AIR FORCE CREW ALLOCATION AND SCHEDULING PROBLEM

by

Minerva Rios Perez

Thesis submitted to the Faculty of the

Virginia Polytechnic Institute and State University

in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

in

Industrial Engineering and Operations Research

APPROVED:

Hanif D. Sherali

Robert D. Foley

Timothy J. Greene

July, 1982
Blacksburg, Virginia

ACKNOWLEDGEMENTS

The author is grateful to Dr. Hanif D. Sherali for serving as chairman of her committee. The guidance provided by Dr. Sherali throughout the research effort is deeply appreciated.

The author is also grateful to Dr. Robert D. Foley and Dr. Timothy J. Greene for serving on her thesis committee.

Special thanks are extended to Ms. Donna S. Lovern for her assistance in the typing of this manuscript.

Finally, the author would like to express her appreciation to Dr. Sherwood Samn at the Brooks Air Force Base, Texas, for presenting the details of the problem and providing the data used in the computations in this thesis.

CONTENTS

Chapter I

INTRODUCTION

Scheduling Problems represent a broad class of mathematical programming problems in which a number of tasks are to be performed over time with a limited amount of resources. The objective of these problems is to determine the sequence, starting times, and manner in which the tasks are to be performed so that the available resources are most efficiently used to perform the specified tasks. Depending on the nature of the scheduling problem, "efficiency" may be measured in terms of various criteria: for example, cost or time.

In this thesis, a special case of the general scheduling problem defined above is studied. The problem under consideration is an airline crew allocation and scheduling problem faced by certain divisions of the United States Air Force. In this problem, a task consists of visiting a set of bases in a pre-chosen order where the starting and ending base is the same (home base). The resources are the planes and the crews. This problem differs from most scheduling problems in that part of the available resources, namely, air force crews, must be allocated among various entities (air force bases) before the actual scheduling process occurs. The problem also differs from the usual airline scheduling problems in that it uses as an input, as opposed to generating, the routes and the number of times each route must be serviced over a certain time period. Due to these unique features, the specific problem under study cannot be solved, heuristically or otherwise, using procedures found in the literature.

As mentioned above, the purpose of this thesis is to study in detail an airline crew allocation and scheduling problem presented by the United States Air Force. Based on the characteristics inherent to this problem, two heuristic solution methodologies are proposed. Although these solution methodologies may be used to solve related problems, they are designed for the specific problem of concern in an effort to obtain as specialized and efficient a solution procedure as possible. Computational experience is provided for these two procedures in an effort to compare their relative merits.

The thesis is organized as follows. Chapter I presents a detailed statement of the problem and its underlying assumptions. A literature review of crew scheduling problems is found in Chapter II. In Chapter III, two solution methodologies, based on the problem structure, are proposed. Finally, Chapter IV presents a comparison of the relative merits of these procedures as well as computational experience for the two proposed solution methodologies using input data available from the Brooks Air Force Base.

1.1 PROBLEM STATEMENT

The problem under study involves n bases, designated B1,B2,...,Bn. Base B1 is chosen as the home base. There exist m distinct routes, each representing some mission, threading through subsets of these bases. Moreover, each route originates at the home Base B1 and is closed in the sense that it also terminates at the home Base B1. Although resource limitations may not permit it, it is desirable to service route j some r_j times per planning horizon, for each j=1,...,m. The resources are in terms of available planes and crews and have the following specifications.

1. Planes

   a) There exist p homogeneous planes in the system. Initially, all p planes are located at the home base B1.

   b) For each leg of the flight, say from base h to base k, on each route j, the expected time t_{hjk} of loading at base h, flying from base h to base k and unloading at base k is specified, for $j=1,\ldots,m$, and $h, k \in \{1,\ldots,n\}$.

c) It is required that a plane, once assigned to a given mission, remains devoted to that mission until its completion. However, the single crew which is required to service each leg of the journey may change from base to base during the flight of a given mission.

d) With regard to the r_j flights of a given mission type j, it is required that there be an interval of at least T_j time units, between the commencement of any two consecutive flights, for $j=1,...,m$. This quantity T_j is typically required to be positive for all missions j if a uniform distribution of missions over the planning horizon is desired. Note that the case $T_j=0$, $j=1,...,m$, essentially makes this assumption inactive.

e) Each plane requires periodic maintenance. This involves a down time of q hours every 45 days for minor maintenance and a down time of Q hours every 90 days for major maintenance. Furthermore, only the home base B1 has the facility for providing maintenance work. The down time of q hours is specified according to a discrete probability distribution. A probability p_1, p_2, p_3 is assigned to different durations q_1, q_2, q_3, hours of down time respectively.

2. Crews

a) There are $c > p$ homogeneous crews in the system. Initially, the c crews are to be distributed among the n bases in some to be determined "optimal" manner. Note that if c_k denotes the number of crews initially located at base k, $k=1,\ldots,n$, then due to assumption (a) on planes, the number of crews at base k at any given time is

$$c_1 - mp(t) + p_1(t) \qquad \text{for base } k=1 \qquad (1.1)$$

$$c_k + p_k(t) \qquad \text{for } k=2,\ldots,n \qquad (1.2)$$

where $p_k(t)$, $k=1,\ldots,n$, denotes the number of planes at base k at time t, and $mp(t)$ denotes the number of jobs being flown at time t. Moreover, it is stipulated that initially $c_1 \geq m$ so that the maximum number of planes can be utilized given that they are all operational.

b) It is assumed that crews do not have home bases to which they have to periodically return. In other words, the initial distribution of crews does not appoint home bases for crews. Realistically, since each crew has some home base, then this assumption implies that whenever a crew goes off duty to return to its home base, it is replaced by some other crew in a manner such that the distribution of crews in the system is determined only by the flights.

c) Each leg (flight from base i to base j) of the journey on each route is assumed to involve a sufficient time duration so as to warrant a rest period for the crew. Crews are hence said to have to "stage at each base". The rest period is taken to be 14 hours, regardless of the time of day when it commences. Moreover if short legs are flown, these are combined so that a route is defined in terms of bases where the crews must rest. Note that in connection with assumption (a) for crews, we assume that the work period of a crew includes both air and ground times. Also, the possible substitution of one crew by another as indicated in assumption (b) for crews is assumed to have no effect on rest periods.

d) Crews are assumed to be integral units and transfer of members from one crew to another is not permitted. Further, the transfer of a crew from one base to another is permitted only if this crew is working on a flight servicing that leg. That is, neither "deadheading" nor any external transfer of crews is allowed. However, a crew may transfer from one mission to another at any base, subject to assumption (c) for planes.

e) The following workload restriction is also imposed on each crew. Namely, the federal regulations do not allow more than 8 flight hours in any given 24 hours time period. Note, for example, that a rest period of 14 hours (as in assumption (c) for crews) implies a maximum of 8 flight hours for a ground time of at least 2.00 hours.

Scheduling Objective

The problem seeks a judicious distribution of crews among the bases and a subsequent scheduling of the missions so as to maximize the expected overall utilization rate for the planes over some period of time. This overall utilization rate is defined to be the average operation time (flight and ground time) per plane per day over the period under consideration. Note that a period of time T_j required between the commencement of any consecutive flights of any given mission j implies a possible idle time for planes while waiting for available missions. Equivalently, one may minimize the total idle time spent by the planes over some period of time.

Although the main interest is the maximization of the utilization rate of the planes, the minimum number of crews and planes required in the system for a given completion time (makespan) is valuable information in this problem.

Different variants of the problem are considered according to the time, T_j, required between the starting time of any two consecutive operations of the same mission j. Three variants are considered, T_j equal to zero, T_j different from zero and calculated during the simulation process, and T_j different from zero and fixed at the beginning of the algorithm. Even though assumption (d) for planes and assumption (a) for crews differ on each variant with T_j, assumptions (a), (b), (c), and (e) for planes, as well as, assumptions (b), (c), (d), and (e) for crews remains the same in the three variants of the problem discussed below.

1.1.1 Problem 1

In this problem, the utilization rate of planes is maximized, given that there exists no restrictions on the starting time of the jobs. This means that for a given number of planes, crews and missions, the number of missions to be performed at any time depends only on the availability of resources (planes and crews) in the system. Therefore, the number of days necessary to complete all the jobs when the number of crews, as well as the number of planes is fixed, is minimized.

In this case, the total time required for the completion of the jobs depends on the allocation of crews among bases as well as on the number of planes in the system since the missions can be flown as long as these two resources are available.

Since the completion time of the jobs depends on the number of planes available in the system, we also seek to determined the minimum number of crews which will yield the completion time.

1.1.2 Problem 2

Here, a restriction of a period of time T_j between the starting time of any two consecutive operations i-1 and i, of mission j is imposed on the problem. Therefore, in this problem, the availability of the missions is also restricted. Moreover, a limited number of planes, equal to the number of missions that can be flown at the same time, is effectively used in the system. Therefore assumption (a) for crews stipulates a minimum of m crews initially placed at Base B1.

The value of T_j is fixed at the beginning of the problem in order to spread the missions over the period of time desired in the problem.

Since this problem in some way restricts the overall completion time, then after a (maximum) utilization rate for the problem is found, it is of interest to determine the minimum number of crews and planes in the system which will yield this maximum utilization rate solution. Obviously it must be expected that a smaller number of crews and planes is required in the system by this problem than by Problem 1, since fewer missions are being flown at the same time.

1.1.3 Problem 3

As in Problem 2 a restriction on the time between the starting time of any two consecutive operations of any mission j, is imposed on this problem. This problem does not allow more than one operation of mission j to be performed at the same time. Therefore, the imposed time restriction T_{ij} depends on the processing time of the operation i-1 as well as on the idle time (if any) incurred by this operation. T_{ij} is calculated for each operation i of mission j using the equation

$$T_{ij} = t_j + It_{(i-1)j} \tag{1.3}$$

where t_j is the processing time of mission j, $It_{(i-1)j}$ is the known idle time of operation i-1 of mission j, and i-1 and i are consecutive operations of mission j. The value

It$_{ij}$ is generated during the simulation process. Hence, T_{ij} is not a prespecified quantity, but is determined by the solution procedure. Since only m planes can be flown at any time, assumption (a) for crews requires a minimum of m crews initially placed at Base B1.

1.2 PROBLEM FORMULATION FOR A RELAXED VERSION

According to the Problem Statement, we break down the formulation below into two cases, namely, $T_j = 0$, $j=1,\ldots,m$ and $T_j \neq 0$, $j=1,\ldots,m$. For both cases, a linear mixed integer programming formulation for an associated relaxed problem is presented. This relaxation of the actual problem considers only assumptions (a), (b), (c), and (d) for the planes while ignoring the assumptions for the crews. The difference in the complexity of the formulations of these two problems is indicative of their relative difficulties. These formulations illustrate the difficulty of the problems. Also one of the proposed methods uses this relaxed version of the problem to find a sequence of missions to be flown by each plane, as well as an upper bound for the number of crews required in the system. However, an exact solution of these relaxed problems is not computationally feasible for all sized problems, and hence, we alternatively propose certain heuristics to obtain good quality solutions.

1.2.1 Case (i) $T_j=0$, $j=1,...,m$

Observe that in this case, the problem basically reduces to a parallel processor scheduling problem [5], with p identical machines (planes) and r_j operations of type j each requiring a processing time t_j units for $j=1,...,m$.

Here, t_j, equals the sum of flight and ground times required to perform mission j. No preemption is permitted. Defining

$$x_{ijk}=\begin{cases} 1 \text{ if machine i processes operation k of mission j} \\ \\ 0 \text{ otherwise} \end{cases} \quad (1.4)$$

for $i=1,...,p$, $j=1,...,m$, $k=1,...,r_j$, the problem of minimizing the makespan T may be formulated as follows

P1: Minimize T

$$\text{Subject to } T\geq \sum_{j=1}^{m}\sum_{k=1}^{r_j} t_j x_{ijk} \quad \text{for each } i=1,...,p \quad (1.5)$$

$$\sum_{i=1}^{p} x_{ijk}=1 \quad \text{for each } k=1,...,r_j \text{ and each } j=1,...,m \quad (1.6)$$

$$x_{ijk} \quad \text{binary for each } i,j,k$$

Constraints (1.5) determine the makespan T and constraints (1.6) ensure that each operation is performed on some machine. The objective of minimizing the makespan is equivalent to minimizing the total machine-idle-time.

Denoting $n = \sum_{j=1}^{m} r_j$, Problem P1 is a mixed-integer program with np binary variables, one continuous variable and n+p constraints. Since no computationally feasible solution procedure for Problem P1 has yet been developed, a heuristic procedure based on the "LPT dispatching rule" [5] is a viable alternative for approximately solving P1.

1.2.2 Case (ii) $T_j \neq 0$, $j=1,\ldots,m$

This case is more complicated than Case (i), since now, not only is the assignment of jobs to machines relevant, but also their sequencing thereon. A relaxation of this problem is considered to illustrate the relative complexity of this case. Defining variables x_{ijk} as in Equation (1.4), one may stipulate that constraints (1.5) and (1.6) must hold as before. Now, define variables

$$y_{jk} = \text{start time for operation } k \text{ of mission } j \tag{1.7}$$

for $k=1,\ldots,r_j$ and $j=1,\ldots,m$. Since it is possible for idle times to occur on machines between the processing of consecutive operations, the makespan T is determined in this case through the inequalities

$$T \geq y_{jk} + t_j + M\,(x_{ijk} - 1) \quad \text{for each } i=1,\ldots,p$$

$$k=1,\dots,r_j \qquad (1.8)$$
$$j=1,\dots,m$$

where M is a suitably large number, e.g., $M = \sum_{j=1}^{m} r_j(\, t_j + T_j \,)$.

Next, for each machine i=1,...,p we need to ensure that there is no overlap between any two operations being processed on that machine. In other words, $x_{irs} + x_{ijk} = 2$ must imply that either $y_{jk} \geq y_{rs} + t_r$ or that $y_{rs} \geq y_{jk} + t_j$ for any job pairs (j,k) and (r,s), for each i=1,...,p. This condition can be enforced by the inequalities

$$y_{jk} \geq y_{rs} + t_r + 2M\,(x_{irs} + x_{ijk} - 2) - M\,\sigma_{jk,rs}$$

$$(1.9)$$

$$y_{rs} \geq y_{jk} + t_j + 2M\,(x_{irs} + x_{ijk} - 2) - M\,(1 - \sigma_{jk,rs})$$

for each pair (j,k) and (r,s), for i=1,...,p where $\sigma_{jk,rs}$ are dummy binary variables.

Finally for each j=1,...,m and for any job pairs (j,r) and (j,s) of type j, one needs to enforce that either $y_{ik} \geq y_{rs} + T_r$ or that $y_{rs} \geq y_{jk} + T_j$. Similar to (1.8), this restriction may be enforced through the inequalities

$$y_{js} \geq y_{jr} + T_j - M\,(\gamma_{jrs})$$

$$(1.10)$$

$$y_{jr} \geq y_{js} + T_j - M\,(1 - \gamma_{jrs})$$

for each pair of jobs (j,r), (j,s) of type j, $j=1,\ldots,m$, where again, γ_{jrs} are suitable dummy variables. Thus, the problem addressed may be formulated as

P2: Minimize T

subject to

subject to	constraint (1.4)	(n)
	constraint (1.6)	(np)
	constraint (1.7)	$(np(n-1))$
	constraint (1.8)	$(\sum_{j=1}^{m} r_j(r_j-1))$
	$T,\ y > 0$ and x, σ, γ binary	

where $n=\sum_{j=1}^{m} r_j$ and where there are (np) x-variables, $(n(n-1))$ σ-variables, $(\sum_{j=1}^{m} (r_j(r_j-1))/2)$ γ-variables, (n) y-variables and T denotes the makespan.

Since the relaxation of P2 itself is a large linear programming problem which requires considerable computational effort, it is clear that except for relatively small problems, solving even the linear relaxation of P2 in order to obtain an approximate solution may be an arduous task. Furthermore, a continuous valued solution may not be useful because a feasible integer solution is not usually available therefrom. Consequently, a heuristic solution procedure may again be preferable.

1.3 OUTLINE OF SOLUTION STRATEGY

Two solution procedures are proposed in this thesis for the given problem. These are referred to as the Two-Phase Method and the One-Phase Method. The principal thrust of these techniques is to obtain a judicious allocation of crews to bases in such a fashion that good quality solutions for the resultant scheduling problem may be obtained.

In the Two-Phase Method, the first phase (Phase I) is an initialization step in which a relaxed version of Problem P2 is solved. The solution to this problem provides three important pieces of information. First, it yields a lower bound on the time period required to schedule exactly r missions of type j, j=1,...,m as well as an upper bound on the number of crews which can be effectively used. Second, it provides a natural starting distribution of crews among bases. Finally, and most important, it suggests which missions should be flown by which planes and in what sequence.

The second phase of the Two-Phase Method (Phase II) is essentially a simulation process repetitiously executed. Given the initial distribution of crews, and the sequence of operations to be performed by each plane as prescribed in Phase I, a repeated simulation of the problem is conducted with different (revised) crew allocations for each simulation, to provide a schedule of missions. More specifically, given a sequence of missions flown by each plane, and an initial allocation of crews to bases (from Phase I), Phase II first simulates the problem to determine a schedule for planes and crews. Based upon the results of the simulation, a re-allocation of crews to bases is then made. Next, the simulation is repeated. The procedure iterates between the allocation of crews to bases and the subsequent simulation phase until some termination criterion is satisfied.

The One-Phase Method is basically a simulation process executed several times. Initially, this method distributes crews among the bases in proportion to the number of times the bases are visited in the problem. The planes are assigned to missions during the simulation process according to certain specified rules. At the end of the simulation, a schedule of missions to planes is obtained. Based on the output of this process, an improvement in the objective fuction is sought through a re-distribution of crews among the bases. Similar to the Two-Phase Method this problem then iterates between the allocation of crews among bases and the simulation process until a suitable stopping criterion is satisfied.

It is important to note here the difference between the Two-Phase Method and One-Phase Method. While in the former method the sequence of missions to be flown by each plane is fixed before the simulation iterations commence, in the latter, the assignment of missions is performed concurrently with the execution of the simulation. There exist other differences between these two methodologies such as the criteria used for the distribution of crews among bases and the determination of bounds for the number of crews required in the system. As will be seen later, these differences in solution methodologies yield significant differences in computational efficiency.

Chapter II

LITERATURE REVIEW

The routing and scheduling of vehicles and crews is an important problem area in Operations Research. A large amount of literature has been written on this topic in the last twenty years. The literature can be classified into three main categories: routing of vehicles, scheduling of vehicles and crews, and routing and scheduling of vehicles and crews. These categories reflect problems of increasing difficulty as more features are considered.

Divisions within the categories exist, based on other problem characteristics and solution strategies applied to the problem. For example, the problem may seek to fulfill different objectives, such as maximizing some index of safety or customer satisfaction or as is usually the case in most routing and scheduling problems, minimizing the total cost. An outline of the ten major considerations in the classification of routing and scheduling problems is given by Bodin and Golden [13].

The solution approach to these problems is another consideration in the classification of the routing and scheduling problems. As discussed by Bodin et al. [14], the main characteristic utilized to classify these problems is whether or not the solution procedure is "polynomially bounded". A solution procedure is considered polynomially-bounded if the computational effort increases only polynomially with the problem size in the worst case. A problem which belongs to this class (P) can usually be solved to optimality in an efficient manner. Generally, whenever an algorithm for the solution to a problem belongs to the class P, the set of procedures employed for the solution of the same type of problem are also P-bounded.

An NP-hard problem is a problem which algorithms with polynomial difficulty cannot solve. The problem that requires a solution effort which increases exponentially with the size of the problem in the worst case is NP-hard. Some scheduling problems have been classified NP-hard. However, there exist problems in the literature that have not as yet been classified as either NP-hard or P-bounded problems.

In this chapter we intend to describe each category separately. The categories will be discussed in order of difficulty. A vehicle routing problem is solved when a sequence of tasks to be performed by these vehicles is obtained. In this problem there is neither a specific time when the tasks must be performed, nor a specific relationship between the tasks (there is no order in which the tasks must be serviced). The four most important problems within this category are: the Traveling Salesman Problem, the Chinese Postman Problem, the Single Main Depot, Multiple Vehicle Routing Problem, and the Multiple Main Depot, Multiple Vehicle Routing Problem.

Many solution procedures have been developed for the solution of the vehicle routing problem. Some of them intend to solve the problem to optimality while others provide good solutions via heuristic procedures. The former are mostly employed in the Single Main Depot, Multiple Vehicle Problem, Multiple Depot, Multiple Vehicle Problem and some variants of the Chinese Postman Problem. The latter are used for the Traveling Salesman Problem and some variants of the Chinese Postman Problem since these are considered NP-hard. These solution procedures will be presented in the next section.

A vehicle scheduling problem is one which seeks a sequence of activities which are restricted to be performed at specific times. The crew sheduling problem requires a sequence of tasks to be performed by each crew when these tasks are restricted to some specific time. The crew scheduling problem will be discussed in this chapter within three categories: Scheduling of workers at a fixed location, Mass Transit Crew/Vehicle Scheduling problem and the Air Crew Scheduling problem.

