

# Optimal Scheduling of Fighter Aircraft Maintenance

by

Philip Y Cho


B.S. Economics and Operations Research  
United States Air Force Academy, 2009

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL  
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

**ARCHIVES**

MASTER OF SCIENCE IN OPERATIONS RESEARCH  
AT THE  
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2011



©2011 Massachusetts Institute of Technology 2011. All rights reserved.

Signature of Author: \_\_\_\_\_  
Sloan School of Management  
May 20, 2011

Certified by: \_\_\_\_\_  
Thomas Magnanti  
Institute Professor, Office of the Institute Professors  
Thesis Supervisor

Certified by: \_\_\_\_\_  
Retsef Levi  
J. Spencer Standish Professor of Management  
Associate Professor of Operations Management, Sloan School of Management  
Thesis Supervisor

Accepted by: \_\_\_\_\_  
Patrick Jaillet  
Dugald C. Jackson Professor, Department of Electrical Engineering and Computer Science  
Co-director, Operations Research Center



# Optimal Scheduling of Fighter Aircraft Maintenance

by

Philip Y Cho

SUBMITTED TO THE OPERATIONS RESEARCH CENTER IN PARTIAL  
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN OPERATIONS RESEARCH

## **Abstract**

The effective scheduling of fighter aircraft maintenance in the Air Force is crucial to overall mission accomplishment. An effective maintenance scheduling policy maximizes the use of maintenance resources and aircraft availability. Currently, maintenance scheduling is a time consuming process that is carried out by airmen whose sole responsibility is to manually generate a maintenance schedule that balances maintenance requirements and flying requirements. In this thesis, we seek to represent the maintenance scheduling process using a mathematical model that ultimately generates an optimal maintenance schedule.

First, we address the scheduling of phase maintenance, the most significant preventative maintenance action, for fighter aircraft. We use a mixed integer program (MIP) to model the phase maintenance scheduling process. The MIP generates a daily maintenance and flying schedule that ensures that the maintenance workload is evenly distributed across the planning horizon. We find that the computational performance of the MIP formulation is less than desirable for large instances of real-world data. Motivated by the need for improved computational performance, we develop an alternative formulation that disaggregates the original MIP into two subproblems that are solved sequentially. The two-stage formulation of the phase maintenance scheduling problem has significantly better computational performance while generating a feasible daily maintenance and flying schedule.

We then address the maintenance scheduling process that is unique to aircraft with low-observable (LO) capabilities. The LO capabilities of an aircraft degrade over time according to a stochastic process and require continuous maintenance attention. We show that the characteristics of the LO maintenance process allow it to be modeled as a variant of the multi-armed bandit (MAB) problem. We then present a variant of the heuristic proposed by Whittle that has been shown to provide near optimal solutions for MAB problems. Applying Whittle's heuristic to the LO maintenance scheduling problem, we generate a simple index policy that can be used to schedule aircraft for LO maintenance. We then compare the index policy to alternate policies and show by simulation that the index policy leads to relatively better fully mission capable (FMC) rates, a common measure of overall fleet health.

**DISCLAIMER CLAUSE:** The views expressed in this article are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government

Thesis Supervisor: Retsef Levi

Title: J. Spencer Standish Professor of Management



ACKNOWLEDGMENTS. This thesis comes as the result of a great deal of support and encouragement. I would first like to thank Professors Retsef Levi and Tom Magnanti for their continual guidance and support. They never ceased to challenge me and encouraged me at every step in my research. I would also like to particularly thank Maj. Eric Zarybnisky for his invaluable assistance in completing this thesis. I sincerely appreciate the hours of patient explanation and discussion of topics regarding my research, as well as the countless hours dedicated towards reviewing my work. I would also like to thank Professor Vivek Farias for his invaluable input and suggestions.

In addition, I would like to thank the many men and women of the United States Air Force who volunteered their time and knowledge in support of my research. Thank you for hosting my visits and answering my countless questions regarding the aircraft maintenance process. My research would have truly been impossible without your support.

Finally, I would like to thank my family and friends. To all my Boston friends, thanks for a great two years. I have enjoyed our time together and sincerely appreciate all of your friendships. You will be missed. And to Mom and Dad, thanks for all your encouragement and support. You are truly appreciated. Thank you all.

THIS PAGE INTENTIONALLY LEFT BLANK

## Contents

List of Figures	9
List of Tables	10
Chapter 1. Introduction	11
1.1. Literature Review	12
1.2. Organization and Outline	15
Chapter 2. Air Force Fighter Aircraft Maintenance and Scheduling	17
2.1. Air Force Flying Hours Program (FHP)	17
2.2. Operational and Maintenance Requirements	18
2.3. Current Aircraft Maintenance Scheduling Process	21
2.4. Summary	26
Chapter 3. MIP Formulation for Phase Maintenance Scheduling	29
3.1. Modeling Framework for the Air Force maintenance scheduling process	29
3.2. MIP Formulation of Aircraft Phase Maintenance Scheduling	30
3.3. Implementation and Practical Considerations	40
3.4. Special Cases	44
3.5. Summary	58
Chapter 4. Two-Stage MIP for Phase Maintenance Scheduling	59
4.1. Framework for Disaggregation of Phase Maintenance Scheduling	59
4.2. Formulations of Subproblems	60
4.3. Discussion of the Disaggregation Models	67
4.4. Implementation of the Two-Stage Model	68
4.5. Summary	71
Chapter 5. Low-Observable Maintenance Scheduling	73
5.1. Low-Observable Maintenance Process	73
5.2. Modeling Framework for the LO Scheduling Problem	80
5.3. Multi-Armed Bandit Problems	82
5.4. Application of Whittle's Index to the SAS Redux Scheduling Problem	84
5.5. Implementation of the DP	90

	CONTENTS	8
5.6.	Simulations	92
5.7.	Expanding the DP Formulation	95
5.8.	Summary	95
Chapter 6.	Conclusion and Recommendations for Future Works	97
6.1.	Future Work	97
Appendix		99
PM-MIP	Formulation	99
Network	Formulation for Special Case PM-MIP	102
M-Sub	Formulation	103
F-Sub	Formulation	105
Bibliography		106



## List of Figures

2.3. Sample monthly flying schedule taken from a long-term fiscal year calendar	23
2.3. Sample checkerboard for a one week period	25
2.3. Ideal Phase Flow	26
2.4. The steps required to generate a detailed maintenance schedule	27
3.1. Functional View of the Fighter Aircraft Phase Maintenance Scheduling Process	31
3.2. Comparison of solve times for setting $\bar{M} = 0$ and $\bar{M} = \left[ Z_{LP}^{\bar{M}=0*} \right]$ , represented by the left and right columns, respectively, for 60 random instances.	40
3.3. Optimal solution shown as a plot where each line represents an aircraft's phase hours over the planning horizon. The jumps in phase hours correspond to an aircraft entering maintenance.	42
3.4. Conceptual framework for network reformulation	53
3.4. Example Network Formulation Instance	57
4.4. Phase hours plot from the optimal solution to the M-Sub	70
4.4. Optimal flying schedule from one instance of the F-Sub	71
5.1. Probability distribution of daily SAS increases	78
5.5. Index values associated with each possible SAS state	91
5.5. Decisions associated with possible SAS state	92
5.6. FMC rates for 5 trials. Each trial had a planning horizon of 1000 periods.	93
5.6. FMC rates for three different methods of selecting which aircraft will fly on a given day.	94

## List of Tables

1	Example Instance Data	41
1	Data from PM-MIP implementation adapted for the two subproblems	69
2	Results for each instance of the F-Sub; one for each period of the M-Sub	72
1	SAS Increase Probabilities	77
2	SAS Increase Correlation Coefficients	79
3	SAS redux downward transition percentages	88

## CHAPTER 1

### Introduction

The ability of the United States Air Force to effectively maintain its fleet of aircraft is crucial to mission accomplishment. A key component of maintenance operations is maintenance scheduling. Without the proper scheduling of preventative maintenance, the Air Force cannot support the flying operations necessary to attain and maintain combat readiness. Therefore, it is necessary to have an effective maintenance scheduling process that maximizes the use of maintenance resources while also ensuring aircraft availability.

