for the simulated shop. It will, of course, vary from one job shop to the next. In this case an arrival frequency of 35 orders per day is being assumed (card 29).

This, however, still does not establish the time pattern of arrival within each day. It is possible that, for all practical purposes orders arrive uniformly throughout the day. This is to say that there is an approximately equal time interval between arrivals. On the other hand it is possible that orders arrive at varying intervals throughout the day. In this eventuality, it is possible to have the program simulate the random arrival of orders. A routine has been built in to accomplish this. Since this is an option which may be accepted or rejected, the choice made must be specified on an input card. The 0 in column 12 of card 30 signifies that orders should not be made to arrive at random and are to be considered as arriving uniformly. A 1 punched in this column would indicate that random arrival is desired. In the case of random arrival the specified frequency per day (in this case 35) is treated as an average. Some days are likely to have more than this average and some less.

The simulator should also be advised regarding the amount of computer storage space which will be required to accommodate shop orders. This will depend on the number of orders in the shop at any given time and the number of computer words required for each order record. For illustrative purposes, the assumption will be made here that a total of 1,000 words of storage will be needed for this purpose. This information is given to the simulator on input card 31.

Order Release Pattern

Given the arrival pattern of orders in the shop, a question still remains regarding the manner of their release for processing. In some shops it may be the practice to release orders for processing one at a time or according to the tempo of their arrival from customers. On the other hand the actual procedure in a given shop may be to accumulate orders and release them in batches at varying time intervals. The simulator must be advised of this, lest it assign a machine tool to an order earlier than the real shop would have released it for such assignment. It is possible to stipulate to the simulator that orders should be released once, twice, three times, or more frequently during the day in correspondence with actual shop procedure. The card illustrated specifies a release frequency of once a day (card 32). A zero punched in this card signifies that orders are to be released as they arrive.

Orders Already in Process

Over and above such immediate considerations as the arrival and release frequencies of orders, the work which a job shop does comes to it in a continuous stream. For this reason when we speak of a particular set of orders we are somewhat arbitrarily abstracting a section of this stream for special attention. Unless a real shop was, so to speak, born yesterday it is probably involved at any given time with work in process. This work in process does not vanish when a new set of orders arrives in the shop. Commitment to a realistic model of the shop's operation requires, therefore, an acknowledgment of this fact to the simulator. If this acknowledgment is not made, the simulator will have to turn a work-empty shop to the processing of a set of new orders. Under these conditions it will take some time before the simulated operation of the shop accurately reflects the work level at which the real shop normally operates. The consequence would be that some of the results shown by the simulator would be misleading.

The number of orders in process in a shop may, of course, vary from time to time. With respect to this variation, what we will generally be interested in, for the purposes of simulation, is a portrait of the shop as it is in its typical condition. An estimate must therefore be made of the average number of partially completed orders which may be expected to be in the shop at any given time. Let us assume that knowledge of the shop in question yields an estimate of 40 such orders (card 33).

In addition to the input card which establishes this "initialization" of the shop for the simulator, the same number of orders must be added to the order input file. The simulator will pick up this additional increment of orders and "place" them in the shop.

The purpose of initialization would clearly not be served if the in-process orders were merely allowed to precede the regular body of orders. This would be tantamount to increasing the volume of orders with respect to which the simulation is being conducted. To say that there are orders in process in the shop implies that these orders are in varying states of completion. The program may therefore be called upon to establish their likely states of completion on a random basis.

The program routine involved will generate a random number between 0 and 1. This number may be interpreted as representing the percentage of total processing time plus lead time which has been completed for an order. It is multiplied by the expected total processing time plus lead time. The number of hours of processing which may be assumed to have already been completed is thereby arrived at. Routing information in the order record may then be used to determine at which machine group the order is supposedly in process at this time. This computation is made for each of the initialization orders and each is placed at some machine group along its work route.

In card 34, DEC 1 indicates that random initialization is to be performed. DEC 0 would signal the program that randomness is not desired.

If the random option is not employed, the simulator will treat the initial orders in the same way that it treats the remaining orders on the order tape. An alternative method of initialization would be to abbreviate the initial orders. The effect of partial completion for any of these orders might be simulated by simply deleting some of the early steps in the order's routing. Thus, if it is desired to place a particular order in the shop half way through its third operation, it is necessary only to delete the first two operations and reduce the specified processing time of the third by half. Such modifications may be made to all of the order records intended for initialization purposes.

Late Order Releases

Regardless of how efficiently a shop may be run, practical exigencies will be responsible sometimes for the late release of orders. This is to say that some orders are likely to be released to the shop for processing after their scheduled start dates.

If the scheduling procedure of the application is employed and the backward scheduling option is chosen, the above circumstance will probably be fulfilled in the schedule. In the case of some orders, calculated start dates will fall into the past.

Where forward scheduling is employed, however, start dates are given and constitute the base of the calculation. In this case it is possible to have the simulator simulate some late releases. This can be done by means of two input cards. The first of these will state a late arrival frequency as a percentage of orders (35). This card defines the percentage of orders to be considered as late arrivals for simulation purposes. The second of these cards declares the amount of lateness in days which is to be assumed (36). As may be seen from the cards shown in the appendix, in the present case it has been presumed that there are no late arrivals.

Origin of Orders for Simulation

It is clear that the validity of a simulation model depends upon the accuracy with which it corresponds to the actual situation it is presumed to mirror. To the extent that there is a failure of congruence between the descriptive model and the reality it represents, the results are questionable. It is essential therefore that the orders used for the simulator be realistic. This does not mean that they must be real in the sense of being actual customer orders. The simulator is basically a research instrument and when it is used its objective is to facilitate the improvement of a job shop in the near future. It is too late to do anything about the past or present. In view of this the problem of constructing the order input for the simulator involves the problem of deciding what kind of order mix will confront the real shop in the immediate future.

Where historical orders satisfy this requirement they may be used to good advantage. Possibly it will be desirable to alter them slightly if this can bring them into closer conformance with an anticipated future order mix. If orders are not expected to change significantly, then certainly the use of historical orders is eminently appropriate.

If historical orders are not available or may no longer be considered representative, it will be necessary to develop a set of "fictitious" orders. These will be fictitious only in the sense that they are not actual current or historical customer orders. The whole point is that they must be "true to life."

The program package provided with the job shop simulator makes available another alternative. A separate program which has been designated the Synthetic Order Tape Generator is provided and may be used to create a set of orders. This program must be thought of as merely a convenient device for generating a set of fictitious orders. It may be advantageously employed if historical orders are no longer pertinent and if the formulae in terms of which it generates fictitious orders are as predictive of actual future orders in the shop as a set of fictitious orders constructed by hand would be. The use of the Synthetic Order Tape Generator requires the specification of such information as the average processing time for each machine group in the shop, the average raw material cost for the order mix, etc. The way in which the Generator uses this information to create an order tape will be discussed in Appendix II, where the program is described.

In briefest outline all the information pertaining to the customer order file which has thus far been discussed in this section will be handled by the simulator in somewhat the following manner:

1. Orders will be read according to the sequence in which they appear on the order tape and at the arrival rate which has been specified.
2. They will be released and moved from one machine to another until they are completed according to the machine sequences and machine times indicated in the order record.
3. Throughout simulation time such information as how many orders are waiting in each machine group waiting line, what their dollar values are, etc., will be established and retained.

B: The Shop's Resources

The orders which enter a shop represent the immediate task confronting the shop. The character of these orders affects the way in which the shop functions.

A full description of a job shop, however, must include information not only on the orders which make it function, but also on the resources with which it functions. These resources constitute the facilities of the shop which are brought into play to accomplish its work.

Time

One of the most basic resources for the accomplishment of any task is time. The simulator will therefore be concerned with the expenditure of time. A preliminary decision must be made regarding how much time should be simulated. The total amount of time through which the simulator is expected to track the shop functioning should be directly stipulated as input. Information in the program's output will describe what happened in the shop during this specified period of time. Now it may also be desirable to use this information to learn about conditions in the shop at intervals within the total time period covered by the simulator. For example, if the total time period to be simulated is 20 working days, it might be helpful to obtain reports which portray the shop at the end of each successive five-day period. This can be done. Cards 37 and 38 indicate respectively the number of days per period and the number of reporting periods desired.

In addition to determining the total time period to be simulated, it is necessary to locate this period more specifically in time. Just when does it begin? An input card is necessary to establish this. It could, of course, be considered "zero" day. If this were done in the present case, the simulation would end on day 20.

However, the beginning point must also be adjusted to scheduling calculations. If zero is used as a beginning point, it might turn out that the scheduled start date for an order would be negative. For example, if a due date of 5 were given for an order and scheduling information predicts 7 days in the shop for this order, backward scheduling would determine a -2 start date for the order. This could also be the result in the case of forward scheduling if late arrivals are specified (see page 23).

The program, however, will not handle negative numbers in this connection. For this reason the beginning point will normally be set at some time later than day "zero." Day 64 has been used in the illustrated card (39).

Once the total number of days to be simulated has been decided, a question which reflects the nature of time as a resource arises. How much time each day is to be allocated to work? It is necessary to advise the simulator of the work time scheme to be utilized in the shop. Cards 40 and 41 signify for a given shop that it operates during three shifts: eight hours on the first, six on the second and four on the third.

The simulation program does not, of course, really operate in terms of either days or hours as such, but merely in terms of time units. Any concrete significance which these units may have is accorded to them by the perspective of a user of the program. A time unit may be interpreted as a month or a day or an hour, etc., whichever is relevant. However, when both larger and smaller time units are involved, the relation between the two must be given to the program. Since days and hours are pertinent here, the fact that there are 24 hours to each day should be indicated (42). On the basis of this item of input, the program will assume that a longer time unit has elapsed for every 24 of the smaller ones. If the input card illustrated specified the number 60, then the longer time unit could be construed as an hour and the smaller as a minute.

As a resource time is a means to an end, and the important thing about it is what is accomplished with it. From this standpoint, all working hours are not necessarily equal.

The hour in which a shift ends and a new one begins tends to have a different value in this sense. For example, if a machine operation which normally took six hours would, by taking that amount of time, be completed half an hour before the end of a shift, there would be a tendency to prolong it so that it would end with the shift. The same tendency might prevent the beginning of a new job when the previous one is completed close to the end of shift time. Conversely, a job which would normally reach completion half an hour after the end of a shift would probably be speeded up so that its completion coincides with the end of the shift.

Such end-of-shift "psychology" can be accorded recognition in setting up the job shop model. The present IBM 704 program uses an "end-of-shift adjustment" card for this purpose. For example, if .5 hours were specified on this input card, any job which would normally be finished between half an hour before and half an hour after the end of a shift would actually be made to end at the end of the shift.

As in the case of all other descriptive input information, the point is correspondence with actual shop conditions. If it is felt that conditions in a given shop do not warrant such allowances for end-of-shift times, a zero may be punched in this input card. This has been done in the illustration shown (43).

Still another real condition which affects the value of the time resource, from the standpoint of what is accomplished with it, may exist in the shop. The speed with which machine operations are performed is not likely to be constant. This is especially true where comparatively new personnel and new operations are involved. If the user of the present application would like to have the simulator make an allowance for learning time when a man is beginning an operation, this can be specified in hours on a single card. As in the case of an end-of-shift allowance, zero hours may be stipulated (card 44).

Continuing to consider time in the perspective of what is accomplished with it, it is evident that as orders move through the job shop the segments of time during which machine operations are taking place should be distinguished from routing or transit time as such. When an operation called for in an order is completed at one machine group, the order must be physically moved to the next machine group in its work route. At least this much time must be allowed by the simulator between machine assignments. Transit times for each of the machine groups in the shop are designated on input cards. They represent, for each group, the average time in hours required to physically move a job to that group. The illustrated card specifies transit times (45).

