

This blank page was inserted to preserve pagination.

A HEURISTIC APPROACH TO ALTERNATE ROUTING
IN A JOB SHOP

by

FRANCIS JOHN RUSSO

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE
DEGREES OF
BACHELOR OF SCIENCE
and
MASTER OF SCIENCE
at the
MASSACHUSETTS INSTITUTE OF
TECHNOLOGY
June, 1965

Signature of Author . . . *Francis John Russo*
Alfred P. Sloan School of Management, May 12, 1965
Certified by . . . *Donald C. Carroll*
Thesis Supervisor
Accepted by . . . *John J.*
Chairman, Departmental Committee
on Graduate Students

"Work reported herein was supported (in part) by Project MAC, an M.I.T. research program sponsored by the Advanced Research Projects Agency, Department of Defense, under Office of Naval Research Contract Number Nonr-4102(01). Reproduction in whole or in part is permitted for any purpose of the United States Government."

*This empty page was substituted for a
blank page in the original document.*

A HEURISTIC APPROACH TO ALTERNATE ROUTING
IN A JOB SHOP

by

FRANCIS JOHN RUSSO

Submitted to the Alfred P. Sloan School of Management on May 21, 1965 in partial fulfillment of the requirements for the degrees of Bachelor of Science and Master of Science.

ABSTRACT

The research reported here investigates the use of heuristics for selecting from several alternate routes resulting from partially ordered tasks in a job shop order file. The experimental vehicle employed was digital simulation.

The concept of the "alternate string" has been developed to generalize the existence of partially ordered operations. That term is defined as a concatenation of operations that can be performed in any order, with the additional specification that all within the string must be completed before any operation past the string can be attempted. The presence of alternate strings with two or more members gives rise to the alternate routing problem, whose solution is approached by heuristic methods.

Choosing from among several alternate routes constitutes a three level decision problem. At the lowest level, routes can be chosen when the order enters the shop. This is equivalent to fixed routing. At a higher level, alternates can be selected at the time of transition from one work station to another. The third decision level occurs at operation time, when one of the alternate operations is placed on a machine. Heuristics were tested at the latter two levels.

There were two prior assertions that this thesis set out to prove. The first was that alternate routing at the highest decision level would produce significant reductions in the mean tardiness of orders completed past their designated due dates, the improvement being both relative to fixed routing and to alternate routing heuristics implemented at lower decision levels. Secondly, the contention was made that the improvement would be of such a magnitude that on-line, real-time systems become economically justifiable as a means of mitigating the attendant control problems caused by non-deterministic paths through the queuing network.

The methodology employed here was to conduct two passes of simulated shop runs. The first, with two artificially high levels of alternate incidence, tested the efficiency of five different alternate routing heuristics in reducing mean tardiness. The second pass consisted of runs with the best heuristic developed during the first experimental phase applied to a realistic length and frequency of alternate strings.

The results of the experiments strongly support the assertions made at the outset of the thesis. The performance characteristics of the different heuristics are discussed at length. In addition, some implications are drawn of the computational nature of alternate routing and the difficulties encountered in implementing alternate routing heuristics at operation time.

Thesis Advisor: Donald C. Carroll
Title: Assistant Professor of Industrial Management

ACKNOWLEDGEMENT

I wish to take this opportunity to express my thanks to the following people for providing aid and succor in the preparation of this thesis:

To Professor Donald C. Carroll, of the Sloan School of Management, whose inspiration provided the primary motivation for this thesis and whose aid and advice were generously given despite important commitments of his own;

To Professor James C. Emery, also of the Sloan School, who took the time from a busy schedule to serve on this thesis committee;

To Paul Clermont, Tom Johnson, Lalit Kanodia, and Allan Scherr, whose good spirits and timely advice successfully counteracted my occasionally damaged psyche and invincible ignorance;

To the staff of Project MAC, too numerous to list here, for their patient attempts to help me use a system I shall never fully understand;

To Miss Carole Robinson for typing the manuscript and otherwise getting it into presentable form.

On a more personal basis, my deepest appreciation to my fiancée for good humoredly accommodating herself to my bizarre hours and bearish disposition and to my parents, without whose lifelong support none of this would have been possible.

TABLE OF CONTENTS

	<u>Page</u>
CHAPTER ONE: INTRODUCTION	1
1. Research in Job Shop Dispatching	1
2. The Concept of Alternate Routing	1
3. A Practical Consideration and Prior Assertions	4
CHAPTER TWO: HEURISTICS FOR ALTERNATE ROUTING	6
1. A Three Level Decision Problem	6
2. The Role of Heuristics in Problem Solving	8
3. Alternate Routing Heuristics	9
CHAPTER THREE: THE EXPERIMENTS	14
1. The Model	14
2. The Method	15
CHAPTER FOUR: RESULTS AND CONCLUSIONS	19
1. General Considerations	19
2. First Pass Results	19
2.1 Comparison of the Alternate Routing Heuristics	19
2.2 The Effect of Alternate Routing on Shop Performance	21
2.3 Interaction of Priority Dispatching Rules	22
2.4 Computational Considerations	23
3. Second Pass Results	25
4. Global Conclusions	26
5. Suggestions for Further Research	30
APPENDICES:	32
A. First Pass Results	32
B. Pass One: Percent Reduction in Mean Tardiness	34
C. Second Pass Results	36

TABLE OF CONTENTS (Cont'd)

	<u>Page</u>
APPENDICES (Cont'd)	36
D. The Alternate Index: A Measure of Potential	37
E. Statistical Analysis of Results	39
1. Alternate Routing and Mean Tardiness	39
2. The Best Alternate Routing Heuristic	41
3. Pass Two ALLQ Performance	42
BIBLIOGRAPHY	43

LIST OF DIAGRAMS

	<u>Page</u>
Figure I.2 Tasks in Fixed Order	2
Figure I.2 Partially Ordered Tasks	2
Figure IV.1 Reduction in Mean Tardiness vs the Alternate Index . . .	28

*This empty page was substituted for a
blank page in the original document.*

Chapter One

INTRODUCTION

1. Research in Job Shop Dispatching

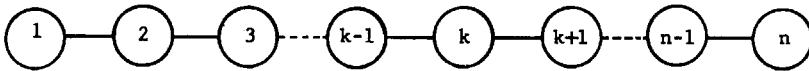
Over the past several years, considerable attention has been given to job shop dispatching. The primary, and thus far the most fruitful, vehicle for such studies has been digital simulation. Researchers have devoted their attention to date principally to the development of effective sequencing or priority rules for resolving conflicts at work stations caused by several jobs competing for machine time.

The work reported in this thesis on quite another aspect of job shop dispatching is likewise a simulation. The model itself is largely an extension of previous work done by Carroll.¹ As such, only a bare minimum of discussion on the underlying structure of the research reported here will be included. The attention of this investigation is focussed on an as yet untouched aspect of job shop dispatching: the alternate route.

2. The Concept of Alternate Routing

A common assumption of job shop models has been that of fixed routing. That is to say, all jobs are presumed to pass through the shop with their operations performed in a predetermined order. The path or route that the job is to follow is fixed. Figure I.1 is a graphic representation of job operations performed in fixed order. For simplicity, one can assume that

¹Carroll, Donald C., "Heuristic Sequencing of Single and Multiple Component Orders," (unpublished Ph.D. dissertation, Sloan School of Management, Massachusetts Institute of Technology, 1965). As of this writing, this work was not in final form. Consequently, no page citations can be included in subsequent footnotes.

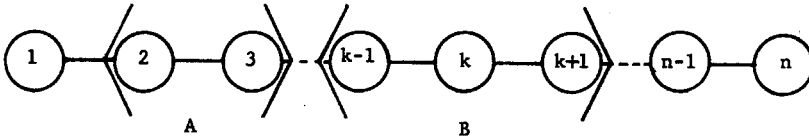


Tasks in Fixed Order

Figure I.1

each numbered circle corresponds to a task to be performed at some machine in the shop. More than one of the tasks may be performed on the same machine, though the restriction imposed is that two consecutively numbered operations must be performed at different work stations. Thus operations would be performed in order, from the first through the nth.

A more realistic situation would be that of partial ordering, in which several operations may be performed interchangeably. Figure I.2 shows how the same job might be depicted with some partial ordering.