The first category deals with the schedule of workers at a fixed location over a certain number of time intervals in order to obtain the required number of workers for each interval. This problem can be formulated as a network flow problem.

The Mass Transit Crew/Vehicle Scheduling Problem attempts to schedule the crews and vehicles in order to minimize the cost of the completion of jobs. Two well known approaches, Set Partitioning based and Run Cutting methods, are presented for the solution of this problem.

The Air Crew Scheduling Problem requires a sequence of tasks to be performed by each crew as well as for each plane. This is a very difficult problem since both aspects of the problem, crews and planes, are considered at the same time. A vast number of techniques have been developed for the solution of this problem. These techniques can be classified under two main groups: mathematical programming approaches and heuristic approaches. The details of the problem as well as the procedures employed for the solution of the problem are discussed in Section 2.2.

The combined Routing and Scheduling Problem considers both the Routing Problem, and the Scheduling problem restricted by precedence relationships, and with respect to intervals of time (time windows) within which the tasks must be performed. Heuristic procedures have been developed since no exact mathematical techniques have been successful for these problems.

A comparison of the objective and solution procedures proposed in this thesis for the problem on hand, with those in the literature is given in Section 2.4 of this chapter.

2.1 ROUTING PROBLEMS

A basic routing problem involves a set of nodes and/or arcs to be serviced by a fleet of vehicles, wherein no temporal restrictions (period of time in which the service must be performed) or precedence constraints exist. Only maximum route length constraints may exist. The problem is to find a set of routes for each vehicle for which the objective fuction is optimized and the constraints are satisfied. "A vehicle route is a sequence of pick up and delivery points which the vehicle must traverse in order, starting and ending at a same main depot or domicile" [14]. Different objectives of the problem and a large variety of constraints produce a vast number of possible variations on the standard problem, although a sequence of routes to be performed by each vehicle is the output in all of them.

The most well known vehicle routing problems are:

1. The Traveling Salesman Problem

2. The Chinese Postman Problem
3. The Single Main Depot, Multiple Vehicle, Routing Problem
4. The Multiple Main Depot, Multiple Vehicle, Routing Problem

The Traveling Salesman Problem requires the determination of the minimum distance in a cycle (Hamiltonian circuit) within which each city is visited exactly once. This is, perhaps, the most discussed problem in the literature of routing problems. Bellmore and Nemhauser [7] present a survey of this problem. Christofides [19] presents a detailed study of this problem.

Different versions of this problem have been studied in the literature. For instance, the M-Traveling Salesman Problem is a generalization of this problem wherein 1) M salesman, rather than one, trace Hamiltonian circuits among a subset of the nodes or cities, and 2) each city is visited at least once. Formulations of this problem can be found in Bellmore and Hong [8], and Orloff [38].

Even though many algorithms have been proposed for solving the Traveling Salesman Problem, problems involving more than 300 cities remain unsolved. Karp [31] has shown that the problem is an NP-hard problem. Hence, different heuristic procedures have been developed. An analytical comparison of these procedures is found in Rosenkrantz et al. [39].

The Chinese Postman Problem deals with the task of finding a minimal distance route in a network where every arc must be traversed at least once. There are three versions of the problem. In the first version, undirected arcs are assumed. Solution procedures for this version have been proposed by Glover [25], Edmond and Johnson [22], Christofides [19] and Bellman and Cooke [6]. The first three papers present P-bounded algorithms which can be employed only in the solution of small size problems.

The second version studies the case where the graph is directed. The solution of this problem is based on the notion of "a continuous path through a network that traverses each arc exactly once" [14]. An algorithm which obtains optimal solutions has been developed by Beltrami and Bodin [9]. The third version considers a mixed graph with both, directed and undirected arcs. The first two versions have been solved to optimality whereas the the third version is an NP-hard problem [14]; there exists no exact procedure for the solution of this version.

The Single Main Depot, Multiple Vehicles Routing problem is a classical routing problem in the literature. Given a set of routes and vehicles in one main depot or garage a demand at each city or depot must be satisfied with a minimum distance traveled. Mathematically, the generalized vehicle routing problem formulation as given in Bodin et al. [14] is :

$$\text{Minimize} \quad \sum_{i=1}^{n} \sum_{j=1}^{m} \sum_{k=1}^{nv} c_{ij}\, x_{ijk} \qquad (2.1)$$

$$\text{Subject to} \quad \sum_{k=1}^{nv} \sum_{j=1}^{n} x_{ijk} = d_i \qquad i=2,\ldots,n \qquad (2.2)$$

$$\sum_{i=1}^{n} x_{ihk} - \sum_{j=1}^{n} x_{hjk} = 0 \qquad k=1,\ldots,nv;\ h=1,\ldots,n \qquad (2.3)$$

$$\sum_{i=1}^{n} d_i \left(\sum_{j=1}^{n} x_{ijk}\right) \leq K_k \qquad k=1,\ldots,nv \qquad (2.4)$$

$$\sum_{i=1}^{n} t_{ik} \sum_{j=1}^{n} x_{ijk} + \sum_{i=1}^{n} \sum_{j=1}^{n} t'_{ijk} x_{ijk} \leq T_k \qquad k=1,\ldots,nv \qquad (2.5)$$

$$\sum_{j=1}^{n} \sum_{i=1}^{n} x_{ijk} \leq 1 \qquad k=1,..,nv \qquad (2.6)$$

$$x_{ijk} \ \text{binary} \qquad i=1,\ldots,n;\ j=1,\ldots,n;\ k=1,\ldots,nv \qquad (2.7)$$

where n is the number of depots, nv the number of vehicles in the system, T_k the maximum time allowed for a route of vehicle k, d_i the demand of depot i (d_1=0), t_{ik} the time required for vehicle k to serve depot i, (t_{kk}=0 for all k), t'_{ijk} the time required for vehicle k to travel from depot i to depot j ($t'_{ikk} = \infty$ for all i), c_{ij} the cost of travel from depot i to depot j, and x_{ijk} equals 1 if the distance from depot i to depot j is traversed by vehicle k, and is 0 otherwise.

The objective function (2.1) to be minimized, represents the total distance traversed. Equation (2.2) ensures that the demand is met at each depot. Equation (2.3) represents the continuity of the routes. The capacity and availability constraints of the vehicles are represented by inequalities (2.4) and (2.6) respectively. The route time constraint is given by inequality (2.5).

Simple modifications in the objective function or the constraints result in different versions of the general problem formulation. For example, if we incorporate in this model a restriction on the exact times when the tasks are to be performed, as it will be seen later in this chapter, a model of a scheduling problem results.

Many solution procedures have been proposed for the solution of this problem. The majority of these procedures can be classsified in five categories.

1. Route-cluster
2. Insertion
3. Improvement/exchange
4. Mathematical programming procedures
5. Exact procedures

Route-cluster is a two step procedure approach, which can be classified in two different procedure groups. The difference between them is the order in which the steps are performed. In one of these procedure groups, Step 1 produces a large route where all the entities are served. This route is usually infeasible. Step 2 breaks the routes down into smaller but feasible routes. In Bodin and Berman [12] an application of this approach to the school bus routing problem, is found. In the second procedure group a small set of entities are produced in the first step. The second step finds routes within each set that minimize the cost. Applications of this procedure are given in Gillett and Johnson [24], and Karp [30]. Also an algorithm for the solution of street-sweeper routing problems using the latter procedure has been proposed by Bodin and Kursh [16].

Insertion procedures have been developed for single main depot routing problems. These procedures compare an alternative solution with the current solution for the problem, where both solutions may be infeasible, but the alternative solution either reduces the objective fuction value or satisfies one of the routes that is not covered by the current solution. The procedure concludes when a feasible solution is at hand. This procedure is used by Hinson and Mulkerkan [27] in a solution method proposed for plane routing problems.

Improvement procedures attempt to reduce the total cost in the problem by exchanging branches of the current solution. Feasibility of solutions is always maintained. The procedure continues until no more improvement via prescibed exchanges are possible. Heuristic rules for this type of procedure have been developed by Baker et al. [4].

Mathematical programming approaches have been devised for routing problems. Most of these solution techniques are heuristic procedures where special structures of the problem are exploited so that formal mathematical procedures such as dynamic programming or Lagrangian relaxation can be incorporated in the (approximate) solution technique for that specific problem. Also, subjective assessments based on previous experience or knowledge of the problem can be adapted to optimization models in order to obtain a more accurate solution procedure for the problem.

Exact procedures such as branch and bound and cutting plane techniques have been applied to routing problems. These techniques provide good bounds for the problems, even though the optimal solution is not obtained for other than small sized problems.

The Multiple Main Depot, Multiple Vehicle routing problem is a generalization of the above problem. This routing problem keeps the same constraints and objective function as the previous one, but now there exist more than one main depot (garage). The solution procedures employed for this type of problem are extensions of the procedures discussed above.

Most of the solution procedures which have been developed for this problem are heuristic procedures. Gillett and Johnson [24], Tillman and Cain [47], and Golden et al. [26] have proposed heuristic algorithms for this problem.

2.2 CREW/VEHICLE SCHEDULING PROBLEMS

Crew/Vehicle Scheduling Problems can be thought of as routing problems with temporal restrictions (specific period of time in which the task must be performed) [14]. The vehicle and crew scheduling problem depends on the sequence of activities to be performed as well as on the specific time that each activity must be performed. In this sense both crew scheduling and vehicle scheduling are similar, although in the former case more restrictions are involved, such as federal regulations, union regulations, contract specifications, etc.

Whenever both crew and vehicle scheduling are required in a problem, the two aspects interact with each other. The ideal situation would be to solve both problems simultaneously. However, the union of both problems generally yields a model too complicated to be solved. Hence the approach usually followed is to use iterative algorithms where the crew or vehicle scheduling problem is alternately solved, given a schedule for the other entity.

A basic crew/vehicle scheduling problem can be formulated as a set partitioning problem [14]

$$\text{Minimize} \quad \sum_{i=1}^{n} c_i x_i \tag{2.8}$$

$$\text{Subject to} \quad \sum_{i=1}^{n} a_{ji} x_i = 1 \qquad j=1,\ldots,m \tag{2.9}$$

$$x_i \quad \text{binary} \qquad i=1,\ldots,n$$

where m is the number of task or jobs to be completed and n the possible schedules covering the jobs. a_{ji} is a binary coefficient taking the value of 1 if schedule i covers job j and a value 0 otherwise. c_i is the cost of schedule i, for each i=1,...,n.

This formulation is modified to obtain different versions of the scheduling problem. For instance, if more than one schedule is allowed to cover the same job, a set covering problem is now obtained where Equation (2.9) is replaced by $\sum_{i=1}^{n} a_{ji} x_i \geq 1$.

A more complicated problem results from the basic crew/ vehicle scheduling problem formulation when additional features are added to it. These features are usually:

1. Restriction on the maximum length of the path allowed

2. Existence of multiple vehicle types

3. Existence of multiple main depot vehicles

The crew scheduling problem can be divided into three main groups of interest:

1. Workers at a fixed location Scheduling Problem

2. Mass Transit Crew and Vehicle Scheduling problem

3. Air Crew Scheduling Problem

The first group is the simplest version of this problem. For a given set of time intervals of a day and a specific demand of workers required in each time interval, a schedule for each worker is sought in the problem. Telephone operator scheduling [43] is an example of this problem. This problem has a network formulation. Algorithms based on this formulation have been proposed by Bennett and Potts [11], Bodin et al. [17] and others.

The Mass Transit Crew and Vehicle Scheduling Problem is a more complicated version of the problem since both aspects of the scheduling problem are considered. A set of tasks is assigned to each crew and vehicle for a one day period. Thereafter, a number of days (with differing schedules) are combined to obtain a weekly schedule. These schedules are combined until a monthly schedule is obtained for each crew and vehicle. In many algorithms an exchange of schedules among the personnel is required in order to ensure equal difficulty of tasks to all the personnel. This last restriction complicates the problem. Thus, an optimal solution is not usually tangible.

The state of art of these problems is presented by Wren [49]. The traditional approaches utilized in the solution of this problem are the Set Partitioning Based Approach and the Run Cutting Approach [14].

Two algorithms based on the Set Partitioning/Covering Approach are presented in Bodin et al. [14]. The first algorithm attempts to decompose the problem based on the duration of the tasks. Several groups are formed according to this criterion and within each group a set of columns is generated. A set partitioning problem is then solved for each group. Parker and Smith [37] develop a second algorithm in which all the tasks are considered at the same time but a set of restrictions is imposed on the set of columns generated in order to reduce the total number of columns. An exact solution procedure for the Mass Transit Crew Scheduling Problem has been developed by Edmonds [21]. This method is based on the set partitioning problem where no more than two ones per column exist.

The Run Cutting Approach is generally used in heuristic methods based on procedures employed by manual schedulers. These methods attempt to incorporate heuristic rules used in manual scheduling for specific problems into computerized techniques. Algorithms based on this approach are given in [32], [33], and [42].

The Air Crew Scheduling problem is one of the most complicated versions of the scheduling problem. Generally, a schedule of planes is conducted concurrently with the scheduling of crew. Many times, a certain frequency of the missions over a period of time is required. Therefore, the jobs must be completed within a given interval of time (time window). Hence a combination of routing and scheduling problems result.

Generally, the solution techniques applied in the solution of the Air Crew Scheduling Problem are also used for the solution of Workers at a fixed location Scheduling Problem and Mass Transit Crew Scheduling Problem. Since our attention is focused on the Air Crew Scheduling problem, a more detailed discussion of the techniques used for the solution of crew scheduling problems will be given within the framework of the problem under consideration.

The solution procedure for the Air Crew Scheduling problem can be divided into two main categories: those which use a mathematical programming approach, and those which use heuristic methods. Within these categories subdivisions are found.

2.2.1 Mathematical Programming Approach

Most Crew Scheduling Problems can be formulated as mixed integer linear programming problems, even though only a small number of them can be completely solved by this approach.

In the Air Crew Scheduling Problem a leg i is defined as the smallest feasible unit of work for a crew, that is a flight from node A to node B without intermediate stops. The planning unit for crews is called a rotation. A rotation is a set of legs flown by the crews which starts and ends at some home base (round trip). For these problems a matrix A is generated where the rows represent legs to be flown and the columns are the rotations. The entry a_{ij} of the generated matrix takes a value of 1 if leg i is flown in rotation j and 0 otherwise. For different cases of the Air Crew Scheduling Problem, this matrix can be either generated or be an input for the problem.

From the matrix A a set of rotations is chosen to be flown. x_j takes value of 1 if rotation j is chosen and 0 otherwise. Associated with each rotation j there exists a cost c_j . Also a specific number of crews is required for each leg i; this number of crews is represented by b_i .

Based on the matrix A the cost c_j for each rotation j, and the required number of crews b_i, for each leg i, a simple formulation of the Crew Scheduling Problem used by many airlines is:

$$\text{Minimize P3:} \quad \sum_{j=1}^{n} c_j x_j \qquad (2.10)$$

$$\text{Subject to} \quad \sum_{j=1}^{n} a_{ij} x_j = b_i \qquad i=1,\ldots,m \qquad (2.11)$$

$$x_j = 0,1$$

Based on Arabeyre et al. [2], the solution procedure for the above formulation can be divided into 3 steps: Generation, Reduction, and Optimization.

1. Generation

The first step is the construction of the matrix A, each column of which corresponds to a feasible rotation. The information used in the development of the matrix, therefore, varies with the specific problem. Hence, each Air Crew Scheduling Problem needs to develop its own matrix based on its specific restrictions.

The general information required for the construction of the matrix is [2]:

a) Crew Regulations

b) Aircraft schedule for each aircraft

c) Definition of the smallest production unit that a crew can fly

Most airlines are concerned with union regulations, federal regulations, and safety. However, these regulations change with the airline and the home country as well as with the type of contract the crew has with the company.

For the construction of this matrix, the legs are listed. Then the rotations are constructed adding legs one by one to it until the addition of a new leg makes the rotation no longer feasible. This process continues until no more rotations are possible.

2. Reduction

The generated A matrix is the matrix considered in a set partitioning problem, which is usually far too large to be solved by any optimization technique. Hence a reduction of the matrix is required.

Many airlines use the most common Integer Programming rules (row dominance, column dominance, etc.) to reduce the matrix A [41]. One of the most useful rules in the reduction process follows.

> For the minimization problem P3, where b_i equals 1 for all i, consider two rows A_i and A_j such that $A_i \subset A_j$. This means that any pairing which contains leg i also contains leg j. Therefore, row A_j may be eliminated and any rotation that does not belong to the intersection $A_i \cap A_j$ can be eliminated since it is infeasible.

Other procedures are also used. For instance Rubin [40] uses an upper bound for the total cost to reduce the dimension of the matrix. Whenever the addition of a rotation to the rotations already chosen produces a cost greater than the given upper bound for it, this rotation is eliminated. The procedure considers all possible combinations of rotations until no reduction is possible.

The reduction step could also be included in the matrix generation process. The combination of these two steps can result in a more efficient method, thus reducing execution time.

3. Optimization

Many Integer Programming solution procedures have been used to solve the Crew Scheduling Problem. The most popular are branch and bound techniques and cutting plane techniques.

United Air Lines, for instance, developed an algorithm based on Gomory's cutting plane technique [34]. This proceeds as follows. After the continuous optimum is computed using linear programming, additional constraints are added to the problem to restrict the continuous solution. The algorithm solves a finite number of linear programming problems and it stops when the continuous solution is an integer optimum. Although in theory there is no limit to the size of the problem that can be solved via this algorithm, in practice, a problem with matrix A consisting of more than 4000 columns is computationally intractable.

Specialized branch and bound techniques developed for the specific problem under consideration, have been sucessfully used by many airlines. A list of airlines which have developed solution procedures based on these techniques is given in Arabeyre, et al. [2]. A brief outline of the solution procedures discussed by Arabeyre et al. [2], is presented here.

Air Canada's Branch and Bound Procedure consists of an implicit enumeration scheme combined with cost projection, relative costing, and subset reduction techniques applied to certain subproblems. This method arranges the columns of the matrix A in order of increasing cost. The uncovered legs are added to the rotations producing a higher cost. Hence, a rearrangement of the matrix is performed. If the total cost exceeds the current bound, a backtrack occurs, and the process is repeated. This process continuous until all legs are covered with the least possible cost [23].

Air France constructs two matrices, matrix A where the columns are the chosen rotations, and matrix B consisting of the remaining rotations. The columns (rotations) of the matrix A are rearranged in an increasing order according to the cost ratio (c_{ij}/n_j), where n_j is the number of nonzero entries in column j. Thereafter the Branch and Bound Technique is applied to the arranged matrix. The branching rule employed is: among the rotations that do not intersect with the elements of the rotations already chosen, the one corresponding to the smallest cost is selected. The matrix with the residual rotations is updated. The process continues until the dimension of the matrix of residuals is less than some given dimensions N_0, M_0. Another version of this algorithm solves the continuous linear program. Certain number of variables with fractional values are set to 1 and the above process is repeated. The solution obtained is compared to the LP lower bound. If the difference is less than a given percentage, the algorithm stops; otherwise, some variables previously fixed at 1 are set free and the process is repeated [1].

The approach used by American Airlines is based on a pre-ordered matrix followed by a branch and bound search [2]. The order of the matrix is based on the time in which each leg flown as well as the cost of each column. Several bounding criteria have been implemented in this algorithm. An initial feasible solution is obtained; this solution is progressively improved upon until either the optimal solution is found or some stopping criteria is satisfied.

Techniques for the solution of network problems have also been applied to Crew Scheduling Problem. Nicoletti [36] presents Alitalia Airlines' Crew Scheduling problem as a network model.

2.2.2 Heuristic Methods

Even though most Crew Scheduling Problems can be formulated as mixed integer linear programming problems, the majority of crew scheduling problems have been solved using heuristic approaches rather than rigorous mathematical procedures. Most of these heuristic methods are based on integer programming solution procedures like branch and bound techniques and Balas' Additive algorithm. Also, network flow algorithms have been used. These techniques are combined with some rules designed for the specific problem under consideration to obtain an optimal or near optimal solutions.

Whenever generation and reduction steps are required in the procedure, these steps are similar to those used in the mathematical programming approach. However the optimization step differs in the procedure to be followed and the answer to be obtained.

Swiss Air has developed a direct search technique based upon heuristic rules which gives an incomplete enumeration of feasible solutions [45]. BEA's Set Covering Algorithm finds a feasible solution by heuristic rules and improves it by adding new constraints according to these rules. This algorithm was proposed by House, Nelson and Rado [28].

Armstrong and Cook[3] develop a two-step algorithm to schedule a limited fleet of aircraft and assign a set of tasks. The algorithm minimizes total flying time while the set of heuristic rules gives attention to priority, customer satisfaction, and maintenance efficiency.

Baker et al. [4] develops a heuristic procedure to solve the Crew Scheduling Problem of the Federal Express Corporation. This problem differs from the ones already presented in the crew payment procedure and the number of home bases. This corporation pays their personnel according to operating hours in a specific rotation rather than a fixed salary. There is only one home base for all the crews, namely, at Memphis. Hence all rotations must originate and terminate at this base.