The importance of proper maintenance scheduling cannot be overlooked in today's Air Force. Over the past decade, the Air Force has been reducing its manpower levels to cut costs and shift funding towards new technologies. In addition, the transition to newer, more advanced aircraft has led to a reduced fleet size due to the higher unit costs associated with these aircraft [15]. Yet the reduction in personnel and fleet size have come in the face of increased military operations throughout the world. Therefore, given the reductions in manpower and fleet size and increased operations tempo, it has become more important than ever to maximize the use of available resources, both manpower and aircraft.

Historically, aircraft maintenance has been known to be one of the most demanding career fields within the Air Force. This is due to perpetual undermanning in conjunction with the direct impact that aircraft maintenance has on flying operations. Each fiscal year, the Air Force determines the manpower levels required to support the forecasted flying operations for all career fields, including aircraft maintenance. Past studies have shown that the methodology used to determine manpower levels underestimates the necessary maintenance manpower[5]. Furthermore, in the majority of cases, the actual number of assigned personnel is less than the manpower level that had been determined to be necessary. Improving the maintenance scheduling process can help sustain a high rate of aircraft availability while also easing the maintenance workload for maintenance personnel as much as possible.

Within the Air Force, maintenance scheduling is done by airmen whose entire careers are devoted to learning and mastering the art of scheduling. These airmen learn to balance the demands of preventative maintenance and the need to meet the flying requirements. Although these airmen have proven to be more than capable, the current maintenance scheduling process is labor intensive and time consuming. Schedulers manually generate maintenance schedules that specify which aircraft will fly on a given day and which will undergo maintenance.

In this thesis, we seek to represent the maintenance scheduling process using mathematical models. These models are intended for two primary purpose. First, the models are intended to provide schedulers with a tool to easily generate maintenance schedules in a non labor-intensive manner. Second, the models are intended to be used as an analysis tool for decision makers to measure the tradeoff between various alternatives (i.e., increased maintenance capacity versus decreased flying operations). Ultimately, we strongly believe that the use of mathematical models in maintenance scheduling can result in significant improvements in Air Force operations.

The first maintenance scheduling problem we consider is phase maintenance which is a preventative maintenance check performed on fighter aircraft. Phase maintenance is a time and labor intensive process that removes an aircraft from the fleet until it is complete. The interval between successive phase maintenance inspections of the same aircraft is driven by the number of flight hours accumulated on the aircraft. Accordingly, this model must consider not only maintenance scheduling but also the assignment of aircraft to sorties. Computational results show that the model is implementable in practice and can provide both maintenance schedules and policy makers with valuable insight.

A second maintenance scheduling problem we consider is directly related to new low observable technologies meant to aid in radar evasion. The low observable properties of an aircraft degrade over time. Based on historical data, we model the degradation and maintenance process as a stochastic process. Due to maintenance capacity constraints, directly solving the model is intractable and we propose an intuitive and easily implemented scheduling heuristic. In addition, we provide computational results that show that the heuristic is close to optimal for a number of scenarios. In addition, we use the stochastic model to explore different sortie assignment rules and provide recommendations on how maintenance schedulers can balance the competing demands of maintaining the low observable capability while also meeting flying requirements.

### 1.1. Literature Review

In this section we begin by providing a brief review of the literature pertaining to Air Force maintenance policy and procedure. We then review existing models regarding flight and maintenance scheduling.

The Air Force has several published documents that outline the fundamentals of aircraft maintenance. Air Force Tactics, Techniques, and Procedures (AFTTP) 3-3 addresses all aspects of aircraft maintenance, including flying and maintenance scheduling. This extensive document develops the a wide range of concepts associated with aircraft maintenance, from maintenance leadership to financial management. Regarding maintenance scheduling in particular, the document provides detailed explanations of scheduling practices that effectively

balance preventative maintenance requirements and flying requirements. Furthermore, the document contains guidance on maintenance scheduling for all aircraft types [8]. It is important to note, however, that the document does not provide any guidance on low-observable maintenance.

Another Air Force document that provides extensive information regarding aircraft maintenance is Air Force Instruction (AFI) 21-101. AFI 21-101 covers many of the same topics presented in AFTTP 3-3, but does not provide many details regarding maintenance scheduling in particular [3]. Therefore, to effectively implement the maintenance concepts presented in AFI 21-101, Air Combat Command (ACC) Instruction 21-165 has been published. ACC Instruction 21-165 deals specifically with the scheduling procedure for fighter aircraft in accordance with the maintenance requirements laid out in AFI 21-101[1].

Lastly, Air Force Materiel Command (AFMC) Instruction 21-165 lays out the planning cycle for aircraft maintenance scheduling. AFMC Instruction 21-165 provides a timeline for proper maintenance scheduling and it explains the approval authorities necessary to generate a final maintenance schedule. In addition, it provides details regarding the measures of performance that used to gauge the effectiveness of a maintenance schedule[2].

In general, the Air Force documents stress the importance of balancing preventative maintenance requirements and the flying requirements. A key concept stressed in each of the documents is the need for long-term planning and foresight. The importance of proper planning to aircraft availability and effective maintenance operations is highlighted throughout. Another common theme across all the documents is the need for maintenance schedulers to maximize aircraft availability by consolidating maintenance actions into a single downtime as much as possible.

We now shift our attention to the existing research on flight and maintenance scheduling. As expected, much of the existing work pertains to commercial airline scheduling. However, due to the distinct differences in military and civilian flight operations, especially maintenance scheduling for fighter operations, much of the existing work is not applicable to the military flying environment. In particular, the military flying environment deals with a far different objective than its commercial counterpart. While commercial airlines focus on profitability and cost savings, military flying operations revolve around combat readiness. This results in a different set of constraints and incentives. Another major difference is that commercial scheduling focuses heavily on route selection and assignments. Fighter aircraft operations are usually fixed at a given location and, therefore, do not involve any decisions regarding routes. Keeping these differences in mind, we present the existing research on flight and maintenance scheduling.

There has been extensive research done regarding commercial airline schedule and route planning. The literature pertaining to commercial airline scheduling is vast, but we present

these works to illustrate a few of the areas of research pertaining to commercial airlines. Early on, Levin (1971) formulated the scheduling and fleet routing problem as an integer program. Since then, work has continued in the area of scheduling and fleet route planning [10]. More recently, Lohatepanont et al. (2004) focused on fleet assignment and schedule planning while maximizing profitability [12]. Another area of research pertaining to commercial airline schedule is crew scheduling. Cohn et al. (2003) presented an tractable approach to the crew scheduling problem that maximizes revenue [4].