Partially Ordered Tasks

Figure I.2

Of the operations shown in the diagram, the following ordering exists:

- No. 1 must be performed first;
- Nos. 2 and 3 can be completed in either order, though both must be finished before No. 4 can be started;
- Nos. k-1 through k+1 may likewise be done interchangeably, though the three must be completed before the job may proceed farther;
- Nos. n-1 and n must be performed in their fixed order; that is, n-1 must be performed before n.

The existence of partial ordering, or alternate operations,

gives rise to the possibility of choosing alternate routes or paths for the job whenever the facility to interchange the order of tasks will result in better performance. It is convenient to conceptualize interchangeable operations as belonging to alternate strings. An alternate string is a concatenation of operations that can be performed in any order, with the additional specification that all within the string must be completed before any operation past the string can be attempted. Referring for a moment back to Figure I.2, it is evident that A is a two-member string, each of its members having a single alternate; B is a three-member string, with each of its members having two alternates. The remaining single operations having no alternates may also be considered to belong to one-member strings. Generalizing the concept of the alternate string, then, one may thus consider the job shop order as made up of a distribution of alternate strings from one to n members in length; the operations of the order are correspondingly members of these strings and have from zero to n-1 alternates.

Comparison of several order files with different distributions of string lengths and frequencies becomes difficult without some common measure. An alternate index is proposed to serve this purpose in Appendix D. With a behavior mathematically equivalent to the arithmetic mean, the alternate index, I_a , suppresses differences in string distributions, providing a basis for comparing string distributions with widely varying characteristics. Such a measure will prove useful later on, when an attempt will be made to draw some conclusions of the effect of alternate incidence on job shop performance.

It is immediately evident that the presence of alternate routes increases the combinatorial nature of the job shop dispatching process.

In addition to the need for choosing from among queue members which job is to receive priority, alternate routing now poses an additional decision problem to the dispatcher: which route among all those possible will result in the best shop performance? This thesis will attempt to approach an answer to this question.

3. A Practical Consideration and Prior Assertions

The combinatorial complexities² of alternate routing probably constitutes the main factor militating against the systematic exploitation of its benefits in industry. There seems to be little doubt that some techniques of alternate routing are applied to expediate special orders or to extricate shops from unconscionably late deliveries. But the attendant control problems from the increased entropy of the non-deterministic queuing network has kept managers from implementing alternate routing in any systematic fashion.

If the control problems are so great that practical benefit can not be derived from alternate routing, then this work is at an impasse--the expansion of the simulation into this additional dimension of decision making is barren of practical benefit. Concomitant to the already specified goal of this thesis to find an effective decision process to choose alternate routes consequently lies the additional burden of demonstrating

²The mathematics is as follows:

- a. The presence of an n member alternate string generates n! alternate routes.
- b. If m = no. of orders in process at any point in time
 l_j = no. of uncompleted alternate strings in the j th job
 n_{ij} = no. of unfinished members of the i th string of the j th job

R_t = total number of possible alternate routes in shop at time t

then,
$$R_t = \sum_{j=1}^m \sum_{i=1}^{l_j} n_{ij}!$$

the feasibility of implementing such a process. The contention at this point is that the payoffs to be derived by a systematic approach to alternate routing in a job shop are of such a magnitude that they justify expensive, sophisticated control systems. The obvious candidate: on-line, real-time information processing systems.

The conjectures of this investigation are thus:

1. That an effective method for systematically routing partially ordered operations will produce a significant improvement in shop performance;
2. That the inherent complexities of alternate routing will be mitigated by payoffs of such a magnitude that control by real-time computation systems will approach economic feasibility.

Chapter Two

HEURISTICS FOR ALTERNATE ROUTING

1. A Three Level Decision Problem

Choosing from among several alternate routes constitutes a three level decision problem, the levels being defined both by logic and by the simulation itself. A brief description of the dispatching mechanism will suffice as support for this contention.

Orders are released into the shop at a preset mean rate sufficient to maintain a theoretical load. Once in the shop, the job is dispatched in the following manner:

a. Upon entering the shop, if the machine designated to perform the first operation is idle, the job is immediately put on that machine; otherwise, it enters the queue waiting for that machine.

b. When a machine finishes an operation, the job is sent to the queue waiting for the machine that is to perform its next operation.

c. When a machine becomes free, its queue is searched and the highest priority queue member is selected to receive machine time. Which one is, of course, determined by the dispatching rule used.³

The three levels at which alternates can be chosen are thus, paral-

³It is appropriate to inject two comments at this point. The first is that obviously a lot more is going on than this paragraph accounts for. The precedent is hereby established that only details with relevance to the task at hand will be included in these pages.

The second will, hopefully, avoid semantic confusion later in the discussion. The convention used throughout this report is that a high-priority job is one with a low-priority index. When the program decides that a job has the highest priority in queue, it does so because that job had the lowest priority index. For example, a job with an index of 1 is of higher priority than one with a priority index of 100.

leling the above structure:

Level I: At Entry

This constitutes the fixed route situation discussed in the preceding chapter. In effect, a choice is made of the sequence in which the operations are to be performed either before or just when the job enters the shop. A route is chosen at a time so remote to the point where information is available that the decision is without much foundation. In a shop with a backlog of work-in-process, a consideration of alternate routes at this point in time would be so remote from real-time conditions that little would probably be gained. For all practical purposes, and for the purposes of this investigation, a Level I decision constitutes fixed routing.

Level II: At Transition

When an operation is completed and the job is sent to the next work station, an opportunity exists to choose which work station if several subsequent tasks can be performed interchangeably. At this transition time, the nature of the queues gives some information which will indicate which of the alternates should be chosen. An alternate headed for an empty queue is a good choice. Beyond this, picking alternates at the transitional level, while involving a higher order decision than in the previous instance, involves only a partial search of the solution space when all alternate queues have jobs waiting to be processed. The decision must be based upon some measure of the relative desirability of the possible alternates at the point of transition, when there is no guarantee that in the next incremental time unit (before the machine becomes free) some new arrivals won't sufficiently alter the situation to render the choice ineffective.

Regardless of the shortcomings, transitional analysis of alternates

promises at least the potential for making a reasonable choice, with minimal amounts of calculation and information handling.

Level III: At Operation Time

At the highest probable level, the choice of the best alternate can be made at operation time, when the job is placed on a machine by the shop itself under its own priority dispatching discipline. It implies that the job be allowed to compete for time at all work stations until one of the alternate operations reaches the highest priority at that machine queue. The further implication is that the alternate is chosen with "perfect information" since by becoming the highest priority operation in its queue that operation simultaneously becomes the best choice, as none of the job's other alternates have become as important within their own queues. It is the case of a local "optimum" constituting a global one. Because the solution space is more exhaustively investigated, an attendant cost--in the form of greater information handling--is incurred over a level II type of decision. If past experience in simulation is of any value, however, one would expect that higher level decision making produces better solutions with minor increases in computational effort. Besides, the "space" involved here, though large in terms of human capability, is a small morsel for a high speed digital computer.⁴

2. The Role of Heuristics in Problem Solving

It might be relevant to digress for a moment here and briefly

⁴Consider, for example, an instance where 100 orders in process each contain five alternate strings, of three operations each. The total number of alternate routes that might be searched are: (see Footnote 2)

$$\sum_{i=1}^{100} \sum_{j=1}^5 (3)! = 3,000, \text{ not a very large magnitude.}$$

consider the role of heuristics in problem solving. When confronted with a problem that has no relevance to an analytic or algorithmic model of easy manipulability, one is forced to accept some "reasonable" rules to limit the search of all possible solutions to only those that, through past experience or intuition, promise to be fruitful. The only alternative is to enumerate all possible answers, a course that for large systems requires prohibitive amounts of time and space (whether paper surface or, in this case, core storage). To surmount this difficulty, the job shop model upon which this research is based utilizes heuristics, "rules of thumb" which selectively limit the solution space. The basic model applies heuristics to the problem of selecting from several jobs waiting to go on a particular machine. The additional dimension of choosing from among several alternate routes possible because of partially ordered operations is attacked in this investigation in a similar fashion. Several heuristics have been tested which select alternate routes according to criteria which seemed either through intuition or experience gained during the course of the research to promise effective improvement in shop performance. All are necessary limitations on the search space imposed by the complexity of the problem. Though none can be claimed "optimum" in any mathematical sense, each is a reasonable balance of a "good" solution and the effort expended to obtain it.

