OPERATIONS RESEARCH

DEPARTMENT

OPERATIONS RESEARCH DEPARTMENT

SCHOOL OF MANAGEMENT

CASE WESTERN RESERVE UNIVERSITY

UNIVERSITY CIRCLE • CLEVELAND, OHIO 44106

28

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Research and Development						
Project Network						
Multiple Project						
Limited Resources						
Heuristic Programming						
Scheduling						

Unclassified
Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Department of Operations Research School of Management Case Western Reserve University, Cleveland, Ohio	2a. REPORT SECURITY CLASSIFICATION Unclassified 2b. GROUP

3. REPORT TITLE

MULTIPLE R AND D PROJECT SCHEDULING WITH LIMITED RESOURCES

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

5. AUTHOR(S) (First name, middle initial, last name)

Corwin, Burton D.

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
September, 1968	24	10

8a. CONTRACT OR GRANT NO. DAHC 19-68-C-0007 b. PROJECT NO. NONR-1141(19) NR277-019 542-3120-2104 c. 542-3120-2510 d.	9a. ORIGINATOR'S REPORT NUMBER(S) Technical Memorandum No. 122 9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)

10. DISTRIBUTION STATEMENT

Distribution of this document is unlimited

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
	Army Research Office Office of Naval Research

13. ABSTRACT

This article considers the application of a heuristic scheduling model for multiple project scheduling with limited resources to R and D projects. The basic operation of the model and its underlying heuristics are described. The use of this scheduling model for multiple R and D project scheduling is discussed.

DD FORM 1 NOV 65 1473

Unclassified
Security Classification

MULTIPLE R AND D PROJECT SCHEDULING
WITH LIMITED RESOURCES

Burton D. Corwin

Technical Memorandum No. 122

September, 1968

This report was prepared as part of the activities of the Department of Operations Research, School of Management, Case Western Reserve University (under Contract Number DAHC 19-68-C-0007 with the Army Research Office and NONR-1141(19) NR277-019 with the Office of Naval Research).

ACKNOWLEDGMENT

I wish to express my thanks to Mr. J. David Cartlidge, Research Operations Analyst, Research and Development Division, the Babcock and Wilcox Company, for the assistance gained through many discussions with him and for discussions with members of the Operations Research Department at Case Western Reserve University.

INTRODUCTION

This research effort has been directed toward an investigation of improved methods of planning and controlling R and D projects in development-oriented governmental laboratories. Specifically, computer-oriented techniques for scheduling single and multiple projects with limited availability of certain key resources have been considered. For purposes of this study, a key resource is one whose limited availability will have a significant effect on the project scheduler's ability to schedule one or more activities requiring this particular resource in some given scheduling period. For the most part, the key resources with limited availability employed by R and D projects will be certain manpower skill classifications.

In their conventional form, such single project planning and scheduling techniques as PERT and CPM fail to consider the constraints imposed on activity scheduling when resource availabilities are limited. Also, these techniques are single project oriented, and as such, do not explicitly handle the multi-project problem normally faced by the R and D laboratroy. An activity associated with any one of the laboratory's many current projects may require the use of one or more types of manpower. In each period of the scheduling horizon, there is a limit on the availability of each of these types of manpower. Thus, an investigation of the applicability of currently available computer-oriented models capable of scheduling multiple projects within resource limits seems appropriate. In this study, primary attention will be given to one such model called SPAR-1 (Scheduling Program for Allocation of Resources) because it is the one most comprehensively oovered in the literature.

DISCUSSION OF HEURISTIC SCHEDULING MODELS FOR LARGE PROJECTS WITH LIMITED RESOURCES

THE SCHEDULING PROBLEM UNDER CONSIDERATION

We will begin this discussion with a description of the scheduling problem being considered. Assume that one is given a number of projects at various stages of completion. Each consists of many separate activities or jobs, at least some of which must be done in a given sequence for technological reasons. A project network diagram in which each activity appears as an arrow and the connections of arrows indicate the predecessor-successor relationships, i.e., technological constraints, between the individual activities can be constructed for each project. Note that in the case of a partially completed project, one is only concerned with those activities that are not yet completed, and only these need to appear in the project network diagram. Then a single amalgamated network is formed from the individual project networks by appending to the beginning and end of each, a dummy activity and a pseudo activity respectively. This single amalgamated project is treated as one large project.

Each of the activities in this large project normally requires one or more resource types and a given time for completion. This time for completion, i.e., the activity duration, may vary with the rate at which resources are applied. Obviously, for a multi-resource activity, the level of resource application must be the same for each of the involved resource types, The activities in the large project that do not require the use of any of the key resources will be either pseudo activities or dummy activities. Pseudo activities show dependency relationships and have an associated activity duration, while dummy activities show dependency relationships only.