Ultimately, much of the work pertaining to the commercial airline industry has little relevance to fighter aircraft scheduling, however, there is one exception. Siriam et al. (2003) specifically address the scheduling of commercial aircraft maintenance. Fighter aircraft have maintenance requirements that are based on the total flight hours accrued. Similarly, commercial aircraft must be maintained after a certain number of flight hours. Siriam et al. formulates an IP that models this maintenance requirement, but the IP is found intractable. Furthermore, the IP does not completely model the flight hours accrued by each aircraft, but rather, it assumes that a given aircraft will simply require maintenance after every  $x$  number of calendar days. Since the flying operations of commercial aircraft remain relatively constant and the flight hours between consecutive maintenance actions remains constant, they argue that the maintenance requirements based on flight hours can be accurately estimated with fixed time intervals. Lastly, Siriam et al. assumes there is no capacity constraint. These assumptions in the IP formulated by Siriam et al. render the model inapplicable to fighter aircraft maintenance scheduling problem [16].

Therefore, the research presented in this thesis regarding phase maintenance for fighter aircraft is unique and separate from previous work done on commercial airlines.

We once again shift our focus from maintenance scheduling to multi-armed bandit (MAB) problems. We will give a brief summary of the literature pertaining to the MAB, but a full treatment of the relevant literature can be found in “The Irrevocable Multi-Armed Bandit Problem” by Farias et al. (2008).

In a classic MAB problem, there are  $n$  projects and the state of each project  $i = 1, 2, \dots, n$  is known to be  $s_1, s_2, \dots, s_n$ . In each time period, a decision must be made to operate one of the  $n$  projects. If project  $i$  is operated, then an immediate reward of  $g(i)$  is gained. Furthermore, project  $i$  transitions from state  $s_i(t)$  to  $s_i(t + 1)$  according to a Markov rule that is project and state dependent. The unoperated projects in a given time period do not yield a reward and they do not transition. The objective is to maximize the expected discounted reward over an infinite planning horizon.

For the classic MAB problem, Gittins and Jones developed an index policy that was shown to be optimal [9]. For a variant of the problem known as the restless MAB, where

projects transition even when unoperated and multiple projects can be operated simultaneously, Whittle proposed a simple index policy. Unlike the classic problem, the restless MAB with a capacity constraint is intractable. To address this issue, Whittle's heuristic is based on the relaxation of the capacity constraint on the number of projects that can be operated at once. The heuristic has been shown to give near optimal results in empirical studies [11]. In addition, it is computationally tractable as the resulting dynamic programs are significantly smaller than the original and can be solved in parallel.

## 1.2. Organization and Outline

The remainder of the work presented in this thesis is structured as follows:

**Chapter 2.** This chapter has three primary objectives. The first is to introduce the Air Force Flying Hours Program and its importance to the mission of the Air Force. The second is to discuss the process in which the requirements from the Flying Hours Program are balanced with various maintenance requirements in order to generate a long-term annual flying plan. Lastly, it presents a detailed description of the current weekly maintenance scheduling and planning process.

**Chapter 3.** The purpose of this chapter is to develop a mixed integer program (MIP) formulation that models the scheduling of phase maintenance for fighter aircraft. First, to clearly identify the inputs and outputs that are relevant to the subsequent mathematical formulations, we present a functional analysis of the maintenance scheduling process. Second, we develop, discuss, and implement the MIP formulation. The computational results motivate the need for an alternate formulation with better computational behavior. Lastly, we present a special case of the MIP formulation in which a simple heuristic can be used to find the optimal solution under a general set of assumptions.

**Chapter 4.** In this chapter, we present alternative methods of modeling the phase maintenance scheduling process. Motivated by the need for a model with better computational behavior, we disaggregate the PM-MIP (from Chapter 3) into two interrelated subproblems that are solved sequentially. We begin by presenting the framework under which we disaggregate the PM-MIP. The subproblems are then developed, discussed, and compared to the PM-MIP. Lastly, we implement the subproblems and compare the solutions with those from the PM-MIP.

**Chapter 5.** In this chapter, we shift our focus to maintenance scheduling issues that are unique to the Air Force's newest generation of low-observable (LO) aircraft. We begin this chapter by presenting a detailed explanation of the LO maintenance process and the dynamics associated with it. We then characterize the process as a restless MAB problem and present a simple index policy that can be used to schedule LO maintenance. Next,

we develop and discuss the dynamic programming (DP) formulation used to generate the desired index policy. Lastly, we simulate the index policy under a range of conditions to quantify the effectiveness of the policy.

**Chapter 6.** We close by summarizing the work presented in this thesis. In addition, we present opportunities for future work in areas pertaining to fighter aircraft maintenance scheduling. We suggest work that builds on the scheduling of preventative maintenance. In particular, we discuss the potential benefits of expanding the index policy presented in Chapter 5.



## CHAPTER 2

### **Air Force Fighter Aircraft Maintenance and Scheduling**

This chapter has three primary objectives. The first is to introduce the Air Force Flying Hours Program and its importance to the mission of the Air Force. The second is to discuss the process in which the requirements from the Flying Hours Program are balanced with various maintenance requirements in order to generate a long-term annual flying plan. Lastly, it presents a detailed description of the current weekly maintenance scheduling and planning process.

#### **2.1. Air Force Flying Hours Program (FHP)**

The Air Force Flying Hours Program (FHP) incorporates inputs from a wide range of documents and guidelines in order to determine the total number of flying hours that must be flown each year. The FHP ensures that aircrews receive the necessary training and experience to attain and maintain combat readiness. Therefore, the FHP is central to the Air Force mission and is the primary factor that drives both flying and maintenance operations [3].

To better understand the FHP and how it is managed within the Air Force, it is necessary to have a basic understanding of the organizational structure of the Air Force. Headquarters United States Air Force (HQ USAF) provides the senior leadership for the entire Air Force. HQ USAF consists of the Secretary of the Air Force (SECAF) and the Chief of Staff of the Air Force (CSAF), and their respective staffs. HQ USAF reports directly to the Secretary of Defense (SecDef) and frequently interacts with policymakers in the federal government. The major subdivisions beneath HQ USAF are the major commands (MAJCOMs). MAJCOMs can be associated either with a specific mission or they can be a geographical MAJCOM. For example, AF Space Command is responsible for all of the Air Force's operations in space while USAF Europe (USAFE) MAJCOM is responsible for all Air Force operations in Europe. Within each MAJCOM, the next major subdivision is the wing. Each wing has a distinct mission and consists of both operational and support units. Each base within the Air Force has one wing assigned to it. Therefore, the terms base level and wing level can be used interchangeably. The organizational structure of the Air Force includes many further subdivisions and special cases, but for the scope of this paper, we highlight only HQ USAF, MAJCOMs, and wings.

Each MAJCOM within the Air Force is required to build their own FHP based on its unique mission, personnel, and resources of the respective MAJCOM. Documents such as the

Joint Mission Essential Task List and Air Force task list outline the roles and responsibilities of the Air Force that have been determined to be crucial to the warfighting capability of the United States. Therefore, each MAJCOM must ensure that their FHP provides their aircrews with the proper training necessary to meet these required roles and responsibilities. While each MAJCOM develops their own individual FHP, the collective ability of the MAJCOMs to successfully execute the FHP determines the overall combat readiness of the Air Force and the United States.