The procedure developed for the Federal Express Corporation problem can be broken into three major steps: rotation contruction, rotation improvement, and composite method. The first step enumerates all possible combinations of legs to obtain the columns (rotations) for the scheduling matrix.

This rotation construction utilizes some criteria based upon the unique properties of the problem. First, it considers each base departure as a seed to a rotation. Each unassigned leg is selected to be assigned to a partially formed rotation following some tie-breaking rules, if this assignment will not violate any regulation. The tie-breaking rules used in this step are the following [4]:

a) First feasible assignment.

b) The next flight leg is assigned to the first feasible pairing for which the aircraft flight number of the partially formed rotation and the flight leg are the same.

c) The crew with the longest rest period takes the first flight leg out.

d) The flight leg is assigned to the rotation with the arriving time closest to its departure time. This rule minimizes the crew idle time.

e) The flight leg is assigned to the rotation for which the additional cost due to this leg is a minimum over all rotations considered.

A leg can be assigned to more than one rotation. When all the assignments of legs to rotations, are completed the more efficient assignments are selected according to tie-breaking rules.

The second step is the rotation improvement for which vehicle routing rules are used. The 2-opt, 3-opt, or generally, K-opt procedures employed in this step attempt to reduce the cost of a route by exchanging any K legs in the route for any K legs not in the route. When all the possible exchanges have been performed, a K-opt solution results.

Another improvement procedure is known as the Merge Technique. This procedure compares pairs of solutions in order to obtain a better merged solution with a reduced cost. The process is repeated until all feasible solutions have been compared.

The last improvement procedure employed by the Federal Express Corporation is called hubturning. The procedure attempts to reduce the cost by combining two pairings starting and ending in the same base.

The third step, composite method, starts with an initial feasible solution, and improves it via the improvement procedures discussed above, until a near-optimal solution is obtained.

2.3 ROUTING AND SCHEDULING PROBLEMS

The routing and scheduling problem requires the design of a schedule for a vehicle or a crew, given time window restrictions and precedence relationships. Time windows restrict a job to be completed within some interval period of time [s,t]. There exists a two-sided window if both extremes of the interval are specified while a one-sided window exists when only one extreme is given. The precedence relationship places a restriction on the order in which the jobs must be performed. When these two constraints are combined, the problem becomes extremely complicated. Further complications may result when more than one origin or destination (main) depot are considered. Hence, no exact procedures based on mathematical programming techniques have been successfully developed; only heuristic algorithms have been implemented. Moreover, this problem has proved to be NP-hard.

The routing and scheduling problem can be broken down into four major groups:

1. School bus routing and scheduling
2. Routing and scheduling of street sweepers
3. Airplane scheduling
4. Dial-a-ride routing and scheduling

The school bus routing and scheduling problem deals with a set of tasks (pick up and delivery) to be performed within a time interval given precedence relations. The objective of the problem is to minimize the number of vehicles required in the system while providing satisfactory service. Usually this problem involves more than one school to be serviced and a different time window for each school.

The problem can be complicated when the transportation cost depends on several factors besides the number of students. These factors can be for example the number of entities which provide the resources, or criteria on which the availability of resources are based, etc. For example in

New York city, the state, subsidizes only those students who live further than 1.5 miles from school [14]. Also different factors are involved in the problem depending on the area where the service is required.

A general procedure is followed by many algorithms proposed for the solution of this problem. First a set of routes for each school is constructed. Second a schedule for the buses is obtained where the schools are ranked according to the time that they begin services. Finally, an improvement of these schedules is sought. Algorithms have been proposed by Newton and Thomas [35], Bodin and Berman [12], and Bernett and Gazis [10] for the solution of this problem.

The street sweeper routing and scheduling problem is an adaptation of the Chinese Postman problem. The problem is stated as follows. A set of streets must be serviced in a certain period of time using a minimum number of vehicles. The basic approach to the solution of this problem is either of the variants of the route-cluster procedures.

A key constraint for this routing and scheduling problem arises from the parking regulations on the streets to be served, since sweeping can only take place when the vehicle comes close enough to the curb [14]. There could exist more than seventy possible parking regulations in one city, increasing the difficulty of the problem. No successful solution procedure has been developed for the solution of this problem, although the Chinese Postman algorithms have been employed for the solution of this problem (see Bodin and Kursh [15], Edmonds and Johnson [22]).

The airplane scheduling problem is a very complicated problem where the time table for crews and planes needs to be generated. Several factors such as frequency of the flights, number of passengers, etc. must be taken into account when the time table is constructed. Procedures to solve this problem were discussed in the previous section.

The dial-a-ride routing and scheduling problem deals with providing service to customers at any desired time. The problem can either be thought of as a dynamic problem wherein the calls are being received while the scheduling is being performed, or as a static problem if the customer demands are known before the scheduling is carried out. Both problems involve precedence relationships and time window restrictions. The dial-a-ride problem has several applications such as package delivery, bank delivery, vehicle for social services, etc. Many algorithms have been developed for the solution of this problem. For example the algorithms presented by Bodin and Sexton [18], and Sexton and Bodin [44]. These algorithms are generally more effective for one-sided window restrictions, than for two-sided window restrictions.

2.4 COMPARISON OF THE PROPOSED PROBLEM WITH THOSE IN THE LITERATURE

The problem under consideration in this thesis as was illustrated in Chapter I can be formulated based on the simple vehicle scheduling problem formulation, although additional constraints such as precedence relationships have been incorporated increasing the difficulty of the problem.

Due to the unique nature of the problem, there does not exist a compatible solution procedure in the literature. The assumptions of this problem, its objective and the requirement to pre-allocate crews among bases are features on which the subsequent scheduling of crews and planes depends, and which give this problem unique characteristics not previously considered in the literature of routing and scheduling problems.

This problem as many others in the literature, becomes intractable when the size of the problem increases. Therefore a formal mathematical programming solution procedure is not computationally feasible. Hence heuristic methods are proposed in this thesis for the solution of the problem.

The solution methods proposed in Chapter III for the Air Force Crew Scheduling problem, as many others which deal with this type of problem, exploit the unique features of the problem in order to obtain a solution procedure. Another similarity of these proposed methods with those in the literature, is the simulation process executed repeatedly in the procedure.

The structure of the Two-Phase Method, is similar in philosophy to the routing and scheduling techniques in that a sequence of missions on planes is obtained before the schedule of crews is considered. The One-Phase Method, however, attempts to schedule both, crews and planes simultaneously.

Chapter III

PROPOSED METHODOLOGY

This chapter deals with the two algorithms proposed for the solution of the Air Force Crew Allocation and Scheduling Problem. These algorithms attempt to schedule missions as well as initially distribute the crews among bases and thereafter schedule them in an appropriate manner. Both algorithms are shown to produce good quality solutions for the problem under consideration.

The major difference between these algorithms concerns the sequence of missions flown by each plane in the system. In the first algorithm, the Two-Phase method, the sequence of missions is fixed according to the output of the first phase; whereas in the One-Phase method, the sequence of missions is obtained concurrently while the missions are being performed. Incidentally, an upper bound for the number of crews required in the system is obtained as a byproduct of the Two-Phase method.

The objective of this problem is to schedule the tasks (missions to be flown) on machines (planes) so as to minimize the makespan T, where the makespan is defined to be the time taken to complete work on all the jobs. This is equivalent to minimizing the total idle time on all the machines, that is minimizing

$$pT - \sum_{j=1}^{m} t_j r_j \qquad (3.1)$$

where p is the number of planes in the system, t_j is the processing time of mission j, and r_j is the number of operations to be performed of mission j.

Equivalently, the objective can be stated as one of maximizing the utilization rate. This utilization rate is the ratio of the total duty hours spend by all the planes, (namely, flight and ground loading and unloading times), to complete all the jobs, and the total time taken for the completion of all the jobs. Mathematically, the utilization rate of planes is given by

$$\frac{\sum_{j=1}^{m} t_j r_j}{pT} \qquad (3.2)$$

Maximizing (3.2) is therefore equivalent to minimizing the makespan T.

The algorithms proposed here are also designed to study the sensitivity of and hence to suggest, the number of crews and /or aircraft needed to complete the missions with minimum idle time.

In the problem under consideration a set of missions is required to be completed over a given period of time. A restriction can be imposed on the time between the commencement of any two consecutive operations of any mission. This constraint implies a period of time T_j ($T_j \neq 0$), between the starting time of any two consecutive operations of mission j. Also a restriction can be imposed on the time between the completion of operation i and the start of operation i+1, of mission j, where i and i+1 are consecutive operations. This constraint yields another variant of the initial problem.

The heuristic dispatching rules developed for this problem are similar for both algorithms since both algorithms have the same assumptions.

A more detailed explanation of each algorithm will be given in the next two sections of this chapter.

3.1 TWO-PHASE METHOD

In this section, a two phase heuristic method is presented for the Air Force Crew Allocation and Scheduling problem. The first step in Phase I of this method deals with a uniform spread of missions over the desired period of time within which the missions must be completed, if such a spread is required by the problem (only one variant, namely, Problem 2, requires this feature). A uniform spreading of missions is performed to ensure that all the r_j jobs of mission j are not performed at the same time. The minimum value T_j between the starting times of any two consecutive operations of mission j, is calculated. The value T_j can be either given or calculated during the simulation process, producing the other two variants of the problem.

The spread of missions, when required by the problem, proceeds as follows.

1. For all missions j, calculate

   $$R_j = r_j t_j \qquad j=1,...,m \qquad (3.3)$$

   where t_j is the processing time (flight and ground time), and r_j is the number of operations of mission j to be performed.

2. Take the difference between the desired completion time (planning horizon) for all the missions, CT, and R_j.

   $$D_j = CT - R_j \qquad j=1,...,m \qquad (3.4)$$

3. Spread the difference for mission j over the total number of operations of mission j.

   $$SR_j = D_j / r_j \qquad j=1,...,m \qquad (3.5)$$

4. In order to ensure the completion of the jobs within the desired period of time, CT, let us consider the factors α and β, where α and β are real values. For positive values of D_j, a value between [0,1] is used for α in order to spread the missions over a certain period of time. A value of 0 would make the spread inactive while a value of 1 would force the last op-

eration of mission j to be finished at the end of the period. In order to ensure the completion of jobs within the planning horizon, even though idle time is incurred.

$$S_j = \alpha SR_j \qquad \text{if } D \geq 0 \qquad j=1,\ldots,m \qquad (3.6)$$

For negative values of D_j, a value β between [1,2] ensures the completion of the r_j operations of mission j within the desired period of time. A value of 1 spreads the missions such that operation r_j will be completed at the end of the period. In order to ensure the completion of jobs within the planning horizon, even though idle time was incurred, a value between (1,2) is recommended for β.

$$S_j = \beta SR_j \qquad \text{if } D < 0 \qquad j=1,\ldots,m$$

5. The starting time of any job i is given by,

$$T_{ij} \geq T_{(i-1)j} + S_j + t_j \qquad i=1,\ldots,r_j \quad j=1,\ldots,m \qquad (3.7)$$

where T_{ij} is the starting time of operation i of mission j, S_j is defined in Equation (3.6), and t_j is the processing time of mission j.

Thereafter, Phase I of this method sets up and solves a relaxed problem as a modified parallel processor problem with precedence constrained operations. In the relaxed problem assumptions (a), (b), (c), and (d) for planes, and no assumptions for crews are involved. Hence, Phase I obtains a lower bound on the time required to schedule all the missions, an upper bound $\bar{c}$ on the number of crews required, and the manner in which these $\bar{c}$ crews may be distributed among the n bases to yield a non-delay schedule for the relaxed problem. In addition, this phase determines for each plane the sequence of operations it must perform. Also Phase I distributes the available $\bar{c}$ crews in direct proportion to the above distribution of the $\bar{c}$ crews, subject to $c_1 \geq p$ for Problem 1 and $c_1 \geq m$ for Problems 2 and 3, where p is the number of planes in the system and c_1 denotes the number of crews placed initially at Base B1.

Phase II can be divided into two steps. Given an initial distribution of crews among the bases, Step 1 simulates the conditions of the problem, using suitable tie-breaking rules whenever necessary, to obtain the resultant schedule for the missions using the sequence from Phase I. Also this phase computes the utilization rate and the total idle time spent by the planes at each base.

Step 2 determines whether the simulation has been executed a pre-chosen number of times, in which case the process is stopped. Otherwise, a re-distribution of the c crews among the n bases is performed in approximately direct proportion to the idle time of the planes at the n bases as determined in Step 1. If a stopping criteria is satisfied, the process is terminated. Otherwise it returns to Step 1.

At termination, the distribution of the crews corresponding to the maximum utilization rate of planes is the one recommended for implementation. A schema of the algorithm is given in Figure 1.

3.1.1 Phase I

In this phase, we consider a relaxed version of the problem described in Chapter I. Specifically, we assume that only stipulations (a), (b), (c), and (d) need hold for the planes and no restrictions need hold for the crews. In other words, we initially assume that whenever a plane is ready for take-off, a rested crew is available to accompany it. Indeed, this assumption will permit us to determine an upper bound on the number of crews required in the system. Having solved the relaxed problem, we invoke the restrictions (a), (b), (c), (d) and (e) on crews in order to determine a good starting allocation of crews among the bases.

Now, under assumptions (a), (b), (c), and (d) on the aircraft, and assuming the ready availability of crews, Problem 3 may be cast into the framework of a modified parallel processor job scheduling problem as follows. Consider a situation in which there are m missions, the j^{th} mission having

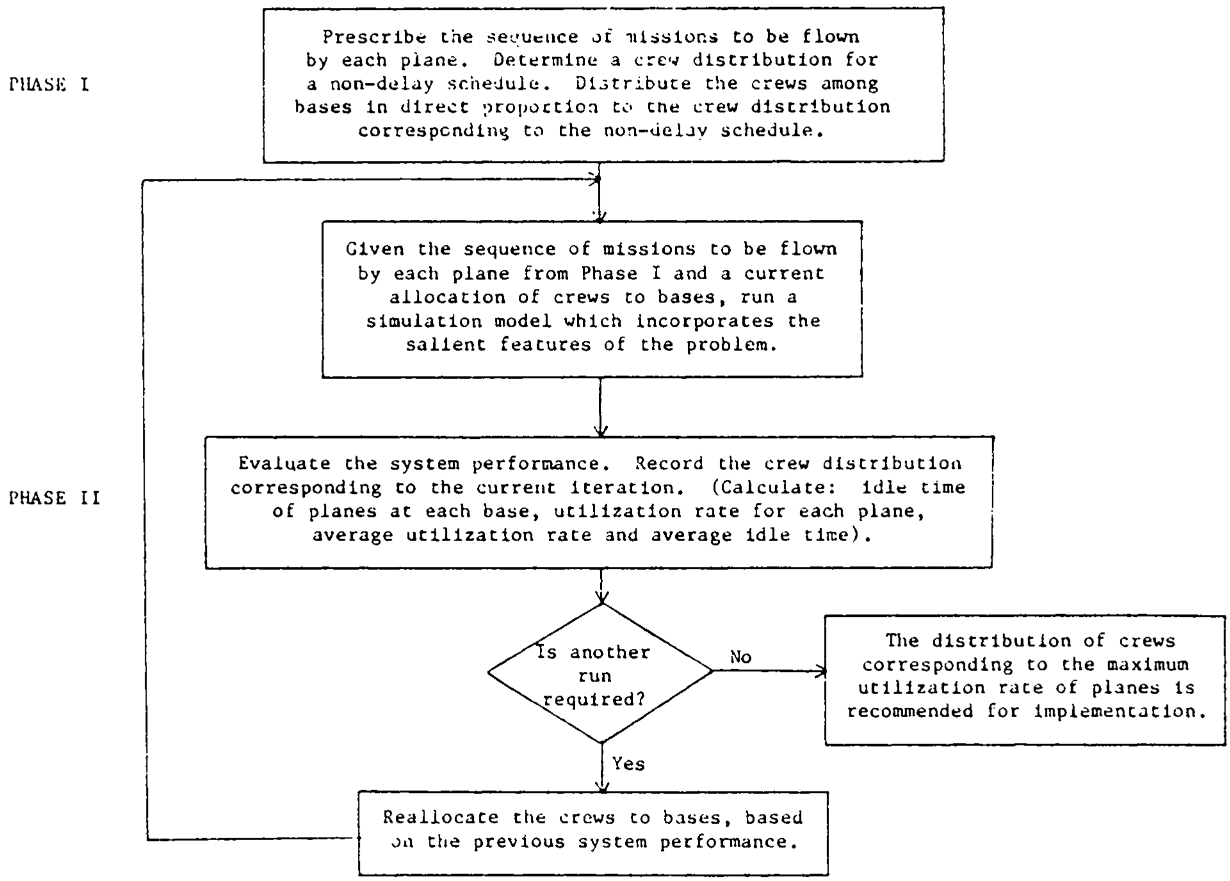

Figure 1: Flowchart of the Two-Phase Method

r_j so called real-operations $j_1,\ldots,j_{r_j}$ and r_j-1 so called virtual-operations $v_{j1},\ldots,v_{j(r_j-1)}$ for each $j=1,\ldots m$. Each real operation of mission j is assumed to take t_j time units whereas each virtual operation of job j is assumed to take T_{ij} time units (where T_{ij} is the time between the completion of operation i-1 and the commencement of operation i of mission j). Further, these operations are constrained to be performed in the sequence

$$j_1, v_{j1}, j_2, v_{j2}, \ldots, j_{(r_j-1)}, v_{j(r_j-1)}, j_{r_j}.$$

Note that in our case, each mission corresponds to a particular type of mission and each real-operation of a mission corresponds to the actual processing time of the corrsponding route. Hence t_j denotes the total non-delay time (processing time) for route j, $j=1,\ldots,m$. A virtual-operation is inserted between consecutive operations to account for the turnover time T_{ij} specified in assumption (d) for planes.

Continuing with the problem description, let there be p identical real-machines and $\sum_{j=1}^{m} r_j - m$ identical virtual-machines. Any real-operation can be performed on any real-machine and any virtual-operation on any virtual-machine. Thus the real machines correspond to the planes in our case and the virtual-machine serves as an artifice in accomodating the minimum turnover time constraints since the number of virtual-operations equal the number of virtual-machines.

The objective is to schedule the jobs on machines so to minimize the makespan T, or equivalently, maximize the utilization rate of the real-machines.

The precedence relationship between the various operations of the missions are illustrated in Figure 2. Here, each node represents an operation and arcs represent direct precedences [5]. The number at each node denotes its operation time, with node NT representing an artificial terminal node. The precedence structure illustrated in Figure 2 is known as an assembly tree (see [5], for example). In general, an assembly tree structure is one in which there exists a unique path leading from any node to the terminal node. Hu [30] has developed a procedure for scheduling single-operation jobs with assembly tree precedence structures and with equal processing times on identical parallel processors. Baker [5] has extended this to the case where the jobs may have non-identical operation times, as in our case. A statement of this scheme specialized to the precedence structure of our problem is given below. Note that this is a heuristic scheme.

Labeling Step

Label each node with the sum of the processing times on the nodes lying on the (unique) path from it (and including it) to the terminal node.

Scheduling Step

Select at most p operations without predecessors, choosing those with the largest labels first. Schedule these operations by assigning them in order of decreasing labels, each time selecting that machine with the least load on it thus far. Remove the scheduled operations from the precedence network and repeat this step. Continue until all operations are scheduled.

Although this procedure produces good solutions for problems where $T_{ij} \neq 0$ as in Problem 3, this is no longer valid for Problems 1 and 2 defined in Chapter I. Therefore, for the sake of uniformity, we will use certain heuristic tie-breaking rules for all cases based on the concepts of the above assembly-tree algorithm.