The amount of time that it actually takes to move work depends both on the position of the machine group from which it is moved and on the position of the group to which it is moved. Since work routes vary, the average time specified actually represents a distribution of times around this average. It is possible to have the simulator select a time at random from this distribution in each instance of transit. If this is desired, a 1 should be punched in a Random Transit Card; if not, the code will be zero (46).

Machine Tools

While the job shop is operating, its machine tool resources are active. Machine tool facilities in a job shop consist of a definite number of machine groups distinguished from each other by type. Each machine group contains a certain number of machine tools of the same type. With regard to these facts there must be a one-to-one correspondence between the simulation model and the constitution of the shop being tested.

Let us assume, for example, that the shop to be simulated contains five machine tool groups — that is, five groups, each containing a number of machine tools which may be considered interchangeable from the standpoint of the operations they can perform. Group 1 might consist of ten lathes, group 2 might contain four drill presses, group 3 fourteen milling machines, group 4 ten grinders, and group 5 eight broaching machines.

The nature of the operations (grinding, drilling, etc.) is not significant for any output information expected from the simulator. Thus it is just the number of machine groups and machines which must be entered as input (47, 48).

In addition, the total number of machines, irrespective of the machine groups with which they are associated, must be specified (49). There cannot be less than one machine in each machine group.

In many instances these figures will precisely reflect the actually existing shop. Sometimes, however, they will be hypothetical. In these cases the point of the simulation will be, at least partially, to test the effect of modifying the actual machine tool constitution of the shop.

As has been emphasized previously, the simulation model is designed as a description of a functioning shop, whether this be a proposed shop or an existing shop. For this reason the static information regarding the number of machines it physically contains should be supplemented by a stipulation regarding how many are actually to be utilized on each shift of the shop's operation. It is possible that not all the machines of the shop are to be active on every shift.

With respect to this decision a separate input card must be provided for each of the machine groups of the shop. The number of entries of each card should correspond to the number of shifts of shop operation, and each entry specifies the number of machines operating for that machine group during that shift (50-54).

Just as this input information acknowledges that machines operate in time, the following two items of information derive from the fact that they also operate in space. The significant thing here is not the physical space occupied by the machines themselves but rather the amount of space available for accommodating waiting lines. As orders enter the shop and are moved from one machine group to another, the simulator places them in waiting lines associated with each machine group. A limit must be placed on the number of orders which each waiting line can support. If this is not done the simulator may create waiting lines which are unrealistic in that they exceed the space available for them. Card 55 establishes a maximum waiting line length for each machine group.

It is conceivable that in the course of simulation the arrival of an order will exceed the specified maximum waiting line for one or more machine groups. The program therefore provides for an overflow area. The largest number of orders which can be contained in this overflow area is determined in the assembly of the program. The simulator is apprised of the capacity of this auxiliary area indirectly. Input card 56 stipulates the maximum number of orders which can be accommodated on a waiting basis. This includes orders on waiting lines and also orders in the overflow area. The difference between the sum of the machine group waiting line capacities (given in the previous input card) and the number given in card 56 establishes the capacity of the overflow area serving all machine groups.

Manpower

Essentially the same type of information as is provided with respect to machine tools must be given for machine tool operators. Operating personnel are ordinarily classified according to machine tool skills. In the first place the number of different labor classes involved must be specified to the simulator (57). In addition, these labor classes must be associated with the machine groups which their respective skills enable them to operate. This information is necessary if the simulator is to be able to play foreman and assign men to the proper machines.

In order to establish the permissible associations, each of the machine groups of the shop must first be identified. Thus far, only the number of them has been specified. They may be identified directly by name (e.g., lathes, drills, etc.) or by convenient number designations which represent names. The latter alternative has been followed in the Machine Group Code Table card shown in Appendix I (58). Each of the numbers (1-5) represents the *name* of a machine group. In the simulation output the first name on this card will be referred to as machine group 1, the second as machine group 2, etc. Thus if the first name on the card were the number 4, this would be referred to as machine group 1 in the output reports.

Labor classes may now be associated with these machine designations. Let us assume for our particular example that the men in labor class 1 are capable of operating machine tools in machine groups 1 and 2; that labor class 2 can operate machines in machine groups 3 and 4; and labor class 3 operates only machines in machine group 5. These facts are conveniently stipulated in the illustrated card (59). The 2 punched in column 12 of this card signifies that labor class 1 is capable of operating the first two machine groups specified in the Machine Group Code Table card. The 2 punched in column 14 indicates that labor class 2 operates the second two machine groups specified in the Table card, etc. While the program provides for assigning a labor class to more than one machine group, the converse is not possible. Thus, two labor classes may not be assigned to the same machine group.

Just as in the case of machine groups, another important factor with respect to labor classes is the number of men from each class who are available for work during each shift of the shop's operation. A separate card is needed for each of the labor classes. Each card establishes for that class the number of operators available for each shift (60-62).

Cost of Resources

At this point the simulator has been provided with relevant descriptive information regarding time, machine tools and labor in the shop to be simulated. Fundamentally, what is given in this description is a pattern for the allocation of resources. It is, however, not the only pattern possible and the simulator may be required to explore the feasibility of others as well. Clearly, one of the fundamental criteria for evaluating alternative possibilities of resource allocation in a shop will be their cost relative to each other.

Since a significant point at which some of these costs may be measured is in connection with the increasing dollar value of work which is in progress, this dollar value must be accumulated for each order as it proceeds through its work route in the shop. Whenever there is an expenditure of time, labor and machine tools, the dollar value of each order is increased beyond its original raw material cost to the extent that these resources are costly. It is convenient to express these cost increments in terms of hourly-machine-labor rate. An hourly-machine-labor rate merely tells us how much it costs to have an operator operate a given machine for one hour. The rate, therefore, sums up the price of all three basic types of resource in a single increment of expenditure. It gives the simulator one relevant reference point for keeping track of what happens when the shop is functioning.

For the purposes of the present illustration, it will be assumed that machine labor rates are $10.00, $9.00, $12.00, $15.00 and $8.53 per hour for machine groups 1 through 5 respectively (63).

The simulator will apply these units of expenditure to each order as it progresses through the shop in simulated time along the work route indicated in the order record. The calculation involved may be illustrated in connection with order 1 of the sample order records shown on page 16.

Before any processing of order 1 has taken place, its dollar value is stipulated by the cost of the raw material involved. In the present case this is $310. The first machine tool operation to which this order is subject takes place in machine group 3. Since the hourly machine labor cost for this machine group, as given above, is $12, and the setup plus machining time specified in the order record is 7 hours, the value of the order at the completion of operation 1 (and therefore at the start of operation 2) will be $394 (that is, $7 \times \$12 + \$310$). Similarly, after operation 2, its dollar value will be $479 (that is, $8.5 \times \$10 + \$394$).

By performing these dollar-value calculations throughout the simulated time periods, the simulator can ultimately provide information on the amount of money tied up with in-process inventory in the shop. This amount is of course identical at any particular time to the full value of the products in the shop. It must, however, be clearly distinguished from the *cost* of in-process inventory. To speak of cost in this context implies the conception of money as a resource of the shop. If money were not tied up by in-process inventory, or indeed finished inventory, it could be put to some alternative use. The cost of having it locked in inventory is the value which would be gained if it were put to such alternative use.

The hypothetical problem of cost, therefore, is that of placing an estimate on the value which would be gained. This can be made a percentage of the total value which is tied up. In the present case an estimate of 10% has been made. This information is given to the simulator (64). The program is thereby in a position to provide output information on the cost of inventory.

The problem of estimating value which would be gained may be considered "hypothetical" only in the sense that for the most part there is no real alternative to tying up funds in inventory. Manufacturing takes time and during this time money will be tied up. What is significant, however, is that the amount tied up can vary depending on other conditions in the shop. The real virtue of obtaining information on inventory cost is that it establishes one convenient scale for measuring the advantage or disadvantage of variation in these other conditions. Wherever in-process inventory can be reduced by such a variation, money is indeed made available for alternative use.

Let us return now to the matter of hourly-machine-labor rates. In addition to enabling inventory cost calculations, once provision for calculating the cumulative value of orders has been made, other distinctions become feasible. It is possible now to distinguish between orders which have a relatively high, low or intermediate dollar value. Output information could then tell us something of the history of these three types as they progress through the shop.

Is it primarily high-value orders which are responsible for a heavy in-process inventory cost? Are comparatively high-value orders characteristically late? Are these the orders which tend to get unduly tied up in waiting lines? And so forth.

If the simulator is to distinguish between high-value, medium-value and low-value orders, it is necessary that the limits of these value categories be defined. This is to say that the user must specify the dollar transition point from a low to a medium to a high-value order. As will be seen, it is desirable to do this for each machine group separately.

It is not necessary to use precisely three value classes. Either one, two or three may be specified. Three represents the maximum with which the simulator can work. Three value classes have been selected for our example (65).

The illustration has also assumed that for machine group 1 any order which is at least $750 in value constitutes a high-value order; an order which is at least $450 and less than $750 will be considered a medium-value order; and any order with an accumulated value of between $1.00 and $450 will be a low-value order. Dollar transition points have been set up in precisely the same format for all five machine groups (66-70).

An additional reason for establishing dollar-value classes for orders may be a desire to consider accumulated value in a dispatching procedure. Priority of assignment to machine tools can then be made to depend, partly or entirely, upon the dollar-value class within which an order falls at the time of its assignment to a particular machine.

The possibility of using dollar-value class as a criterion for dispatching decisions points up the virtue of devising different value class limits for each machine group. Just why is this desirable?

Although, in a job shop, work order routes are typically varied, they are not necessarily randomly diverse. It is true enough that a given machine tool may be the third operational step for one order and the fifth step for another order. However, by reason of their nature, some machine tool operations are more likely to be called upon during either the earlier, middle or later phases of a work sequence. For example, a polishing operation will probably play its role in a final phase. This is because polishing, if it is done at all, is usually done last. Now, inasmuch as the dollar value of an order depends in part on how many machine tool operations it has already "digested," the average dollar value of orders arriving at machine group stations will fluctuate. It will vary with the probable position of the group in the general work route sequence. If we are dealing with a group of machines whose function it is to buff or polish, the average dollar value of arriving orders will, in all probability, be comparatively high.

Now under these conditions setting up dollar-value class limits independently of machine group would result in an unbalanced proportion of orders being classified as high by the time they arrive at the polishing operation. The limiting case would exist if all orders now fell within the bounds of the high-value class. To the degree that the limiting case is approached, the effectiveness of a dispatch rule which relies on the distribution of orders into dollar-value classes is lost. The problem can be eliminated when different value class limits are established for each machine group.

C: Dispatching Rules — Organization of Resources in Relation to Work

When work orders and shop resources have been described, two of the three fundamental elements necessary for a descriptive model of a functioning shop have been provided. The third major element consists of the controlling principle in terms of which facilities are applied to work. This governing principle has been termed the dispatch rule.

Throughout a shop's operation the problem of which customer order to assign to an available machine tool arises. The dispatch rule contains the criteria in terms of which such operating decisions are made. In this sense it exercises a "flow" discipline over orders which are eligible for access to a machine.

Since "desirable" shop conditions vary from company to company, it is not possible to recommend an optional dispatch rule in abstract. For this reason, the simulator has been designed to permit the use of any special dispatching formula which the user considers appropriate and advantageous for his shop situation. Either the one actually in use in a shop or some alternative possible rule may be provided for the simulator to work with. The mechanics of incorporating such rules into the 704 simulation program are fairly simple since the simulator was designed with this in mind. Information on how this can be done is available along with the program package. It is expected that a user will take full advantage of this important feature of the simulation program in order to investigate the effect different rules have on operating efficiency in the job shop.

Some discussion will, however, be devoted here to ten dispatch rules for which routines have been programmed and which have already been built into the 704 job shop simulation program. The discussion is abstract. The actual consequences of any of these rules for a particular shop are dependent upon many conditions in the shop.

The purpose of this discussion is primarily to illustrate briefly the kind of reasoning which enters into the construction of a rule and to suggest some sense of the range of alternatives. However, any one of the rules discussed may be employed by a user for the purposes of a particular simulation run. A single input card which specifies the DRN (Dispatch Rule Name) for the rule chosen is sufficient.