After each MAJCOM has developed their FHP, the flying hours are then distributed amongst the various wings within it. The FHP requirements are specific to each mission design series (MDS), or unique aircraft type. Therefore, at the wing level, commanders are responsible for ensuring that their fleet of aircraft collectively meet the flying hours requirements outlined by the MAJCOM. In practice, the flying units within the Air Force strictly follow the flying hour requirements. If a given wing reaches its flying hours requirement earlier than planned, to avoid overflying the unit will often cease all flying activity. On the other hand, if a wing is behind schedule and is unlikely to achieve its flying hours requirement, other wings within the same MAJCOM will be notified and required to fly additional hours to ensure that the total flying hours for the MAJCOM are in line with the requirements of the FHP. For example, suppose that wing *X* is assigned 5,000 flight hours in a fiscal year and wing *Y* is assigned 8,000 flight hours. Assume both are part of the same MAJCOM. With a few months remaining, wing *Y* projects that it will fall short of its 8,000 flight hours by 100 hrs and it notifies MAJCOM headquarters. Personnel at MAJCOM headquarters will then notify wing *X* to increase its flying operations and complete 100 additional flight hours.

In summary, the FHP is central to the Air Force's readiness and combat capability. As a result, the FHP is monitored closely by Air Force leaders as well as legislators in Congress who provide funding based on the FHP. Given the centrality of the FHP, flying and maintenance operations within the Air Force must always revolve around the FHP requirements.

## **2.2. Operational and Maintenance Requirements**

Based on the FHP for each MAJCOM, each wing within the MAJCOM is given a flying hours requirement at the start of the fiscal year. To meet these flying hours requirements at the wing level, the operations and maintenance personnel must interact closely to develop a long-term annual flying plan. This plan dictates the total number of sorties, or flights, that will be scheduled in the coming year as well the aircraft configuration requirements associated with each sortie. A feasible plan must have enough sorties to meet flying hours requirements, but cannot require more sorties than maintenance resources can support. Therefore, operations and maintenance schedulers must interact closely to balance operational requirements and maintenance capabilities.

Operations schedulers must account for several key factors when developing the annual flying plan. First, they must balance the number of sorties with the average sortie duration (ASD). Although the flying hours program is based on total flying hours and not the number of sorties, operations schedulers must ensure there are enough sorties over the course of the year for all aircrew members to maintain combat readiness as defined by the ready aircrew program (RAP). Therefore, it is critical to plan for the proper number of sorties. They must also plan for various sortie types, such as weapons training sorties and night sorties, that are constrained by aircraft configurations or time of day. Lastly, operations schedulers must account for historical attrition data when generating annual flying plans. To ensure that the wing does not fall short of FHP requirements due to random aircraft failures or inclement weather, every wing schedules additional sorties based on historical attrition rates. By scheduling additional sorties to compensate for expected attrition, schedulers do not have to regenerate the flying plan each time an aircraft breaks or sorties are canceled due to weather.

Maintenance schedulers are responsible for properly allocating limited maintenance resources to ensure all aircraft are maintained for combat readiness. Maintenance schedulers must ensure that preventative maintenance for each aircraft is completed in accordance with published guidelines for each MDS. Preventative maintenance requirements are specific to each MDS and specify numerous maintenance actions that must be completed at prescribed intervals. Maintenance schedulers must consider different types of preventative maintenance requirements (see Section 2.2.1 for more details). In determining the annual flying plan, maintenance schedulers must work with the operations schedulers to determine the extent to which maintenance resources can support sortie requirements while also completing all required preventative maintenance.

**2.2.1. Preventative Maintenance Requirements.** A thorough understanding of preventative maintenance requirements is essential to the development of a feasible annual flying plan. Based on the preventative maintenance requirements, maintenance and operations schedulers must estimate the amount of aircraft downtime needed for maintenance and the resulting impact on sortie generation. Preventative maintenance requirements are specific to each MDS and are laid out in a published document referred to as the *-6 (dash 6)*. The *-6* lists all required preventative maintenance actions for a given MDS. It is updated as needed. There are two primary categories of preventative maintenance requirements: calendar based requirements and usage based requirements.

Usage based requirements are maintenance actions that are driven by flying hours. For each usage based requirement, the *-6* dictates the maximum number of flying hours that can be accrued between consecutive maintenance actions. Of all usage based maintenance requirements, *phase maintenance inspections* are the most significant. Phase maintenance, often simply referred to as *phase*, is an extensive inspection that requires the aircraft to be

inactive for a week or longer, depending on the MDS. During this downtime, maintenance personnel inspect an aircraft in fine detail, inspecting almost every system and part on the aircraft. If they discover any discrepancies, the proper action is taken to repair the aircraft. Phase maintenance is carried out at the base level, meaning that fighter aircraft undergo phase maintenance at whichever base they are stationed rather than at a centralized location. In most cases, dedicated docks or hangars that are used solely for phase maintenance. While there might be other minor usage based maintenance actions, phase maintenance is the consolidation of the vast majority of usage based maintenance actions.

As an example, the F-16 used to undergo phase maintenance every 300 flying hours meaning that 300 flying hours was the maximum number of flying hours allowed on an aircraft between consecutive phase maintenance actions. In 2003, the F-16 phase maintenance interval was extended to 400 hours. The change in the phase maintenance interval led to an estimated 20% reduction in F-16 maintenance workload and resulted in annual savings of millions of dollars. This illustrates the importance of phase maintenance and the impact it has on maintenance workloads. During a typical seven day phase inspection, maintenance personnel examine different areas of the aircraft each day following a standard predetermined inspection schedule. During a typical phase inspection, approximately 200-300 discrepancies are fixed. Upon completion of the inspection, the F-16 is then restored, reassembled, and flight tested.

Calendar based maintenance is treated similarly. The -6 for a given MDS specifies the maximum number of calendar days that are allowed between consecutive calendar based maintenance actions, independent of the number of flying hours that are accrued by an aircraft. Calendar based maintenance requirements can range from intervals of once every 30 days to once every several years. In most cases, calendar based maintenance does not require the extended aircraft downtime that is needed for phase maintenance. In many cases, these maintenance requirements can be completed relatively quickly between sorties and therefore do not require any actual scheduled downtime. An example of a calendar based maintenance action is an ejection seat inspection that must be completed once every 18 months.

In addition to the two primary categories of preventative maintenance, there are a few other types of preventative maintenance worth mentioning. Depot level maintenance is a form of preventative maintenance that occurs every several years and is completed at a single, centralized location for each MDS. Depot maintenance requires that an aircraft be flown to an MDS-specific location where highly specialized maintenance personnel conduct a complete overhaul of the aircraft. They conduct any upgrades that are necessary and can perform extensive maintenance that is beyond the capabilities of the maintenance personnel at the base level. Since depot level maintenance is carried out at a single consolidated location

for a given MDS, operations and maintenance schedulers at the base level have no input as to when an aircraft is scheduled for depot maintenance. Instead, base level schedulers are provided with specific dates on which a given aircraft must be delivered to depot maintenance. Although depot maintenance can lead to months of downtime for an aircraft, since it is not scheduled at the base level, it is not something that can be accounted for when determining a feasible annual flying plan. Another category of preventative maintenance actions is time compliance technical orders (TCTOs). TCTOs are often one time requirements that have been deemed necessary for an MDS and they must be completed by a certain date.