3. Alternate Routing Heuristics

Five major alternate routing heuristics were tested, four at the transitional level and the fifth at operation time. At transition from one machine queue to the queue at the machine where the next operation is to be performed, a route (that is to say, an operation) was chosen by picking the task destined for the queue having the:

1. Lowest average priority--LOAP;
2. Maximum average slack per remaining operation--MASRO;
3. Maximum average critical start date--MACSD;
4. Maximum average processing time--MAPT.

At the operational level, a rule was programmed which might be conveniently called:

5. ALLQ--put all alternates in their respective queues, allowing the job to compete for time in as many queues as it has alternate operations. When the job receives priority at one machine, let that be the alternate chosen, and take the remaining tasks out of their respective queues until the operation picked is completed.

A brief description of each of the above rules might be in order at this point.⁵

The lowest average priority heuristic (LOAP) operates in the following manner. For all alternates possible at transition time, calculate the average priorities of the respective queues by summing the priorities of the queue members of each queue and dividing by the number in that queue. The alternate route is chosen by selecting the operation that will enter the queue with the lowest average priority. If any operation will enter an empty queue, choose it as the best alternate.⁶ The net effect of this rule is to defer tasks that would be competing for machine time with high priority jobs, first performing alternate operations that will enter either

⁵The intention here is to provide the reader interested in programming one or more of the heuristics for his own model with all the information needed to duplicate exactly what has been done here. As such, the descriptions that follow are completely general, and assume only that those who might choose to do so have data bases and source languages sufficient to the task.

⁶An empty queue was selected without further search in all the alternate choosing heuristics.

empty queues or those that contain jobs with a low level of urgency. Averaging the priority measures is demonstrably reasonable; for example, consider two alternates destined for the following queues.

<u>Queue</u>	<u>No. in Q</u>	<u>Priority Index</u> ⁷	<u>Sum</u>	<u>Average</u>
A	1	10	10	10
B	2	5	10	5
		5		

Clearly, the best alternate is the one that will enter A, since A's member is less urgent than either of B's. If the sum were used as a criterion, one would be indifferent. The more discriminating averaging mechanism is thus more effective than a simple sum as a measure of how far behind a machine is. This rationale is behind the use of an average for all the transition level heuristics.

The lowest average priority, as thus described, was used to pick alternate routes for all dispatching rules but one. The exception was the COVERT family, which divides queue members into three classes: critical, late, and early jobs. For this rule the sum was taken as the measure for the first two categories, while the average used for the third. Queue comparisons were class by class, lower categories used to break ties in those higher. A class by class comparison was also applied in the other multi-class dispatching rule (see Chapter Three).

The second alternate picking heuristic--maximum average slack per remaining operation (MASRO)--was largely motivated by preliminary tests of the slack per remaining operation dispatching rule with the previously mentioned alternate choosing rule. Improvement in performance was so marked in that instance that it was decided to pair the slack based alter-

⁷A high priority index implies a low priority job, and vice versa; (see Footnote 3.)

nate choosing heuristic with the rest of the dispatching rules. It is actually a subset of the lowest average priority method detailed on the previous page, to which explanation may be referred for further detail.

Also at the transitional level, the maximum average critical start date rule (MACSD) chooses the route by selecting the operation that will enter the queue whose members have the latest average critical start date.⁸ The calculations involved are of the same nature as those discussed previously: at transition time, calculate the average critical start date for all queues to which alternates may be sent, choosing as the alternate the operation whose queue has the largest value.

The last transition level heuristic considered was that based on processing time: select the alternate operation that will enter the queue whose members have the largest average processing time (MAPT). This rule will postpone alternates that will compete for time with short operations while choosing those that will go on longer-task filled queues. It thus should have performance effects analogous to the shortest operation priority rule, which is particularly effective in reducing wait times.

At the operational level just one heuristic was considered--ALLQ, which has been generally described previously. It works in the following manner. At transition, put all members of the alternate string not yet completed in their respective queues. Some method of recording that the job has alternates in queue and for later removing rejected operations

⁸The critical start date is the last date an operation may be started without the order becoming tardy.

from their queues is necessary.⁹ When one of the alternates becomes the highest priority operation in its queue, the remaining members of the alternate string are removed from their queues until the chosen operation is completed. The choice of the alternate route by this heuristic is intimately dependent upon the priority dispatching rule employed, since it is the priority rule, by setting one of the alternates highest in its respective queue, that determines that the rejected alternates are in queues which have other more pressing commitments. The rule is patently crude. Conflicts caused by two or more alternates receiving priority on their respective queues simultaneously are resolved by a random choice. Nevertheless, it seems a good first blush attempt to choosing alternate routes at operation time.

⁹The discussion on this point is necessarily vague, as the methodology involved depends upon the source language and the data structure employed in the model. To indicate that other alternates must be removed from queues poses little problem. Tracking down the rejected queue entries is another matter. A list processing technique seems most economical both to the programmer and the computer.

Chapter Three

THE EXPERIMENTS

1. The Model

Since the basic job shop model employed in this research was developed and reported elsewhere,¹⁰ discussion of its assumptions and structure will be kept to a minimum. The programs for simulating the shop were all written in Fortran Assembly Programming (FAP) for the IBM 7094. The system consists of a monitoring main program and several subroutines which initialize the shop for either empty or an initial load, release orders into the shop at a rate sufficient to maintain a predetermined load, remove completed operations from machines, place the next operation in the proper queue, assign jobs to empty machines from the waiting lines, and compile statistics on shop performance along with a histogram of job lateness. The priority rule employed is determined by loading an appropriate subroutine. Additional programmed facilities are provided for output of periodic snapshots of shop status and for dumping terminal shop status for later use as an initial load.

Internal data organization employs a threaded list structure, with pointers utilized for reference among the several information centers within the shop. The order header, for example, has pointers to the order file location of the current operation and the queue location of that operation. Each queue entry contains the order file address of the particular entry; likewise each machine word contains the location of the order

¹⁰ Again the reader is directed to Carroll, op. cit..

being processed.

This particular internal data structure proved particularly effective in implementing the higher level alternate routing rules, particularly ALLQ. These heuristics, themselves programmed in FAP, were made considerably easier to implement because of the thoughtful data organization employed in the original model.

The shop was run with several priority dispatching rules, as briefly outlined below:

1. FCS--first come, first served;
2. EARS--earliest start date;
3. SHTOP--shortest operation;
4. TRSIO--a two class truncated shortest operation rule that for critical jobs (i.e., negative slack) takes the shortest operation within the critical class; if none are critical, the shortest operation is taken;
5. SLKROP--slack per remaining operation, with slack defined as critical start date minus current time;
6. COVERT--a family of cost \div time rules that assumes that all costs are tardiness costs and tardiness cost is a linear function of the probability of lateness; for critical jobs (i.e., already late) chooses the shortest operation, for late jobs (i.e., behind schedule, though not irretrievably late) the one with the largest cost over processing time, for only early jobs in queue chooses the shortest operation.¹¹

2. The Method

The general experimental method consisted of two passes. The first series of runs was used to obtain some information on the relative perfor-

¹¹This highly effective priority dispatching rule has not been yet reported in the literature. Carroll (1) discusses it thoroughly in his work cited. Since the work reported here is not directly concerned with dispatching rules, it was considered inappropriate by the author to include any more detail. COVERT need not exclude other costs (e.g. carrying); tardiness cost was the only one included here. TRSIO, actually a less sophisticated COVERT, is also Carroll's, op. cit.

mance of the different alternate routing heuristics amongst each other and with fixed routing. Pass two used the best of the rules tested in the first pass to make some realistic comparisons between alternate and fixed routing. The details of both phases are discussed below.