Given the large project described above and the constraints in the form of limited resource availabilities in the successive scheduling periods, one must find a schedule that satisfies the technological and resource availability constraints and minimizes the overall duration of the single large project. By proper selection of the durations of the pseudo activities appended to the end of each project this objective can be closely related to the one of meeting specified individual project completion dates or at least minimizing their overruns. This procedure will be discussed in more detail later. Finally, it is important to observe that finding a schedule is equivalent to determining when each activity will be started and at what level of resource application it will be maintained during each period it is active.

THE BASIC OPERATION OF THE HEURISTIC SCHEDULING MODEL

The available scheduling techniques for the problem described above are based on heuristic programming. Weist (3) defines a heuristic scheduling program as a procedure for generating a project schedule on the basis of one or more scheduling rules. Each scheduling rule used is a heuristic of the rule-of-thumb variety. Insight into the basic approach employed by these available scheduling techniques is provided by Weist's description of the basic operation of his heuristic scheduling model SPAR-1. This description is given as follows:

In its basic approach, the model focuses on available resources, which it serially allocates, period by period, to jobs listed in order of their early start times. Jobs are schedules, starting with the first period, by selecting from the list of those currently available and ordered according to their total slack (which is based on technological constraints only and normal resource assignments). The jobs in this list are scheduled sequentially, starting with the first, and as many jobs are scheduled as available resources permit. If an available job fails to be scheduled in that period, an attempt is made to schedule it the next period. Eventually all jobs so postponed become critical and move to the top of the priority list of available jobs (3). The basic flow diagram of SPAR-1 is shown in Figure 1.

The actual scheduling process performed in each successive period requires the period by period maintenance and appropriate updating of essentially two separate lists of activities. The first is a list identifying the activities continued from the previous period. For each activity in this list, information concerning its current resource assignment level, its current scheduled completion times, and its current total slack value is kept. From one's knowledge of the current resource assignment level for each of these activities, together with information regarding the type of resource required by the activity and the quantity of this resource associated with each resource assignment level provided by a master list of all activities in the large project, one can readily determine the amount of each resource type available to be allocated at the beginning of the present period. The second list identifies the activities currently available to be scheduled. Information concerning the type of resource required, along with the quantity of resource and activity duration associated with each level of resource assignment, and the current total slack value would be included for each activity in this list. The activities in this list are ordered according to increasing current total slack value.