Rule 1: First Come, First Served

This is perhaps the simplest dispatch rule which can be used (71). It has already been discussed to some extent in Part I of this manual and will therefore only be identified here. Where it is used, the priority of an order is determined exclusively by the sequence of its arrival in a machine group waiting line. The order which arrives earliest in a particular waiting line is the first one assigned when one of the machines in the group has completed any previous work assigned to it.

Rule 2: First Come, First Served, within Dollar-Value Classes

This rule uses the same principle, but in a slightly more sophisticated way (72). It presupposes that the defining limits of dollar-value classes have been established for the simulator (see page 28). When this dispatch rule is employed and there are three value classes, high-value orders are accorded priority over all other orders, and medium-value orders are always assigned before low-value orders. Within each of these three categories, however, the first come, first served rule is followed.

Let us assume, for example, that all three value classes are represented in the orders awaiting assignment at a particular machine group dispatching station. The first to gain access to a machine tool by this rule will be that high-dollar-value order which arrived first in the waiting line for that machine group.

A rule such as this one might be of considerable advantage if comparatively high-value orders are heavily involved in in-process inventory. It will have the economically soothing effect of increasing the flow rate of the orders largely responsible for tying up funds in the shop.

There are, of course, many factors which in practice influence the speed with which an order moves or flows through a shop. Everything that happens to it takes a certain amount of time. Work moves from one machine station to another in order batches. This transit takes time. It takes time for an operator to remove a piece from a batch container and apply it to a machine tool. These expenditures of time are not affected by the present dispatch rule. Only time invested in waiting upon the availability of a machine tool is manipulated.

Taking the job shop as a whole, a certain amount of time must be expended in waiting. The attempt to create a machine tool and labor bank which eliminates any need for waiting is, as was pointed out in Part I, uneconomic. In this context time is being used as an alternative resource of the shop. It is an alternative to machine tools and labor. To use a dispatch rule which focuses on the dollar value of an order is in effect to choose a way in which the time resource is to be allocated. It is this rule which has been used in our example.

Rule 3: Minimum Slack Time per Remaining Operation

Another criterion for determining the relative priorities of waiting orders is the amount of "slack" time (i.e. nonproductive time) per operation, which remains between the "present" time and the due date on each order (73). It will be recalled that the scheduling procedure, in attaching dates to orders, took into consideration not only processing time for each operation, but also an estimated waiting time and possibly a transit time for each operation. As a result of the immediately preceding shop operations, some orders might be behind schedule and others ahead. This might be due to variations in the schedule resulting from variability of processing times, transit times and expected waiting times. In reaching a decision on order priorities, this dispatch rule simply compares the amount of remaining nonproductive time which has been allocated per operation for each order, and gives the highest priorities to those competing orders which have the least amount of this slack time per remaining operation.

For example, let us suppose that the relative priority of two orders is in question and that both have the same due date. Now if this date is April 30 and the date of the calculation is April 28, then 48 hours of a three-shift operation remain before the scheduled completion of both orders (due date – present time). In addition, let us assume that the processing time (machining + setup time) for all remaining operations of the first order is 8 hours. This means that, for this order, as much as 40 hours can be afforded or expended in waiting without preventing its actual completion by due date (due date – present time – remaining processing time). The significance of this 40 hours still depends on how many operations remain before completion. If there are very many, 40 hours may not be much, whereas if there are very few remaining, 40 hours may be a great deal. For this reason it is desirable to divide expendable time by the number of remaining machine tool operations.

The complete decision formula which yields the expendable time per remaining operation is therefore as follows: (due date – present time – remaining processing time) ÷ number of remaining operations.

Let us assume that this formula has been applied to both of the orders in our example, and that in the first case the expendable time per remaining operation is two hours whereas in the second case it is three hours. In this situation the dispatch rule will give priority of assignment to the first order.

The rule which we have been discussing might conceivably be made more sophisticated. It is evident upon closer examination, for example, that the amount of waiting time which an order can afford in the above sense is not really a sufficient criterion for determining the extent to which it may safely be kept waiting. The latter also depends on conditions in the shop at that time. In terms of these conditions even a comparatively great deal of time may not be sufficient.

For example, if, in the illustration above, the effective waiting lines for all subsequent machine tool operations for order 2 are very long compared with those for order 1, it might be well to give assignment precedence to order 2. Although it can afford more time than the first order, it will be obliged to follow a more difficult path and may therefore need more time. If this is true it is in the relatively worse position. The dispatching rule could be programmed to perform some such additional calculation and adjust its priority assignments accordingly.

It should be noted that this is the only rule which uses the "present" time as a factor in calculating priorities. This means that every time a new order is to be placed on a machine, the program automatically recalculates a priority code for every competing order. In the other nine rules a priority code is computed for each order once per operation when the order arrives on the waiting line, and subsequent events do not change this code.

Rule 4: Shortest Processing Time for Present Operation

Still another rule has already been built into the 704 simulation program (74). This type of decision results in the assignment of highest priority to the order with the shortest expected processing time (setup plus machining time) for the operation being currently considered. Conversely, the greatest amount of waiting time is allocated to those particular orders which have the longest expected processing time. These are, after all, the orders whose processing keeps other orders waiting longest. Since the orders which require the *least* amount of time on a machine tool are being moved through first, this procedure is likely to have the effect of minimizing the *number* of orders kept waiting.

This dispatch rule does not, of course, alter the final amount of processing time which must be invested in the total order mix. It simply institutes one possible strategy for expending this time. The scheme dictates that, during the earlier phases of shop work on a particular order mix, the tempo of passage through machine tool stations is high. The part is not an index to the whole, however, for by the same strategy the tempo will be comparatively slow during the latter phases of work. Those orders requiring a relatively great amount of processing time must from a practical standpoint ultimately also be sent through.

To mention only two considerations pertinent to the possible choice of this rule, the questions of in-process inventory and due date commitments are involved. It may be, for example, that the orders held in waiting the longest are precisely those with the highest dollar value. Indeed there is a certain likelihood of this, since orders having the highest dollar value are probably the orders which entail the greatest amounts of processing time. Their high dollar value has come about by virtue of the accumulation of hourly-machine-labor charges. If long processing jobs are at the same time high-dollar-value jobs, the danger exists that in-process inventory cost will mount beyond desirable limits. Similarly, these orders may not arrive at completion until well past their due date. In part, the issue depends upon whether these orders have a high raw material cost. It also hinges upon whether they require a relatively great number of machine tool operations or only a few, and whether only a few or most of these are long processing operations. A dispatching rule could be programmed to consider these and other factors.

Rule 5: Longest Processing Time for Present Operation

One way of fully exploring the consequences, for a particular shop, of a dispatching procedure such as the one just discussed is to simulate the same shop with its opposite rule. Rule 5 (the reverse of Rule 4) assigns the order with the longest expected processing time to the current operation before other orders (75).

By virtue of the dispatching decisions made through this rule, the orders which could move through machine stations fastest (shortest processing time) are kept waiting longest. Thus, barring the influence of factors which may act in a counter direction, the machine group stations would probably display a tendency to build up waiting lines quickly and, beyond a certain point, to whittle them down quickly. As has been indicated, it is possible that the character of the orders in the shop is such that the orders which demand the longest processing times are, with a relevant degree of consistency, the ones which accumulate the highest dollar value. In this situation the present rule may be considered a proper antidote to the previous one. Getting the high-dollar-value orders through quickly may have the effect of reducing in-process inventory.

There is no reason why the two opposite rules could not be combined in a single, more worldly-wise rule. The diplomacy of such an approach could be made to revolve about dollar-value class and proximity to due date. Feedback regarding these variables could trigger the switching between the two opposites. The shortest processing time rules would be operative wherever in-process inventory and lateness are not prohibitive. As soon as the latter problems become formidable, the program could shift to a longest processing time rule.

Rule 6: Earliest Planned Start Date

A fairly straightforward attack upon the specific problem of due date commitments, in so far as it can be handled by a dispatch rule, may be made in terms of the planned start dates of an order. The scheduling procedure discussed on pages 16-20 calculated a theoretical start date for each machine tool operation through which an order must proceed. Rule 6 simply examines the planned start dates of waiting orders and selects the order with the earliest one for assignment (76). If there is more than one order with the same start date for this operation, it selects from these at random.

In principle this rule elevates the shop schedule to the status of a blueprint. The schedule functions as a law of the shop and the rule operates in the direction of enforcing it. If one of the purposes of the simulation is to gain information with which to develop a more accurate basis for scheduling, the rule can be used to advantage. Inasmuch as it specifically attempts to adjust flow speeds in the shop to the schedule, if the final results still do not conform to it very closely the schedule is clearly unrealistic.

Rule 7: Earliest Planned Start Date within Dollar-Value Classes

The seventh preprogrammed dispatch rule is much the same as the sixth except that it incorporates an additional concern for the accumulated dollar value of an order (77). As in the case of Rule 2, all high-value orders will precede both medium- and low-value orders, and those in the medium-value class will precede the low-value orders. Within each of the dollar-value classes, the order with the earliest planned operation start date will be assigned first.

While the sixth rule occupies itself with meeting delivery commitments, it does not *insure* that they will be met. When due dates are unrealistic, orders may still be late. The seventh rule attempts to limit the extent to which funds are tied up in such late orders.

Rule 8: Earliest Due Date

This is a simple rule which incorporates a concern for meeting delivery commitments (78). It will be recalled that the due date of an order may have been specified originally in the order record, or it may have been calculated by the scheduling procedure. According to this rule the order which has the earliest indicated due date will be the first assigned to an available machine tool.

The theory of this approach leaves little to imagination. The closer an order is to its scheduled completion date, the less it should be held up.

One of the things the rule does not take into account is how close the order is to being completed already. If it has an earlier due date than any other order, it is assigned first regardless of whether it is now being considered in connection with the first operation it requires or the last. Thus an order which is due May 2 will be given precedence over one due May 4, despite the fact that it needs only two more hours of processing compared with two hundred for the May 4 order. Depending on the constitution of the order mix in the shop, the rule could easily fill the warehouse with finished stock inventory and turn the shop itself into a warehouse for in-process inventory.

Rule 9: Dollar Value and Start Date

The ninth dispatch rule represents a somewhat more subtle approach to the problems of both meeting delivery commitments and minimizing in-process inventory (79). Its first concern is to hasten the flow of orders behind schedule; the second is to avoid keeping relatively high-value orders waiting in line too long. The construction of the rule admits of a sensitive control of the relations between these two objectives.

The earliness or lateness of an order relative to its predetermined schedule is gaged by comparing the planned start date for the current operation with the current date (SD-T). Thus, if the planned start date is the 24th and the present date is the 23rd, the order has arrived earlier than expected at the dispatching station in question. If the present date is the 25th, the order is one day late. In this latter instance the result of the subtraction (SD-T) will be negative. The dispatch rule assures a positive result by adding a constant $(K + SD - T)$.

Basically, the accumulated dollar value of an order (VAL) is added to the result of the above computation, and assignment precedence is given to that order which then has the maximum index. The complete formula for establishing priorities is:

$$(A \cdot VAL)^B + [C(K + SD\text{-}T)]^D$$

A, B, C and D in this formula represent weights which can be assigned respectively to the variables of dollar value and schedule deviation. They afford considerable flexibility for decision regarding the relative importance of these two factors. Just exactly what these weights should be in the case of any given shop at any given time is a matter to be determined by simulation experiment and an examination of the alternative consequences.

In using this rule the user must specify that there are four variable factors. This is done on input card 81. This card might be:

```
ND PAR     DEC     4
```

In addition the four values involved must be given. Thus card 82 might be:

```
DPAR     DEC     .5,  1.,  .5,  1.
```

These are the values of A, B, C and D in this case. This is the only one of the ten built-in dispatch rules which requires this type of information.