The first pass of experimental runs was undertaken to test the relative performance of the different alternate routing heuristics under rather high levels of alternate incidence. Two levels of incidence were tested during this phase:

- Case I: each operation is a member of a two alternate string;
i.e., each operation has one alternate;
- Case II: each operation is a member of a four alternate string;
i.e., each operations has three alternates.¹²

Two assumptions underlie the method used. The first assumption made was that alternates can exist among all operations regardless of the machine on which the tasks are to be performed. That is to say, alternate strings involving any particular combination of machines are no more likely than those having any other combination. The second assumption applies particularly to the ALLQ heuristic. Somewhat less defensibly, it is assumed that no additional transit time is involved by allowing a job to compete for time on more than one machine. It assumes that some central buffer storage area exists that is equally accessible to all machines. Some realism might have been added to the model by adding some amount to the processing time to provide for increased transit times because of additional handling. This,

¹²Strictly speaking these distributions will hold only for orders having a number of operations evenly divisible by the lengths of the strings. Thus, "end effects" which will produce some one member strings in Case I, and one, two, and three member strings in Case II have been ignored.

however, was not done, partly because so doing requires additional assumptions about when and how much time should be added. The assumption is somewhat supported by the consideration that, in the "real world", the "dispatcher" could anticipate the imminent selection of one of the alternate routes and start the job moving to the proper machine.

Experimental conditions peculiar to the first pass included:

1. An initial load of 544 orders dispatched under the FCFS priority rule with no alternate routes;
2. Test runs of 512 orders, approximately 6,000 tasks;
3. A shop load of 80 o/o.

The second pass was undertaken to appraise the performance of the best alternate routing heuristic found in pass one over longer runs with a more realistic incidence and distribution of alternate strings. Just one case of alternate incidence was tested in pass two:

- a. 85 o/o of the alternate strings contain but one member; i.e., 71 o/o of all the operations have no alternates;
- b. 10 o/o of the strings have two members; i.e., 17 o/o of the operations have one alternate route;
- c. 5 o/o of the strings have three members; i.e., 12 o/o of the operations have two alternates.¹³

No additional assumptions behind those discussed in the first pass were made, other than the obvious one that the above distribution of alternate strings is a "realistic" one.

Experimental conditions peculiar to the second pass were:

¹³The same point made about end effects made for the pass one incidence levels also applies here.

1. An initial load of 1,152 orders with no alternates dispatched with FCFS;
2. Test runs of 3,072 orders, approximately 37,000 tasks;
3. A shop load of 80 o/o.

Common to both passes were orders with exponentially distributed holding and wait times, eight machine groups with single server queues, a pure job shop (i.e., equal transition probability from each machine to any other machine), and equal loads at all machines. Only single component orders were considered, that is, orders without assembly operations.

The above figures on the number of orders and tasks only include data orders, for which shop performance data was collected. In both passes an initial run of non-data orders (in addition to those in the initial load) was performed to "settle down" the shop. For the first pass 100 orders were used, for the second, 200. For both passes the non-data orders averaged 12 tasks per order.

Chapter Four

RESULTS AND CONCLUSIONS

1. General Considerations

Since due dates are exogenous inputs and the model assumes no early delivery of orders, the only variable cost is that associated with tardiness of orders beyond the given due dates. The variable of interest upon which the major conclusions of this research are based is thus: the mean tardiness of the orders completed past their designated due dates, as the assumption being made is that the only relevant cost is that incurred from late delivery to customers. Consequently, the best alternate routing heuristic is that which produces the smallest mean tardiness. Wait times will receive some attention, however, though only as a secondary consideration.

2. First Pass Results

First pass results confirmed a priori feelings of the effect of alternate routing on job shop performance, though there were a few surprises. Appendix A contains a tabulation of mean tardiness and mean wait time for all priority rules and alternate routing heuristics for both cases of alternate incidence tested in this series of runs. Appendix B presents the same output in different form. The percentage reductions in mean tardiness over fixed routing are tabulated for all combinations of priority rules and alternate routing heuristics for both two-member and four-member strings.

2.1 Comparison of the Alternate Routing Heuristics

Without benefit of statistical support, some important generalizations can be drawn about the behavior of the alternate routing heuristics.

The first concerns the relative behavior of the slack based transitional rules, MASRO and MACSD, and MAPT, which is based on processing time. The former pair behave in a manner similar to analogous slack based dispatching rules. Their main effect is to reduce mean tardiness, by reducing the tail of the order lateness distribution. The reduction in wait time is comparatively small, as these two act to trade-off tardiness and wait time, sacrificing the latter to improve the former. MAPT, however, has a smaller effect on tardiness than either of the previous two. This one would expect, since it acts in a similar manner as the shortest operation priority rule. But unexpectedly, MAPT has even a smaller effect on wait time than either of the slack based alternate routing heuristics. True, MAPT combined with its sister dispatching rule SHTOP did produce the lowest wait times of any pair of transition level alternate heuristics and priority rules for both cases of alternate incidence. In general, however, MAPT did not do as well as the other rules in reducing wait times. The general implication to be drawn, then, seems to be that processing time-based alternate routing heuristics applied at the transitional level are inferior to slack-based rules in both reducing tardiness and wait time.

Because the lowest average priority heuristic (LOAP) is actually composed of six basic rules, one for each of the six basic dispatching rules tested, little can be said of its performance in general. Its primary role was to provide some insight into possibly effective alternate routing rules and compare the behavior of the various priority disciplines both as dispatching rules and alternate routing heuristics. The most surprising performer in this set was the LOAP-SLKROP combination, which led to the programming of MASRO for use with all dispatching rules.

The relation of MACSD to MASRO in reducing mean tardiness was another

unanticipated phenomenon. Both heuristics are slack based in nature. MASRO, in addition, has a time dimension which makes it a dynamic rule, one whose measure changes with time. As such, one would expect it to be more effective in reducing tardiness than MACSD. The results, however, do not substantiate this. MACSD is marginally more effective, though not substantially so. More investigation is necessary, therefore, before any implications can be drawn of the relative efficacy of static and dynamic alternate routing heuristics at the transitional level.

In advancing one of the alternate routing heuristics as being the best of those tested, some statistical analysis was employed.¹⁴ As predicted, the highest level heuristic, ALLQ, produced a reduction in mean tardiness significantly greater than any of the transition level rules. In postponing the selection of the alternate until operation time, ALLQ has the advantage of choosing on the basis of superior information relative to that possessed by the other rules. In addition to the time dimension of its information, ALLQ searches a larger portion of the solution space. The evidence thus supports the prior contention that the higher level rule would outperform those at lower decision levels. ALLQ was thus chosen as the alternate routing heuristic to be implemented in the next phase of the experiment.

2.2 The Effect of Alternate Routing on Shop Performance

The prior contention of this research was that alternate routing would produce substantial improvements in shop performance, particularly in the paramount measure, mean tardiness. The primary support for this

¹⁴The statistical analyses upon which this and other conclusions so noted are based are performed in Appendix E.

assertion was the general assumption that the more degrees of freedom given to a decision-maker the more effective the decisions that are made. This is assuming, of course, that the responsible party possesses an efficient decision-making mechanism. The statistical evidence provided by the first pass results overwhelmingly supports prior judgment. The poorest alternate routing heuristic, MAPT, for the two-member string case produced significantly lower mean tardiness with all the dispatching rules than did the shop run without alternate routes. By implication, then, ALLQ, which was far superior to MAPT at the more realistic two-member string level, can claim a massive capability in reducing order tardiness. How well it performs in a more realistic situation will be determined in the second pass. The results thus far, however, provide some hint of what can be expected.

The previous chapter stated that the partial aim of this phase of the experiments was to gain some feel for the effect of alternate incidence on mean tardiness---thus the two-member and four-member strings. The alternate index, I_a , previously referred to can serve a useful role here. In going from fixed routing to two-member to four-member strings, I_a has gone from 0 to 1 to 3. A cursory glance at the tardiness figures in Appendix B reveals that decreasing marginal reductions in mean tardiness result with increases in I_a . With such large reductions in tardiness with the two-member strings, there is therefore good reason to expect significant improvement in tardiness at lower, more realistic levels of alternates with fractional values of I_a . Indeed, had there not been such significant improvement at these rather high levels of alternates, pass two, at a lower level of incidence, would have been completely unnecessary.