In trying to meet the requirements of the FHP, operations and maintenance schedulers must ensure that the annual flying plan allows for the completion of all necessary preventative maintenance actions.

### 2.3. Current Aircraft Maintenance Scheduling Process

To balance the demands of the FHP and preventative maintenance, the Air Force must carefully schedule aircraft preventative maintenance. This responsibility falls on the shoulders of the operations and maintenance schedulers assigned to each flying squadron. Separate from aircraft maintenance personnel, the maintenance schedulers' careers within the Air Force are solely devoted to learning and mastering the aircraft maintenance scheduling process. Currently, aircraft maintenance schedules are generated manually using simple spreadsheets and rules of thumb. This results in different approaches to maintenance scheduling that vary depending on the personality of the schedulers. Hence, the process of aircraft maintenance scheduling is often referred to as an "art" that requires years of experience to master. The remainder of this section will provide an in-depth discussion of the current process by which sortie requirements are generated as well as the subsequent maintenance scheduling process for fighter aircraft. It is important to note, however, that much of the aircraft maintenance scheduling process is not standardized across the Air Force and, in many cases, fighter squadrons will have differing planning processes. Although details regarding the maintenance scheduling process may differ amongst squadrons, this section will provide the reader with a foundational understanding of the steps involved in generating an operational aircraft maintenance schedule.

**2.3.1. Building Flying Schedules and Sortie Requirements.** As explained earlier, the aircraft maintenance scheduling process for fighter aircraft begins with the Flying Hour Program (FHP). Based on the flying hours requirements that are dictated by the FHP for a given year, each operational flying squadron generates an estimated long-term annual flying plan. The long-term flying plan uses the FHP and the training requirements for the aircrews to determine how many sorties of various types need to be flown over the course of the year to meet the FHP for the upcoming year. When creating this long-term flying plan, factors

such as weather and attrition are taken into account. This flying plan does not specify daily sortie requirements, but rather, it is high-level plan used to determine the total number of sorties that need to be completed in the coming year.

The long-term annual flying plan is used to generate a more detailed flying schedule that specifies the number of various types of sorties that are to be flown on each calendar day. Figure 2.3.1 shows a sample monthly schedule from a yearly flying schedule. For each day, the schedule shows the number of sorties that need to be completed. In this case, the calendar includes the flying schedule for two squadrons represented by the two rows of numbers in each day. The morning and afternoon sortie requirements are separated by an “x”. The “p” refers to a *pit*, which means an aircraft that has just landed is refueled with the engines running to allow for a quick turnaround time. As an example, “9p9x8” would represent the following sequence of events:

- (1) Nine aircraft takeoff in the morning.
- (2) The nine aircraft land, refuel with engines running, and takeoff again that same morning.
- (3) The nine aircraft land and are shut down.
- (4) Eight aircraft take off in the afternoon.

In general, there are several common practices used when it comes to generating detailed a flying schedule at the group level. As can be seen in the example monthly schedule in Figure 2.3.1, flying operations are almost exclusively conducted during weekdays. Only in rare instances are sorties scheduled for weekends or holidays. Another common practice is to schedule “pits” only in the morning. Furthermore, pits are used more often earlier in the week to increase the proportion of sorties that are flown during the first half of the week. Since aircraft are likely to break during the week, fewer sorties are scheduled for the latter half of the week to hedge against the possibility of reduced aircraft availability.

Along the same line of reasoning, the number of sorties scheduled in the morning is always greater than the number of sorties scheduled in the afternoon. There is a significant amount of preflight work necessary to prepare an aircraft to fly on a given day. Therefore, to minimize the total preflight preparation for each day, the same aircraft are used to complete all morning and afternoon sorties. Given this policy, there is a high likelihood that at least one aircraft will be unavailable to fly in the afternoon due to a maintenance issue after the morning sorties. Hence, the number of afternoon sorties is always less than the number of morning sorties.

In summary, the FHP leads to a flying plan that specifies the total number of sorties that need to be flown in a year. Next, the flying plan is used to build a more detailed flying schedule that specifies the number of sorties to be flown each day. This flying schedule is generated so that the total number of sorties scheduled over the course of the year meets

		Squadron XX Monthly Flying						
		SUNDAY	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY
27 94	X/C		1 9p9x8 9x8	2 26 9p9x8 17 9p9x8	3 26 9x8 26 9p9x8	4 17 9x8 26 9x8	5 17 9x9 17 9p9	6 18 18
		7	8 9p9x8 9x8	9 26 9p9x8 17 9p9x8	10 26 9x8 26 9p9x8	11 17 9x8 26 9x8	12 17 9 17 9	13 9 4x3 9 4x3
27 94	X/C	14	15 9p9x8 9x8	16 26 9p9x8 17 9p9x8	17 26 9x8 26 9p9x8	18 17 9x8 26 9x8	19 17 9p9 17 9p9	20 18 18
		21	22 9p9x8 9x8	23 26 9p9x8 17 9p9x8	24 26 9x8 26 9p9x8	25 17 9x8 26 9x8	26 17 9x9 17 9	27 18 9
27 94	X/C	28	29 MX TNG DAY No Fly	30 9p9x8 9p9x8	31 26 9x8 26 9p9x8			
		<b>Total Sorties</b>	<b>457/ 457 / 914</b>	<b>(27 / 94 / Total)</b>	<b>Attrn: 18.9</b>	<b>Contract</b>	<b>371/ 371 / 742</b>	

FIGURE 2.3.1. Sample monthly flying schedule taken from a long-term fiscal year calendar

the FHP requirements taking into account a historical rate of attrition. Furthermore, some common practices are used to guide the development of the flying schedule. The flying schedule ultimately provides what will be referred to as the sortie requirements for a given fleet of aircraft.

**2.3.2. Assigning Aircraft to Sortie Requirements.** After the flying schedule has been generated using the FHP requirements, the next step is to develop a detailed schedule that assigns specific aircraft to each sortie requirement while also scheduling downtime for

required preventative maintenance. Maintenance schedulers are solely responsible for deciding which aircraft will be flown each day in support of the sortie requirements and when each aircraft will undergo preventative maintenance. In making these decisions, the maintenance schedulers must be keenly aware of the state of each aircraft and, in particular, must understand the distribution of phase hours remaining on each aircraft. The assignments of aircraft to sortie requirements and scheduled downtimes for preventative maintenance will be collectively referred to as the maintenance schedule.

The maintenance schedule is often presented in a format known as the checkerboard. The checkerboard lists all aircraft in a given unit and specifies what each aircraft will be assigned to do on every day of the week. In addition, the checkerboard contains information regarding the phase hours remaining on each aircraft as well the sortie requirements on each day.

Figure 2.3.2 shows a sample checkerboard for a fighter squadron. This particular checkerboard covers a one week period and 20 aircraft. The first column lists the tail number of each aircraft in the squadron. The second column shows the phase hours remaining on each aircraft as well the current configuration of the aircraft. For each day of the week, an aircraft is assigned to a certain activity. An “F” means that the aircraft is scheduled to fly on that day. All other symbols are defined in the legend. The final row, entitled “Daily Turn”, shows the sortie requirements for each day, which come from the flying schedule discussed in the previous section. Aircraft that are not assigned to fly, could be assigned to numerous other activities, many of which are preventative maintenance. If an aircraft is not in need of preventative maintenance and is not assigned to fly, it is commonly assigned to USM, or unscheduled maintenance. Being assigned to USM does not mean that an aircraft will actually be maintained, but it is not flying and maintenance personnel might work on the aircraft if they feel there is a need.