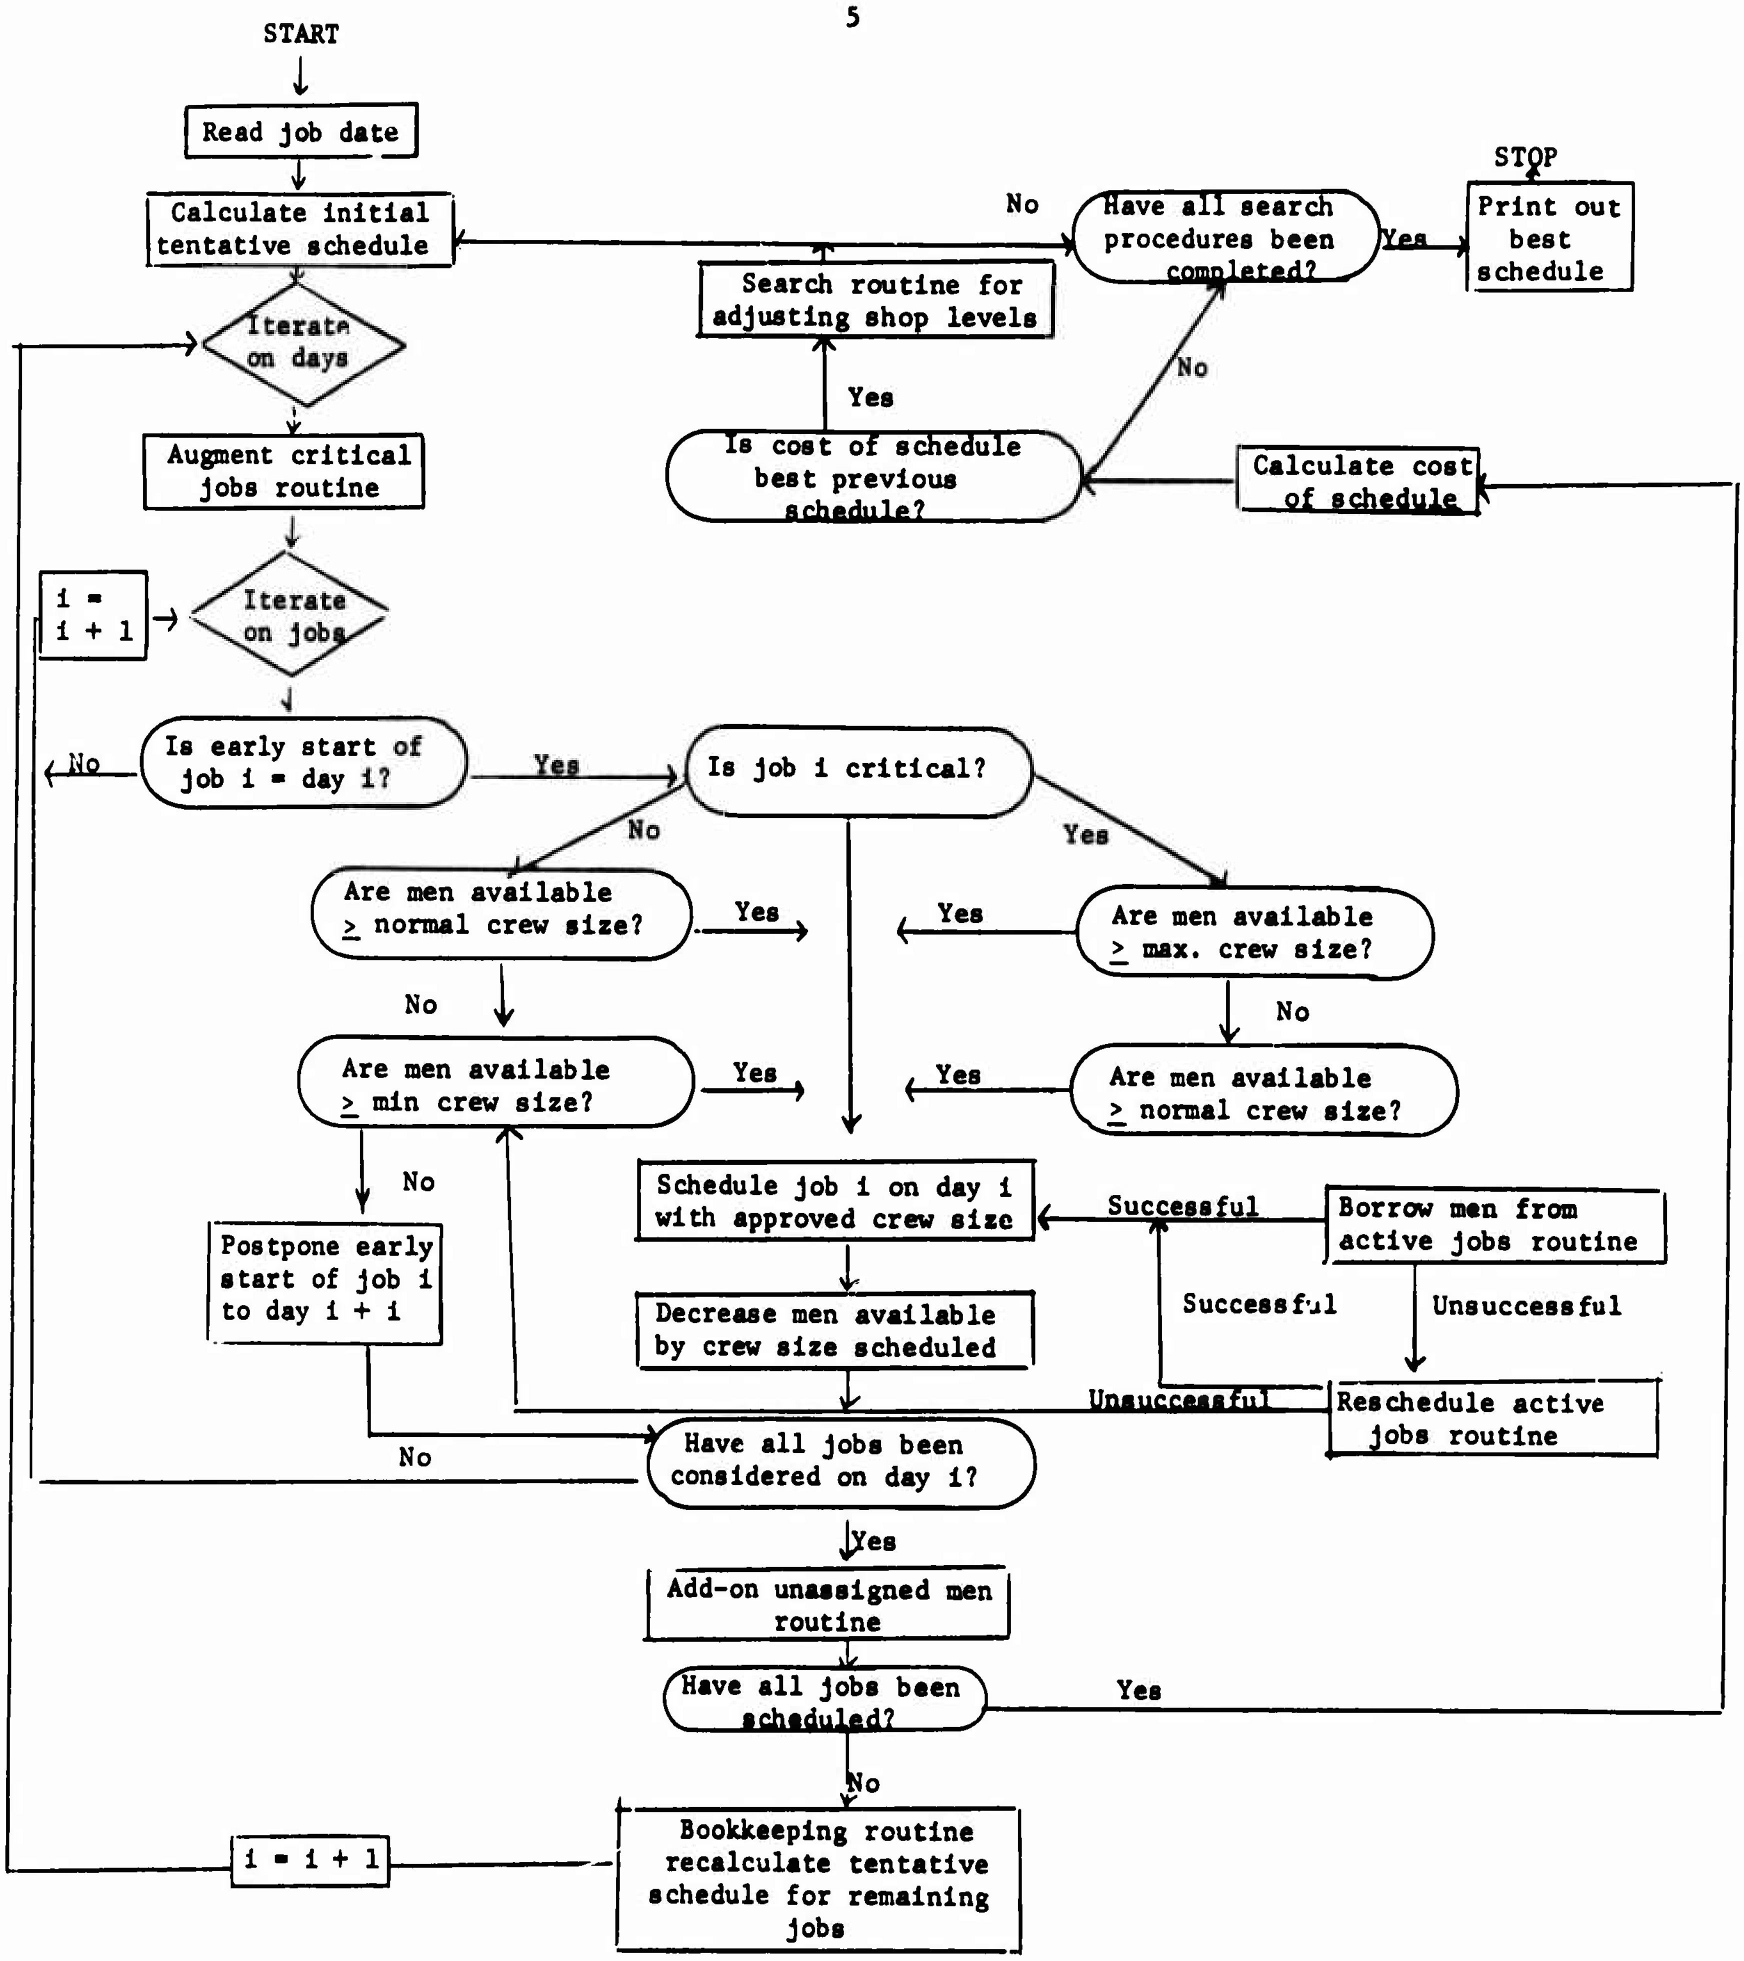

Figure 1

Flow diagram for SPAR-1 from J. D. Wiest, "A Heuristic Model for Scheduling Large Projects with Limited Resources", Management Science, Vol. 13, No. 6 (February, 1967), pp. 359-77.

At the beginning of the present period, the list of activities currently available for scheduling is scanned sequentially. Activites are scheduled until either the list is exhausted or until the limits for each resource type are reached. If a sufficient quantity of resource is not available to schedule a currently critical activity at its normal resource assignment level, the Borrow from Active Activities and/or the Reschedule Active Activities scheduling heuristics described below are brought into play. The list of activities continued from the previous period provides the candidate activities for both of these scheduling heuristics. It is important to note that if an activity with a current total slack value of zero is delayed from starting at the beginning of the present period, the overall duration of both its associated project and the single large project will increase by one.