2.3 Interaction of Priority Dispatching Rules

The first pass experiments also show that all priority dispatching

rules do not benefit equally from alternate routing. The results do not allow for generalizations, but the change in the relative performance of the dispatching rules with and without alternate routing is evident. Some of the poorer priority rules manifested better performance (relative to the better dispatching rules) with alternate routing than they were able to muster under fixed routing. There is certainly greater room for improvement with the less effective dispatching rules. Probably the primary cause for this phenomenon is the greater number of empty queues that occur in a shop operating with a less effective priority rule. Such rules display the characteristic of allowing long queues to build up at certain work stations while other machines stand idle. The presence of an empty queue in alternate routing situations offers a special opportunity for any of the heuristics, since by choosing an operation that will enter an empty queue, it is assured that the wait time for that task will be zero. The performance of the shortest operation dispatching rule provides perverse support for this interpretation. SIO has the characteristic of providing short wait times at work stations, keeping queues short and of approximately equal length. Its consistent refusal to obtain as much improvement as the other dispatching rules from alternate routing thus acts as a strong indication that priority disciplines that display the opposite characteristics (as the poorer rules indeed do) have build-in potentials for improvement.

2.4 Computational Considerations

The experience gained during this first phase provides a basis for making some comments on the computational nature of alternate routing in a simulated job shop. The highest level heuristic, ALLQ, involved the largest programming effort, chiefly because of the threaded list manipulations needed to keep track of the queue locations of the alternate opera-

tions and remove the rejected tasks from queue. This particular consideration is closely related to the programming language used. Although in this research FAP was used exclusively, the generalization that a higher level heuristic will involve more complex programming is probably a valid one. As one mitigating factor, although ALLQ necessitated greater information handling (even to the extent of requiring additional COMMON storage allocation), less actual calculation was involved than in the lower level heuristics. On the other hand, more program interaction was necessary than in the lower level rules. In the final analysis, however, computational and information processing requirements notwithstanding, the higher level heuristic provides its own mechanism for overcoming its drawbacks. ALLQ's effectiveness in reducing work in process by reducing order throughput time releases enough core data storage to offset the rule's need for larger program storage. A concomitant of this reduction was that computer execution time to complete the simulation was not adversely affected by the additional processing needed to implement alternate routing.¹⁵

Not parenthetically, the importance of data organization deserves some comment here. The development of the original job shop model provided for a rather involved information structure. Considerable setup cost was incurred because of this; a goodly portion of the simulator's instruction set is devoted to maintaining the data base and keeping the threaded lists

¹⁵These last two statements are made without supporting data. The first is motivated largely by impressions gained via the core dumps taken while testing and debugging the programs, and by queue snapshots outputted by the shop (not reported here). The second is difficult to precisely support because of the systems changes made during the course of this work on Project MAC's IBM 7094. The largest increases in execution time over the basic model were on the order of 10 o/o.

up to date as the shop status changes over time. In a sense the basic model has a good deal of "over-kill" in its data structure; more simply structured it would still function as basic job shop. But in extensions into higher level problems such as alternate routing¹⁶ such prior provision proved invaluable. The relative ease with which alternate routing could be programmed into the basic shop was largely due to the effective data organization. If such had not been done, the conclusions of the previous paragraph might have been very different.

3. Second Pass Results

It will be recalled that the aim of the second set of runs was to test the performance of the best heuristic developed during the first experimental phase under more realistic alternate incidence. ALLQ was thus applied to an order file with an alternate index of 0.2 and the results compared with the same file, only with fixed routing.¹⁷

As the tabulation in Appendix C shows, there is a considerable improvement in mean tardiness to be derived from alternate routing at realistic levels of alternate incidence. Statistically, the percent reductions in tardiness obtained by the ALLQ heuristic are significant. Even without benefit of test, however, the magnitude of the reductions are impressive enough to have broad implications on the feasibility of

¹⁶ Another extension of the basic model, provision for machine substitutability, is currently under investigation.

¹⁷ See section III.2 for a more detailed presentation of experimental conditions. The fixed routing runs with which the pass two results are compared were performed by Carroll (1). Pass two thus replicates the single channel, single component, 80 o/o loaded shop subset of his experimental runs.

$$\text{For the second pass, } I_a = 0(.85) + 1(.10) + 2(.05) = 0.20.$$

real-time control systems for job shop operation.

Deferring the more global discussion for the moment, some time might well be spent just looking at the nature of the pass two results. As predicted from the experience gained during the first experimental phase on the effect of alternate incidence (quantified as I_a), there was indeed a "surge" in mean tardiness reduction at low I_a 's. Percentage reductions in mean tardiness ranged from 6.2 for SHTOP to 52.0 for the best of the COVERT dispatching rules. The range is of minor importance, however. What is significant is that all classes of slack based priority disciplines enjoyed improvements of over 40 o/o. The best overall performer was COVERT SPM.5 which also received the largest reduction in tardiness with the ALLQ heuristic for alternate routing. This pair represents the most computationally difficult priority discipline and alternate routing heuristic tested in this research. The wait times, which reflect flow time, were reduced on the order of 10 o/o for all dispatching rules. The more global conclusions which must be drawn from these results are of major importance to the current state of the art of managerial control of job shop operation at the dispatching level.

4. Global Conclusions

The global conclusions that must follow as a result of the work reported here bring this thesis full circle. Two assertions were tentatively offered in the introductory section of this report. The first was that alternate routing, approached in an effective, systematic fashion, provides a potential for significantly reducing tardy completion of customer orders. The results reported thus far confirm this contention. The second prior hypothesis was that attendant control problems would be largely mitigated by such large reductions in tardiness that on-line, real-

time control systems would become economically feasible. The number of considerations that warrant discussion prevent any conclusive analysis here. Some tentative remarks on this point, however, can be attempted.

The benefit to be gained from an on-line control system that can facilitate the implementation of alternate routing is functionally related to a number of variables not considered in this research. The size of the shop, the dollar volume of its business, and the exact nature of its cost-tardiness relationship are immediately evident factors. In addition, one must consider the possibility of time-sharing the job shop control system with other real-time and batch processing applications that might be implemented on an already operating computing facility, together with the availability of suitable man-machine interfaces for real-time interaction. The broad spectrum of conditions that exist in different industrial facilities thus preclude any definitive statement on the economic justification of such systems.

Within the scope of this research, however, some conclusions can be ventured. Assuming for the purposes of discussion that the typical job shop operates under an average of 80 o/o utilization with an earliest start date (EARS D) priority dispatching rule, one can expect reductions in mean tardiness with the ALLQ alternate routing heuristic something on the order of the curve in Figure IV.1 on the following page.¹⁸ The point made here is that an on-line control system could effect a reduction in

¹⁸The graph was prepared by comparing the mean tardiness of ALLQ--EARS D with the results of EARS D with fixed routing. At $I_a = 0.2$, the pass two results were compared; pass one provided the remaining figures.