Toward this end, consider the following allocation of a sequence of specific operations to each plane. This sequence plays an important role in Phase II. A value p_j is first assigned to each mission j to denote the product of the processsing time t_j and the number r_j of operations of mission j to be performed. An ordered list of missions is constructed according to p_j, Using the LPT rule [3]. The first mission in the list is assigned to the first aircraft available. Now, the number of visited bases in common bet-

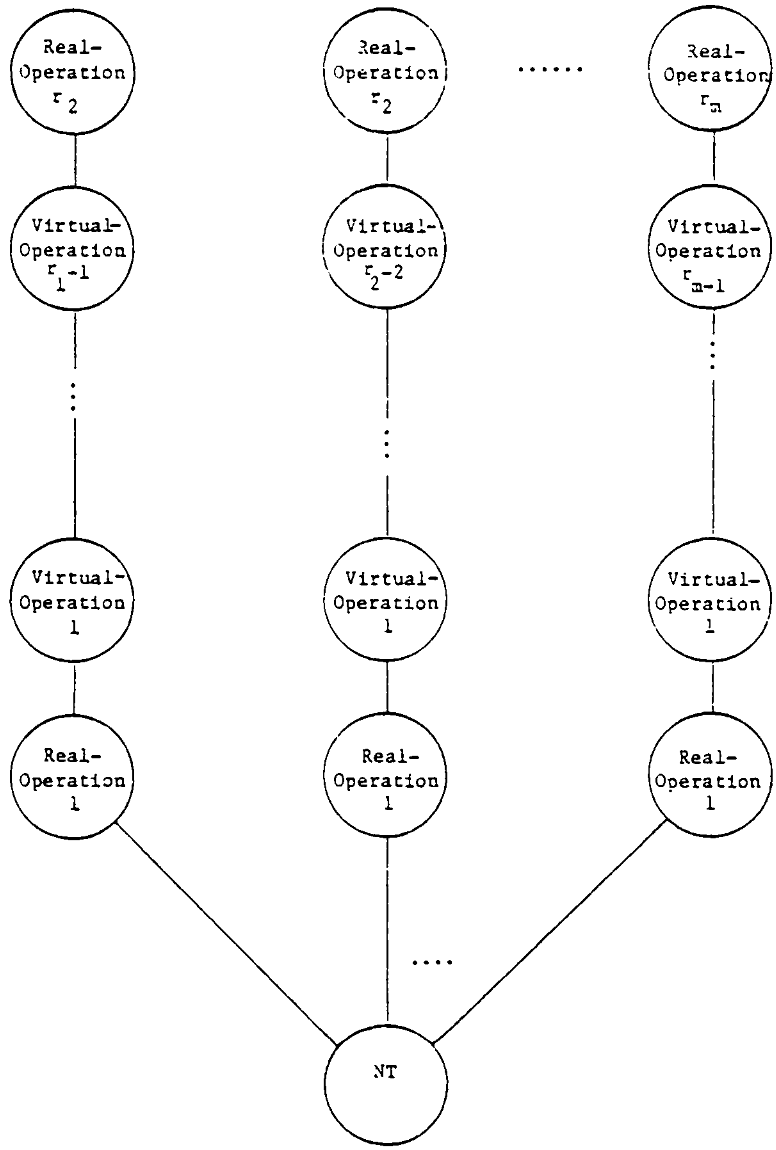

Figure 2: An Assembly Tree

ween mission j and the mission being served is calculated for each mission in the list. Let v_j denote this value. The mission with the minimum v_j is assigned to the next plane. In case of a tie wherein more than one mission has the same value of v_j, the mission with the largest p_j is chosen. The value p_j for the assigned mission, as well as the list of the remaining missions, is updated. v_j is re-calculated for the updated list and another mission is assigned continuing the process until there are no more missions or aircraft available. Then whenever a plane returns to home Base B1 or a mission is available, the availability of mission in the former case or plane in the latter case, is tested. If there exists such a feasible assignment, it is performed using the following tie breaking rule if necessary.

> If at home base more than one plane is ready for take-off, the plane which has the fewest flight hours at this point is chosen for take-off.

Otherwise the plane or mission enters in a waiting line until an assignment is possible. Notice that the number of missions in the list will decrease with each assignment. When all the missions have been assigned and all the planes return to Base B1, the process stops. A Flowchart of the rule described above is given in Figure 3.

Now, given the sequence of missions for each plane obtained above, a simulation process for a non-delay schedule is conducted to obtain an upper bound on the number of crews required at each base. Assumptions (b), (c), (d), and (e) for crews are now involved.

This simulation process assumes an infinite number of crews in the system. Initially p crews, equal to the number of planes in the system, are placed at Base B1. Subsequently, whenever a plane is ready for take-off at any base i, a crew is placed at that base i if there is no crew available according to assumption (d) for crews. The process continues, increasing the number of crews at each base in the system when necessary, until the scheduling of all the jobs is over, i.e., all the missions are completed. In this manner the number of crews required at each base for a non-delay schedule is obtained. Therefore, an upper bound for the total number of crews required in the system is obtained.

If the resultant schedule is optimal, then clearly its makespan is a lower bound on the time required to schedule r_j operations of type j, j=1,...,m.

Further, if the solution is optimal, the number of crews determined in the manner described above is an upper bound on the number of crews which the system may use effectively.

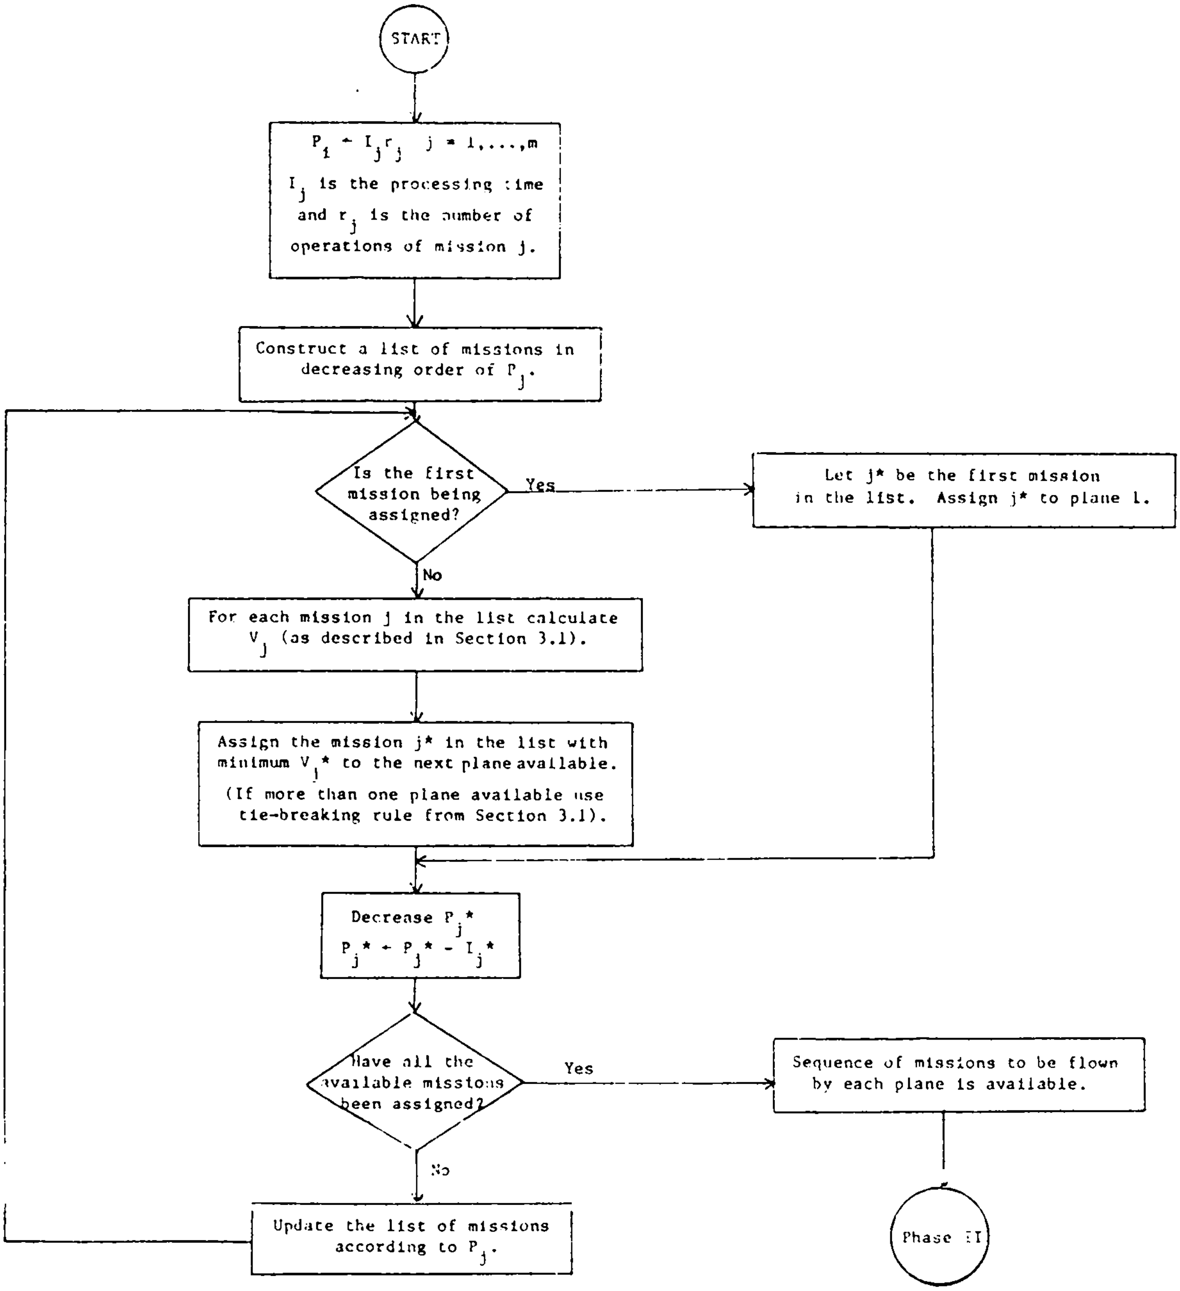

Figure 3: Missions Assignment Tie-breaking Rule

Next, we invoke assumption (a) on the crews. That is, we determine an initial distribution for the c available crews among the bases by allocating them in direct proportion to the distribution obtained for the crews required for the above non-delay schedule. Specifically,

$$c_i = \left\langle \frac{\tilde{c}_i \cdot c}{\sum_{i=1}^{n} \tilde{c}_i} \right\rangle \qquad i=1,\ldots,n \qquad (3.8)$$

where c_i is the number of crews assigned to base i, $\tilde{c}_i$ is the number of crews at base i for the non-delay schedule, c is the total number of crews in the system and where $\langle u \rangle$ is the integer part of u. Also at least one crew is required at those bases visited more than NV times (NV=2 in our computations). If $\sum_{i=1}^{n} c_i$ becomes greater (less) than c due to integrality of c_i and the restriction of at least one crew at the required bases, the number crews in $\left| \sum_{i=1}^{n} c_j - c \right|$ bases with the minimum number of visits is decreased (increased) by one, while maintaining at least p crews for Problem 1 and m crews for Problem 2 and 3, at Base B1, and at least one crew at each base visited at least NV times. A schema of Phase I is given in Figure 4.

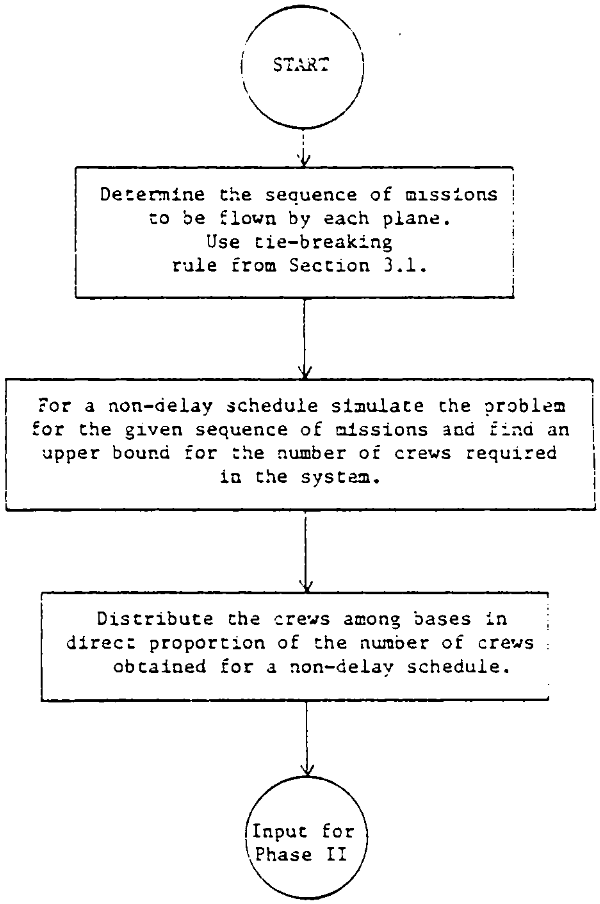

Figure 4: Flowchart of Phase I

3.1.2 Phase II

Given the sequence (not the schedule) in which each plane flies its assigned missions as prescribed in Phase I, an iterative simulation phase is set up. This phase incorporates the salient features of the problem as specified by assumptions (a), (b), (c), (d), and (e) for planes and assumptions (a), (b), (c), (d) and (e) for crews. The simulation process is essentially a Gantt Chart or a schedule construction routine. Hence, this procedure must be performed just once for each distribution of crews among bases. At the end of each iteration the crews are re-distributed among bases approximately in direct proportion to the idle time of the aircraft at each base in the current iteration.

The simulation process proceeds as follows. For the initial distribution of crews among bases and the sequence of missions to be performed by each plane as obtained from Phase I, the scheduling of the assigned missions is performed subject to the availability of mission-resource pairings. The simulation is performed on the basis of next-event incrementing, wherein whenever an event occurs, all the parameters in the system are updated. The events are:

1. arrival of a plane at any base r
2. availability of planes at any base k
3. availability of crew at any base k

4. availability of a mission at Base B1

If the event is the availability of a mission or of one of the resources, the availability of the other two factors is verified in order to incorporate another operation in the simulation process. If one of the factors is not available the others factors wait until they are all available while the next event is selected. Otherwise the following factors are updated for mission j being flown by plane i from base k to base r (k=B1)

1. list of plane available at base k
2. time for the availability of plane i at base r
3. next base to be visited by plane i (base r)
4. arrival time of plane i at the next base r
5. list of crews available at base k
6. time for the availability of the next crew at base k if one exists
7. total plane idle time at base k
8. mission j being flown by plane i
9. time for the availability of mission j
10. number of hours already flown by plane i
11. number of missions already flown by plane i

If the next event is the availability of a plane or a crew at any base k other than Base B1, factors (1) through (7) are updated.

For the arrival of plane i at any base r the following information is updated:

1. availability of plane i at base r
2. number of planes at base r
3. number of crews at base r
4. availability of the arriving crew at base r
5. availability of the next crew (existing or arriving) at base r

If the arrival occurs at Base B1, assumption (e) for plane is also verified. This process continues until all the missions are completed.

In order to obtain the desired information for the problem at the end of the process, the following information must be maintained and updated when necessary.

1. idle plane time at each base
2. total flight hours for each plane
3. number of times that each mission has been performed
4. distribution of crews for each iteration
5. number of iterations already executed

Assumption (e) for planes is incorporated using some stochastic features. That is, every 45 days, a minor maintenance for planes is required in our problem. The length of the down time is determined via a discrete distribution fuction. A random number is generated assuming a uniform distribution. According to the value obtained for the random number, one of three possible down times is chosen:

1. 24.00 hours are required for a routine minor maintenance with .60 probability of occurence
2. 60.00 hours are required for minor repair with .28 probability of occurence
3. 96.00 hours are required for major repair with .12 probability of occurence

In this manner, a schedule is obtained for the planes and for each crew over some period. After the schedule for the missions is thus obtained, the utilization rate for each plane is calculated by,

$$UR_i = \frac{\sum_{j=1}^{q} t_{ij}}{TD} \qquad i=1,\ldots,p \qquad (3.9)$$

and the overall utilization rate of planes is calculated by

$$UR = \frac{\sum_{j=1}^{q} \sum_{i=1}^{p} t_{ij}}{\sum_{j=1}^{q} \sum_{i=1}^{p} t_{ij} + \sum_{i=1}^{m} It_i} \qquad (3.10)$$

where t_{ij} is the total processing time of mission j on plane i including flight and ground (loading and unloading) time, TD is the total time required to complete all the jobs, UR_i is the utilization rate for plane i, $q=\sum_{j=1}^{m} r_j$ and It_i is the total idle time of the planes at base i. The total idle time of the planes at each base is computed by the following equation

$$It_i = \sum_{j=1}^{p} ip_{ij} \qquad i=1,\ldots,n \quad (3.11)$$

where It_i is the total idle time of planes at base i, and ip_{ij} is the idle time of plane j at base i.

Another iteration, if necessary, is initiated as follows. The available c crews are re-distributed among the n bases according to the rules discussed below. A test of the stopping criteria given below is performed. If one of the criteria is satisfied the process is terminated; otherwise, Phase II is executed again. Then, the crew distribution corresponding to that schedule which has the maximum utilization rate is recommended for implementation. A schema of this Phase is shown in Figure 5.

Tie-breaking rules were developed in this algorithm for the assignment of planes, crews and routes. These rules are based upon the unique properties of the problem. The heuristic rules along with the stopping criteria are listed below:

1. Aircraft Assignment Rules

   a) If at home Base B1, more than one plane is ready for take-off and there exists only one crew which can accompany it (feasible to the constraints), then the aircraft which has logged the fewest

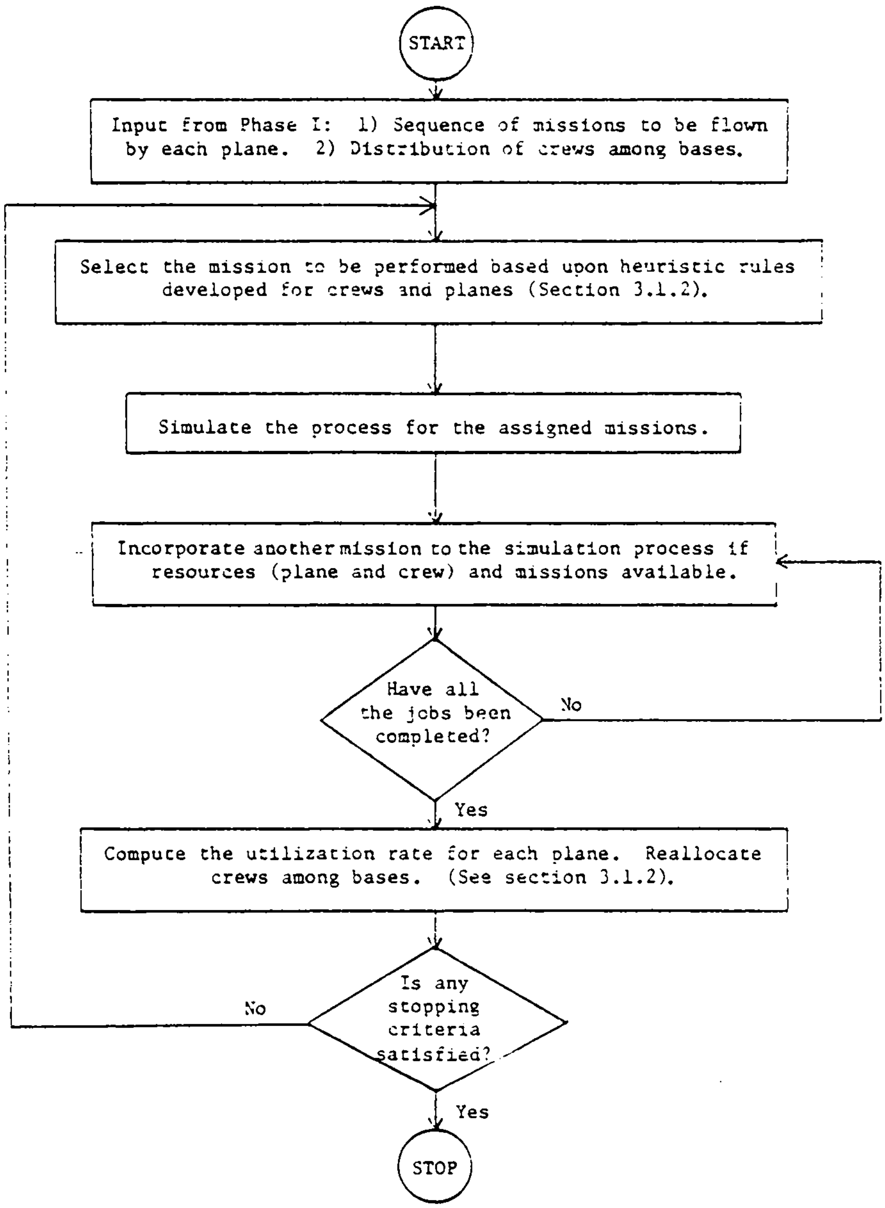

Figure 5: Flowchart for Phase II

flight hours in the process is chosen for take-off. This assignment of aircraft is performed one plane at the time.

b) If at any base other than home Base B1, a rested crew is available and there exists at that base more than one flight which this crew can accompany, then the flight associated with that mission j is selected for which the sum of the number of crews assigned to the remaining bases to be visited by that mission is a maximum.

2. Crew Assignment Rule

a) If at any base a plane is ready for take-off and there exists more than one crew which can accompany it (feasible to the constraints), then that crew which has had the longest rest period at that base is assigned to the plane.

3. Crew Re-distribution Rule

a) At the end of each iteration, the redistribution of crews in the system is based on the idle time of the planes at each base in the current iteration. An exchange of k crews is performed from the k bases with minimum idle time to the k bases with maximum idle time. A maximum difference of one crew per base is allowed between two consecutive iterations. The redistribution of crews is as follows:

i) Construct a list of n bases in decreasing order of the total idle time spent by the planes in the current iteration at each base.

ii) Increase the number of crews for the first K bases (for the problem to be discussed in Chapter IV, K=2) in the list by one.

$$c_i = c_i + 1 \qquad i=1,\ldots,K \qquad (3.12)$$

iii) Decrease the number of crews by one for the last K bases in the list, while $c_i \geq 1$ is satisfied for those bases visited more than NV times (in our computations NV=2). Note that more than K bases may be considered in the decreasing process due to the restriction of $c_i \geq 1$.

$$c_i = c_i - 1 \qquad i=n-K+1,\ldots,n \quad (3.13)$$

iv) Keep the same number of crews in the remaining bases.

$$c_i = c_i \qquad i=K+1,\ldots,n-K \quad (3.14)$$

4. Stopping Criteria

a) If the process has already been executed a pre-chosen number of times, a termination occurs.

b) If the resulting distribution has been generated at a previous iteration, the process stops.

c) If there is no idle time for planes at all the bases, the process stops.