As a result of the scheduling process for the present period, each activity in the list of activities continued from the previous period will be (1) continued at the same level of resource application during the present period, (2) continued at a higher level of resource application during the present period as a result of either the heuristic for augmenting critical activities or the heuristic for adding-on unused resources, both of which are described below, (3) continued at a lower level of resource application during the present period as a result of the Borrow from Active Activities routine, or (4) de-scheduled for the present period as a result of the Reschedule Active Activities routine. Note that de-scheduled activities are available for scheduling at the beginning of the next period. Some of the activities continued in the present period may be completed at the end of the present period. The completion of an activity at the end of the present period will make one or more immediately succeeding activities available for scheduling at the beginning of the next period if the other immediate predecessor activities for each of these activities are also completed. Also, some of the activities in the list of activities currently available for scheduling may not be scheduled during the present period. Thus, at the end of the present period, both lists must be appropriately updated to reflect the changes in both the activity content of each list and the relevant information recorded for each activity. The current total slack value of each activity not completed by the beginning of the next period is updated by performing a standard CPM analysis using the remaining activity duration for each activity. For partially completed activities these remaining durations are based on the current resource application level associated with the activity, whereas for those activities not yet started, they are based on the assumption of normal resource application levels. The updated versions of both lists form the input for the scheduling process in the next period. The scheduling process thus proceeds from period to period until all activities have been scheduled.

THE SCHEDULING HEURISTICS INCORPORATED IN THE MODEL

The scheduling problem under consideration is essentially a limited resource allocation problem. Whenever the availability of a certain resource type is less than the total requirement of both currently scheduled and currently available activities, a choice in allocating it must be made. Thus the heuristic scheduling model must incorporate rules of priority and policy, i.e., scheduling heuristics, to enable the available resources of each type to be allocated efficiently. There are essentially two heuristics underlying the basic approach described above. First, resources are allocated serially in time. That is, the model starts on the first day and schedules all jobs possible, then does the same for the second day, and so on. Secondly, when several activities are competing for the same resource, preference is given to these activities with the least total slack. Thus, total slack is used to confer relative priorities between competing activities. If the durations of the pseudo activities are selected properly, the total slack of an activity is actually measured with respect to the desired completion date of its associated project.

Pseudo activities can be employed to establish relative priorities between individual projects, and thus between activities in different projects. Assume that one has K individual projects in various stages of completion at schedule time zero, a project due date d_i for each individual project $i = 1,2,\ldots,K$, and a scheduling horizon $d \geq d_i \ \forall \ i = 1,2,\ldots,K$. The d_i values could be arbitrarily specified by management. One should observe, however, that in most cases it would be impractical to assign a d_i value smaller than the remaining overall project duration calculated by performing a standard CPM analysis using the remaining activity durations associated with a normal level of resource application for each activity in project i. The value $(d-d_i)$ is esentially the total slack value of project i with respect to the other $(K-1)$ projects. If d_i is obtained by means of the CPM analysis just described, the total slack for each activity in project i at schedule time zero would be its within project i total slack, which is measured with respect to d_i, plus $(d-d_i)$. Irregardless of how d_i is obtained, one can append a pseudo activity with duration less than or equal to $(d-d_i)$ to project i in order to decrease the total slack of each activity in project i with respect to time d in any scheduling period by this same amount. Thus one can assign a total slack value to each project i with respect to time d in order to reflect the priority of its activities relative to the activities of the other (K-1) projects. If the duration of the pseudo activity appended to project i is set equal to $(d-d_i)$, $i = 1,2,\ldots,K$, the scheduling model applied to the single large project tends to enforce the due dates d_i of the individual projects i.

The operation of the basic scheduling model defined by the two heuristics described above can be modified by a number of additional scheduling heuristics designed to increase the use of available resources and/or decrease the overall duration of the single large project. These additional heuristics essentially represent the rules of policy that are incorporated in the heuristic scheduling model. Several of the ones employed by Weist in his SPAR-1 scheduling model (3) will be described briefly. In the discussion that follows, we will refer to the level of application of a resource type to an activity as the resource level.

The first modifying heuristic to be described refers to the selection of a resource level for an activity to be scheduled. Associated with each activity are three different resource levels, namely normal, maximal and minimal, and their corresponding activity durations. A normal resource level denotes the number of units of the required resource normally assigned to an activity. A maximal resource level refers to the number of units of the required resource needed for crashing the activity, while a minimal resource level refers to the smallest number of units of the required resource that can be assigned to it. The rules for resource level selection are given as follows. If an activity to be scheduled is currently critical and a sufficient amount of the required resource is available, the activity is scheduled at its maximum resource level. If an insufficient amount of the required resource is available to do so, or even to schedule it at its normal resource level, an attempt is made to obtain the additional units of resource needed to reach its normal resource level by means of the Borrow from Active Activities and Reschedule Active Activities routines described below. In case the activity can not be scheduled at its normal, or even its minimal resource level despite the use of these two routines, its start date is delayed one period. As a result, since the activity is currently critical, the overall duration of both its associated project and the single large project is increased by one period. This activity will remain critical at the beginning of the next period, and the model will attempt to schedule it then. Each currently non-critical activity is scheduled at its normal resource level if resource availabilities permit. If the amount of the required resource available is not sufficient for scheduling the activity even at its minimum resource level, the activity is delayed for consideration until the next period.