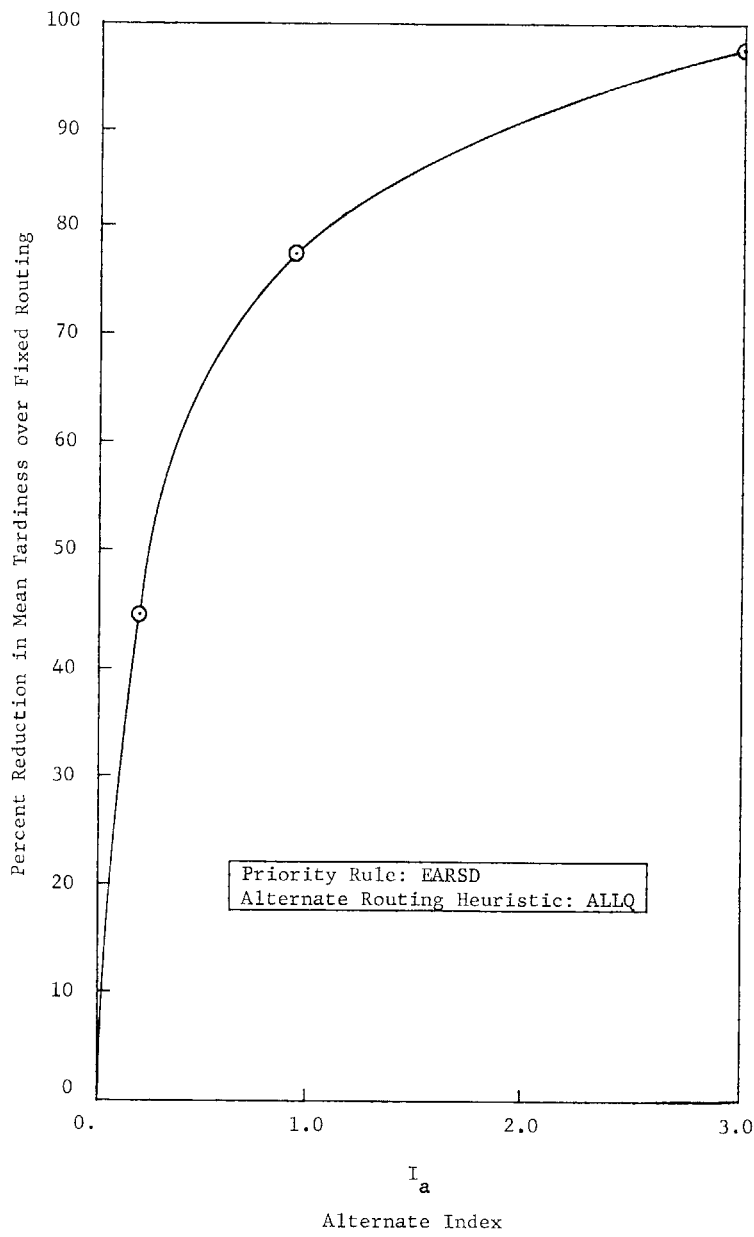


Figure IV.1

tardiness costs on the order of 50 o/o¹⁹ at comparatively low alternate indices with the ALLQ heuristic over a "typical" shop with an EARS D dispatching rule. Moreover, by combining a COVERT priority rule with the ALLQ alternate routing rule, tardiness reductions approach 100 o/o (that is, mean tardiness approaches zero) over a fixed route, EARS D dispatching scheme. The case for real-time computer control of job shops is hereby rested. With demonstrable performance improvements such as have been reported here, the defensibility of the economic feasibility assertion seems assured.

For smaller shops, as well as for large installations with insurmountable objections (economic or otherwise) to computer control, implementation of alternate routing with a transition level heuristic seems eminently reasonable for a manual control system. The extent of the decisions involved and rather limited control problems would well justify a manually executed, systematic approach to alternate routing. Through none of these rules has been tested here at the low level of alternate incidence, reductions of 30 to 40 percent in mean tardiness can be projected from the results obtained elsewhere in this research. After the initial throes of changeover, there seems no reason to expect significantly larger control problems with alternate routing than would exist with fixed routing, sporadically modified by the various expediting techniques now in use. This presupposes, of course, that a well thought out, well administered dispatching and routing control mechanism is used.

¹⁹This is in accordance with the assumption of a linear relationship between tardiness and cost ventured during the brief discussion on the COVERT priority discipline.

5. Suggestions for Further Research

Though one hesitates to call this thesis "ground-breaking," the fact nevertheless remains that alternate routing in job shops has received little or no attention to date. Consequently, much additional research effort is needed to both confirm (or disprove) what is reported here and to advance the present state of knowledge. The work done here has demonstrated the rich harvest that can be expected from additional cultivation of this field.

As an initial suggestion, it would seem that the evidence indicates that the application of alternate routing heuristics at what has here been called the operational level has more promise than transitional rules. It is hoped, therefore, that further research will concentrate on developing more sophisticated heuristics than ALLQ that will operate on a comparable decision level. The opinion is offered, however, that decreasing returns to scale are to be expected beyond the level of sophistication achieved here. What seems immediately of interest is a method of further analyzing the other queues containing alternate operations once an ALLQ type rule allows the priority discipline to select one of the operations as the highest priority within its queue. In addition, some look ahead might be worth investigating.

A more immediate extension of this thesis might be to attempt what has, because of time limitations, been left undone. The effect of alternate routing on differently loaded shops is an obvious candidate for further study. Here only an 80 o/o load has been tested. Other levels of utilization should also receive attention. Higher loads particularly offer considerable experimental interest. Another item left open is an attempt to program alternate routing in conjunction with look ahead heuristics of

the "hold-off" and "sneak-in" variety that have been attached to some priority disciplines. The obvious difficulty is that heuristics of these types look at impending arrivals at machine queues before assigning an operation to a machine. Where alternate routing exists, however, it becomes impossible to predict with certainty which orders are going to arrive within the period of look ahead. Some probabilistic model, perhaps utilizing a Bayesian approach, might make such compatibility possible.

Alternate routing in job shops holds the promise of providing an interesting area for the research efforts of a considerable number of people in the near future. As a closing note to this thesis, a word of caution is offered to potential investigators. As previously ventured, data organization is of paramount importance in any attempt to implement a higher level alternate heuristic, of which ALLQ is a representative example. The researcher who embarks upon a voyage in this area without providing his model with an effective referencing mechanism for a central file is due for some frustrating moments. It was extremely fortuitous that such a problem was not encountered here.

APPENDIX A


First Pass Results (Arbitrary Time Units)

Alternate Routing Heuristic	Priority Dispatching Rule	No Alternates		All 2-Member Strings		All 4-Member Strings	
		Mean Tardiness	Mean Wait Time	Mean Tardiness	Mean Wait Time	Mean Tardiness	Mean Wait Time
Lowest Average Priority (LOAP)	FCFS	41.9	15.4	18.1	12.5	4.02	9.31
	SHTOP 1	13.1	7.73	12.6	6.82	8.59	5.99
	EARS	33.5	15.3	11.9	12.6	2.73	9.59
	TRSIO 2	6.13	8.86	4.30	7.90	2.13	6.42
	SLKROP	19.5	14.4	6.23	12.5	0.0957	10.1
	COVERT SPM.5	1.97	9.32	0.963	8.56	0.443	6.99
	COVERT SPP0	3.06	11.3	2.00	10.5	0.680	8.69
	COVERT SPP.5	2.89	11.2	2.17	10.6	1.32	9.38
	COVERT SPP1	4.15	10.8	2.82	10.1	1.11	8.86
	Maximum Average Slack Per Remaining Operation (MASRO)	FCFS	41.9	15.4	22.8	12.9	7.85
SHTOP		13.1	7.73	10.8	7.18	5.03	6.22
EARS		33.5	15.3	11.7	12.7	3.09	9.82
TRSIO 2		6.13	8.86	3.12	7.93	1.22	6.36
SLKROP		19.5	14.4	6.23	12.5	0.0957	10.1
COVERT SPM.5		1.97	9.32	1.14	8.70	0.383	7.09
COVERT SPP0		3.06	11.3	1.28	9.98	0.242	7.98
COVERT SPP.5		2.89	11.2	1.36	10.5	0.242	8.70
COVERT SPP1		4.15	10.8	2.39	10.0	0.437	8.65
Maximum Average Processing Time (MAPT)		FCFS	41.9	15.4	29.8	13.92	14.7
	SHTOP 1	13.1	7.73	12.6	6.82	8.59	5.99
	EARS	33.5	15.3	17.1	13.22	4.56	10.62
	TRSIO	6.13	8.86	4.03	7.80	2.68	6.67
	SLKROP	19.5	14.4	7.58	12.91	0.799	10.07
	COVERT SPM.5	1.97	9.32	1.44	8.85	0.652	7.31
	COVERT SPP0	3.06	11.3	1.54	9.79	0.875	8.64
	COVERT SPP.5	2.89	11.2	1.54	10.39	0.807	8.43
	COVERT SPP1	4.15	10.8	1.61	9.96	0.848	8.50

APPENDIX A (Continued)