The dispatch rule used in our example requires no "dispatch parameters" such as those discussed above. It is still, however, necessary to enter input cards 81 and 82. Card 81 in the Appendix states that there is "one" dispatch parameter. Card 82 defines this "one" parameter as zero.

Rule 10: Random Selection

This rule can be used as a convenient reference point for assessing the consequences of other rules (80). No special priorities are assigned to orders in this case. The order which is to be assigned to an available machine tool is selected at random from those waiting in line for that machine group.

In effect, this rule dispatches orders as though there were no coherent pattern apart from chance governing machine tool assignments. All orders are equal in that each has just as good a chance of being assigned first as any other. The approach does not distinguish between relevant and irrelevant characteristics of an order. In the light of this impartiality, the effect of this rule may be used as a basis for comparison with other dispatching disciplines. By comparing results, the advantages and disadvantages of the various rules show up in relief. Essentially, therefore, the Random Selection rule may be considered an exploratory device.

PART III

The Development of Simulation Information for Decision

Once an adequate description has been made of the orders entering a shop, the resources of the shop, and the dispatching principle in terms of which the latter are applied to the former, the job shop model is complete. At this point the simulator can proceed to simulate the job shop under the control of the descriptive model.

The Simulation Procedure

In the course of simulation, the simulator program operates in a manner parallel to the activities in an actual shop. This process will be described briefly.

First of all the simulator brings into the high speed storage of the computer those orders which have been designated for "initialization" purposes. These are the orders which are presumed to be in-process at the very outset, and for which a relative state of completion may have to be calculated. The procedure for accomplishing this has been discussed previously (pages 22-23). Whether the state of completion is calculated or already given on the order tape, the next step is to place these orders at the indicated machine groups.

Work orders corresponding to the daily releases to the shop may now be brought in. These orders may be presumed to have already been processed by the scheduling portion of the program. A general procedure for each set of such releases can then be followed. On the basis of the routing information contained in each order record, pertinent information regarding the order record is moved to a storage location which corresponds to the waiting line associated with a particular machine group.

The program then proceeds to examine each machine group as well as the storage locations assigned to manpower. In effect it is, at this point, looking for idle men and machine tools. When this combination is found, the priority ratings of the orders in the waiting line for that machine group are compared. The order with the highest priority is assigned to the idle man and machine. Priorities are determined on the basis of the dispatch rule in use. If an idle man can be assigned to more than one machine group and a machine is available in each, the high-priority orders in each waiting line are evaluated to determine the one with the highest priority.

An order which has been assigned to a machine tool is assumed to remain there for the number of hours specified as processing time in the order record. If processing times are to be adjusted, it is the adjusted times which are used here. Variability around the adjusted times may also be taken into account. The number of hours involved is placed in the storage positions allotted to the machine tool.

Repetitive tests are made to determine whether the operation is finished. When it has been completed, the machine is available for further assignment. The order previously assigned to it must now be moved to the waiting line associated with the next machine group on its route.

Throughout the assignment procedure the release pattern specified in the program input is considered in determining which orders are eligible for assignment. In addition, specified transit times are taken into account for establishing the moment of eligibility.

When in the course of time a counter indicates that a work shift has come to an end, all orders are removed from the machines. The end-of-shift allowance is considered (see page 25). The orders are placed into the waiting lines for their machine groups along with any other orders which may have been waiting. Orders which were in process of being machined are assigned the very highest priority to distinguish them from the orders which have been waiting.

New orders may be brought into the system and placed into the proper waiting lines. They will not be eligible for assignment until the proper transit times have elapsed. Orders remaining from a previous shift may, however, be assigned immediately on the basis of their priority. Those which were being machined at the close of a previous shift are thereby assigned first. Remaining orders are subject to the flow discipline of the dispatch rule.

Information respecting the state of the system is accumulated and stored continuously. This is the information which will be contained in the output reports. Among other things it will include such relevant matters as the number of orders completed and the time of their completion relative to their due dates. It will include information on idle labor and/or machines. The current dollar value of orders in the shop, which has been calculated by applying hourly-machine-labor rates to processing times, will be recorded. In addition, the average amount of waiting time for orders in each of the machine group waiting lines will be noted.

As was pointed out at the beginning of this section, in its processing the simulator enacts the behavior of a real shop in all respects deemed relevant to management decision regarding a real shop. The output of the simulator summarizes the conditions which have existed in the shop during the simulated periods. It is to an examination of this output that we must turn now.

Simulation Output

Generally speaking, the output of the simulation application confronts management with the consequences which follow from having set up a shop in a specific way. The consequences are those which are pertinent to the range of decisions open to management regarding the shop. In view of this it is possible to compare the consequences of one possible set of decisions with those which follow when other decisions are made. It is in terms of an evaluation of consequences that a specific course of action in a real shop can be decided upon.

The evaluation of consequences, the rating of them as good or bad, appalling or wonderful, the balancing of some of the consequences against others, will presumably be made in terms of the goals which management is striving to achieve. The job shop simulation application does not establish what these goals are. It does not determine the purposes of an organization. This remains a management responsibility.

For example, simulation output will stipulate how many orders are late and just how late they are. However, it will not state whether these consequences fall above or below some threshold of acceptability. The latter judgment draws upon management's knowledge of its business.

In this sense management's prerogative to establish policy is not pre-empted. Indeed it is precisely this responsibility which is emphasized. An optimum job shop is not created for management by simulation. What simulation does do is to provide management with the information it needs to create such a shop. As the first section of this manual emphasized, the basic stumbling block has always been the difficulty of obtaining precise and sufficient information.

Output Reports

All together the simulation program produces eight reports. Each of these will be discussed below. The specific instances of each report shown and discussed were produced by a simulation run. This simulation was controlled by the specific input information shown in Appendix I.

1. Identification Report. The first report shown below requires very little in the way of explanation. It simply consists of (1) a printout of identifying information, and (2) a summary of certain basic information describing the shop. This is merely some of the data which has been given to the simulator on in-

ABC CORPORATION
JOB SHOP SIMULATION
SAMPLE ANALYSIS BY IBM 704

DATE			AUG 30, 1959
SIMULATION IDENTIFICATION NUMBER			SIM1
ORDER SCHEDULING METHOD			SR01
SHOP DISPATCHING METHOD			FCFSV
NUMBER OF REPORTING PERIODS			4
NUMBER OF DAYS PER PERIOD			5
ORDER TAPE IDENTIFICATION			SYNTH
INITIAL SHOP LOAD (ORDERS)			40
AVERAGE DAILY SHOP LOAD (ORDERS)			35.00
NUMBER OF SHIFTS			3
WORKING HOURS PER SHIFT	8.00	6.00	4.00

LABOR CLASS INFORMATION

CLASS	NUMBER OF MEN PER SHIFT		
1	14	10	8
2	24	20	18
3	8	6	4

MACHINE GROUP INFORMATION

CLASS	GROUP	OPERATING MACHINES PER SHIFT		
1	1	10	8	6
1	2	4	4	3
2	3	14	12	12
2	4	10	10	10
3	5	8	6	4

put cards. Among other things indicated are the number of periods for which reports will be printed, the number of days in each of these periods, the initial load of orders in the shop and the average daily shop load, machine group and labor information, etc.

The basic information describing the shop has already been provided to the simulator for the purposes of simulation. There remains, however, an additional category of purely identifying information which has not as yet been described. This information is not essential for the simulation as such but is provided only so that it may appear on the first output report. It is convenient for distinguishing the set of reports acquired as a result of one simulation experiment from those associated with other such experiments.

Six items of information belonging in this category are shown in Appendix I (page 43). Each requires a separate input card. The first three items shown are merely name cards. Whatever is punched in columns 13-72 of these cards will be printed on the first output report exactly as punched. It will make up the first three heading lines of this report. Any information identifying the shop or the specific purpose of the simulation run may be punched in these cards.

The next card makes provision for the date of the simulation run. The final two cards shown identify respectively the order tape which has been used in the simulation and the specific deck of input cards utilized.

The first report containing all of this identifying information is printed only once for the entire simulation. The remaining seven reports are printed for each reporting period. Those shown on the following pages represent the reports for the second period. Reports for all periods differ in content, of course, but are the same with respect to structure and the type of information they contain.

2. *Load Analysis Report.* The second report consists of an analysis of the hours of machine loads which are to be *expected* in the shop within the time period covered by the report. The load expectations expressed in the report derive from the schedule plan which has been developed for the orders prior to simulation. They do *not* indicate, therefore, what has actually happened during the simulation. It only informs us that in terms of the time schedule which has been set up beforehand, such and such machine groups may be expected to be called into play for X number of hours during, say, the second five days of processing. It is possible that the number of hours required by the schedule will exceed the actual capacity of the shop.

As may be seen from the example of this report shown below, it contains not only information regarding a particular reporting period but also year-to-date information. The latter consists of cumulative information for this period plus all previous periods. In the present illustration it therefore displays machine loads over the time covered by the first two reporting periods.

The first line of figures in the report contains information about machine group 1. The number of hours scheduled for this machine group is shown here as divided between high (H), medium (M) and low (L) value orders. In order to differentiate machine loads by dollar-value classes, the program has, of course, had to trace the increasing dollar value of each scheduled order according to the way it will proceed along its processing route. The first figure on the machine group 1 line means that we may expect 144 hours of processing time to be needed on this machine group. This time will be needed during the period covered by the report for orders which at that time may be expected to fall within the specified limits of the high-dollar-value class. Similarly, 111 is the number of scheduled hours for medium-value

PERIODIC OUTPUT

LOAD ANALYSIS

PERIOD 2

	THIS PERIOD						YEAR TO DATE					
MACH GROUP	LOAD H HRS	LOAD M HRS	LOAD L HRS	LOAD T HRS	AVAIL CAPAC HRS	LOAD CAP	LOAD H HRS	LOAD M HRS	LOAD L HRS	LOAD T HRS	AVAIL CAPAC HRS	LOAD CAP
1	144	111	339	594	760	.782	213	171	671	1055	1520	.694
2	35	61	190	287	340	.844	63	114	393	570	680	.838
3	194	406	423	1023	1160	.882	295	686	869	1851	2320	.798
4	91	223	491	805	900	.895	144	471	901	1517	1800	.843
5	58	81	277	415	580	.716	110	147	642	899	1160	.775
	523	882	1720	3125	3740	.836	825	1589	3477	5890	7480	.787

orders during period 2 on this machine group, and so on for low-dollar-value orders. The corresponding figures for the remaining machine groups are given in the lines below the first.

The column of figures under "T" shows the total number of hours of processing load for each machine group irrespective of the dollar-value consideration. These figures may be compared with the shop's machine-load capacity which is given for each machine group in the column immediately to the right. The capacity figures are derived from input information regarding the number of machines in each machine group and the number of hours in each shift.

The results of a comparison are in fact expressed in the next column. This "Load Cap" column is arrived at by dividing the total scheduled load ("T") by the available capacity ("Capacity Hours"). The figures in this column therefore constitute indices of the extent to which the scheduled load has fallen below or exceeded the shop's available capacity.

3. Shop Performance Report. The next report is illustrated below. The previous report has clarified the implications of the schedule plan with respect to the processing time loads. The present report, on the other hand, records the actually utilized capacity of each machine group *after* the simulation has been performed. If the lead times and processing times which have been allowed for scheduling purposes closely correspond to the waiting times and processing times which actually develop in the shop, the two reports may resemble each other quite closely.

As in the previous case the Shop Performance report provides both year-to-date and specific time period figures. Within each of these divisions of the report, available machine capacity (the same as in the load analysis report) is contrasted with actually utilized capacity. The available capacity for each machine group is shown in the first column and the actually utilized capacity is indicated in the second column. What this information implies in the way of idle time and percentage of capacity utilized is spelled out in the next two columns.

Unlike the case of load analysis, it is evident that in this report the number of hours of actual processing time on a given machine group cannot exceed the number of possible or available hours. Essentially what this report will tell us, therefore, is how close we have come to taking full advantage of the shop's machine capacity. Was the flow of actual work in the shop such that some machines were reduced to idleness? If so, which machine areas were involved, and during which time intervals were they involved?