A weekly checkerboard such as the one shown in Figure 2.3.2 is often developed two weeks beforehand by maintenance schedulers who account for a wide range of factors. To choose which aircraft that will be assigned to fly on a given day, schedulers must consider the phase hours remaining on each aircraft as well as any upcoming preventative maintenance requirements.

Schedulers cannot focus on an individual aircraft’s hours remaining until phase maintenance. Rather, they must look at the distribution of phase hours across the entire fleet of aircraft, or the *phase flow*. Managing the phase flow is intended to prevent and avoid a situation where multiple aircraft come due for phase maintenance nearly simultaneously. Performing phase maintenance on multiple aircraft at the same time, either requires maintenance personnel to work extra hours or will lead to extended aircraft downtime due to backlogging. Therefore, to avoid these scenarios, schedulers seek to keep the phase hours



22 FS WEEKLY UTILIZATION AND MAINTENANCE SCHEDULE WEEK OF:										17 - 23 JUN 2002							
ACFT	PHS HRS	MON	TUE	WED	THU	FRI	SAT	SUN	ACFT	PHS HRS	MON	TUE	WED	THU	FRI	SAT	SUN
	CONFIG	17	18	19	20	21	22	23		CONFIG	17	18	19	20	21	22	23
90-0813	176.5 T2B	F	SI	F	SI	F			91-1339	48.3 T2B	CANN	CANN	CANN	CANN	CANN		
90-0916	55.7 T2B	F	DR	F	SI	F			91-1340	66.9 T2B	SI	SP	DR	SI	USM		
90-0827 PACER	187.3 T2B	SI	F	WP	F	DR			91-1341	266.1 T2B	SP	DR	F	SI	F		
90-0828	114.4 T2B	SI	USM	USM	USM	SI			91-1342	121.2 T2B	SI	USM	USM	SI	USM		
90-0829 SQ	25.8 T2B	F	F	F	SI	F			91-1343	118.3 T2B	PD	SP	SP	SP	SP		
90-0831	215.5 T2B	F	F	DR SI	SI	F			91-1344	126.3 T2B	USM	USM	USM	SI	USM		
90-0833	196.2 T2B	USM	USM	USM	USM	USM			91-1351	66.2 T2B	USM	USM	USM	F	F		
91-1336	35.4 T-7	USM	USM	USM	SI	DR			90-0843	96.7 T-6	DR	SI	USM	SI	F		
91-1337	113.2 T2B	F	F	SI	SI	PD			91-1464	191.3 T-6	F	F	F	SI	SI		
91-1338	146.6 T2B	USM	USM	SI	DR SI	USM			91-1352	72.4 WG T2B	DR	F	F	SI	SP	SP	
DAILY TURN		8P8X6	8P8X6	8P8X6	8P8X6	10P10	DAILY				8P8X6	8P8X6	8P8X6	8P8X6	10P10		
A/R - AS REQUIRED		DEMO-DEMONSTRATION			FDI - FIRE DEPT TRNG			PD - PREPOST OOCK			SP-SPARE FLYER			USM - UNSCHEDULED MAINT			
A/SP - AIR SPARE		DR - DDC REVIEW			FTD - FIELD TRNG DET			RECON-RECONFIGURE ACFT			T/F - TAIL FLASH			WLT - WEAPONS LOAD TRNG			
AMQP-ACFT MX QUAL PRGM		EC - ENGINE CHANGE			FXC-CROSS COUNTRY			RF-RED FLAG PRIME			TQ-TIME CHANGE ITEM			WA - AM WASH			
CAC - CHART "A" CHECK		EOD- EOD TRAINING			G/SP - GROUND SPARE			RFS-RED FLAG SPARE			TCTO-TIME COMP TECH ORDER			WP - PM WASH			
CT-COMBAT TURN		F - FLYING			LTF - LOG TRAINING FLT			RTWS-RADAR THREAT WRNG SVS			TDY - DEPLOYED						
DD-DELATED DISC		F OBB - FLY OUT & BACK			IWS- IWS MOD			SD - STATIC DISPLAY			TM - TIME MANAGEMENT						

FIGURE 2.3.2. Sample checkerboard for a one week period

remaining on each aircraft evenly spaced. If schedulers are successful in keeping the aircraft evenly spaced the resulting phase flow will resemble the an ideal state phase flow illustrated in Figure 2.3.3 in which all aircraft have well spaced phase hours. When determining which aircraft will fly on a given day, schedulers must be aware of the phase flow of the fleet of aircraft and seek to make flying assignments that help achieve or maintain a reasonable phase flow. Schedulers must carefully manage the utilization rates of each aircraft, especially for aircraft that are nearing phase maintenance, to ensure that phase maintenance inspections come due at the proper times.

For calendar based maintenance requirements, the utilization rates are irrelevant since the rate of flying has no impact on when an aircraft must be maintained. However, schedulers must determine what days an aircraft can be down for calendar based maintenance and must ensure that, as a fleet, there are enough aircraft available to meet the sortie requirements.

Due to the vast number of preventative maintenance requirements that exist for a given MDS, schedulers must always seek opportunities to consolidate maintenance actions into a single aircraft downtime. In many cases multiple maintenance actions can be completed simultaneously since they each address separate systems on the aircraft. For example, the ejection seat can be maintained at the same time as the gun. This allows for consolidation in the sense that multiple maintenance actions are scheduled for a single downtime. Without consolidation, it is unlikely that a squadron would ever have enough aircraft to meet the sortie requirements. Therefore, schedulers must analyze all upcoming maintenance actions and see if there might be a way to consolidate maintenance actions by doing certain maintenance actions earlier than required.




FIGURE 2.3.3. Ideal Phase Flow

After taking into account all these factors, schedulers generate a checkerboard for each week by weighing the trade offs of possible decisions. For a fleet of 20 aircraft or more, being able to come up with a feasible checkerboard is a nontrivial task. Furthermore, since checkerboards are generated on a week by week basis, schedulers cannot fully weigh many of the long term consequences that might come about as a result of the decisions made in the current period. For example, consider a case in which there is a large increase in sortie requirements in the coming weeks. Schedulers might not consider this increase in developing the current week's checkerboard and therefore not attempt to complete as much preventative maintenance as possible in preparation for the coming spike in demand. Hence, in the coming weeks, maintenance personnel might be required to work extended hours to increase the number of available aircraft even though it might have been possible to complete the maintenance in earlier weeks. Ultimately, the Air Force relies on experienced schedulers, who have mastered the "art" of scheduling over long careers, to generate maintenance schedules that properly balance FHP requirements and preventative maintenance requirements.