Alternate Routing Heuristic	Priority Dispatching Rule	No Alternates		All 2-Member Strings		All 4-Member Strings	
		Mean Tardiness	Mean Wait Time	Mean Tardiness	Mean Wait Time	Mean Tardiness	Mean Wait Time
Maximum Average Critical Start Date (MACSD)	FCFS	41.9	15.4	29.1	13.76	7.38	9.80
	SHTOP	13.1	7.73	9.62	7.16	2.91	5.93
	EARSD	33.5	15.3	12.4	12.61	2.51	9.68
	TRSTO	6.13	8.86	2.81	7.74	1.16	6.23
	SLKNOP	19.5	14.4	6.21	12.34	0.135	9.69
	COVERT SPM.5	1.97	9.32	.801	8.56	0.145	6.76
	COVERT SPPO	3.06	11.3	1.15	9.64	0.174	7.30
	COVERT SPP.5	2.89	11.2	1.57	10.35	0.506	8.53
	COVERT SPPI	4.15	10.8	1.85	10.12	0.758	8.45
	ALLQ						
	FCFS	41.9	15.4	14.3	11.8	1.50	7.87
	SHTOP	13.1	7.73	9.66	6.25	9.26	5.23
	EARSD	33.5	15.3	7.34	11.6	0.877	8.77
	TRSTO	6.13	8.86	2.49	7.24	0.912	6.13
	SLKNOP	19.5	14.4	1.45	11.7	0.0254	9.83
	COVERT SPM.5	1.97	9.32	0.359	7.78	0.0391	5.89
	COVERT SPPO	4.14	11.6	0.217	8.83	0.0488	6.82
	COVERT SPP.5	4.23	11.4	0.885	9.82	0.0957	7.90
	COVERT SPPI	3.35	11.1	0.908	9.70	0.2188	8.10

1. Both cells tagged "1" are equivalent.
2. Both cells tagged "2" are equivalent.



APPENDIX B (Continued)

Alternate Routing Heuristic	Priority Dispatching Rule	All 2-Member Strings		All 4-Member Strings	
		% Reduction in Tardiness		% Reduction in Tardiness	
Maximum Average Processing Time (MAPT) (Cont'd)	SLKROP	61.1		95.9	
	COVERT SPM.5	26.9		66.9	
	COVERT SPP0	49.6		71.4	
	COVERT SPP.5	46.7		72.1	
	COVERT SPP1	61.2		79.6	
	Average	40.1 %		69.8 %	
Maximum Average Critical Start Date (MACSD)	FCFS	30.5		82.4	
	SHTOP	26.5		77.8	
	EARS	62.9		92.5	
	TRSI0	54.1		81.0	
	SLKROP	68.2		99.3	
	COVERT SPM.5	59.3		92.6	
	COVERT SPP0	62.4		94.3	
	COVERT SPP.5	45.6		82.5	
	COVERT SPP1	55.4		81.7	
		Average	51.7 %		87.1 %
ALIQ	FCFS	65.8		96.4	
	SHTOP	26.3		29.3	
	EARS	78.1		97.4	
	TRSI0	59.3		85.1	
	SLKROP	92.5		99.9	
	COVERT SPM.5	81.8		98.0	
	COVERT SPP0	95.0		98.9	
	COVERT SPP.5	79.1		97.7	
	COVERT SPP1	73.0		90.5	
		Average	72.3 %		88.1 %

APPENDIX C

Second Pass Results

Alternate Routing Heuristic: ALLQ

Alternate Index: $I_a = 0.20$

Distribution of Strings: 85% one-member
10% two-member
5% three-member

Priority Dispatching Rule	No Alternates		$I_a = 0.20$		Percent Reduction in Tardiness	Percent Reduction in Wait Time
	Mean Tardiness	Mean Wait Time	Mean Tardiness	Mean Wait Time		
FCFS	36.58	14.40	22.23	12.51	39.2	13.1
SHTOP	11.32	7.00	10.62	6.54	6.2	6.6
EARS	24.71	14.20	13.75	12.57	44.3	11.5
TRIO	4.62	7.96	2.90	7.11	37.2	10.6
SLKROP	16.22	13.90	7.53	12.69	53.5	8.7
COVERT SPM.5	1.430	8.78	0.686	7.79	52.0	11.3
COVERT SPP0	2.536	10.18	1.370	9.51	46.4	6.6
COVERT SPP.5	2.588	10.40	1.658	9.67	34.8	7.0
COVERT SPP1	3.461	10.12	2.577	9.59	25.5	5.2

APPENDIX D

THE ALTERNATE INDEX: A MEASURE OF POTENTIAL

The allowance for alternate routing, by adding a dimension of decision making to the dispatching process, contains a potential for improving job shop performance by reducing order tardiness. The amount of this potential is directly related to the length and frequency distribution of the alternate strings that can be allowed. The types of distributions possible are legion. In some shops, the strings might be short and rather frequent in occurrence. Others may have few long strings of alternates. The point is that a common dimension is desirable for measuring alternate incidence that is directly related to the potential for improvement that exists. This measure should allow for comparisons of distributions of alternate strings that masks out the differences in string length and frequency distribution while preserving the one feature of partially ordered operations that is of interest: the potential for reducing tardiness.

The role assigned to the measure desired is thus equivalent to that fulfilled by the familiar statistical entity the arithmetic mean. The measure proposed here, called the alternate index, displays the required properties. It has the following physical property: it is the average number of alternates existing for each operation. This quantity loses physical significance at low levels of incidence when it becomes a fraction. Thus, the adoption of the term "index".

The alternate index I_a is defined as follows:

Let n_i = length of all strings of length n , $n = 1, 2, \dots, \infty$

f_i = fraction of strings in the order file of length n

k = number of different string lengths in the order file.

$$\text{Then, } I_a = \sum_{i=1}^k (n_i - 1) f_i .$$

APPENDIX E

STATISTICAL ANALYSIS OF EXPERIMENTAL RESULTS

1. Alternate Routing and Mean Tardiness

The first hypothesis to be tested is whether or not there are significant reductions in mean tardiness through alternate routing. To test the hypothesis, the Wilcoxon test²⁰ was applied to the data set with the smallest percent reduction in mean tardiness: MAPT with two-member alternate strings. If the poorest case can claim statistically lower tardiness times, so can those better.

The formulation is thus:

H_0 : The mean tardiness with two-member strings routed with MAPT is no better than mean tardiness with fixed routing.

H_1 : Mean Tardiness with MAPT and two-member strings is lower than with fixed routing (a one-tail test).

PRIORITY RULE	TARDINESS NO ALTS	TARDINESS 2-MEMBER STRINGS	RANK OF d	RANK WITH LESS FREQUENT SIGN
FCFS	41.9	29.8	12.1	8
SHTOP	13.1	12.6	0.5	1
EARSD	33.5	17.1	16.4	9
TRSIO	6.13	4.03	2.10	5
SLKROP	19.5	7.58	11.92	7
COVERT SPM.5	1.97	1.44	0.53	2
COVERT SPPO	3.06	1.54	1.52	4

²⁰Siegel, S., Nonparametric Statistics, McGraw-Hill, New York, 1956 contains a discussion of the Wilcoxon test for paired samples on pp. 75-83. Critical values of T for different sample sizes are given on page 254.

PRIORITY RULE	TARDINESS NO ALTS	TARDINESS 2-MEMBER STRINGS	d	RANK OF d	RANK WITH LESS FREQUENT SIGN
COVERT SPP.5	2.89	1.54	1.35	3	
COVERT SPP1	4.15	1.61	2.54	6	

T = 0
N = 9

At a significance level $\alpha = 0.005$ the critical value of T is 2. Therefore, the null hypothesis must be rejected. Mean tardiness with alternate routing is significantly lower for MAPT, the poorest performer and, by deduction even more so for the better heuristics.

2. The Best Alternate Routing Heuristic

The concern is now whether anyone of the alternate routing heuristics tested can be said to be superior to the others. The Wilcoxon test was applied to the two heuristics with the largest percent reduction in mean tardiness for the two member strings (the more realistic of the two incidence cases): ALLQ and MACSD. Their paired percentage reductions in tardiness, as tabulated in Appendix B, will be compared.

H_0 : ALLQ and MACSD produce equivalent reductions in mean tardiness.

H_1 : ALLQ produces greater reductions than MACSD (a one-tail test).

PRIORITY RULE	ALLQ: PERCENT TARDINESS REDUCTION	MACSD: PERCENT TARDINESS REDUCTION		RANK OF d	RANK WITH LESS FREQUENT SIGN
FCFS	65.8	30.5	35.3	9	
SHTOP	26.3	26.5	- 0.2	-1	1
EARSD	78.1	62.9	15.2	3	
TRSIO	59.3	54.1	5.2	2	
SLKROP	92.5	68.2	24.3	6	
COVERT SPM.5	81.8	59.3	22.5	5	
COVERT SPPO	95.0	62.4	32.6	7	
COVERT SPP.5	79.1	45.6	33.5	8	
COVERT SPP1	73.0	55.4	17.6	4	
					T = 1 N = 9

At a significance level $\alpha = 0.005$ the critical value of T is 2.