It was pointed out at the very beginning of this manual that one of management's goals is the best use of the shop's machine resources. Information in the Shop Performance report establishes the extent to which this goal would be approximated by one possible shop.

To the extent that the shop falls short of what, in a given case, management might consider adequate performance, certain hypotheses might be entertained. Perhaps machines in a particular machine-group area have been unduly idle because of an unbalanced waiting line situation in some other machine area. This in turn might conceivably be a result of the number of machines available in this other area. It might be the effect of the particular dispatch rule employed.

Hypotheses such as these may be confirmed or refuted as the figures for different time periods are examined and compared. In addition they may be borne out by information contained in other reports. If an improvement in shop performance is of practical importance in a given case, subsequent simulation runs can always provide the test of any hypothesis.

SHOP PERFORMANCE

PERIOD 2

			THIS PERIOD				YEAR TO DATE			
LABOR CLASS	MACH GROUP	MACHINE GROUP DESCRIPTION	AVAIL CAPAC HRS	UTIL CAPAC HRS	IDLE TOTAL HRS	UTIL CAPAC	AVAIL CAPAC HRS	UTIL CAPAC HRS	IDLE TOTAL HRS	UTIL CAPAC
1	1		760	564	196	.743	1520	1077	443	.703
1	2		340	283	57	.833	680	590	90	.867
2	3		1160	973	187	.839	2320	1893	427	.816
2	4		900	786	114	.873	1800	1511	289	.839
3	5		580	421	159	.726	1160	932	228	.803
			3740	3028	712	.810	7480	6003	1477	.803

LABOR UTILIZATION

PERIOD 2

CLASS	THIS PERIOD AVAIL MAN HRS	UTILIZ MAN HRS	UTILIZ	YEAR TO DATE AVAIL MAN HRS	UTILIZ MAN HRS	UTILIZ
1	1020	851	.834	2040	1652	.810
2	1920	1789	.932	3840	3454	.899
3	580	388	.669	1160	897	.773
	3520	3028	.860	7040	6003	.853

4. Labor Utilization Report. This report, shown above, provides essentially the same type of information with respect to the manpower resources as the Shop Performance report does for machine resources. The total number of available man-hours in each labor class is indicated in the first column. The number of these hours actually used and the percentage of utilization is recorded in columns 2 and 3 respectively. The information is given on a period as well as a year-to-date basis.

Essentially it is the adequacy of the shop's manpower resource which is revealed in this report. It may turn out that the available man-hours in a specific labor class were completely or almost completely used. This might give rise to the suspicion that an intermittent unavailability of man-hours has been responsible for some idle machine time. Checking the relationship between labor class and machine group in the previously described Shop Performance report may support this hypothesis.

5., 6. Analysis of Queues. The next two reports present an analysis of the waiting line situations which have developed in the simulated shop. Both reports contain essentially the same information, except that the first covers a specific reporting period whereas the second provides the year-to-date figures.

The basic categories of these reports are the dollar-value classes of orders and the various machine groups

ANALYSIS OF QUEUES

PERIOD 2 — THIS PERIOD

MACH GROUP	ARRIV H NO	ARRIV M NO	ARRIV L NO	ARRIV T NO	DEPART H NO	DEPART M NO	DEPART L NO	DEPART T NO	Q-TIME H DAYS	Q-TIME M DAYS	Q-TIME L DAYS	Q-TIME T DAYS	AVER Q H NO	AVER Q M NO	AVER Q L NO	AVER Q T NO
1	22	19	58	99	20	19	57	96	.016	.054	.045	.041	.1	.3	.8	1.1
2	37	54	155	246	37	52	150	239	.033	.037	.136	.099	.3	.4	5.4	6.2
3	34	83	83	200	33	80	84	197	.022	.024	.066	.041	.2	.5	1.2	1.9
4	30	57	127	214	32	57	119	208	.022	.030	.112	.076	.1	.4	4.4	4.9
5	38	37	105	180	38	40	113	191	.031	.034	.061	.049	.3	.2	1.3	1.8
	161	250	528	939	160	248	523	931	.026	.032	.093	.065	.9	1.8	13.1	15.9

ANALYSIS OF QUEUES

PERIOD 2 — YEAR TO DATE

MACH GROUP	ARRIV H NO	ARRIV M NO	ARRIV L NO	ARRIV T NO	DEPART H NO	DEPART M NO	DEPART L NO	DEPART T NO	Q-TIME H DAYS	Q-TIME M DAYS	Q-TIME L DAYS	Q-TIME T DAYS	AVER Q H NO	AVER Q M NO	AVER Q L NO	AVER Q T NO
1	37	40	131	208	34	38	129	201	.014	.032	.029	.027	.1	.1	.5	.7
2	62	107	333	502	61	105	327	493	.032	.035	.125	.094	.2	.4	5.4	6.0
3	57	158	189	404	55	152	185	392	.019	.021	.044	.032	.1	.4	1.1	1.6
4	50	131	254	435	50	130	240	420	.018	.024	.077	.053	.1	.4	2.8	3.3
5	61	73	257	391	61	73	255	389	.032	.039	.133	.100	.2	.3	4.4	5.0
	267	509	1164	1940	261	498	1136	1895	.024	.028	.092	.066	.7	1.7	14.2	16.6

on which they are processed. The first eight columns of the report show the number of orders arriving and departing from the waiting line associated with each machine group. Both arrivals and departures are differentiated by high-, medium-, and low-dollar-value class. The total number of orders arriving and departing, irrespective of value class, is also shown in columns 4 and 8 respectively.

The remaining columns of the report are devoted to recording the *amount* of time which orders have had to wait in each "Q" or waiting line, and the *length* of the line within which they have had to wait. First, the average amount of waiting time in 24-hour days is given. This applies to orders which have left the machine group involved during the time period covered. Thus, high-value orders have waited an average of .016 days in the waiting line for machine group 1 before leaving. In this waiting line .1 has been the average *number* of high-value orders.

There is one difference between the period report and the year-to-date report. In addition to including previous periods, the latter also includes in each waiting line, on the arrivals side, the orders which were presumed to be in progress in the shop at the time the simulation began.

In effect, the above reports represent views of the patterns by which orders may be expected to flow through the shop toward their completion. The occurrences recorded are immediately relevant to such concerns as the amount of money tied up by in-process inventory. These funds are embodied in orders kept waiting for access to machines. The waiting lines are also in a cause-and-effect relation to the order completion date picture. If the shop cannot meet its delivery date commitments, the reasons may be revealed in the waiting line analysis.

A point of interest in the latter connection is a comparison between average waiting times for each machine group and the lead times which were allowed for scheduling purposes. It was the function of lead-time estimates to predict actual waiting times. The correspondence of actual completion dates with scheduled completion dates depends, to a considerable degree, on how accurately the prediction of waiting times was made.

If judgment concludes that the simulated shop is a good one, as revealed and assessed in terms of all of its consequences, then the waiting times shown in this report may be used to revise lead-time estimates in scheduling. It is quite possible, however, that at some points the waiting lines shown in the report are unacceptable for the real shop. In this case the report may provide clues regarding the source of the difficulty.

It is possible that the length of a given waiting line consistently increases through the successive time periods covered in the reports. A machine tool-operator combination which is consistently overloaded will develop waiting lines which grow indefinitely. The Shop Performance and Labor Utilization reports can be compared on this point. In this situation the imbalance might be corrected through the addition of machine tools or manpower. Alternatively, an examination of the waiting line situation with respect to dollar-value class might justify consideration of a dispatch rule which would allocate waiting time with reference to this factor. In point of fact, once the time and place of the problem is located, any number of solutions may suggest themselves. Later experiments with simulation would provide the test of validity.

7. *Tabulation of Completions.* The ability of the simulated shop to meet customer delivery commitments is clearly displayed in the next report shown. This report is again divided on the basis of current period and cumulative year-to-date information. The latter category includes all orders completed since the beginning of the simulation. Within each of these categories a distinction is made between orders which were released for processing "on time" and those which were released after their scheduled start date. In each of these subcategories the number of orders completed is shown in relation to their scheduled dates of completion as well as in relation to their dollar-value class. Thus, the number of high-value orders completed on the proposed day may be compared with the number completed earlier and later.

Various patterns of completion may show up in this report. It could, for example, turn out that almost all orders are completed on or before their scheduled due date. Sixty percent of these might, however, be over ten days early. This constitutes either a delightful result or a serious problem, depending partly on whether or not the organization must store early completions until the originally promised delivery date. If it is in fact under this costly necessity, a completion pattern might be preferable wherein fewer orders are completed on time but the rest are early or late only by a comparatively low number of days.

When orders are late the problem arises of assessing cost in loss of customer confidence, with the bearing this has on future sales. Judgment regarding these matters is dependent upon a knowledge of the specific context. In this connection it is management's task not only to balance these consequences against others for the sake of evaluating the total shop, but also to decide just what the ultimate consequences are.

TABULATION OF COMPLETIONS

PERIOD 2	THIS PERIOD												YEAR TO DATE											
EARLY OR LATE	ON TIME RELEASED				LATE RELEASED				ALL ORDERS				ON TIME RELEASED				LATE RELEASED				ALL ORDERS			
	H	M	L	T	H	M	L	T	H	M	L	T	H	M	L	T	H	M	L	T	H	M	L	T
	NUMBER OF ORDERS COMPLETED												NUMBER OF ORDERS COMPLETED											
E31+UP																								
E26-30																								
E21-25																								
E16-20																								
E11-15																								
E 6-10														1		1						1		1
E 5														1		1						1		1
E 4																								
E 3	2			2					2			2	4	4		8					4	4		8
E 2	3	8	4	15					3	8	4	15	5	14	10	29					5	14	10	29
E 1	17	22	38	77					17	22	38	77	26	51	82	159					26	51	82	159
0	4	13	56	73					4	13	56	73	12	25	110	147					12	25	110	147
L 1																								
L 2																								
L 3																								
L 4																								
L 5																								
L 6-10																								
L11-15																								
L16-20																								
L21-25																								
L26-30																								
L31+UP																								

If final evaluation discards the simulated shop as unsatisfactory, the information in the completions report may suggest corrective measures. In particular, the separation of completed orders by dollar value category affords insight into the possible virtue of modifying a dispatch rule. Perhaps a disproportionate number of late orders have been of high value. Such a result would yield to an altered priority assignment strategy. Still another alternative which might be considered is the possibility of changing the dollar-value-class limits. This would not change the dispatching principle, but it would change its results.

8. Inventory Carrying Cost Report. Between the time orders enter a shop and the time they emerge as completed, they are increasingly acquiring value. This value is, however, merely residual; it is not active value in the sense of functioning to accomplish anything. It can yield no return to the organization until the orders have been completed.

INVENTORY CARRYING COST IN $ PER ANNUM
INTEREST RATE 10.0 PERCENT

PERIOD 2	THIS PERIOD						YEAR TO DATE					
GROUP	Q(H)	Q(M)	Q(L)	Q(T)	M(T)	RATIO	Q(H)	Q(M)	Q(L)	Q(T)	M(T)	RATIO
1	11	17	15	44	360	.121	8	10	12	29	312	.094
2	43	34	121	199	148	1.338	34	32	124	189	148	1.280
3	25	32	24	82	679	.120	18	27	22	67	600	.112
4	24	29	89	143	462	.309	16	25	59	100	421	.238
5	41	18	36	95	243	.392	32	24	114	170	238	.714
TOTAL	145	131	287	562	1891	.297	108	118	330	557	1719	.324
EARLY JOBS	722	475	286	1483	1891	.784	577	690	297	1564	1719	.910
GRAND TOTAL	867	605	572	2044	1891	1.081	685	808	627	2120	1719	1.234

In principle, the cost of this investment consists in what might be gained if the capital temporarily tied up in in-process inventory were active. It is convenient to calculate this cost in terms of an interest rate. For this case it has been assumed that the cost of inventory, while it is in process within the shop, may be gaged by applying a 10% per annum interest rate. The Inventory Carrying Cost Report presents the results of this calculation. It is therefore the fundamental basis for deciding whether in-process inventory cost has remained within acceptable limits.