Many rules based upon the distribution of crews at previous iterations were tried to re-distribute the crews among the bases. Although these rules were not succesful for our problem, some of them are mentioned here in order to discuss their relative merits as well as to compare them with the redistribution rule being used in this algorithm. Three of these rules are explained below:

1. The distribution of crews in the next iteration depends on the number of crews at each base and the idle time of planes at the current iteration. The number of crews at each base is given by the following equation.

$$c_i = \left\langle \frac{(c_i^{\alpha} w_i)\ c}{\sum_{i=1}^{N} c_i^{\alpha}\ w_i} \right\rangle \qquad i=1,\ldots,n \qquad (3.15)$$

where c_i is the number of crews at base i in the current iteration, c is the total number of crews in the system, w_i is the idle time at base i in the current iteration. α is an exponent factor, which attempts to weight c_i differently than w_i, i.e., c_i is weighted more than w_i if α is a positive integer and is weighted less than w_i if α is a positive fraction between 0 and 1. Finally, $\langle u \rangle$ denotes the integer part of u. The problem requires a minimum of one crew at each base i, visited more than certain number of times, namely NV, and a minimum of p crews at Base B1, where p is the number of planes in the system. If the total number of crews in the system becomes greater (less) than c, c_i is decreased (increased) by one for the $|c-\sum_{i=1}^{n} c_i|$ bases with minimum (maximum) idle time in the current iteration, subject to p crews at Base B1.

2. Previous iterations are weighted in order to obtain a good distribution of crews among the n bases. The formula designed for this rule follows:

$$\left\langle \frac{\sum_{j=1}^{I} (c_{ij}^{\alpha} w_{ij})\, c}{\sum_{j=1}^{I} \sum_{i=1}^{n} (c_{ij}^{\alpha} w_{ij})} \right\rangle \quad i=1,\ldots,n \qquad (3.16)$$

Here, I is the number of previous iterations considered in the equation, c_{ij}, is the number of crews at base i at iteration j, w_{ij} is the idle time of all the planes at base i at iteration j, c is the total number of crews in the system. As in rule 1 the number of crews at any base must be an integer number greater than zero. The number of crews at Base B1 must be at least p and the sum of c_{ij} at iteration j, over all i=1,...,n must be equal to c. If the total number of crews in the system becomes greater (less) than c the process continues as in Rule 1.

3. The number of crews at base i in the next iteration is based upon the difference between the average idle time and the idle time at base i in the current iteration. This rule attempts to uniformally distribute the crews among bases according to the average idle time in the system. The equation employed for this rule is the following.

$$c_i = \left\langle (w_i - av) . sf + ca \right\rangle \qquad i=1,...,n \qquad (3.17)$$

$$sf = \sum_{i=1}^{n} c_i \; / \; \sum_{i=1}^{n} w_i \qquad (3.18)$$

where av is the average idle time of planes in all bases at the current iteration, ca is the average number of crews per base (c/n), sf is a scale factor, w_i is the idle time of planes at base i. As in the previous rules, c_i must be a nonnegative value and c_1 must be at least p.

The rules discussed above present several disadvantages for our problem. First, rule 1 as well as rule 2 base the assignment of crews to base i on the product of the number of crews at base i in the current iteration and the idle time at this base in the same iteration. Hence, whenever there is no idle time (w_i=0) at one base, at most one crew will be assigned to this base in the next iteration (due to the constraint of a minimum of one crew at each base) regardless of the number of crews at this base in the current iteration. As a result, an erroneous distribution of crews among bases may be obtained, and the number of iterations required to obtain a good distribution of crews, if one is obtained, increases significantly. Although rule 2 considers more than one previous iteration, the product w_ic_i, affects the distribution considerably and too many iterations are needed before a good distribution is reached.

Rule 1, as well as, rule 2 use a factor α in order to weight c_i. Although this factor may produce a better distribution of crews among bases, a trial and error procedure is required before an appropriate factor for the problem under consideration is found.

Even though rule 3 produces better solutions than rules 1 and 2, it is not recommended for implementation because too many iterations are required before a good solution is obtained.

In contrast with these rule, the redistribution rule used in this algorithm produces an improved solution usually within two iterations. The average number of iterations required to obtain a good solution is approximately 3 iterations. Also the implemented Re-distribution Rule is simpler than the rules mentioned above and less computational effort is required (less execution time).

3.2 ONE-PHASE METHOD

The solution method proposed in this section for the Air Force Crew Allocation and Scheduling problem is essentially a deterministic simulation process. In this method, instead of fixing a sequence of missions for each plane before finding the schedule, the sequence of operations for each plane is obtained simultaneously while the simulation of missions is being executed.

This proposed method consists of four steps: uniform distribution of missions over the desired completion time (when it is required), initial distribution of crews among the n bases, simulation of jobs, and re-distribution of crews among bases and the test of the stopping criteria. Steps 1 and 2 do not involve assumptions for planes or crews while Steps 3 and 4 involve assumptions (a),(b),(c),(d), and (e) for aircraft as well as assumptions (a),(b),(c),(d), and (e) for crews. Figure 6 presents a schema for the proposed method.

3.2.1 Step 1

This step is required only when $T_j \neq 0$, and T_{ij} is fixed for all mission j at the beginning of the process (Recall that for some other variants of the problem, $T_j \neq 0$ means be calculated during the simulation process itself).

The first step of this method distributes or spreads the missions to be performed over the desired period of time within which all the missions must be completed. The uniform spread of missions is performed to ensure that all the operations of any specific mission will not be performed at the same time. Details of the technique used to perform this step have been described in the discussion of the Two-Phase Method in Section 3.1.

3.2.2 Step 2

The initial distribution of crews among the n bases is based upon the number of times that each base is visited by all the jobs (i.e. the number of times that each base has been visited at the end of the process by all the missions). Let m be the number of missions to be performed, r_j the number of operations of mission j (number of times that mission j must be performed by the end of the period) and v_j the number of times that base k is visited by mission j. Then the total number of times that mission j visits base k is

$$v_{kj} r_j \qquad j=1,\ldots,m \quad k=1,\ldots,n \qquad (3.19)$$

Based on this equation one can calculate the total number of visits to base k by

$$h_k = \sum_{j=1}^{m} v_{kj} r_j \qquad k=1,\ldots,n \qquad (3.20)$$

For a given c, the total number of crews in the system, the number of crews initially assigned to base k is given by

$$c_k = \left\langle \frac{(c - mp)\ h_k}{\sum_{k=2}^{n} h_k} \right\rangle \qquad k=2,\ldots,n \qquad (3.21)$$

where mp is the number of crews at Base B1. As required by the problem, the number of crews initially assigned to Base B1 is p for Problem 1 and m for Problem 2 and 3. Due to the rounding in Equation (3.21) and the restriction $c_i \geq 1$ for those bases with at least NV visits, it is possible to obtain a total number of crews greater (less) than the number of crews available in the system. In this case a crew reduction (increasing) process is conducted. In this reduction (increasing) process, the number of crews in the system for the first $|c - \sum_{i=1}^{n} c_i|$ bases is decreased (increased) by one until there are c crews in the system. When a reduction process is required, it takes place in any $|c - \sum_{i=1}^{n} c_i|$ bases except Base B1 to satisfy the constraint $c_i \geq mp$ while in an increasing process all the bases are considered.

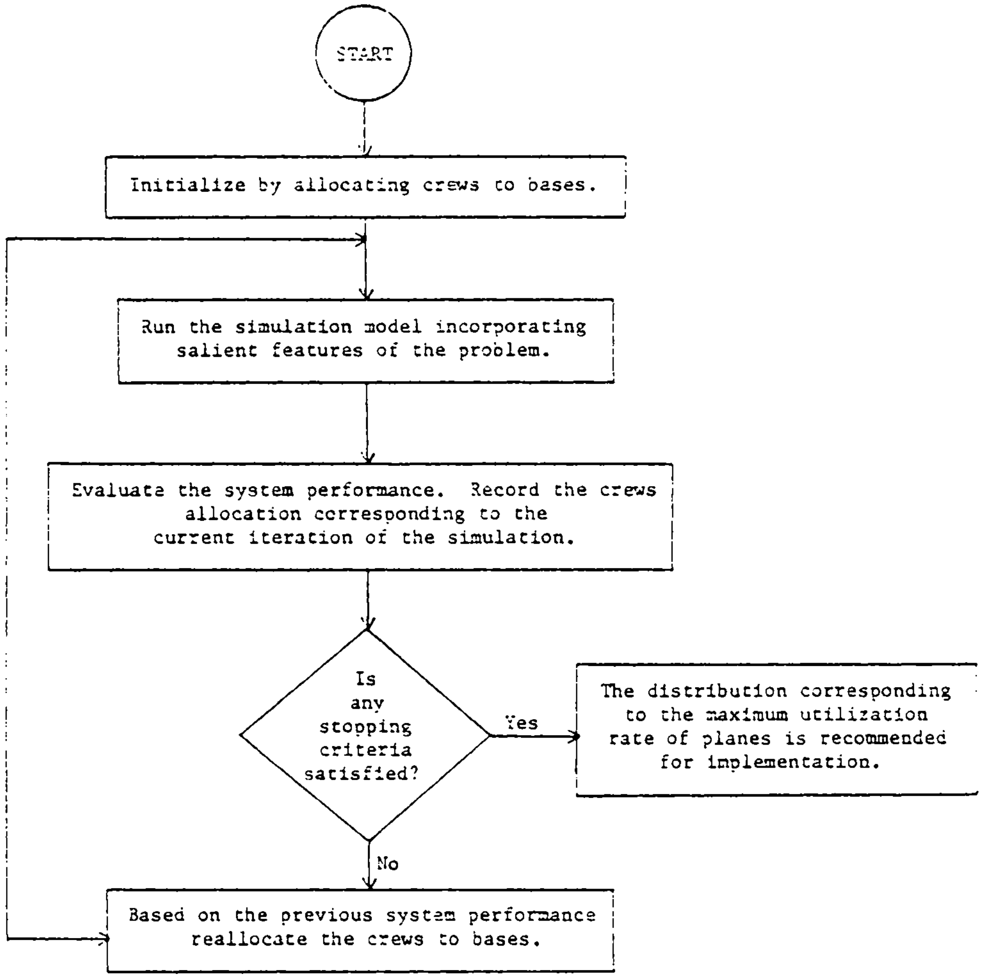

Figure 6: Flowchart for the One-Phase Method

3.2.3 Step 3

Step 2 conducts a simulation process to find a good heuristic schedule of missions. This step incorporates the salient features of the problem as specified by assumptions (a),(b),(c),(d), and (e) for aircraft and assumptions (a),(b),(c),(d), and (e) for crews. The simulation process needs to be performed just once for each distribution of crews among bases.

The first part of the simulation is the assignment of missions to planes where the following three requirements must be satisfied:

1. Availability of aircraft
2. Availability of crew
3. Availability of mission

The simulation process is conducted for the assigned missions producing the following output for each plane on duty:

1. Arrival time at each base visited by the mission in progress
2. Departure time from each base visited by the mission in progress
3. Idle plane time at each base

Simultaneously a test of availability for aircraft, mission, and crew is performed at the home base. During the simulation process four events can occur:

1. availability of crew at any base k
2. availability of plane i at any base k
3. availability of mission j at Base B1
4. arrival of plane i at any base r

Whenever an event occurs some parameters in the system are updated. If any one of the first three events occur a test for the availability of the other two factors is performed. If the availability restrictions are satisfied for a certain assignment, it is incorporated into the simulation process. Otherwise the available factors have to wait until all the factors are available. Then another event is determined.

Whenever one of the first three events occurs at base k the following information is updated:

1. next base to be visited by the mission being flown by plane i
2. list of planes at base k
3. arriving time of plane i at the next base to be visited
4. list of crews at base k
5. time when the next crew will be ready for take-off at base k
6. total plane idle time at base k

If the availability of plane or crew occurs at Base B1 the following information also needs to be updated:

1. mission to be flown by plane i
2. number of times the assigned mission j has been performed
3. time when mission j can be reassigned

When the arrival of plane i at any base r occurs, the following information is updated:

1. departure time for plane i
2. availability time for the arriving crew
3. time when the next crew (existing or arriving) is ready for take-off at base r

If the arrival occurs at Base B1, assumption (e) for planes needs also to be verified.

In order to obtain the desired information at the end of the process, the following information needs to be maintained and updated:

1. idle plane time at each base
2. total flight time for each plane
3. number of times that each mission has been performed
4. distribution of crews at each iteration
5. number of iterations executed

Assumption (e) for planes is incorporated based on some stochastics features. That is, after 45 days, each plane requires a scheduled maintenance. The duration of this down time period is determined via a discrete probability distribution fuction. A uniform random number is generated in order to determine how many days are required for the regular maintenance. According to the random number three case can occur:

1. 24.00 hours are required for minor maintenance with .60 probability of occurence
2. 60.00 hours are required for minor repair with .28 probability of occurence
3. 96.00 hours are required for major repair with .12 probability of occurence

The simulation process continues until all the missions have been completed.

Some heuristic rules have been developed according to the special properties exhibited by this problem in order to assign crews, planes and missions during the simulation process. These rules are as follows:

1. Aircraft Assignment
   a) If at the home Base B1 more than one plane is ready for take-off and there exists a crew which can accompany it on some mission (feasible to the constraints), then the plane which has logged the fewest flight hours so far is chosen for take-off. This assignment of aircraft is performed one plane at the time.

b) If at any base different from the home Base B1, a rested crew is available and there exists at that base more than one flight which can accompany it, then the flight for which the sum of the number of crews assigned to the remaining bases to be visited in that operation is a maximum, is selected.

2. Crew assignment

a) If at any base a plane is ready for take-off and there exists more than one crew which can accompany it on some mission (feasible to the constraints), then that crew which has had the longest rest period at that base, is assigned to the plane.

3. Crew Re-distribution Rule

a) At the end of each iteration a re-distribution of crews is performed based on the total idle time of planes at each base in the current iteration. An exchange of k crews is performed one from each of the k bases with minimum total idle time to the k bases with maximum total idle time. A maximum difference of one crew per base is allowed between two consecutive iterations. The restribution is performed as in the Two-Phase Method (Section 3.1).

4. Route Assignment

a) The first assignment of missions at the begining of each iteration follows these steps :

i) Construct an ordered list of missions according to the product $p_j = t_j r_j$ where t_j is the processing time for mission j, and r_j is the number of operations of mission j.

ii) Assign the mission with maximum p_j to the first plane available.

iii) Update p_j for the assigned mission by subtracting t_j from it.

iv) Update the list eliminating the missions not available.

v) For each mission j in the list compute v_j as the number of bases visited by mission j that are also visited by the missions being flown.

vi) Assign the mission with minimum v_j on the list to the next plane available. If there are more than one mission with the same minimum v_j, the one that corresponds to the maximum p_j on the list, is assigned.

vii) Repeat steps (iii) to (vi) until there is no mission or plane available (the list is empty).

b) If more than one mission is available and there is only one plane ready for take-off and/or there is only one crew to accompany it, the mission to be assigned is chosen following these steps:

i) Construct an ordered list of available missions according to p_j.

ii) Compute v_j, the number of bases visited by mission j that are also visited by the missions being processed. v_j is computed only for those missions available.

iii) The mission with minimum v_j in the list is assigned to the available plane. If there is more than one mission with the same v_j, the one that corresponds to the maximum p_j in the list, is assigned.

iv) Update p_j for the assigned mission.

v) Update the list.

At the end of this step a schedule of operations is obtained for each plane and for each crew over some period. Now Step 4 is executed.

3.2.4 Step 4

For the schedule of missions obtained in Step 3 the utilization rate for each plane is computed by the following equation:

$$UR_i = \sum_{j=1}^{q} t_{ij}/TD \qquad i=1,\ldots,p \qquad (3.22)$$

where UR_i is the utilization rate for plane i, t is the processing time (flight and ground) of mission j assigned to plane i, TD is the completion time of the last mission assigned to plane i, $q = \sum_{j=1}^{m} r_j$.

Also an overall utilization rate is calculated by

$$UR = \frac{\sum_{i=1}^{p} \sum_{j=1}^{q} t_{ij}}{\sum_{i=1}^{p} \sum_{j=1}^{q} t_{ij} + \sum_{i=1}^{n} It_i} \qquad (3.23)$$

where It_i is the total idle time of planes at each base, calculated as follows:

$$It_i = \sum_{j=1}^{p} it_{ij} \qquad i=1,\ldots,n \qquad (3.24)$$

where it_{ij} is the idle time spent by plane j at base i.

According to the above information a re-distribution of crews among bases is performed if the total idle time in the current iteration is different from zero. The rule used to re-distribute the crews in this step was discussed in Section 3.1.2 of this chapter.

The next part of this step is concerned with the stopping criteria. These criteria determine whether another iteration is required or whether the process should be stopped.

The stopping criteria employed in this step are the following:

1. If the process has already been executed a pre-chosen number of times, a termination occurs.
2. If the new crew distribution has been generated at any previous iteration, the process stops.
3. If there is no idle time for planes at any base, the process is terminated.

If none of the above criteria is satisfied, another iteration is performed, and Step 3 and Step 4 are executed again. Otherwise the distribution of crews among the bases and the schedule of planes corresponding to the iteration with minimum makespan i.e. the maximum average utilization rate of planes, is the one recommended for implementation.

3.3 EXAMPLE PROBLEM

In order to illustrate the algorithms proposed in this thesis a small example problem is presented in this section. Lets consider a problem involving 5 planes, 10 crews and 3 missions. The planning horizon is 1.00 day.

The missions are required to be spread over the desired period of time. Therefore Problem 2 is being considered. For this purpose a period of time between the commencement of any two consecutive operations of mission j, T_j is calculated for all missions j at the beginning of the process.

A rest period of 2 hours is required for crews at the end of each leg. Ground time of 1.5 hours is required at each base for loading and unloading. The values for α and β used to spread the missions over a period of 1.0 day are 0.75 and 1.50 respectively.

The data for this problem is presented in Table 1.

TABLE 1

Example Problem Data

Route 1 frequency = 3

base #	1	2	4	1	
ft		2.5	1.5	1.0	

Route 2 frequency = 2

base #	1	3	2	3	1
ft		2.0	2.0	3.0	1.0

Route 3 frequency = 2

base #	1	4	2	3	1
ft		3.5	2.0	3.0	1.0

"ft" is the expected flight time between base i and base j

a ground time of 1.5 is required at the end of each leg

3.3.1 Two-Phase Method

As was described in Section 1 of this chapter, Phase I uniformly spreads the missions over the desired period of time. Using Equation (3.3)

$$R_1=3(9.5)=28.5$$

for $t_1=2.5+1.5+1.5+1.5+1.0+1.5=9.5$ and $r_1=3$. Then from Equation (3.4)

$$D_1=24.00-28.5=-4.5$$

where CT=24.00 hours (1.00 day). This diffence D_1 is spread over the total number of operations of mission j, from Equation (3.5)

$$SR_1=-4.5/3=-1.5$$

Since D_1 has a negative value a factor β is employed in order to obtain S_1 from Equation (3.6),

$$S_1=(1.50)(-1.5)=-2.25$$

In the similar manner S_2 and S_3 are calculated.

S_1	-2.25
S_2	-3.00
S_3	-5.25

Now using Equation (3.7) the starting time of operation i of mission j is obtained. Let us consider operation 2 of mission 1,

$$T_{21}\geq 0.00+9.50-2.25$$

In the similar manner, the starting time of each operation of missions 2 and 3 are obtained. For operation 2 of missions 2 and 3 T_{ij} is calculated.

T_{21}	7.25
T_{22}	11.00
T_{23}	10.25
T_{31}	14.50

The next step of this phase is to determine the sequence of missions to be flown by each plane. For this step the tie-breaking rule from Section 3.1.1 is employed. As described in Figure 3 a list of missions to be assigned is constructed in order of p_j where

p_1=28.5-9.5

p_2=28.0-14.0

p_3=31.0- 15.5

and whenever a plane or mission is available another assignment is performed until no more missions are available. The sequence obtained for each plane is given below.

Plane 1 : Route 3

Plane 2 : Route 2, Route 1

Plane 3 : Route 1, Route 2

Plane 4 : Route 3

Plane 5 : Route 1

Given the sequence of missions to be flown by each plane, the number of crews required at each base is calculated for a non-delay schedule. Initially 3 crews are placed at Base B1. At time 2.5 Plane 2 is ready for take-off at Base B3. Since no crew is available the number of crews at this base is increased by one. At time 7.0 Plane 3 is ready for take-off at Base B2 where there is one crew available, therefore the number of crews placed at that base remains the same. This process continues yielding the following distribution

number of crews at Base B1	7
number of crews at Base B2	3
number of crews at Base B3	3
number of crews at Base B4	2

Based on this distribution for a non-delay schedule, the crews available in the system are distributed among the bases according to Equation (9). The distribution is as follows.