#### 2.4. Summary

The Air Force Flying Hours Program (FHP) is central to the combat readiness of the Air Force. The FHP determines the number of total flying hours that are necessary to attain and maintain a level of combat readiness that will allow the Air Force to fulfill its mission. The FHP is used to determine how many flying hours must be flown by each squadron of the Air

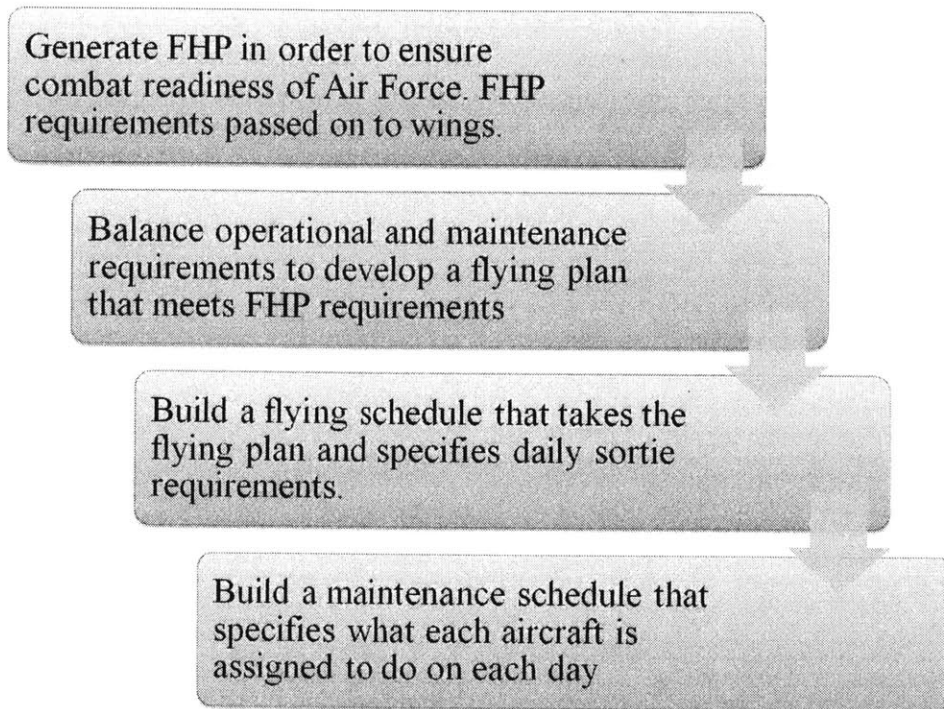


FIGURE 2.4.1. The steps required to generate a detailed maintenance schedule

Force. Based on the flying hours requirements, operations and maintenance schedulers work together to develop a flying plan that specifies the total number of sorties of various types that will be flown. To do this, the operations and maintenance schedulers must balance the FHP requirements, aircrew training requirements, and limited maintenance resources to ensure that all flying and maintenance requirements are met. The flying plan then leads to a flying schedule that specifies the exact number of sorties that are required each day.

Once the sortie requirements have been outlined in the flying schedule, maintenance schedulers are solely responsible for assigning aircraft to the sorties while also ensuring that all preventative maintenance requirements are met. To develop a feasible and sustainable maintenance schedule, schedulers must consider the phase flow of the fleet and the upcoming preventative maintenance requirements. They must also seek to consolidate maintenance actions into a single downtime as much as possible. Using rules of thumb that have been formed from prior experience, they generate a checkerboard for each week. Figure 2.4.1 summarizes the process through which the maintenance schedule is generated.

Ultimately, the Air Force relies on experienced maintenance schedulers to properly balance FHP requirements and preventative maintenance requirements. The remaining chapters will present various mathematical models aimed at capturing the maintenance scheduling process and the trade offs that are involved.

THIS PAGE INTENTIONALLY LEFT BLANK

## CHAPTER 3

### MIP Formulation for Phase Maintenance Scheduling

The purpose of this chapter is to develop a mixed integer program (MIP) formulation that models the scheduling of phase maintenance for fighter aircraft. As explained in Section 2.2.1, phase maintenance inspections are completed at the base level. They are based on the cumulative flying hours accrued by an aircraft since its last inspection. Due to the extended aircraft downtime required to complete phase maintenance, the proper scheduling of phase maintenance is critical to meeting operational requirements. Currently, phase maintenance schedules are generated manually, but given a set of operational requirements, a schedule could be generated by solving a mixed integer program (MIP).

First, to clearly identify the inputs and outputs that are relevant to the subsequent mathematical formulations, we present a modeling framework for the maintenance scheduling process. Second, we develop, discuss, and implement the MIP formulation. The computational results motivate the need for an alternate formulation with better computational behavior. Lastly, we present a special case of the MIP formulation in which a simple heuristic can be used to find the optimal solution under a general set of assumptions.

#### 3.1. Modeling Framework for the Air Force maintenance scheduling process

The aircraft maintenance scheduling process is highly complex due to a myriad of factors that must be simultaneously taken into consideration. Ideally, a mathematical model could be used to represent the entire scheduling process and generate an optimal schedule. However, such an extensive mathematical model would likely be intractable. Beyond tractability issues, some of the factors that impact the maintenance scheduling process are subjective and difficult to represent mathematically. Given these limitations, the scope of our initial model is limited to phase maintenance scheduling. To model the phase maintenance scheduling process, we first present a general modeling framework of the problem that describes the inputs and outputs of the model.

**3.1.1. Inputs to the Phase Maintenance Scheduling Process.** To generate a feasible phase maintenance schedule, several key pieces of data must be known, starting with the sortie requirements over the planning horizon. The sortie requirements come directly from the flying schedule. They are specified for each time period, with a single time period equivalent to half a day. If there are multiple sortie types, such as sorties types of differing

durations, the sortie requirements must specify the number of each type of sortie that must be flown in each time period. Next, the preventative maintenance requirements must be known. In this chapter, we are concerned with the preventative maintenance requirements only for phase maintenance. Phase maintenance is based solely on the cumulative flying hours accrued by each aircraft (unlike many other preventative maintenance actions that are based calendar days). We must know the maximum flying hours allowed between phase maintenance actions as well the associated downtime for completing phase maintenance. Third, in addition to the sortie and maintenance requirements, the number of *primary assigned aircraft* (PAA) and the initial state of each of these aircraft must be known. Since we are concerned only with phase maintenance scheduling, the initial state data for each aircraft only consists of the flying hours remaining until phase maintenance, or phase hours remaining. The number of PAA determines how many aircraft are available to satisfy the sortie requirements. These are the primary inputs that factor into the generation of a phase maintenance schedule.

Given these inputs, we find it appropriate to model the phase maintenance scheduling process as a deterministic model. As discussed earlier, the flying hours program (FHP) requirements are known at the start of each fiscal year and the sortie requirements are built to satisfy the FHP requirements. Therefore, once the sortie requirements are set, they are highly unlikely to change. Also recall that the sortie requirements are determined while taking into account historical attrition rates and random aircraft failures. Therefore, the stochastic nature of aircraft failures is implicitly captured by the sortie requirements.

**3.1.2. Outputs of the Phase Maintenance Scheduling Process.** Given this data, we need to make two primary interrelated sets of decisions: (1) a set of flying decisions that assign each required sortie to a specific aircraft, and (2) a set of maintenance scheduling decisions must that schedule each aircraft for maintenance at a specific time. Since the maintenance scheduling decision for phase maintenance is solely driven by total flying hours, the two sets of decisions are highly interrelated and must be made simultaneously. The flying decisions must be made to ensure that all sortie requirements are met. Concurrently, it is necessary to take into account the effect that the flying decisions have on phase hours and maintenance demand. Figure 3.1.1 summarizes the set of inputs and outputs for the phase maintenance scheduling process.

## 3.2. MIP Formulation of Aircraft Phase Maintenance Scheduling

In this section we present a formulation of the phase maintenance MIP (PM-MIP). We explain the characteristics of the objective function used in the PM-MIP and present a simple use of the LP relaxation to improve the PM-MIP formulation.