Of the time which orders have been in the shop, some part has been taken up with actual machining and some part has been spent in waiting lines. The Inventory Carrying Cost Report distinguishes between the costs included in these respective portions of total simulation time. The first three columns of the report specify carrying costs (dollar value x interest rate) for orders during the time that they have been waiting in line. Dollar-value class has been the only basis for distinction here. The interest rate has been applied to the *average* dollar value of each order in each class during the period covered in the report.

The fourth column of the report states the total carrying cost of all orders for the time they were waiting in line. The fifth column, on the other hand, stipulates carrying costs for orders which were on machines during the period covered. The ratio column expresses the relation between total waiting time carrying cost (column 4) and machining time carrying cost (column 5).

The ratio figures are of special significance in evaluating the effectiveness of the shop. Very little can be done about carrying costs which accrue while orders are being processed. To be sure, the prospect of faster machine tools may always be entertained. But this is a question with many ramifications in the economics of the shop. Given machines with certain processing speed capacities, the problem of carrying costs settles in the area of waiting times. It is in connection with this area that dispatch rules offer a handle for control. All other things being equal, the shop which minimizes the proportion of carrying cost attributable to waiting lines is the more profitable shop. In other words, the smaller the ratio figure, the better.

In considering alternative dispatching procedures, the comparative magnitudes of cost involved in high-, medium- and low-value orders is of course also significant.

Another calculation made for the purpose of this report is the determination of a carrying cost for orders completed early. Unlike late orders, a cost for early orders can be computed on essentially the same basis as for in-process orders. However, by virtue of the fact that some organizations can ship orders immediately upon completion, regardless of due date, there may be no actual carrying costs for early completions. It is for this reason that a separate tabulation has been made for these orders.

The year-to-date information in this report is not cumulative in the same sense as it is for the reports previously discussed. It is simply the same calculation of costs taken over a longer period of time. The cost figures on the year-to-date side might easily be less than those on the period side. This is essentially because orders are continually entering and leaving the shop. Thus, average dollar values, taken over a relatively long period of time, could be less than they are for some portion of this time.

It should be evident at this point that all of the reports discussed above are related to each other. Basically, their combined purpose is to display, in terms relevant to management's goals, the consequences of setting up a job shop in a given way. As has been persistently emphasized throughout this manual, a job shop production system is extremely complex. The output of the job shop simulation application reflects all the interdependencies involved.

The information contained in the above reports provides management with an effective tool for controlling manufacturing processes in the service of management goals. If a new product is introduced or the character of the order mix changes, the consequences of the change may be speedily assessed. The effect of any practical measures which could be taken can similarly be accurately gaged. Alternative possibilities may be compared with each other and decisions for a real shop can be based on the balance of advantages which simulation can show any one of them to have. In this sense the job shop simulation application constitutes a laboratory in which experimentation can be carried out toward an optimum job shop.

Appendix I

List of All Input Cards

A: Input Concerning Orders

CARD NO.	CARD COLUMNS 1-6	8-10	12-72	TEXT PAGES	EXPLANATION
1	LOWMC	DEC	1.00	16	Lowest Material Cost
2	LOWSU	DEC	0.01	16	Lowest Setup Time
3	LOWMT	DEC	0.01	16	Lowest Machine Time
4	SRN	BCD	1SR01	19	Scheduling Rule Name
5	UGDD	DEC	0	19	Use Given Due Date
6	SPAR	DEC	.054, .076, .147	19	Lead Times in 24-Hour Days for Each Machine Group, by Value Categories
7		DEC	.082, .105, .163	19	
8		DEC	.036, .057, .076	19	
9		DEC	.154, .263, .271	19	
10		DEC	.025, .026, .032	19	
11	NSPAR	DEC	15	20	Number of Scheduling Parameters
12	SCHP	DEC	1.0, 0.0	20	Coefficients of the Scheduling Polynomials
13	SCHPD	DEC	1	20	Degree of Scheduling Polynomials
14	NSCHP	DEC	1	20	Number of Scheduling Polynomials
15	SCHPG	DEC	1, 1, 1, 1, 1	20	Assignment of Scheduling Polynomials to Machine Tool Groups
16	SCHPL	DEC	0	21	Scheduling Polynomials with Logarithms
17	SIMP	DEC	1.0, 0.0	21	Coefficients of the Simulation Polynomials
18	SIMPD	DEC	1	21	Degree of Simulation Polynomials
19	NSIMP	DEC	1	21	Number of Simulation Polynomials
20	SIMPG	DEC	1, 1, 1, 1, 1	21	Assignment of Simulation Polynomials to Machine Tool Groups
21	SIMPL	DEC	0	21	Simulation Polynomials with Logarithms
22	RNSCH	DEC	321528735	21	Scheduling Random Number Seed
23	RNSIM	DEC	321528735	21	Simulation Random Number Seed
24	VARP	DEC	0.0, 0.0	22	Coefficients of Variability Polynomials
25	VARPD	DEC	1	22	Degree of Variability Polynomials
26	NVARP	DEC	1	22	Number of Variability Polynomials
27	VARPG	DEC	1, 1, 1, 1, 1	22	Assignment of Variability Polynomials to Machine Tool Groups.
28	VARPL	DEC	0	22	Variability Polynomials with Logarithms
29	ARRFR	DEC	35.0	22	Arrival Frequency of Orders
30	RAO	DEC	0	22	Random Arrival of Orders
31	NF	DEC	1000	22	Dimension of Order File
32	RELFR	DEC	1.0	22	Release Frequency of Orders
33	INSLD	DEC	40.0	23	Initial Shop Order Load
34	RIN	DEC	1	23	Random Initialization
35	LARRF	DEC	0.00	23	Late Arrival Frequency
36	LARRT	DEC	0.00	23	Late Arrival Time

B: Input Concerning Shop Resources

CARD NO.	CARD COLUMNS 1-6	8-10	12-72	TEXT PAGES	EXPLANATION
37	ND	DEC	5	24	Number of Days per Period
38	NP	DEC	4	24	Number of Reporting Periods
39	BEG	DEC	64	24	Beginning of Simulation
40	NS	DEC	3	24	Number of Shifts
41	HS	DEC	8.0, 6.0, 4.0	24	Working Hours per Shift
42	HPD	DEC	24	25	Hours-to-Days Relationship
43	ESA	DEC	0	25	End-of-Shift Adjustment
44	STA	DEC	0	25	Start-up Allowance
45	TRANS	DEC	.2, .5, .3, .2, .6	25	Transit Time for Each Machine Group
46	RTR	DEC	0	25	Random Transit Time
47	NG	DEC	5	25	Number of Machine Groups
48	MG	DEC	10, 4, 14, 10, 8	25	Number of Machines in Each Machine Group
49	NM	DEC	46	25	Number of Machines
50	MGS	DEC	10, 8, 6	26	Number of Machines Operating on Each Shift for Each Machine Group
51		DEC	4, 4, 3	26	
52		DEC	14, 12, 12	26	
53		DEC	10, 10, 10	26	
54		DEC	8, 6, 4	26	
55	QG	DEC	20, 20, 20, 20, 20	26	Queue Length per Machine Group
56	NQ	DEC	400	26	Total Dimension of Queues
57	NC	DEC	3	26	Number of Labor Classes
58	GCTBL	DEC	1, 2, 3, 4, 5	26	Machine Group Code Table
59	GC	DEC	2, 2, 1	26	Machine Groups Handled by Each of the Labor Classes
60	MCS	DEC	14, 10, 8	26	Number of Men Working on Each Shift for Each Labor Class
61		DEC	24, 20, 18	26	
62		DEC	8, 6, 4	26	
63	RG	DEC	10.00, 9.00, 12.00, 15.00, 8.53	27	Hourly-Machine-Labor Rates for Each Machine Group
64	INTER	DEC	10.0	27	Interest Rate
65	NV	DEC	3	28	Number of Value Classes
66	HMLVL	DEC	750.00, 450.00, 1.00	28	Lower Dollar Limits of High-, Medium-, and Low-Value Classes for Each Machine Group
67		DEC	850.00, 500.00, 1.00	28	
68		DEC	900.00, 350.00, 1.00	28	
69		DEC	1000.00, 400.00, 1.00	28	
70		DEC	800.00, 500.00, 1.00	28	

C: Alternative Input Concerning Dispatch Rules

CARD NO.	CARD COLUMNS 1-6	8-10	12-72	TEXT PAGES	EXPLANATION
71	DRN	BCD	1FCFS	29	First Come, First Served
72	DRN	BCD	1FCFSV	29	First Come, First Served, within Dollar Value Classes
73	DRN	BCD	1MINSOP	29	Minimum Slack Time per Remaining Operation
74	DRN	BCD	1MINPRT	30	Shortest Processing Time for Present Operation
75	DRN	BCD	1MAXPRT	30	Longest Processing Time for Present Operation
76	DRN	BCD	1MINSD	31	Minimum Start Date
77	DRN	BCD	1MINSDV	31	Minimum Start Date within Dollar Value Classes
78	DRN	BCD	1MINDD	31	Earliest Due Date
79	DRN	BCD	(VAL + SD)	31	Dollar Value and Start Date
80	DRN	BCD	1RANDOM	32	Random Selection
81	NDPAR	DEC	1	32	Number of Dispatch Parameters
82	DPAR	DEC	0	32	Dispatch Parameters

D: Input Concerning Identification

CARD NO.	CARD COLUMNS 1-6	8-10	12-72	TEXT PAGES	EXPLANATION
83	NAME1	BCD	ABC CORPORATION	35	(Maximum 60 Characters)
84	NAME2	BCD	JOB SHOP SIMULATION	35	(Maximum 60 Characters)
85	NAME3	BCD	SAMPLE ANALYSIS BY IBM 704	35	(Maximum 60 Characters)
86	DATE	BCD	2AUG 30, 1959	35	Date
87	OTID	BCD	1SYNTH	35	Order Tape Identification
88	PAR	BCD	1SIM1	35	Parameter Deck Identification

Appendix II

The Synthetic Order Tape Generator

This appendix will describe the development of an order tape for use in connection with the job shop simulation program. An entirely separate program, the *Synthetic Order Tape Generator,* is available and may be used for this purpose. It is no more nor less than one convenient means for generating a set of fictitious orders when historical shop orders are inappropriate.

From the standpoint of simulating a particular job shop, the specific function of the order generator is to create a set of orders the constitution of which is representative of the order mix which may actually be expected in the shop. This realism must be significantly approximated in all the elements which make up an order. Thus, raw material costs, machine tool routes, setup and machining times should all be typical. The basic point here again is that simulation can be effective only to the extent that the model used contains all the *essential* features of the reality it simulates. What these features are, so far as orders are concerned, and the way in which the generator makes use of them to develop a set of orders, will be discussed below.

The first thing that must be determined for the order generator is the volume of orders which we wish it to create. It will be recalled that a tape containing 740 orders (20 days at the rate of 35 per day, plus an initial load of 40) would be sufficient for our example. However, inasmuch as it may be desirable to run the simulator for more than 20 simulation days on subsequent experiments, a margin for increase may be provided. Input card 1 on the final page of this appendix will insure that the order generator will produce an order tape containing 1,000 orders.

With respect to all the other variables which enter into the construction of a complete order picture, the assumption of the generator is that it is not necessary that they be specified for each individual order, and that a probabalistic approach will sufficiently approximate reality to afford a reliable basis for calculating the real system's behavior. A decision to use the generator, or perhaps to modify it, must be based on a review, in the light of a given shop's real characteristics, of the way in which the generator implements the above assumption.