The second modifying heuristic considered deals with augmenting the resource level of currently critical activities which currently have resource levels less than their maximums. Before any activities in the list of activities currently available are scheduled to start in the present period, the activities in the list of activities continued from the previous period are examined. If any of these activities is currently critical and has a resource level assigned less than its maximum, and a sufficient amount of the required resource is available, its resource level for the present period is increased as much as possible up to the maximum. These additions can be regarded as being only temporary and it can be assumed that the involved activities return to their previous level of resource assignment prior to the start of the scheduling process for the next period.

The third modifying heuristic considered is concerned with the treatment of multi-resource activities. If an activity requires more than one resource type, separate activities are created for each resource type, and these activities are constrained to start the same period with the same level of resource assignment. Thus we can assume that there is only one resource type associated with each activity.

The fourth modifying heuristic to be described is the Borrow from Active Activities routine referred to previously. If a sufficient amount of the required resource is not available to schedule a currently critical activity at its normal resource level, then the model enters a procedure for searching the currently noncritical activities in the list of activities continued from the previous period to see if enough additional units of the required resource can be borrowed from these activities to schedule this currently critical activity at its normal resource level. Decreasing the resource assignment level of a currently active activity by borrowing units of resource from it results in an increase in its remaining duration. Now its remaining duration value is based on the assumption that it will be continued in subsequent periods at this new lower resource level. Units of resource are borrowed from an activity only when the resultant increase in its remaining duration does not cause its recalculated current total slack value to become negative, and thus cause the overall duration of both its associated project and the single large project to increase.

The fifth modifying heuristic to be examined was also referred to previously. This is the "Reschedule Active Activities" routine. If a sufficient amount of the required resource is not available to schedule a currently critical activity at its normal resource level and can not be obtained by means of the Borrow from Active Activities routine, the Reschedule Active Activities routine is brought into play. The model scans the currently noncritical activities in the list of activities continued from the previous period and determines those which use the same resource type as the currently critical activity and could be postponed without becoming currently critical. For our purposes, an activity is considered critical if its total slack value does not exceed zero. When an activity previously scheduled is de-scheduled, its earliest start time becomes the beginning of the present period and its remaining duration becomes its total duration assuming a normal level of resource assignment. Therefore, its current total slack must be recalculated based on these adjustments. If this recalculated current total slack value is non-positive, the activity can not be rescheduled since to do so would cause the overall duration of both its associated project and the single large project to increase by at least one unit.

The final modifying heuristic described by Weist deals with adding on unused resources at the conclusion of the scheduling process for a given period. After as many activities as possible are scheduled in a given period, there still may by unused resources of various types available. The model compiles a list of the activities scheduled during the given period to which these resources might be assigned and arranges them in ascending order of their current total slack. Proceeding down the list, the model increases the resource level of these activities until the unused resources or the list of activities is exhausted. These additions are only temporary and the involved activities return to their previous level of resource assignment at the beginning of the next period.

There are two additional features incorporated in SPAR-1 (3) which can also be utilized to reflect rules of policy and as such are essentially also modifying scheduling heuristics. These will now be considered.

The first deals with the interruption of an activity before its completion, i.e., allowing an activity to be processed in either consecutive or intermittent time periods. Weist's heuristic scheduling model SPAR-1 has the capability of treating each activity contained in each of the individual projects as being either completely interruptable or completely uninterruptable. A completely interruptable activity does not need to be scheduled in successive time periods. If such an activity is scheduled during some given period it does not appear in the subsequent period's list of activities continued from the previous period. Rather, it appears in the list of activities currently available for scheduling with its total slack value and remaining duration at each resource level appropriately updated. Thus, in this subsequent period, this activity competes with all other currently available activities for the resource it requires. On the other hand, a completely uninterruptable activity must be scheduled in consecutive time periods. If such an activity is scheduled during some given period and not completed by the end of this period, it must be either continued at the same, or perhaps at a different resource level in the subsequent period, or postponed in its entirety then by means of the Reschedule Active Activities routine.

The second feature is concerned with resource interrelationships. In some scheduling environments there may exist a large degree of substitutability between the various resource types that must be incorporated into the heuristic scheduling model. Some of the activities contained in each of the individual projects may have the property that they can be completed through the use of one of several resource types. For each different resource type that could be applied to such an activity there would exist a distinct value for the activity duration associated with the quantity of either resource corresponding to each level of resource assignment. In Weist's model SPAR-1, any given activity can be specified as requiring one of two alternate resource types. This is done by assigning a dummy resource type to the activity which serves to identify the two actual resource types involved and the order of preference between them. For such an activity the input values given for the activity duration associated with each resource level are those corresponding to the preferred resource type. The activity durations for the substitute resource type are assumed to be linearly related to those for the preferred resource type. Thus, whenever the activity appears in the list of activities currently available for scheduling, its current total slack value is based on the assumption that it will be assigned to its preferred resource type by the scheduling process. Since the substitute resource type will normally have longer activity durations associated with the respective activity levels, the current total slack figure corresponding to the substitute resource type would be smaller than the value being used. When the activity is considered for scheduling, based on its current total slack values, the SPAR-1 model has the capability of first examining the availability of the required amount of the preferred resource type, and then if this availability is not sufficient, examining the availability of this amount of the substitute resource type. If the activity is scheduled with its substitute resource type, this is noted and its current total slack is adjusted appropriately. Once such an activity is scheduled with one of its two resource types it must continue with that resource type until its completion. Note that if the activity is subsequently postponed by means of the Reschedule Active Activities routine it may be rescheduled using either of the two resource types.