Base B1	7
Base B2	1
Base B3	1
Base B4	1

Now, the information obtained from Phase I serves as an input to Phase II. In this phase the tie-breaking rules discussed in Section 3.1.2 are employed. Initially using rule

(a) for the assignment of planes, Planes 1, 2, and 3 are chosen to fly at time 0.00. Planes 3 and 4 stay at Base B1 since there is no route available to be flown.

For the chosen planes the simulation process is performed. During the simulation process more than one plane can be ready for take-off at any base. In this situation rule (b) for assignment of planes is used to determine the plane to be chosen. For example, consider the situation when the next event is the availability of crew at Base B2 (time=20.50). At this base there are two planes, Plane 3 and Plane 4, ready for take-off. Only one of them can take-off due to the availability of a single crew. Therefore the total number of crews at the bases to be visited in the remaining legs of this mission is calculated for the Planes 3 and 4. The plane with the larger number of crews located at the remaining bases is the one chosen. In this case for both planes the number of crews available is the same. Therefore, the plane that has been at the current Base B2 for a longer period of time is chosen for take-off (here, Plane 3 is chosen). Now the availability time for a crew at Base B2 is updated. Also the arriving time of Plane 3, as well as the next base to be visited, is updated.

At the end of this iteration a re-distribution of crews is performed. The distribution is based on the total idle time for planes at each base. The idle time at each base is as follows

Base B1	0.00
Base B2	3.30
Base B3	2.50
Base B4	1.00

The utilization rate for each plane is calculated from Equation (3.9), and the average utilization rate is 15.00 hour/day per plane.

The re-distribution rule used is discussed in Section 3.1.2. In this iteration Base B2 has the maximum idle time. Using k=2 the number of crews at this base is increased by one. Since Base B4 has only one crew and it is visited more than NV times (NV=2) no crew can be removed from this base. Therefore one crew is removed from Base B1 to be placed at base B2. Now, with k=2, Base B3 must also receive an additional crew if possible. However, no exchange of crew is allowed here because only Base B1 has more than one crew, and a difference of only one crew is allowed at each iteration for any base. The new distribution is hence as follows

number of crews at Base B1	6
number of crews at Base B2	2
number of crews at Base B3	1
number of crews at Base B4	1

Now the stopping criteria are tested. Since no stopping criteria is satisfied (Section 3.1.2) another iteration is performed. In the second iteration a maximum utilization rate of 16.154 hour/day per plane in 1.08 days, is obtained. At this iteration no stopping criteria is satisfied therefore a third iteration is performed. An average utilization rate of 16.00 hour/day per plane is obtained. Thereafter a re-distribution of crews is performed. This distribution turns out to be the same as that obtained in iteration 2, and therefore, the process is terminated.

3.3.2 One-Phase Method

This method consists of three steps. Step 1 as discussed in Section 3.2.1, spreads the missions over the desired period of time as is required by the present variant of the problem. The details of this step are given in Section 3.3.1 for the spread of missions in the Two-Phase Method.

The second step of this algorithm deals with the distribution of crews among the bases. From Equations (3.19) and (3.20) the total number of times that each base is visited, is calculated for each mission. For example, $h_1=2(3)+2(2)+2(2)=14$. Then from Equation (3.21) the initial distribution of crews obtained for this problem is as follows

number of crews at Base B1 5

number of crews at Base B2 2

number of crews at Base B3 2

number of crews at Base B4 1

For the given allocation of crews among bases from Step 2, Step 3 conducts a simulation process. First, the missions are assigned according to the tie-breaking rules discussed in Section 3.2.3. Route 3 is the first mission to be assigned since it has the maximum p_j value calculated from rule (a) for the assignment of missions. Next, Routes 1 and 2 are assigned. These routes have the same value of v_j calculated from Rule (a) for the assignment of missions. Therefore the route with the maximum p_j is assigned next. That is, Route 2 is assigned to the second plane available and Route 3 is the last route to be assigned at time 0.00. Since a period of time T_j is required between the starts of any two consecutive operations of a mission j, only these three jobs are initiated in the simulation process at the beginning of the process. The simulation of the missions is conducted using the tie-breaking rules for the assignment of crews, planes, and missions as discussed in Section 3.2.3.

Step 4 is performed. In this step the average utilization rate of planes, and the idle time of planes at each base is calculated via Equations (3.25) and (3.26) respectively. On this iteration an average utilization rate of 16.471 hours/

day per plane in 1.06 days, and an average idle time of 1.00 hour, is obtained.

At the end of the simulation process (when all the jobs have been completed), a re-distribution of crews among bases is performed using the re-distribution rule presented in Section 3.2.3. From Equations (3.22) through (3.24) a new distribution of missions is obtained as follows.

number of crews at Base B1	7
number of crews at Base B2	1
number of crews at Base B3	1
number of crews at Base B4	1

Also a test of the stopping criteria is performed. No stopping criteria is satisfied, therefore another iteration is required.

Steps 2, 3, and 4 are executed again for iteration 2. Tie-breaking rules are employed, wherever necessary, for the assignment of planes, crews and missions. For example in this iteration at time 11.00 hours, Planes 3 and 5 are ready for take-off at Base B2. Therefore tie-breaking rule (b) for the assignment of planes discussed in Section 3.2.3, is used. In this rule the total number of crews assigned to the remaining bases to be visited by the mission being flown for each plane at that base, is calculated. In this problem both planes have the same number of crews located at the remaining bases. Hence, the plane which has been at that base for the longest period of time is chosen for take-off. In this case, Plane 3 is chosen. At iteration 2 an average utilization rate of 15.00 hours/day per plane in 1.17 days, and an average idle time of 1.688 hours, is obtained. No stopping criteria is satisfied, and therefore, iteration 3 is performed.

Iteration 3 yields an average utilization rate of 1.082 hours/day per plane in 1.08 days, and an average of 0.688 hours of idle time. In Step 4 of this iteration the distribution of crews obtained is the same as the distribution obtained at iteration 2. Hence the algorithm stops, yielding a maximum average utilization rate of 16.471 hour/day per plane at iteration 1.

The algorithms discussed in this chapter were implemented using data from the U.S. Brooks Air Force base. The results obtained, along with a comparison of the algorithms, are presented in the following chapter.

Chapter IV

RESULTS AND CONCLUSIONS

This chapter presents the results of computational experience with the two methodologies, namely, the Two-Phase Method and the One-Phase Method, proposed in Chapter III for the solution of the Air Force Crew Allocation and Scheduling Problem. The (real-world) data used for computational purposes, as well as a comparison of the relative merits of these methods, is also presented in this chapter.

4.1 DATA FROM BROOKS AIR FORCE BASE

In order to study the relative merits of the proposed methods, a set of data from the United States Brooks Air Force base was used, and the proposed heuristics were implemented on an IBM 370/158 computer with coding in FORTRAN IV.

The data concerns eighteen planes, fifty two crews, fourteen bases and eight routes. The routes, flight times of the legs on each route, and frequency of occurence of each route is given in Table 2. Since the exact values for the flight time of the legs are unknown, expected values are used. The planning horizon for the problem is 90 days.

For assumption (e) on planes, a maintenance down time of 24.00 hours with probability of .60 of occurence, 60.00

hours for minor repairs with probability of 0.28 of occurence, and 96.00 hours for major repairs with probability of 0.12 of occurence is used.

Results and graphs for this problem are given for the proposed methods.

Since two different methods have been tested using the data from Table 2, a discussion of the results for each one of the problems discussed in Chapter I, is presented separately for each method.

TABLE 2

Data from Brooks Air Force Base

Route 1 frequency = 10

base #	1	2	3	4	3	2	3	4	3
stage		0	1	1	1	0	1	1	1
ft		0.8	1.3	8.2	9.2	1.8	1.3	8.2	9.2

base #	2	3	4	1
stage	0	1	1	1
ft	1.8	1.3	8.2	9.4

Route 2 frequency = 14

base #	1	5	6	7	4	3	5	6	7
stage		0	1	0	1	1	0	1	0
ft		0.4	5.1	5.6	3.8	9.2	0.4	5.1	5.6

base #	4	3	5	6	7	4	1
stage	1	1	0	1	0	1	1
ft	3.8	9.2	0.4	5.1	5.6	3.8	9.2

Route 3 frequency = 5

base #	1	8	9	10	3	8	9	10	3
stage		1	0	1	1	1	0	1	1
ft		3.6	9.8	1.3	8.2	3.6	9.8	1.3	8.2

base #	8	9	10	1
stage	1	0	1	1
ft	3.6	9.8	1.3	8.2

Route 4 frequency = 6

base #	1	13	6	9	10	14	13	6	9
stage		0	1	0	1	1	0	1	0
ft		1.8	6.5	4.1	1.3	9.4	0.8	6.5	4.1

base #	10	14	13	6	9	10	1
stage	1	1	0	1	0	1	1
ft	1.3	9.4	0.8	6.5	4.1	1.3	8.2

Route 5 frequency = 20

base #	1	11	12	9	10	3	11	12	9
stage		0	1	0	1	1	1	1	0

ft		1.1	6.5	6.1	1.3	8.2	6.1	6.5	6.1

base #	10	3	11	12	9	10	3	1
stage	1	1	1	1	0	1	1	1
ft	1.3	8.2	6.1	6.5	6.1	1.3	8.2	5.8

Route 6 frequency = 15

base #	1	11	3	10	3	11	3	10	3
stage		0	1	1	1	0	1	1	1
ft		1.1	5.3	6.9	8.2	6.1	5.3	6.9	8.2

base #	11	3	10	3	1
stage	0	1	1	1	1
ft	6.1	5.3	6.9	8.2	5.8

Route 7 frequency = 25

base #	1	8	3	9	10	3	8	3	9
stage		0	1	0	1	1	0	1	0
ft		4.0	3.1	7.7	1.3	8.2	3.6	3.1	7.7

base #	10	3	8	3	9	10	3	1
stage	1	1	0	1	0	1	1	1
ft	1.3	8.2	3.6	3.1	7.7	1.3	8.2	5.8

Route 8 frequency = 5

base #	1	9	10	3	11	9	10	3	11
stage		0	1	1	1	0	1	1	1
ft		10.9	1.3	8.2	6.1	10.9	1.3	8.2	6.1

base #	9	10	3	1
stage	0	1	1	1
ft	10.9	1.3	8.2	6.1

stage = 1 if crew needs to rest at that base, 0 otherwise.

ft = expected flight time in hours.

All ground times (loading and unloading) are 2.3 hours at each base.
For computational purposes only bases with stage=1 are considered, since only those may need the allocation of crews. The flight and ground time required on any leg where a stage=0 base is visited, is added to the flight time between the previous base with stage=1 and the next base with stage=1.

4.2 TWO-PHASE METHOD

A computer program designed for the Two-Phase Method, was run for the available data. As in the previous section, results for this method were obtained for three different variants of the problem.

4.2.1 Problem 1

In this problem the minimum number of days (makespan) required for the completion of the jobs where all the restrictions over planes and crews are satisfied, is sought. For a given number of planes (18), the number of crews in the system were varied in order to obtain the maximum utilization rate for this fixed number of planes. The minimum number of days required to complete the jobs for 18 planes and 52 crews is 30.520 days with a corresponding (maximum) utilization rate of 19.952 hour/day per plane.

Next, the program was run for different number of crews in the system. Figure 7 presents a plot of the number of crews versus the utilization rate. The results for this are given below.

Number of crews	Maximum Utilization Rate
52	19.952
49	19.336
45	18.679
40	17.398
35	15.762
26	11.850

The average number of iterations required for this problem is 5.6 iterations with an average of 200.00 seconds of total execution time.

As will be discussed later in this chapter, Problem 1 requires a larger number of iterations than Problems 2 and 3 to obtain a good solution. Hence, the execution time required for this problem is also greater. This problem consumes up to three times the execution time required for Problems 2 and 3 with the same number of resources in the system.

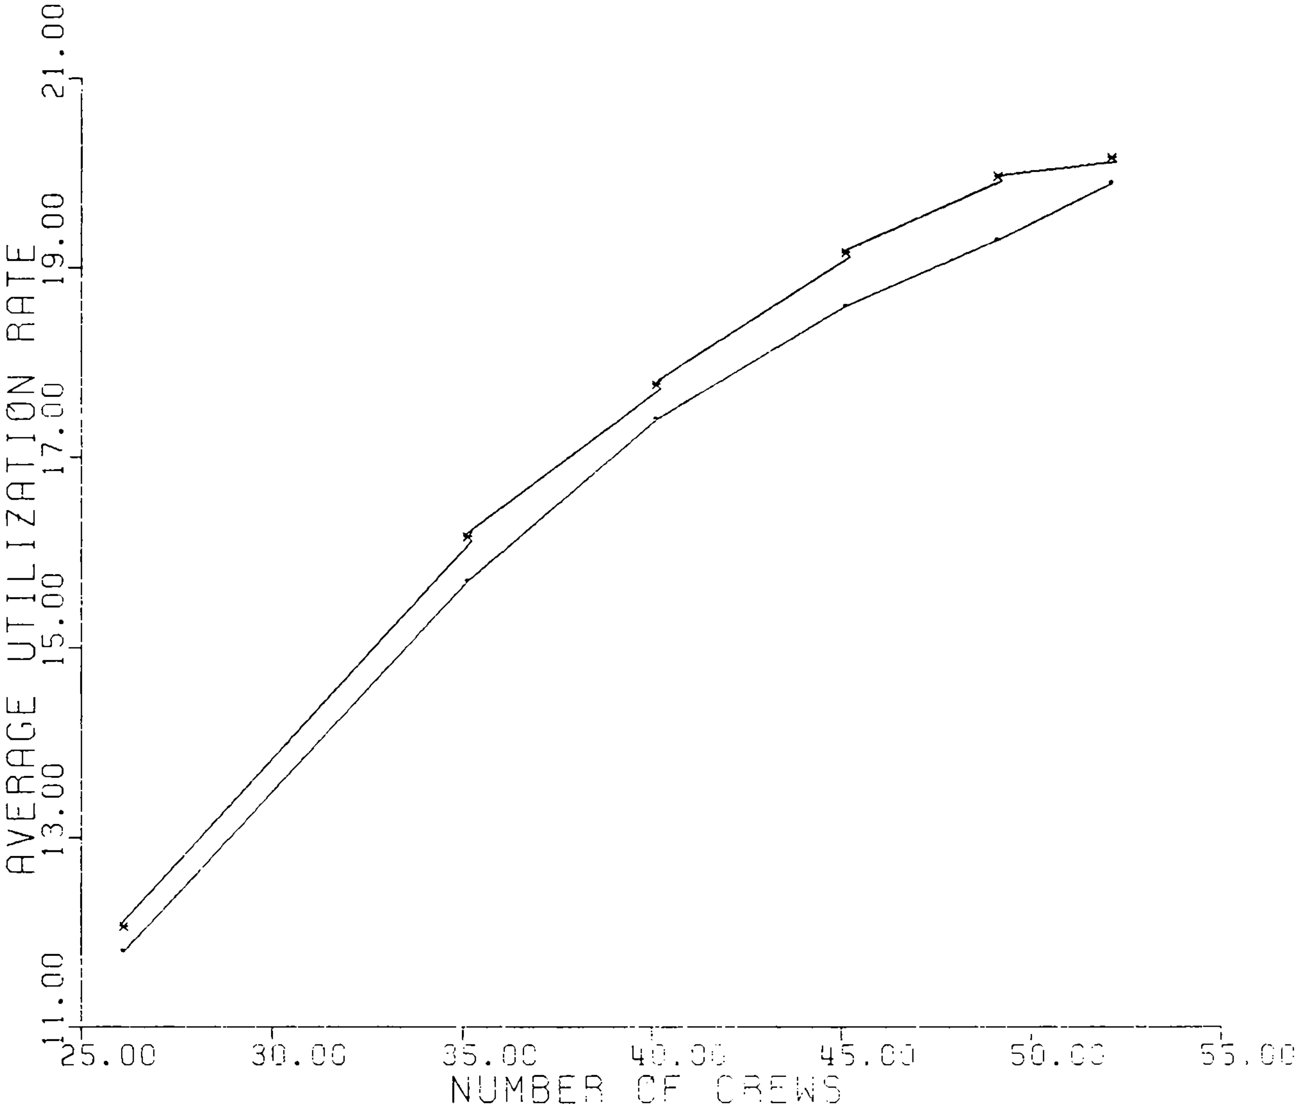

Legend:
. . Two-Phase Method
* * One-Phase Method

Figure 7: Graph for Problem 1 with 18 planes

4.2.2 Problem 2

The objective of this problem is to maximize the utilization rate of the planes when a fixed period of time T_j is required between the commencement of any two consecutive operations of the missions. For these given T_j, j=1,...,n the sequence of missions for each plane is fixed in Phase I. Also, an upper bound for the number of crews required in the system for a non-delay schedule is obtained. A minimum time of 84.875 days, with an average idle time of 0.00 hours, and a maximum average utilization rate 7.174 hour/day per plane, was obtained for this problem with 52 crews and 18 planes available in the system, in 1 iteration and 34.16 seconds of execution time.

For the utilization rate obtained above, the computer program was run for 6 different number of crews and 4 different number of planes in the system. The results are given in Table 3 and Figures 8 through 11. From this information a minimum of 24 crews and 10 planes is required in the system to complete all the jobs within the same number of days that 52 crews and 18 planes need to complete all the jobs, although some idle time is incurred. The actual completion time in this problem is 84.875 days with an average idle time of 43.931 hours. Since fewer planes are available in the system, the average utilization rate increases to a (maximum) value of 12.914 hours/day per plane.

The average number of iterations for a fixed number of crews and different number of planes in the system varies between 1 and 5 iterations. It is of some interest to observe that this average number of iterations becomes larger when the number of crews and planes required in the system is closer to the optimal combination obtained by this algorithm. For the combination of 24 crews and 10 planes in the system, 4 iterations were necessary in 66.68 seconds of execution time. A maximum number of 8 iterations was required for the combination of 24 crews and 18 planes, with the computational effort being 150.01 cpu seconds.