The order generator must provide a *raw materials cost* for each of the orders it creates. It will do this on the basis of information regarding the average cost of raw materials in the shop. Card 2, listed on the final page, specifies an *average* materials cost of $300 for our example.

In assigning a specific materials cost to each order, the generator, as presently constituted, assumes that the range of material costs in the shop's orders are best described by a "negative exponential distribution" with the average as specified on the input card. After setting up this distribution it establishes a cost for each separate order by selecting a specific value from the distribution on a random basis.

By definition, a random selection may be from any point in the distribution. A lower limit is therefore specified to the generator in order to avoid the selection of an unrealistically low material cost. Synthetically generated material costs which are lower than this limit will be made equal to it. Card 3 establishes $1.00 as the lower limit for our example.

The nature of the distribution presupposed here is pertinent to the question of realism. A negative exponential distribution expresses the presumption that the greater quantity of orders in the shop will have a raw materials cost which is less than the average. This is to say that only a comparatively few orders have a high raw materials cost, and that in these instances the cost is sufficiently great to bring the average up to $300.

This assumption may of course not be warranted for a particular shop. In a given case a more realistically descriptive assumption might be that orders are as likely to have a raw materials cost which is higher than the average specified as they are to have one which is lower. If this is the case the random selection of raw material values for each individual order would better be made from a symmetrical distribution because the latter would be more descriptive of the facts. The order generator may be reprogrammed to operate in terms of any type of distribution. The specific decision can be made only in the context of knowledge about a particular shop.

In addition to a raw materials cost, machine tool work routes must be determined for each order. This requires, first of all, that the generator program have information regarding the number of machine groups among which orders must move. Inasmuch as five machine groups have been indicated in our example for the purpose of the simulation itself (p. 25), this information should also be given to the order generator (card 4). Individual operations on the resulting order tape will be identified by the number of the machine group involved. (Clearly, information given to the simulator must also be given to the order generator, and it is important to be sure that it is the same information.)

The order generator establishes the sequence of operations to be performed on a particular order in terms of probabilities. The first problem is to determine the machine group at which any given order will undergo its first operation. What is the probability at the outset that it will be machine group 1 or any one of the other machine groups in the shop? These are not ordinarily random probabilities but depend upon particular types of orders and the specific kinds of operations they involve. For example, the nature of certain operations in relation to the typical orders of a shop may be such that they will be very likely or very unlikely to come first in an operation sequence. Where experience has distinguished between the horse and the cart there is usually no problem about deciding which comes first.

The conclusions of experience with respect to this matter are provided to the generator on an input card (card 5). In the present case this card stipulates that the probability of the first operation occurring on machine group 1 is .111. This means, in other terms, that 11.1% of all orders may be expected to begin at this machine group. Similarly, 26.6% of all orders may be expected to start at machine group 2, and so on through machine group 5, where 21.1% of the orders will probably begin. Needless to say, the total probability of all the alternative possibilities must add up to one. For programming reasons the sixth entry on this card must consist of zeros.

The programmed procedure employed by the order generator involves the initial generation of a random number between the limits 0 and 1 for each order. If this number turns out to be .111 or less, the first machine group in the route which this order must follow is designated as machine group 1. In the event that the random number generated falls between the limits .111 and .377 (.377 = .111 + .266), the starting point for this order becomes machine group 2, etc. The random number generated must fall within the limits of one of the probabilities associated with a machine group.

Once the starting point for each order has been established, the question of the next station on its work route may be dealt with. For the present illustration, it has been assumed that an order will not undergo two successive operations at the same machine group. This being the case, only two basic alternatives remain. The first is that the order will proceed to any one of the other machine groups for its next operation. The other is that it may now be moved out of the shop.

Information regarding the probabilities which attach to all of the alternatives must be provided to the order generator. This is done with five input cards which establish a probability matrix in terms of which all route sequence decisions may be made (cards 6-10). Each card has for its reference one of the machine groups. In the present case there are five. The first card represents machine group 1 and the .000 in the first position on this card signifies that orders at this station will never stay there for their next operation. The succeeding entries on this card indicate that 24.9% of them will go next to machine group 2, 18.8% to machine group 3, 19.9% to machine group 4, 19.9% to machine group 5, and 16.5% of the orders at machine group 1 will be complete after the current operation. The remaining cards provide the same exit probability definitions for orders on each of the other machine groups.

As the order generator proceeds to develop operation sequences for each order in terms of the above information, orders are progressively retired from the shop system. However, since the order completion point is reached only on a probability basis, it is possible that some orders will remain in the shop too long and develop an unrealistic number of operations. The generator should therefore be advised as to a realistic upper limit. Card 11 establishes a limit of thirty operations for any given order in the present case.

For both scheduling and simulation purposes, each order record must contain information regarding the amount of processing time associated with each operation. The order generator will calculate these times on the basis of given information regarding the average processing time for each machine group (card 12). This average depends, of course, upon the nature of the operation performed. It has been assumed for this example that an order which arrives at machine group 1 remains there for processing an average of six hours; at machine group 2, 1.25 hours; and so on. As in the case of raw materials cost, the generator assumes that processing times in the shop are best described by a negative exponential distribution with averages as indicated for each machine group. A random sampling from these distributions is performed to determine the specific processing time assigned to each operation.

The amount of processing time invested in each operation is a sum of two factors, actual machining time and setup time. It is necessary, therefore, to provide the generator with some basis for differentiating the two. In the order generator, as presently constituted, a very general approach has been taken.

An input card is supplied which identifies setup time as a percentage of total processing time. For our example the decision has been made that 20% of processing time will be setup, and 80% will be machining. This fact is recorded in card 13. Cards 14 and 15 establish lower limits for setup time and machining time respectively. This is insurance against the development of unrealistic times.

In the process of selecting specific values for such variables as processing times and raw material costs, etc., it is necessary for the order generator to create random numbers. To do this it must be given a base number with which to work. Card 16 provides a "random number seed" for this purpose.

There is a certain type of formal information which requires specification in the same way for both the job shop simulator and the order generator. This is the information for controlling the format of records on the order tape. Where, in the record, is the designation of material cost, setup time, etc., located? The generator needs this information for setting up the order record in the first place, and the same information tells the simulator how it has been set up.

The first format problem has to do with the number of 704 words which will be assigned to contain information for the order as a whole. Each order, for example, will contain a single raw material cost, a single due date, etc. Our example (card 17) makes provision for five words in this section of the order record. The following represents our illustration of the type of information which might be contained in each of these words:

1. Order Identification Number
2. Number of Units per Order
3. Predetermined Priority of Order
4. Due Date of Order
5. Material Cost of Order

Item 5 is the only one actually supplied by the order generator. Thus, if the generator is employed, the first four words in the above scheme would simply contain zeros.

If information is desired in the remaining four words it may be supplied by reprogramming the generator. As presently constituted, the job shop simulation program makes no use of an order identification number. An inserted identification may therefore be of any kind, so long as it can be accommodated in one word. Similarly, the simulation program uses neither the number of units per order nor any designation of predetermined priority. Such information might be placed in the order records in connection with a programmed dispatch rule which could be set up to consider these factors.

The desired position of raw material cost information within the first five words of the order record is specified in card 18. This again must correspond to the information given the simulator.

Beyond the portion of the order record designed for constant information, record lengths will vary depending upon the number of machine tool operations involved. The basic information associated with each of these operations can be contained in three 704 words. This information consists of the identification of the machine group involved as well as the amount of setup and machining time required. There will therefore be a set of three words containing this type of information for each operation of the order.

The order generator should be advised regarding the total number of words used for each operation and, in addition, of the relative positions of the above information within this set. Cards 19-22 perform this function. The format model established for the three items of information associated with each operation is effective for whatever number of operations may be involved in the order.

Given all the input cards discussed above, the order generator may be put to work developing an order tape. Operating instructions and information for using the order generator are available with the job shop simulation program package.

Input for the Order Generator

CARD NO.	CARD COLUMNS 1-6	8-10	12-72	TEXT PAGES	EXPLANATION
1	NORDR	DEC	1000	45	Number of Orders to be Generated
2	MMC	DEC	300.00	45	Mean Materials Cost
3	LOWMC	DEC	1.00	45	Lowest Material Cost
4	NG	DEC	5	46	Number of Machine Groups
5	MATR	DEC	.111, .266, .200, .212, .211, .000	46	Probability Matrix
6		DEC	.000, .249, .188, .199, .199, .165	46	
7		DEC	.121, .000, .220, .233, .233, .193	46	
8		DEC	.113, .272; .000, .217, .217, .181	46	
9		DEC	.115, .276, .207, .000, .220, .182	46	
10		DEC	.115, .276, .207, .220, .000, .182	46	
11	MAXOP	DEC	30	46	Maximum Number of Operations per Order
12	MPT	DEC	6.0, 1.25, 5.0, 3.75, 2.5	46	Mean Processing Time by Machine Group
13	PERCS	DEC	.20	47	Percent of Setup Time
14	LOWSU	DEC	.01	47	Lowest Setup Time
15	LOWMT	DEC	.01	47	Lowest Machining Time
16	RNGEN	OCT	012345670123	47	Random Number Seed
17	WOW	DEC	5	47	Words Per Order
18	WMC	DEC	5	47	Position of Material Cost
19	WJW	DEC	3	47	Words per Operation
20	WGC	DEC	6	47	Position of Group Code
21	WSU	DEC	7	47	Position of Setup Time
22	WMT	DEC	8	47	Position of Machining Time

Appendix III

Order Analysis Program

During the course of this manual much emphasis has been placed on the multiplicity of factors which have a bearing on the effectiveness of a job shop. It is the purpose of the simulation program to provide information in terms of which this effectiveness, in relation to a given set of orders, may be increased. However, even before simulation attacks this issue, it is possible to determine whether the shop is basically adequate for handling a given set of orders.

It may be that the machine tool resources of a shop are simply not sufficient. A given set of orders might require 1,000 hours of processing on machine group 1 within a certain interval of time, if the orders are to be completed by their designated due dates. In the event that the number of machines available in the shop makes only 800 hours of processing time *possible,* then the shop is simply not equal to the situation. However advantageous a particular dispatch rule may be in increasing the effectiveness of the shop, it cannot overcome this type of limitation. This holds true of any other measure which might be taken, short of increasing the number of available machines.

Information regarding the machine time demands which an order mix will be making upon a shop can be obtained directly from an examination of the orders. An order analysis program has been written which will provide information of this general type to a user. This program, which is entirely distinct from the simulation program, is available with the simulation package. It subjects a set of order records, historical or generated, to a detailed analysis the results of which it presents in report form.

This examination of the characteristics of the order mix in relation to the shop's facilities provides information which is useful in several connections. It can, for example, be helpful in making decisions about effective dispatch rules. It is apparent that what constitutes a good principle for assigning order priorities must be partly dependent upon the characteristics of the order mix.

When orders have been created by the order generator, an additional function of order analysis is to provide a basis for checking on their representativeness. If the generated orders are not capable of creating shop situations comparable to those which real orders would create, they are useless for simulation purposes. Order analysis, by spelling out some of the operating implications of an order mix, makes it possible to more completely compare these orders with what is known about real orders in the shop.

There are two main reference points around which the analysis revolves. The first of these is the amount of processing time which orders are expected to consume, and the second is the dollar value attained by the orders as they progress through the shop. The analysis breaks this information down by machine group to provide more insight into individual operations.

An illustration of the first type of report produced is shown on the top of page 52. On the first line of this report there is an identification of the machine group under consideration. There will be one of these reports for each machine group in the shop. This first line also specifies the average dollar value of orders when they arrive at this machine group. The final item on this line specifies the average amount of processing time which operations on this machine group will consume.

In the body of the report the processing time and dollar-value characteristics of the order mix relative to a machine group are each expressed in terms of ten classes. The ten successive dollar-value classes are enumerated horizontally across the report on the second line following the word "Value." The ten processing time classes are designated vertically on the left margin of the report under the heading "Time."