ILLUSTRATIONS USING VARIOUS SIMPLIFIED VERSIONS OF THE HEURISTIC SCHEDULING MODEL

A detailed example using a single small project and a simplified version of the heuristic scheduling model described above is found in Weist (4). The project used in this example consists of ten activities, each of which requires a certain amount of time and a given number of units of the same resource. For simplicity, it is assumed that for each activity, the quantity of the resource associated with its minimal, normal and maximal resource levels is the same and that the resource limit is constant throughout the scheduling horizon. The heuristic scheduling model employed by this example is based on three of the heuristics described above. The resource is allocated period by period until all activities are scheduled. In each period scheduling preference is given to those activities with the least total slack. The Reschedule Active Activities routine is brought into paly as required. Also, a similar but somewhat larger example is given in Battersby (1). The Reschedule Active Activities routine is not included in this example. Finally, the reader may be interested in knowing that two different heuristic scheduling procedures for scheduling a particular project so that constraints on resource availabilities are satisfied are given in Moder and Phillips (2). Both of these are considerably different than the heuristic scheduling model described in this section. In the opinion of the author they are both considerably less useful than Weist's SPAR-1. They are both primarily suited for a single project and a single resource type.

THE USE OF THE HEURISTIC SCHEDULING MODEL SPAR-1 FOR R AND D PROJECT SCHEDULING

Before we consider using a heuristic scheduling model such as SPAR-1 for a particular R and D multiproject scheduling problem, it is important to note that we are making the following three assumptions. First, we assume that we are dealing with individual R and D projects that can be represented as project network diagrams. Secondly, we assume that an appropriate scheduling period, i.e., a hour, a day, etc., can be defined for the problem. Finally, we assume that the relevant key resources can be identified and that limits can be assigned to the availability of each one in each period of the scheduling horizon. The limits for each resource type may vary from period to period. In order to apply any heuristic scheduling model to the multi-project scheduling problem faced by a particular R and D laboratory, one must first examine the heuristics that form the basis of the model. Then one must determine if these are appropriate for the given problem, i.e., if they adequately reflect the scheduling rules of priority and policy judged as being important for the given problem. If this is not the case, one must decide if the heuristics incorporated in the model can be altered and/or new ones included without extensive modification of the specific scheduling model or its associated computer program in order to satisfy the requirements of the given problem.

If we were considering the application of SPAR-1 to the given problem it seems that we would examine the appropriateness of each of four features of the model. We would like to determine if these features can be applied in the given scheduling environment without modification. If this is not the case, we must decide what specific modificaitons are required and how much difficulty is involved in incorporating them into the existing SPAR-1 computer program.

The first feature to be considered is the model's association of three different resource levels with each activity. We must determine if the management personnel of the given R and D laboratory are able to define, for each activity, values for the quantity of the required resource and the activity duration associated with each ofthe three resource levels. If only the pair of values associated with normal resource level can be obtained for each activity the model can still be applied. However, it is restrained from exercising several of the modifying heuristics designed to increase the utilization of available resources. If these three pairs of values can be determined for each activity, we must next decide on how intermediate resource levels should be handled. That is, we must decide if only the resource quantities corresponding to each of the three resource levels can be assigned to an activity or if any resource quantity between that corresponding to the minimum and maximum resource levels can be assigned. In the latter case, some assumption as to the form of the relationship between the number of units of the required resource assigned to an activity and its time duration must be made. For the purpose of allowing intermediate resource levels an activity can be considered as consisting of a certain total number of resource periods to be completed. At the end of each period in which an activity is scheduled the number of resource periods remaining to be done can be appropriately updated based on the quantity of resource assigned to it in that period. In order to establish the remaining activity duration required by the updating CPM analysis performed at the end of the period, the adjusted number of resource periods remaining is divided by the quantity of the required resource associated with the resource level (minimum, normal, maximum) assigned during the period. Note that the resource level assigned will be one of these three because additional resources assigned by both the augment critical activities heuristic and the add-on unused resources heuristic are assumed to be only temporary and the involved activities are decreased to their previous resource level at the beginning of the next period. SPAR-1 assumes that the form of this relationship is linear for all activities which may or may not be reasonable for the given problem. The computer program could be changed to allow a different form of this relationship for some or all activities. Alternatively, it could be modified so as to restrict the assignment to each activity of its required resource to be at one of the three specified levels only.