Therefore, the null hypothesis must be rejected. ALLQ is significantly better than MACSD in reducing mean tardiness, and, by deduction better than the other alternate routing heuristics tested.

3. Pass Two ALLQ Performance

The question to be considered is whether or not ALLQ with an alternate index of 0.2 produces significantly lower mean tardiness than the shop without alternate routing. Appendix C alone should suffice as proof of a large reduction in tardiness. For the sake of formalism, however, the Wilcoxon test is again applied.

H_0 : Fixed routing and the ALLQ heuristic with $I_a = 0.2$ produce equivalent values of mean tardiness.

H_1 : ALLQ with an alternate index of 0.2 produces lower tardiness than fixed routing (a one-tail test).

PRIORITY RULE	TARDINESS NO ALTS	TARDINESS ALLQ $I_a = 0.2$	d	RANK OF d	RANK WITH LESS FREQUENT SIGN
FCFS	36.58	22.23	14.35	9	
SHTOP	11.32	10.62	0.70	1	
EARS	24.71	13.75	10.96	8	
TRSIO	4.62	2.90	1.72	6	
SLKROP	16.22	7.53	8.69	7	
COVERT SPM.5	1.430	0.686	0.744	2	
COVERT SPFO	2.556	1.370	1.186	5	
COVERT SPP.5	2.588	1.658	0.930	4	
COVERT SPP1	3.461	2.577	0.884	3	

T = 0
N = 9

At a significance level $\alpha = 0.005$ the critical value of T is 2. Therefore, the null hypothesis must be rejected. At an alternate index of 0.2 the ALLQ heuristic is significantly better than fixed routing with respect to mean tardiness.

BIBLIOGRAPHY

1. Carroll, D. C., "Heuristic Sequencing of Single and Multi-Component Orders", an unpublished Ph.D. dissertation submitted at Massachusetts Institute of Technology, June, 1965.

2. Siegel, S., Nonparametric Statistics, McGraw-Hill, New York, 1956.

The following books and articles are suggested as general reference material:

3. Baker, C. T. and Dzielinski, B. P., "Simulation of a Simplified Job Shop", Management Science, (VI, April, 1960), 311-323.
4. Conway, R. W., Johnson, B. M., and Maxwell, W. L., "An Experimental Investigation of Priority Dispatching", Industrial Engineering, (XI, May, 1960), 221-230.
5. Conway, R. W., Johnson, B. M., and Maxwell, W. D., "Some Problems in Digital Systems Simulation", Management Science, (VI, October, 1959), 92-110.
6. Emery, J. C., "An Approach to Job Shop Scheduling Using a Large-Scale Digital Computer", Industrial Management Review, (III, Fall, 1961), 78-96.
7. Gere, W. S., A Heuristic Approach to Job Shop Scheduling, an unpublished Ph.D. thesis submitted at Carnegie Institute of Technology, September, 1962.
8. International Business Machines, General Information Manual: Improved Job Shop Management Through Data Processing, 1960.
9. Minsky, M., "Steps Toward Artificial Intelligence", Proceedings of the IRE, (Vol. 49, No. 1, January, 1961).
10. Muth, J. F. and Thompson, G. L., eds., Industrial Scheduling, Prentice Hall, Englewood Cliffs, N.J., 1963; selected articles:

Allen, M., "The Efficient Utilization of Labor Under Conditions of Fluctuating Demand", 252-276.

Conway, R. W., and Maxwell, W. L., "Network Scheduling by the Shortest-Operation Discipline", 277-299.

Fischer, H. and Thompson, G. L., "Probabilistic Learning Combinations of Local Job-Shop Scheduling Rules", 225-251.

Jackson, J. R., "Queues with Dynamic Priority Discipline", 308-326.

Pounds, W. F., "The Scheduling Environment", 5-12.

BIBLIOGRAPHY (Cont'd)

11. Nanot, Y. R., An Experimental Investigation and Comparative Evaluation of Priority Disciplines in Job Shop-Like Queues, Management Sciences Research Project Report No. 87, Western Management Sciences Institute, University of California, Los Angeles.
12. Rowe, A. J., "Toward a Theory of Scheduling", Industrial Engineering, (XI, March, 1960), 125-136.
13. Simon, H. A., The New Science of Management Decision, Harper and Row, New York, 1960.

CS-TR Scanning Project
Document Control Form

Date : 12/11/95

Report # LCS-TR-19

Each of the following should be identified by a checkmark:
Originating Department:

- Artificial Intelligence Laboratory (AI)
 Laboratory for Computer Science (LCS)

Document Type:

- Technical Report (TR) Technical Memo (TM)
 Other: _____

Document Information

Number of pages: 54(60-images)
Not to include DOD forms, printer instructions, etc... original pages only.

Originals are:

- Single-sided or
 Double-sided

Intended to be printed as :

- Single-sided or
 Double-sided

Print type:

- Typewriter Offset Press Laser Print
 InkJet Printer Unknown Other: _____

Check each if included with document:

- DOD Form Funding Agent Form Cover Page
 Spine Printers Notes Photo negatives
 Other: _____

Page Data:

Blank Pages (by page number): _____

Photographs/Tonal Material (by page number): _____

Other (note description/page number):

Description :	Page Number:
<u>IMAGEMAP: (1-54) UN#ED TITLE & BLANK PAGES, i-iii,</u>	
<u>UN#ED BLANK, V-Vii, UN#ED BLANK,</u>	
<u>1-44</u>	
<u>(55-60) SCANCONTROL, COVER, DOD, TRGT'S (3)</u>	

Scanning Agent Signoff:

Date Received: 12/11/95 Date Scanned: 1/18/96

Date Returned: 1/11/96

Scanning Agent Signature: Michael W. Cook

UNCLASSIFIED
Security Classification

DOCUMENT CONTROL DATA - R&D		
<i>(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)</i>		
1. ORIGINATING ACTIVITY <i>(Corporate author)</i> Massachusetts Institute of Technology Project MAC	2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	2b. GROUP
3. REPORT TITLE A Heuristic Approach to Alternate Routing in a Job Shop		
4. DESCRIPTIVE NOTES <i>(Type of report and inclusive dates)</i> Master's Thesis, Sloan School of Management		
5. AUTHOR(S) <i>(Last name, first name, initial)</i> Russo, Francis John		
6. REPORT DATE June 1965	7a. TOTAL NO. OF PAGES 51	7b. NO. OF REFS 13
8a. CONTRACT OR GRANT NO. Office of Naval Research, Nonr-4102(01)	9a. ORIGINATOR'S REPORT NUMBER(S) MAC-TR-19 (THESIS)	
b. PROJECT NO. Nr-048-189	9b. OTHER REPORT NO(S) <i>(Any other numbers that may be assigned this report)</i>	
c.		
d.		
10. AVAILABILITY/LIMITATION NOTICES Qualified requesters may obtain copies of this report from DDC.		
11. SUPPLEMENTARY NOTES None	12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency 3D-200 Pentagon Washington, D. C. 20301	
13. ABSTRACT <p>The research reported investigates the use of heuristics for selecting from several alternate routes, resulting from partially-ordered tasks in a job-shop order file. The experimental vehicle employed was digital simulation on the Project MAC Compatible Time-Sharing System.</p> <p>The methodology was to conduct two passes of simulated shop runs. The first, with two artificially high levels of alternate incidence, tested the efficiency of five different alternate routing heuristics in reducing mean tardiness. The second pass consisted of runs with the best heuristic developed during the first experimental phase, applied to a realistic length and frequency of alternate strings. The performance characteristics of the different heuristics are discussed at length. In addition, some implications are drawn of the computational nature of alternate routing and the difficulties encountered in implementing alternate routing heuristics at operation time.</p>		
14. KEY WORDS Computer On-line computer systems Machine-aided cognition Real-time computer systems Multiple-access computers Time-sharing Time-shared computer systems		

Scanning Agent Identification Target

Scanning of this document was supported in part by the **Corporation for National Research Initiatives**, using funds from the **Advanced Research Projects Agency of the United States Government** under Grant: **MDA972-92-J1029**.

The scanning agent for this project was the **Document Services** department of the **M.I.T. Libraries**. Technical support for this project was also provided by the **M.I.T. Laboratory for Computer Sciences**.