The upper bounds for the number of crews required when different numbers of planes are available in the system are:

Number of planes	Upper bound
18	51
14	51
10	44
8	44

TABLE 3

Average Utilization Rate per Plane

Number of Crews	Number of Planes: 18	14	10	8
52	7.174 (0.00)*	9.224 (0.00)	12.914 (0.00)	15.686 (0.00)
40	7.174 (1.96)	9.224 (4.73)	12.914 (0.000	15.686 (2.95)
32	7.174 (7.86)	9.224 (11.23)	12.914 (6.18)	15.686 (6.40)
30	7.174 (4.57)	9.224 (11.23)	12.914 (9.83)	15.686 (7.08)
24	7.714 (39.51)	9.222 (42.33)	12.914 (43.93)	15.784 (88.57)
16	5.770 (627.30)	7.497 (530.84)	10.395 (498.47)	12.167 (584.33)

*(·) hours of idle time

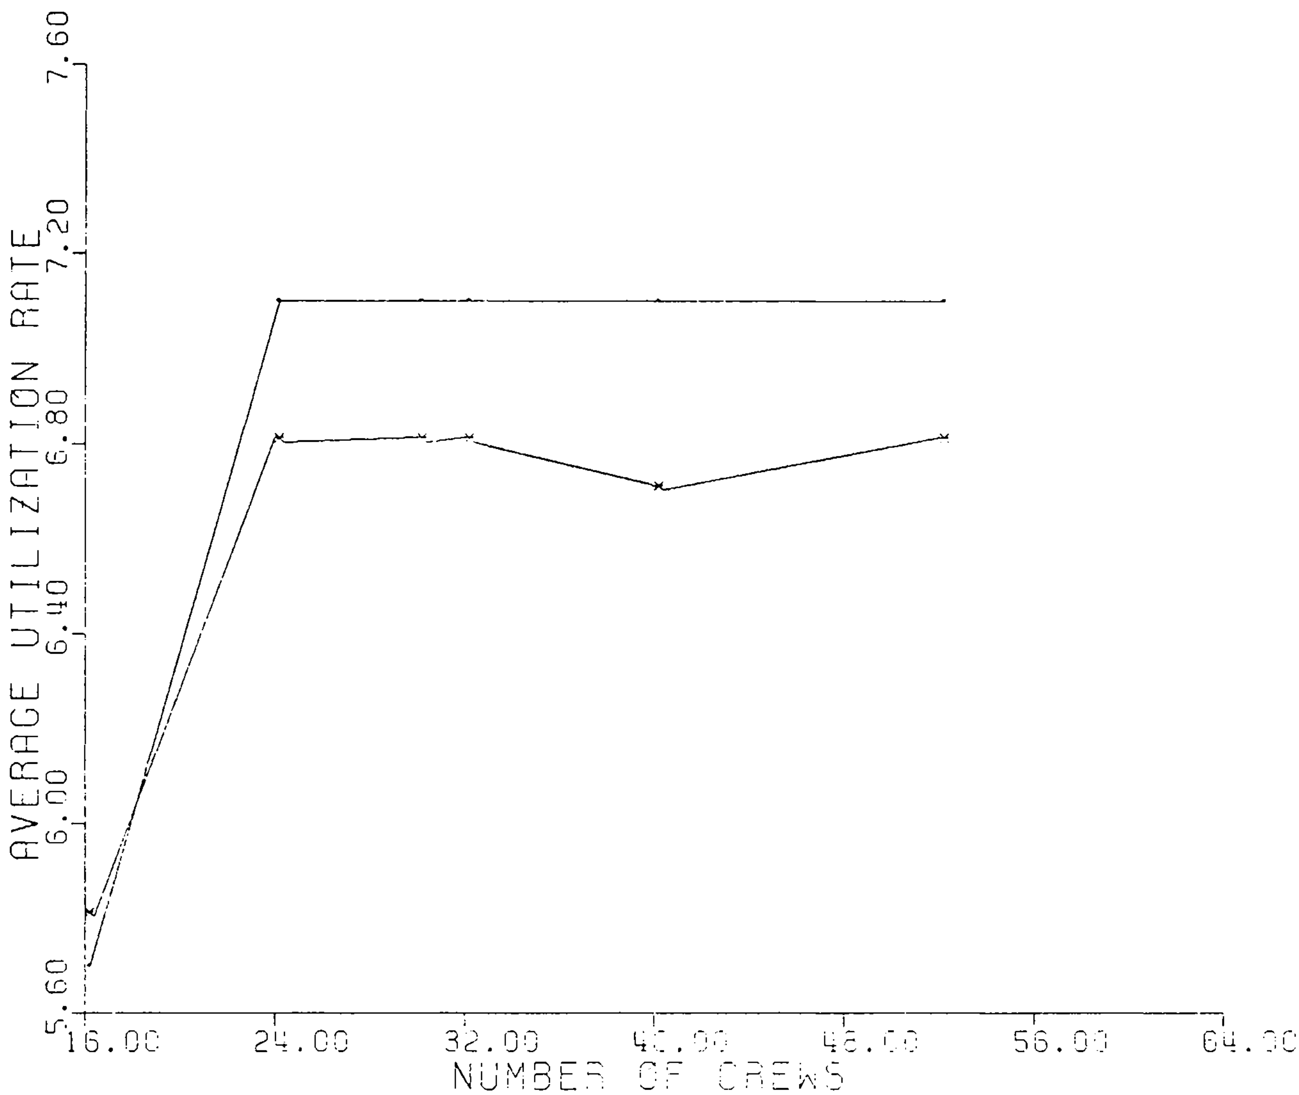

Legend:
. . Two-Phase Method
* * One-Phase Method

Figure 8: Graph for Problem 2 with 18 planes

4.2.3 Problem 3

In this problem the missions are spread so that at most one operation of each mission is being flown at any given time. This problem coresponds to Problem 3 presented in Chapter I.

A minimum idle time of 0.00 hours was obtained for Problem 3 in only one iteration with an average utilization rate of 5.161 hours/day per plane. Since 117.98 was the minimum number of days required to complete the jobs, T_j as defined above cannot complete the missions over the specified planning horizon of 90 days. However, notice that an optimum solution was obtained by the Two-Phase Method within one iteration in 33.86 seconds of execution time. Also an upper bound of 47 crews was obtained for the number of crews required in the system for a non-delay schedule.

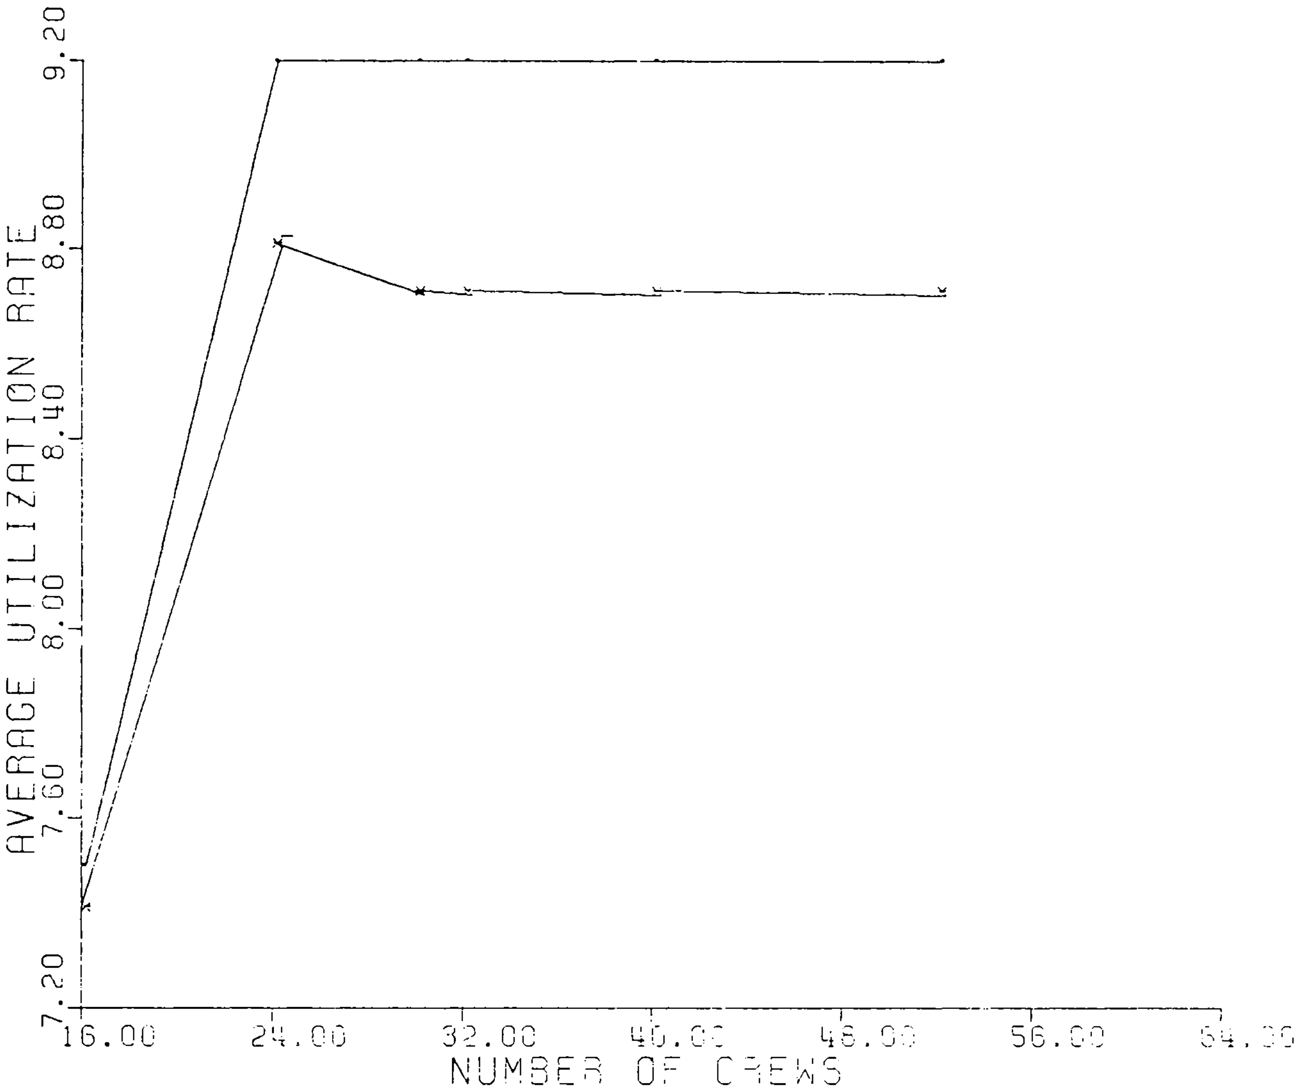

Legend:
. . Two-Phase Method
* * One-Phase Method

Figure 9: Graph for Problem 2 with 14 planes

4.3 ONE-PHASE METHOD

The computer program for this method was run using the data provided by the Brooks Air Force base. The three variants of the problem stated in Chapter I were solved. Each problem is discussed below.

4.3.1 Problem 1

Given the data from Table 1 the maximum utillization rate for the planes in the system is sought. In this problem there exist no restrictions on the completion time of the jobs. This means that the purpose of the problem is to find the minimum number of days necessary to complete all the jobs, when the number of crews, as well as the number of planes in the system, is fixed. For this case a program was run producing a minimum average idle time of 5.04 days in 30.312 days and a maximum average utilization rate of 20.089 hour/day per plane.

The program was run for a different number of crews and a fixed number of planes in the system. Results for this problem are given below. Figure 7 shows the rate of change of the utilization rate of planes for different numbers of

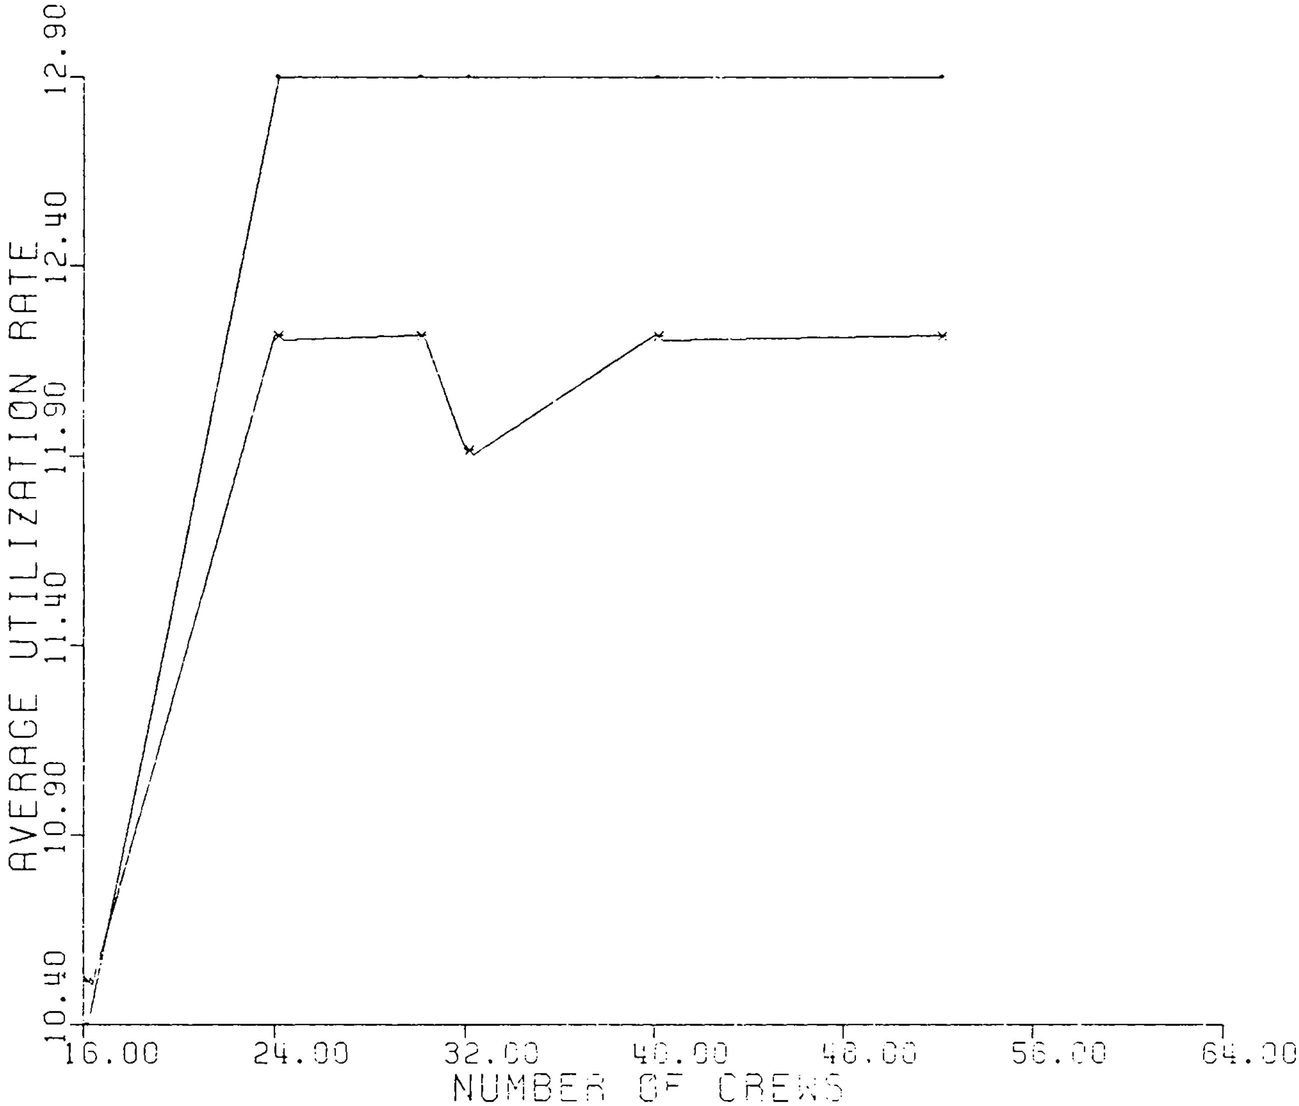

Legend:
. . Two-Phase Method
* * One-Phase Method

Figure 10: Graph for Problem 2 with 10 planes

crews in the system.

Number of crews	Maximum Utilization Rate
52	20.089
49	19.903
45	19.099
40	17.719
35	16.168
26	12.090

An average of 4 iterations were required in order to obtain a good solution for this problem. This number of iterations implies an average of 140.00 cpu seconds of execution time.

4.3.2 Problem 2

For Problem 2 presented in Section 1 the available data was used in order to obtain a maximum utilization rate of planes. For Problem 2 a minimum average of 0.00 hours of idle time was obtained in 89.19 days (a 90 day period is desired) which yields a maximum average utilization rate of 6.827 hour/day per plane. Only one iteration was required in 25.12 seconds of execution time.

For the given completion time obtained in Problem 2, the computer program was run for 6 different number of crews and

4 different number of planes in the system in order to obtain a desirable plane-crew configuration for the problem. Table 4 shows the maximum average utilization rate per plane obtained for each case. Figure 7 presents a plot of the utilization rate versus the number of crews in the system for a fixed number of planes.

From Table 4 and Figure 8 through 11 a non monontone increase in the average utilization rate of planes is observed. These fluctuations in the average utilization rate is a consequence of the initial distribution of crews among bases. When the system has more crews than required in order to obtain a fixed utilization rate, no idle time for planes exists, although a better allocation of crews is possible, producing a better utilization rate due to different sequences of missions that can be flown by each plane. Since the re-distribution of crews depends on the idle time at the previous iteration of the algorithm, the process is terminated and no improvement is obtained. However, if the number of crews in the system is smaller than or closer to the required number of crews in the system for a fixed utilization rate, a larger number of different distributions is attempted, thereby producing better solutions.

The number of iterations required to obtain a good solution for this problem ranges between 1 and 8 with an average

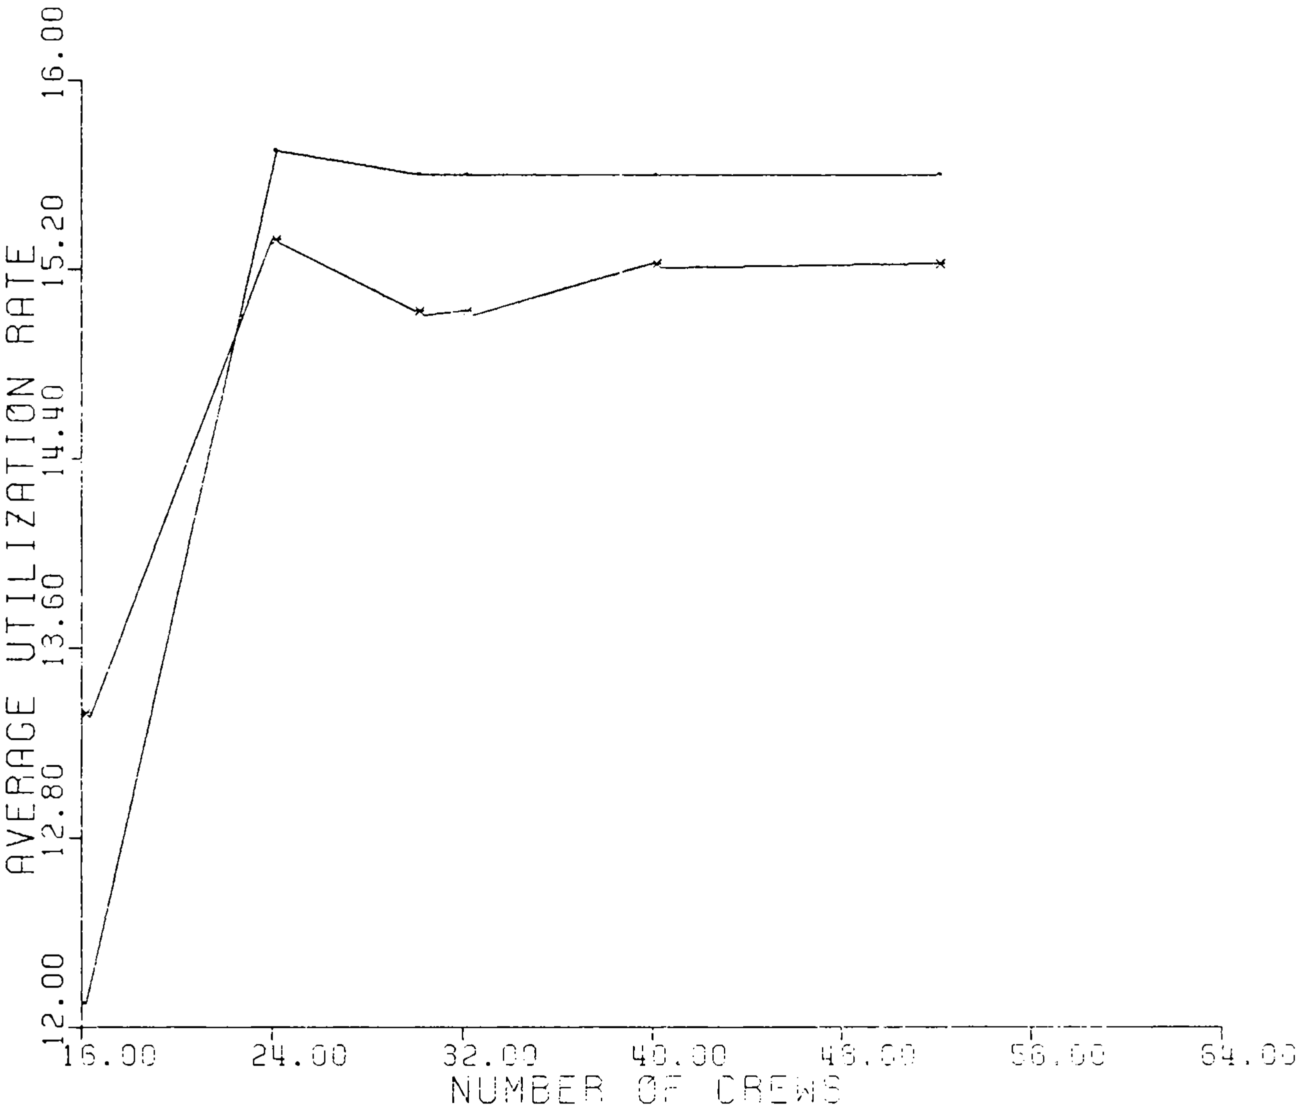

Legend :
. . Two-Phase Method
* * One-Phase Method

Figure 11: Graph for Problem 2 with 8 planes

of 2 iterations. This program was run for a fixed number of crews and different number of planes. It was noticed that the number of iterations was increased when the number of crews and planes in the system were closer to the best combination obtained for this algorithm. Hence the execution time also increased for these combinations. A maximum of 8 iterations were required for a combination of 18 planes and 24 crews consuming a maximum of 180. seconds of execution time.

From the information given in Table 4 and Figure 8 through 11, a minimum makespan of 88.048 days for a maximum average utilization rate of 8.892 hours/day per plane is obtained with 24 crews and 14 planes in the system The solution for this combination was obtained in 1 iteration in 37.08 seconds of execution time.