The specific limits for each decile class are established in the program by dividing the total range of dollar value and processing time for all orders into ten equal parts. (They will therefore differ from one machine group to another.) These specific limits are indicated as lower limits, immediately under the class number in the case of dollar value, and immediately to the right of the class number in the case of processing time. Thus the range of the first dollar class is from $1.20 through $205.07. Similarly, the first time class includes processing times from .03 hours through 3.23 hours. Putting these two classes together, we can read from the report that this order mix will confront machine group 1 with 73 orders, each having an arrival value between $1.20 and $205.07 and each consuming between .03 and 3.23 hours of processing time at this machine group. Similarly, the number of orders whose dollar value upon arriving at machine group 1 will be between $1224.50 and $1428.37 and which will consume between 25.70 and 28.90 hours of processing time will be found at the juncture of column 7 and row 9.

The row of figures following the symbol "LR/AV" contains the result of dividing the lower limit of the dollar-value class (immediately above) by the mean dollar value of all orders ($469.28). Since this is done for each class, it provides a convenient measure of the extent to which the orders in any given class have a dollar value greater or less than the dollar value of the average order arriving at this machine group.

The column under the symbol "LR/AV" has the same type of significance for processing time. It contains the results of dividing the immediately preceding numbers by the mean processing time for all orders at this machine group (5.78).

The total number of operations in each time class and value class is summarized. For value classes this is done in the row beginning with 165, and for time classes it is done in the column beginning with 282. The grand total of all operations passing through machine group 1 is also given (628) at the point where the two total lines meet. Since orders may go through a machine group more than once, this quantity may be greater than the number of orders. Multiplying this figure by the average processing time (5.75 hours) would establish the total machine capacity which these orders will require of this machine group in the shop. Percentage figures for each class immediately follow total figures for both value and time. The final column and row stipulate cumulative percentages.

The second report produced by the order analysis program is also shown on page 52. This is merely a summary table which consolidates the information contained in each of the machine group reports. It therefore provides time and value information for all orders and operations within the shop irrespective of machine group. This table is interpreted in the same fashion as those described above for the individual machine groups. The total given here, however, represents the total number of orders on the order tape.

The final report provided by the order analysis program (on page 52) consists of order information broken down by raw material costs and number of units in the orders. It will be recalled that each order has a certain raw material cost associated with it. In addition, each order may call for a certain number of units to be produced. The example which has been used in this manual has not specified unit quantities.

An order record may, however, spell out the number of units involved and specify raw material cost per unit. In this case the final report will establish the number of units and the value characteristics of the order. The report is read in basically the same way as the others. The unit deciles are shown vertically and the value deciles are shown horizontally. Since, in the illustrated case there has been no division of orders into units, all orders fall within the first unit decile.

It should be evident that if the order analysis program is to produce the reports described above it will require more working information than is contained on the order tape alone. For example, before it can calculate the dollar value of orders as they arrive at each machine group, it must be given information regarding the price of hourly machine-labor usage. It can then keep track of the increasing value of each order as it progresses through its routing from one machine group to another.

First of all, the order analysis program must be given information regarding the number of machine groups referred to in the routings on the order tape. The maximum number of machine groups which may be specified in the program as presently constituted is 18. The program may be recompiled to admit of several hundred. In the present case, five machine groups are involved and the input card will look like the following:

Card Columns:	1–4	
	0005	(Number of Machine Groups)

In addition to stipulating the number of machine groups involved, the name of each must be given. These names must be punched exactly as they appear on the order tape. The Order Analyzer assumes that the machine name on the on-order tape is in binary coded decimal. It also assumes that six IBM card code characters are used for each name on the input card. Since it is possible that there are as many as 18 machine groups, two cards may be necessary and have been made standard for this input information. If there were 18 names, these names would occupy columns 1 through 72 of the first card and 1 through 36 of the second card. In our example there are five machine groups. Their names will therefore occupy columns 1-30 of the first card, and the second card (which must also be entered) will be blank. Since the Synthetic Order Tape Generator was used to create an order tape for our example, the machine groups have merely been identified on the order tape by the numbers 1-5. The first name card will therefore look as follows:

Card Columns:	1–6	7–12	13–18	19–24	25–30
	000001	000002	000003	000004	000005

Another input card must be added to provide the order analysis program with information regarding the format of the order records. As previously indicated, the order record contains information which is constant for each order as well as information which varies for each operation in the order's route. The following input card signifies that there are five words of constant information and three additional words of information for each operation:

Card Columns:	1–4	5–8
	0005	0003

Some of the input information required by the order analysis program varies from one machine group to another. Thus it is necessary to state this information on a separate input card for each of the machine groups involved in the simulation. Since there are five machine groups in our example, five cards will be required. Each contains information about a single machine group.

The first item of information on each of these cards will represent the cost of using a machine – the hourly-machine-labor rate for that machine group. Needless to say, these rates must be the same as those specified to the simulator. This information will be placed within columns 1-13 of each card and punched to the extreme right of the field. Right justification will be the rule in all of the fields on these cards.

The five cards shown below contain the hourly-machine-labor rate information in columns 1-13. If no decimal point is punched in these columns the number is assumed to have two decimal places. The information required for the remaining fields in these cards is described below.

Card Cols:	1–13	14–24	25–35	36–46	47–57	58–68	69–72
	0000000010.00		1.E00				2
	0000000009.00		1.E00				2
	0000000012.00		1.E00				2
	0000000015.00		1.E00				2
	0000000008.53		1.E00				2

In the discussion of order scheduling, a procedure for adjusting the processing times given on the order tape was discussed (pages 20-21). The order analysis program should use these adjusted times. Its analysis of the processing times of all the orders on the order tape can thereby be made in terms of the more realistic adjusted processing times.

In order to do this it requires an identification of the polynomial which was the basis for adjustment. Five fields between columns 14-68 in each of the machine-group cards are reserved for this type of information.

It should be expressed in floating point according to the following scheme:

Columns:		
	14–24	Constant term of scheduling transformation polynomial.
	25–35	First-degree coefficient of scheduling transformation polynomial.
	36–46	Second-degree coefficient of scheduling transformation polynomial.
	47–57	Third-degree coefficient of scheduling transmation polynomial.
	58–68	Fourth-degree coefficient of scheduling transformation polynomial.

In addition to defining any polynomial transformation, the order analysis program should be advised as to whether the log-antilog routine will be utilized (page 21). An additional Logarithm Transformation card is required to fully accomplish this. Any integer may be punched in this card. In the illustration below, the integer 0001 has been used:

Columns:	1–4
	0001

The program will compare this integer with whatever has been punched in columns 69-72 of the five cards discussed above. Wherever the same integer has been placed in these columns, the log-antilog routine will be utilized. If the integer in columns 69-72 is *not* the same, this routine will not be utilized. Thus it may be employed in the case of some machine groups, and not in others.

MACHINE GROUP 1 MEAN VALUE 469.28 MEAN TIME 5.78

DEC.	LRAN.	VALUE	1	2	3	4	5	6	7	8	9	10	TOT.	PCT.	CPCT.
			1.20	205.08	408.97	612.85	816.73	1020.62	1224.50	1428.38	1632.27	1836.15			
TIME		LR/AV.	0.00	0.44	0.87	1.31	1.74	2.17	2.61	3.04	3.48	3.91			
1	0.03	0.00	73.	80.	51.	38.	17.	7.	6.	6.	1.	3.	282.	44.90	100.00
2	3.24	0.56	35.	44.	27.	16.	9.	6.	1.	1.	1.	2.	142.	22.61	55.10
3	6.44	1.12	19.	27.	12.	10.	4.	2.	2.	2.	0.	0.	78.	12.42	32.48
4	9.65	1.67	18.	16.	6.	6.	2.	4.	3.	1.	0.	0.	56.	8.92	20.06
5	12.86	2.23	8.	5.	9.	1.	1.	3.	1.	0.	0.	0.	28.	4.46	11.15
6	16.07	2.78	6.	5.	4.	3.	4.	0.	0.	1.	0.	0.	23.	3.66	6.69
7	19.28	3.34	2.	3.	1.	0.	0.	0.	0.	0.	0.	0.	6.	0.96	3.03
8	22.49	3.89	2.	0.	1.	0.	1.	0.	0.	0.	0.	0.	4.	0.64	2.07
9	25.70	4.45	2.	0.	1.	0.	0.	0.	0.	0.	0.	0.	3.	0.48	1.43
10	28.91	5.00	0.	1.	0.	2.	1.	1.	1.	0.	0.	0.	6.	0.96	0.96
TOT.			165.	181.	112.	76.	39.	23.	14.	11.	2.	5.	628.		
PCT.			26.27	28.82	17.83	12.10	6.21	3.66	2.23	1.75	0.32	0.80			
CPCT.			100.00	73.73	44.90	27.07	14.97	8.76	5.10	2.87	1.11	0.80			

FINAL VALUE VS. TOTAL TIME MEAN VALUE 507.56 MEAN TIME 17.45

DEC.	LRAN.	VALUE	1	2	3	4	5	6	7	8	9	10	TOT.	PCT.	CPCT.
			11.82	261.53	511.23	760.94	1010.64	1260.35	1510.05	1759.75	2009.46	2259.16			
TIME		LR/AV.	0.02	0.52	1.01	1.50	1.99	2.48	2.98	3.47	3.96	4.45			
1	0.03	0.00	270.	180.	76.	36.	17.	9.	3.	0.	1.	0.	592.	59.20	100.00
2	15.35	0.88	15.	122.	46.	26.	9.	9.	3.	1.	1.	1.	233.	23.30	40.80
3	30.67	1.76	0.	30.	53.	19.	4.	2.	1.	1.	0.	0.	110.	11.00	17.50
4	45.99	2.64	0.	0.	13.	11.	8.	3.	0.	0.	0.	0.	35.	3.50	6.50
5	61.31	3.51	0.	0.	1.	6.	3.	3.	4.	1.	0.	1.	19.	1.90	3.00
6	76.64	4.39	0.	0.	0.	1.	0.	0.	0.	0.	0.	1.	2.	0.20	1.10
7	91.96	5.27	0.	0.	0.	0.	2.	2.	0.	0.	1.	0.	5.	0.50	0.90
8	107.28	6.15	0.	0.	0.	0.	0.	1.	0.	0.	0.	0.	1.	0.10	0.40
9	122.60	7.03	0.	0.	0.	0.	0.	0.	1.	1.	0.	0.	2.	0.20	0.30
10	137.92	7.91	0.	0.	0.	0.	0.	0.	0.	0.	1.	0.	1.	0.10	0.10
TOT.			285.	332.	189.	99.	43.	29.	12.	4.	4.	3.	1000.		
PCT.			28.50	33.20	18.90	9.90	4.30	2.90	1.20	0.40	0.40	0.30			
CPCT.			100.00	71.50	38.30	19.40	9.50	5.20	2.30	1.10	0.70	0.30			

MATERIAL COST VS. UNITS MEAN VALUE 308.12 MEAN UNITS 0.

DEC.	LRAN.	VALUE	1	2	3	4	5	6	7	8	9	10	TOT.	PCT.	CPCT.
			1.00	220.31	439.61	658.92	878.22	1097.53	1316.83	1536.14	1755.44	1974.75			
UNITS		LR/AV.	0.00	0.71	1.43	2.14	2.85	3.56	4.27	4.99	5.70	6.41			
1	0.	0.	514.	241.	120.	61.	33.	18.	7.	3.	1.	2.	1000.	100.00	100.00
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
5	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
6	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
7	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
TOT.			514.	241.	120.	61.	33.	18.	7.	3.	1.	2.	1000.		
PCT.			51.40	24.10	12.00	6.10	3.30	1.80	0.70	0.30	0.10	0.20			
CPCT.			100.00	48.60	24.50	12.50	6.40	3.10	1.30	0.60	0.30	0.20			

E20-6071

IBM

International Business Machines Corporation
Data Processing Division
112 East Post Road, White Plains, N. Y. 10601

Printed in U.S.A. E20-6071