The second feature examined is the use of current total slack as the single measure of the relative priorities between activities competing for the same resource. We must determine if this is an adequate measure of priority for the given R and D laboratory. This is equivalent to determining if the objective of primary importance to R and D management personnel is simply to meet individual project due dates, or at least minimize their overruns subject to the given resource limitations. If this is the case, current total slack alone will be sufficient since the proper selection of pseudo activities will tend to enforce individual project due dates by establishing the appropriate relative (current total slack) priorities between the individual projects and thus between their respective activities. However, it may be the case that in the given R and D laboratory the relative priority of an individual project depends on other factors such as its sponsoring agency and its dollar value, as well as its due date. The SPAR-1 scheduling model has no explicit provision for handling these alternate priority measures. However, they could readily be incorporated into the model and its associated computer program by means of one of the following procedures. The management personnel of the given R and D laboratory could classify each of the individual projects into one of a fixed number of different categories according to all relevant alternate priority measures. One procedure for incorporating these alternate priority measures would be to define a different critical slack value K for each category, and then adjust the current total slack values of each activity to a common base of zero. This adjustment would be made by subtracting the value K corresponding to the priority category of its associated project from the activity's current total slack value. The list of activities currently available for scheduling would be ordered according to this adjusted value. An alternative procedure would be to define a common critical value K for all projects, and when two or more currently critical activities or currently noncritical activities with the same current total slack are competing for the same resource, give preference on the basis of the priority categories of the associated projects. By the combined use of these two activity priority indicators, namely the current total slack of the activity and the priority category of its associated project, it seems that a sufficient single measure of the relative priorities between competing activites could be obtained for most R and D laboratory situations.

The third feature to be considered is the model's capability for allowing an activity to be scheduled in non-consecutive time periods. As discussed in the previous section, for the purposes of SPAR-1, activities must be either completely interruptable or completely uninterruptable. Furthermore, the duration of a completely interruptable activity associated with each level of resource application is the same whether or not the activity is scheduled with interruptions. Thus it is assumed that work on an interrupted activity can be resumed exactly where it left off without any loss of time. One or both of these assumptions may not be appropriate for the activities of the individual projects being scheduled by the given R and D laboratory. Based on the size of the chosen scheduling period, it may be necessary to require that an interruptable activity is not completely interruptable, i.e., that each time it is scheduled it must be processed for at least a certain minimum number of consecutive periods. This requirement could be incorporated in the SPAR-1 model without too much modification to its associated computer program.

On the other hand, if the management personnel of the given R and D laboratory decided that the second assumption is not appropriate, extensive modification to the model and its associated computer program would be required. For each interruptable activity the increase in its duration associated with each resource level caused by an interruption would have to be specified and a method for determining its remaining duration for a given resource level would have to be established for purposes of the updating CPM analysis performed at the end of each period. It seems that if an activity was not scheduled with too many interruptions, the assumed preemptive resume nature of these interruptions would be appropriate.

The final feature to be examined is the model's provision for dealing with activities that can be completed through the use of one of two different resource types. We must first decide if it is sufficient to allow only two alternative resources for any activity in the particular R and D scheduling problem. If this is the case, we must then determine if the assumption of the same linear relationship between the values of the activity durations associated with each resource level for the preferred and those for the substitute resource type is appropriate. This assumption may not be valid for the activities of the individual projects being scheduled in the given R and D laboratory if resource type substitution is used extensively. If this assumption is not appropriate, the SPAR-1 model and its associated program could be modified to allow either different relationships for different categories of activities or two sets of activity durations for each activity with an associated preferred and substitute resource type.

In conclusion, if the scheduling environment of a given R and D laboratory is well suited for the application of a computer-oriented heuristic scheduling model such as SPAR-1, it can provide valueable assistance in solving the complex multiple project scheduling problem faced by the laboratory's management personnel.

Specific References

1. Battersby, A., Network Analysis for Planning and Scheduling, St. Martin's Press, New York, 1964.

2. Moder, J. J. and C. R. Phillips, Project Management with CPM and PERT, Reinhold Publishing Co., New York, 1964.

3. Wiest, J. D., "A Heuristic Model for Scheduling Large Projects with Limited Resources", Management Science 13 (February, 1967), 359-377.

4. ____________, "Heuristic Programs for Decision Making", Harvard Business Review 44, (September - October, 1966), 129-143.

General References

1. Buffa, E. S., Production-Inventory Systems: Planning and Control, Irwin, Homewood, Illinois, 1968.

2. Davis, E. W., "Resource Allocation in Project Network Models - A Survey", Journal of Industrial Engineering 17, (April, 1966), 177-188.

3. Federal Electric Corporation, Programmed Introduction to PERT, John Wiley and Sons, Inc., New York, 1963.

4. Kelley, J. E., Jr., "The Critical Path Method: Resource Planning and Scheduling" Chapter 21 in J. F. Muth and G. L. Thompson, Industrial Scheduling, Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1963.

5. McMillan, C. and R. F. Gonzalez, Systems Analysis, Irwin, Howewood, Illinois, 1965.

6. Miller, R. W., Schedule, Cost and Profit Control with PERT, McGraw-Hill Book Company, New York, 1963.