TABLE 4

Average Utilization Rate per Plane

Number of Crews	Number of Planes: 18	14	10	8
52	6.827 (0.00)*	8.711 (0.00)	12.244 (0.00)	15.253 (0.00)
40	6.793 (2.13)	8.711 (0.00)	12.244 (0.00)	15.289 (2.34)
32	6.813 (15.01)	8.727 (14.85)	11.994 (8.63)	15.044 (14.08)
30	6.813 (29.29)	8.710 (22.78)	12.218 (14.85)	15.035 (15.89)
24	6.686 (131.57)	8.892 (150.84)	12.258 (28.53)	15.305 (23.19)
16	5.894 (732.49)	7.476 (727.43)	10.567 (655.64)	13.354 (498.93)

*(·) hours of idle time

4.3.3 Problem 3

The available data was also implemented for Problem 3 with the One-Phase Method. Although in Problem 3, a minimum idle time of 0.00 hours was obtained, 117.98 days were required to complete the jobs, which yields an average utilization rate of 5.161 hour/day per plane. Therefore Problem 3 does not satisfy the 90 day requirement of our problem. However, notice that for Problem 3 a minimum idle time was obtained in only one iteration and 26.83 seconds of execution time.

Since this problem could not be solved within the desired period of time given the maximum number of crews and planes available in the system, the program was not executed for any other combination of crews and planes.

4.4 COMPARISON OF THE PROPOSED METHODS

The Two-Phase Method and the One-Phase Method, proposed in this thesis for the solution of the Air Force Crew Allocation and Scheduling Problem were primarily designed to find the minimum makespan T (overall completion time for all the jobs) for the problem. However, these methodologies can be also employed for different variants of the problem to obtain other valuable information for our problem. Although both methods can be used to obtain essentially the same information, several advantages exist in one method over the other depending on the variant of the problem to be solved. These advantages are discussed below.

The Two-Phase Method provides an upper bound for the number of crews required in the system. This upper bound can reduce the number of times the program needs to be run in order to obtain the minimum number of crews and planes required in the system to attain a given utilization rate.

The initial distribution of crews obtained in the Two-Phase Method gives better initial solutions than the distribution of crews obtained by the One-Phase Method. As was seen before, the initial distribution obtained in the former method is based upon the distribution of crews among bases for a non-delay schedule performed in Phase I of this method. Since this initial distribution is a good starting allocation of crews among bases, and the minimization of idle time depends directly on the allocation of crews in the system, the initial solution obtained by this method is better than the one obtained by the One-Phase Method.

A comparison of these two methods based on the three variants of the problem presented in Chapter I is given below. A slight difference is found in the number of iterations required in these two methods before a solution for Problem 2 is obtained. This number of iterations is directly proportional to the number of crews and planes available in the system. This is obvious, since the number of possible exchanges of missions between planes increases with the number of planes. In the same manner, an increase in the number of crews implies more possible different arrangements of the crews among the bases. In the Two-Phase Method an average of 3 iterations were required to find a solution to the problem while in the One-Phase Method an average of 2 iterations were performed. The difference in the number of iterations is more noticeable, when the available crews in the system becomes closer to the minimum number required for a given utilization rate.

The average number of iterations required in the Two-Phase Method and the One-Phase Method, versus the number of crews in the system for Problem 2 is provided in Figure 12. In reading this figure, it is important to note that when the number of crews available in the system becomes closer to the number of crews necessary in the system for a fixed utilization rate, the number of iterations before a solution is obtained, increases.

Although the number of iterations required by the One-Phase Method is smaller than the Two-Phase Method for the three variants of the problem, the execution time of these methods is comparable.

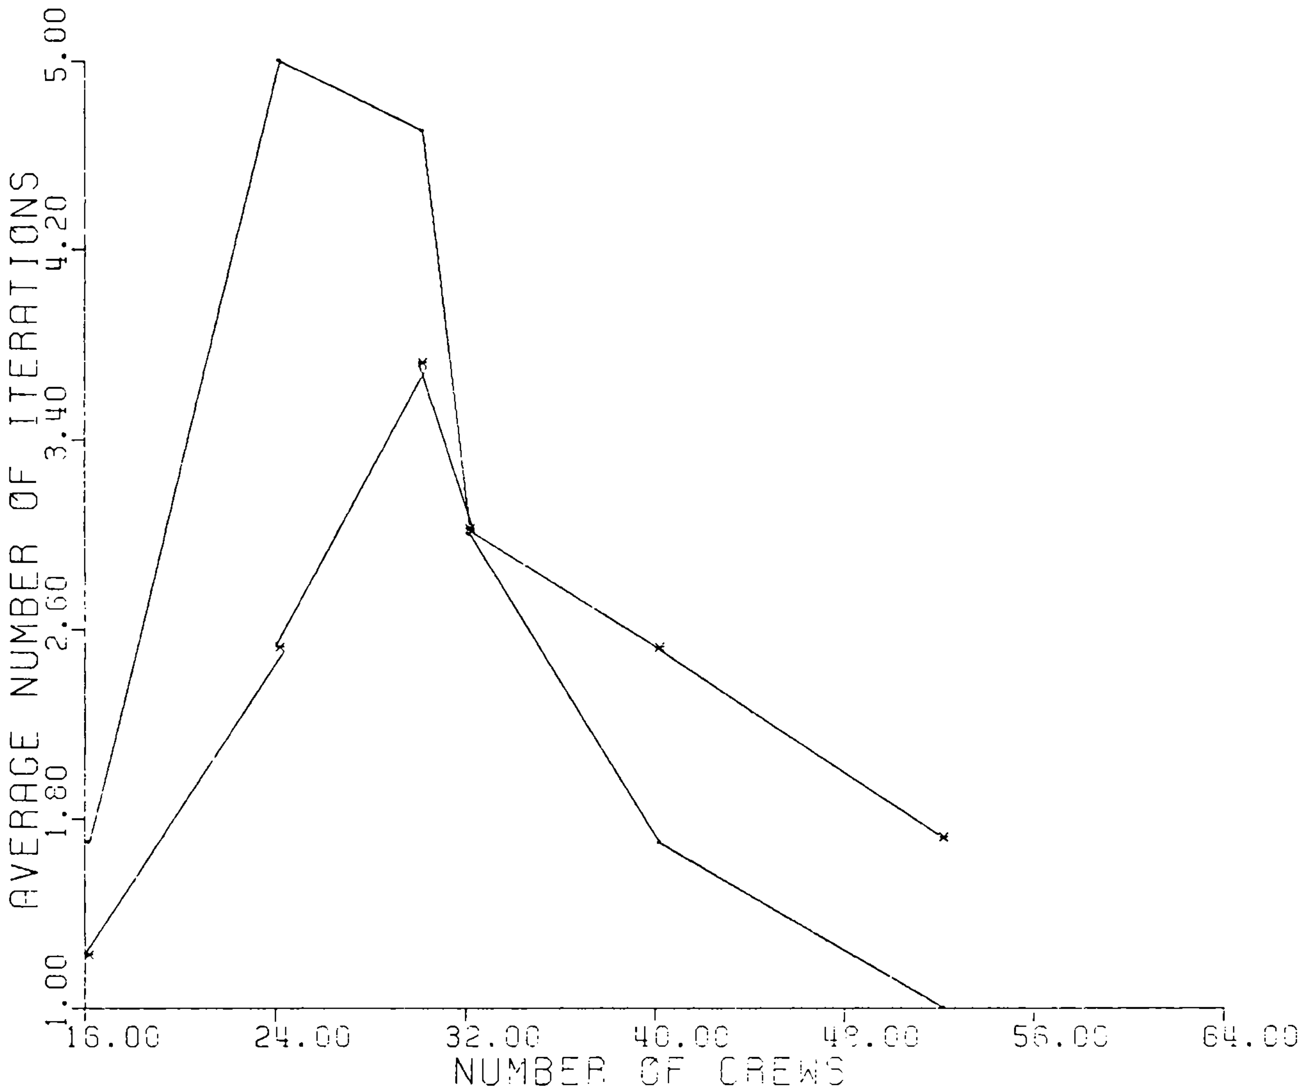

Legend :
. . Two-Phase Method
* * One-Phase Method

Figure 12: Average Number of Iterations for Problem 2

The similarity in the execution time between the proposed methods is a consequence of the execution time required for each iteration in these methods. Since the Two-Phase Method fixes the sequence of missions for each plane in Phase I, less effort is required for the simulation process conducted in this method. The simulation process becomes shorter than the simulation process performed in the One-Phase Method. This remark is explained by the fact that the tie-breaking rules required for the assignment of missions to planes are considered only once in the former method (Phase I is conducted only once in the execution of the program) while in the latter method, each iteration requires an assignment of missions to planes, thereby significantly increasing the simulation effort.

However, in Problem 1 a significant difference exists in the execution time required for each algorithm. An average of 140.00 seconds is required to obtain a good solution in the One-Phase Method while 200.00 seconds of execution time is needed in the Two-Phase Method. This fact is explained by the number of iterations required. The Two-Phase Method requires twice the number of iterations required by the One-Phase Method in order to obtain a good solution. The increase in the number of iterations is a consequence of the sequence of missions which is fixed in Phase I of the method. Therefore no flexibility exists in order to improve the sequence of missions, hence a better solution can be obtained only with a better distribution of crews. In the One-Phase Method more flexibility exists producing better quality solutions. That is, while in the former method, a different sequence of missions can be flown on each iteration, in the latter method this sequence is fixed in the first phase of the method and remains the same during all the iterations.

In general the average utilization rate of planes obtained from the Two-Phase Method is better than the solution obtained from the One-Phase Method in Problem 2. The difference obtained in the solutions of these methods generally remains within an average utilization rate of 0.60 hour/day per plane as was observed in the figures for Problem 2.

However, the solution obtained for Problem 1 by the One-Phase Method is superior to the one obtained by the Two-Phase Method. The difference in the average utilization rate of planes obtained by these two methods is approximately 0.30 hours/day per plane.

Hence, from an overall viewpoint of compromising between computational effort and solution quality, the Two-Phase Method is recommended over the One-Phase Method for Problem 2. Furthermore, when there is a large number of resources in the system, considerably lesser computational effort is required, by the former method. However, for smaller sized problems, good solutions (combinations of planes and crews) may be obtained by the One-Phase Method with the required computational effort being comparable with that for the Two-Phase Method.

However, from the figures, it is observed that the One-Phase Method produces better quality solutions with lesser computational effort than the Two-Phase Method for Problem 1.

Actually, from the observations made in this thesis, it appears that a combination of the methods may be an overall effective alternative. Such a composite method may use Phase I and the first iteration of Phase II of the Two-Phase Method, and then after redistributing the crews among the bases, it may switch over to the One-Phase Method. In this manner, the tendency of the Two-Phase Method to produce good quality initial solutions may be combined with the flexibility of the One-Phase Method in later iterations.

4.5 SUMMARY

This thesis proposes two solution methodologies for the Brooks Air Force Base Crew Allocation and Scheduling Problem. A mathematical formulation of the relaxation of the problem is provided to show the difficulty of the problem. Since a mathematical approach is not computationally feasible for this problem, heuristic solution procedures have been developed.

A review of the available literature concerned with the scheduling aspects of the problem is presented. A comparison of the problem under consideration with the problems in the literature is provided in order to show that no solution procedure exists in the literature that can be directly applied to this problem. Also the similarities of the proposed methodologies with the solution procedures developed for similar problems in the literature have been discussed.

The two proposed methods have been discussed in detail and heuristic scheduling rules have been provided which may be generally incorporated into other heuristic approaches as well. Furthermore, computational experience has been obtained with these methodologies using data from the Brooks Air Force base. Based on the results obtained, the Two-Phase Method is recommended over the One-Phase Method for Problem 2. However for Problem 1 the One-Phase Method is recommended over the Two-Phase Method, since better quality solutions are obtained. For Problem 3 neither method appears superior; therefore no preference exists, similar solutions are obtained with comparable computational effort.

A composite method may be an overall alternative solution. This method combines the good quality initial solution of the Two-Phase Method with the flexibility of the One-Phase Method.

REFERENCES

1. Agard, J., Arabeyre, J.P., Vautier, J., "Generation Automatique de Rotations D'equipages", Revue d'Informatique et de Recherche Operationelle, No.6, 1967.

2. Arabeyre, J.P., Fearnley, J., Steiger, F.C., Teather, W., "The Airline Crew Scheduling Problem: A Survey", Transportation Science. Vol.3, No.2, May 1969, pp. 140-163.

3. Armstrong, Ronald D., Cook, Wade D., "Optimal Linking of Tasks and Schedule Development for Military Air Transportation", INFOR, Vol.14, No.1, February 1976, pp. 13-31.

4. Baker, E. K., Bodin, L. D., Finnegan, W. F., Ponder, R. J., "An Efficient Heuristic Solution to a Airline Crew Scheduling Problem", AIIE Transactions, June 1979, Vol.11, No.2, pp. 79-85.

5. Baker, Kenneth R., Introduction to Sequencing and Scheduling, John Wiley and Sons, Inc., 1974.

6. Bellman, R. and Cooke, K., The Konigsberg Bridges Problem Generalized, Journal of Mathematical Analysis and Applications, 1969, Vol 25, p.1.

7. Bellmore, M., and Nemhauser, G., "The Traveling Salesman Problem: A survey", Operations Research, 1974, Vol.16, pp. 538-558.

8. Bellmore, M. and Hong, S., "Transportation of Multi-Salesman to the Standard Traveling Salesman Problem" Journal of the Association of Computing Machinery, 1974, Vol. 21, pp. 500-504.

9. Beltrami, E. and Bodin, L., "Networks and Vehicle Routing for Municipal Waste Collection", Networks, 1974, Vol. 4, pp. 65-94.

10. Bernett, S. and Gazis, D., "School Bus Routing by Computer", Transportation Research, Dec. 1972, Vol. 6, pp. 317-325.

11. Bernett, B. and Potts, R., "Routing Roster for a Transit System", Transportation Science, Feb. 1968, Vol. 4, No. 1, pp. 14-34.

12. Bodin, L. and Berman, L., "Routing and Scheduling of School Buses by Computer", Transportation Science, 1979, Vol.13, No. 2, pp. 113-129.

13. Bodin, L. and Golden, B., "Classification in Vehicle Routing and Scheduling" ,University of Maryland, 1980, Working Paper MS/S 80-006.

14. Bodin, L., Golden, B., Assad, A., Ball, M., "The State of Art in the Routing and Scheduling of Vehicles and Crews", 1981, University of Maryland, Maryland, Working Paper MS/S 81-035.

15. Bodin, L., Kursh, S., "A Detailed Description of a Sweeper Routing and Scheduling System", Computer and Operations Research, 1979, Vol.6, pp. 181-198.

16. Bodin, L., Kursh, S., "A Computer-Assisted System for the Routing and Scheduling of Street Sweepers", Operations Research, 1978, Vol. 26, No.4, pp. 525-537.

17. Bodin, L. and Rosenfield, D., "Estimation of the Operating Cost of Mass Transit Systems", Report No. nahcups-umta-1-76, College of Urban and Policy Science, State University of New York, Stony Brook, N.Y..

18. Bodin, L. and Sexton, T., "The Subscriber Dial-a-Ride Problem", Report No. umcp-umta-1-79, 1979. College of Business and Management, University of Maryland, College Park, Maryland.

19. Christofides, N., "The Optimum Traversal of a Graph", Omega, 1973, Vol. 9, No. 6, pp. 719-732.

20. Christofides, N., "The Vehicle Routing Problem", Rev. Frans. Res. Oper., 1976, Vol. 10, pp. 55-70.

21. Edmonds, J., "Paths, Trees and Flowers", The Canadian Journal of Mathematics, 1965, Vol. 17, pp. 449-467.

22. Edmonds, J. and Johnson, E., "Matching, Euler Tours, and the Chinese Postman", Mathematical Programming, 1963, Vol. 5, pp. 88-124.

23. Fearnley, J., "Crew Scheduling Development in Air Canada", AGIFORS, Symposium, 1966, Air Canada.

24. Gillet, B. and Johnson, T., "Multi-Terminal Vehicle Dispatch Algorithms", Omega, 1974, Vol. 4, pp. 711-718.

25. Glover, F., "Finding an Optimal Edge-Covering Tours of a Connected Graph", ORC 67-13, Operations Research Center, University of California, Berkeley, California, 1967.

26. Golden, B., Magnanti, T., Nguyen, H., "Implementing Vehicle Routing Algorithms", Networks, 1977, Vol. 7, pp. 113-148.

27. Hinson, J., Mulkerkar, S., "Improvement to the Clarke and Wright Algorithm as Applied to an Airline Scheduling Problem", Technical Report, Federal Express Corp., 1975.

28. House, R.W., Nelson, L.D., Rado, T., "Computer Studies of a Certain Class of Linear Integer Problems", Recent Advances in Optimation Techniques, 1966, A. Lavin and T. Vogl (eds.) Wiley, N.Y.

29. Hu, T. C., "Parallel Sequencing and Assembly Line Problems", Operations Research, November 1961, Vol. 9, No. 6, pp. 841-848.

30. Karp, R., "Reducibility Among Combinatorial Problems", Complexibility of Computer Computations, 1972, Plenum Press, New York, pp. 85-104.

31. Karp, R., "Probabilistic Analysis of Partitioning Algorithms for the Traveling Salesman Problem in the Plane", Mathematics of Operations Research, 1977, Vol. 2, pp. 209-244.

32. Landis, R. "A Perspective on Automated Bus Operator Scheduling: Five Years Experience in Portland, Oregon", Computer Scheduling of Public Trasportation Urban Passeger Vehicle on Crew Scheduling, A. Wren, ed., North Holland Publishing Co., 1980, pp. 67-70.

33. Lassard, R., Rousseau, J., DuPuis D., "Hasters I: A Mathematical Programming Approach to the Bus Driver Scheduling Problem", Computer Scheduling of Public Trasportation Urban Passenger Vehicle Crew Scheduling, A. Wren, ed., North Holland Publishing Co. 1981, pp. 256-268.

34. Martin, G.T., "An Accelerated Euclidean Algorithm for Integer Linear Programming, Recent Advance in Mathematical Programming, R. L. Graves and P. Wolfe (eds.), 1963, McGraw-Hills, pp. 311-318.

35. Newton, R. and Thomas, W., "Bus Routing in a Multi-School System", Computers and Operations Research, 1974, Vol. 1, pp. 213-222.

36. Nicoletti, Bernardo, "Automatic Crew Routing", Transportation Science, February 1975, Vol.9, No. 1, pp. 33-44.

37. Parker, G. and Smith, B. "Two Approaches to Computer Crew Scheduling" Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling Wren, A., ed., North Holland Publishing Co.,1981, pp. 193-222.

38. Orloff, C., "Routing a Fleet of Vehicles to/ from a Central Facility", Networks, 1974, Vol. 4, pp. 147-162.

39. Rosenkrantz, P., Stearns, R., Lewis, P., "An Analysis of Several Heuristics for the Traveling Salesman Problem", SIAM Journal Computing, 1977, Vol. 16, pp. 563-581.

40. Rubin, Jerold, "A Technique for the Solution of Massive Set Covering Problems, with Application to Airline Crew Scheduling", Trasportation Science, February 1973, Vol. 7, No. 1, pp. 34-48.

41. Salkin, Harvey M., Integer Programming, Addison-Wesley Co., 1975.

42. Schmidt, J., Knight, R., "The Status of Computer-Aided Scheduling in North America", Computer Scheduling of Public Transport Urban Passenger Vehicle Crew Scheduling, A. Wren, ed. North Holland Publishing Co. 1981, pp. 17-22.

43. Segal, M., "The Operator-Scheduling Problem: A Network-Flow Approach", Operations Research, 1974, Vol. 22, pp. 808-823.

44. Sexton, T. and Bodin, L., "The Single Vehicle Many to Many Routing and Scheduling Problem with Desired Delivery Times", Work Paper No. 80-014, University of Maryland, 1980.

45. Steiger, F. and Neiderer, M., "Scheduling Air Crews by Integer Programming", Presented at the IFIP Congress 1968, Edinburgh.

46. Soyster, A.L., "Airlift Crew-Ratio Scheduling Task 12", Report submitted to the Texas A & M Research Foundation, November 1979, Project No. 3919, Task 12.

47. Tillman, F. and Cain, T., "An Upper Bound Algorithm for the Single and Multiple Terminal Delivery System", Management Science, 1972, Vol. 18, pp. 664-682.

48. Wong, J., "Computer Aided Scheduling For Wheel-Trans Service", Project No. 3323, Transit System Office, Systems Research and Development Branch, Toronto: Ontario, Ministry of Transportation and Communications, March 1978.

49. Wren, A., ed. Computer Scheduling of Public Transportation Urban Passenger Vehicle and Crew Scheduling, North Holland Publishing Co. Amsterdam, 1981.

The vita has been removed from
the scanned document

AIR FORCE CREW ALLOCATION AND SCHEDULING PROBLEM

by

Minerva Rios Perez

ABSTRACT

This thesis addresses an airline crew allocation and scheduling problem faced by certain divisions of the United States Air Force. Three variants of the problem under consideration were posed by the Brooks U.S. Air Force Base. This thesis reports on experience with two heuristic methods developed, each applicable to the different variants of the problem. Although the problem described herein is peculiar to this situation, the heuristic scheduling and dispatching rules developed have been found to be very effective, and are generally applicable in other related contexts of routing, and crew and vehicle scheduling problems as well. The two algorithms developed have been applied to a coded set of real world data. The results indicate that each one of the two methods is preferable over the other for one of the two variants of the problem, and they are equally effective for the third variant.

The observations made in this study suggest an overall effective composite technique for this class of problems.