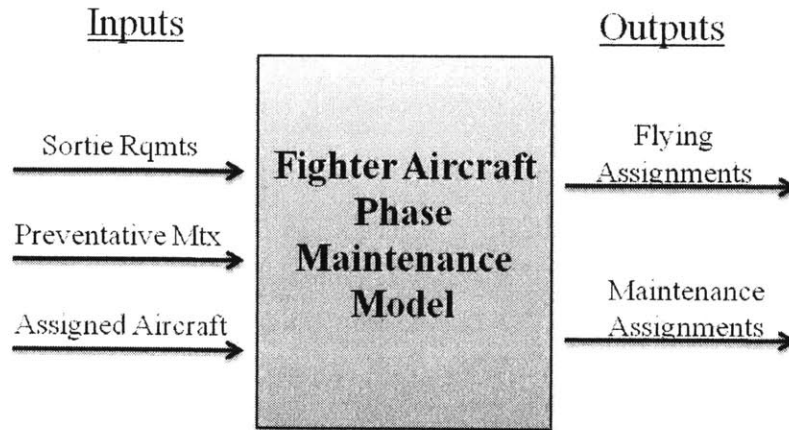


FIGURE 3.1.1. Functional View of the Fighter Aircraft Phase Maintenance Scheduling Process

**3.2.1. Model Ingredients and Formulation.** This section introduces the first mathematical formulation we use to model the phase maintenance scheduling process. At the start of the planning horizon, we assume that we know the phase hours remaining on each aircraft as well as the sortie requirements in each time period. Throughout the planning horizon, the model keeps track of the phase hours remaining for each aircraft as it decides which aircraft are assigned to meet the sortie requirements in each period. By keeping track of the phase hours for each aircraft, the model also decides when to schedule each aircraft for phase maintenance. The model must enter each aircraft into phase maintenance to ensure that the flying hours accrued never exceeds the maximum allow between phase maintenance inspections. In addition, the model cannot enter an aircraft into phase maintenance excessively early since it unnecessarily increases costs. At the end of the horizon, it is crucial that the phase hours remaining for all aircraft be well spaced. Therefore, we include evenly spaced end of horizon targets in the model to ensure that the distribution of phase hours at the end of the planning horizon is close to an ideal state. The model must assign each aircraft to an end of horizon target which then determines the number of phase hours remaining each aircraft must have at the end of the planning horizon.

The objective of the model is to minimize the maximum number of planes in maintenance at any given time during the planning horizon. The min-max objective function is motivated by the desire to minimize the variance in maintenance workload over the planning horizon, rather than simply minimizing the total maintenance over the planning horizon. By using a min-max function, the model attempts to spreadout the phase maintenance actions as much as possible. Ultimately, this MIP produces a feasible flying and maintenance schedule that balances the variability in maintenance demands and the total number of aircraft in phase maintenance at any given time.

**PM-MIP Notation.** We begin by defining all the relevant sets to the PM-MIP. We let  $I$  denote the set of all aircraft and  $|I|$  is the total number of aircraft. The set of all aircraft includes all the aircraft that share the same phase maintenance resources. Many operational bases have multiple flying squadrons that generate independent maintenance schedules, but since they share the same phase maintenance resources, they must be considered together in a phase maintenance scheduling model.

As discussed in Section 3.1.1, the sortie requirements for PM-MIP come from the annual flying schedule. The sortie requirements specify the number of each type of sortie that must flown in each time period. To accurately capture the sortie requirements, we must be able to differentiate between various sortie types. Therefore, we let  $J$  denote the set of all sortie types and  $|J|$  is the number of different sortie types. Examples of different sortie types are night sorties and weapons range sorties.

To ensure that the phase hours remaining on the aircraft at the end of the planning horizon is well spaced, we assign each aircraft to an end of horizon target. This is crucial if the model is to be solved in a rolling horizon setting since it forces aircraft to be well spaced at the end of the planning horizon. The end of horizon targets require that each aircraft have a certain number of phase hours remaining at the end of the horizon so as to prevent multiple aircraft from coming due for phase maintenance simultaneously in the future. The targets are evenly spaced from zero phase hours remaining to the maximum possible phase hours remaining. We let  $Q$  denote the set of all end of horizon targets. In this formulation, we assume that each aircraft is assigned to a unique end of horizon target so  $|Q| = |I|$ . However, if  $|Q| < |I|$ , multiple aircraft could be assigned to the same end of horizon target.

The sets in the PM-MIP formulation are summarized below.

**Sets.**

$I =$	set of all aircraft $i \in I$
$J =$	set of all sortie types $j \in J$
$Q =$	set of all end of horizon targets $q \in Q$
$ I  =$	number of aircraft available in the model
$ J  =$	number of unique sortie types
$ Q  =$	number of end of horizon targets (we assume $ Q  =  I $ unless otherwise noted)

Next, we present the data associated with the PM-MIP. In the formulation,  $T$  is the length of the planning horizon, with each time period  $t = 1, 2, \dots, T$ , representing a half a day. While the formulation could be adapted to allow each period to represent any arbitrary period of time, we segment the planning horizon into half day periods to accurately represent current Air Force operations. In general, fighter aircraft are scheduled for sorties in the morning and afternoon which naturally lends half day periods. In addition, since flying operations and phase maintenance are generally restricted to weekdays, we assume that weekends are not



included in the planning horizon. The data could easily be adapted to include weekends, but at the cost of increasing the size of the model. For the remainder of the chapter, all instances that are solved assume no activity on weekends.

The parameter  $\bar{h}$  represents the phase maintenance interval, that is, the maximum flying hours an aircraft can accrue between phase maintenance actions. This parameter is specific to the aircraft type or mission design series (MDS) being considered in the model. Likewise, the parameter  $k$ , which defines the length of downtime needed to complete a phase maintenance action, is also specific to the MDS. In operational settings, to prevent aircraft from entering phase maintenance excessively early, maintenance scheduler can only schedule an aircraft for phase maintenance if its phase hours remaining is below a certain limit,  $h_{max}$ . This parameter  $h_{max}$  represents the maximum number of phase hours an aircraft can have when entered into phase maintenance.

The parameters  $b_i$  and  $e_q$  describe the states of aircraft at the beginning and end of the planning horizon. At the start of the planning horizon, each aircraft,  $i \in I$ , has a given number of phase hours,  $b_i$ , remaining. At the end of the planning horizon, it is necessary to ensure a distribution of phase hours among the aircraft so that a large number of aircraft will not later come due for phase maintenance at the same time. This is of particular importance if the model is solved in a rolling horizon setting. The parameters  $e_q$  are used to define the end of horizon targets and they are evenly spaced on the interval  $[0, \bar{h}]$ . For each target  $q \in Q$ ,  $e_q$  is the target number of remaining flying hours for any plane assigned to that target at the end of the horizon.

The parameter  $\bar{M}$  is of particular importance due to its relevance for the objective function and the implications it can have on the computational behavior of the model.  $\bar{M}$  is a fixed maintenance capacity that can be used without cost or penalty. We explain the role of the parameter  $\bar{M}$  further in Section 3.2.2.

#### Data.

$T$	length of time horizon, $t = 1, 2, \dots, T$
$l_j$	length of sortie type $j \in J$
$s_j^t$	minimum number of sorties of type $j \in J$ required in period $t$
$\bar{h}$	maximum accrued flying hours between phase maintenance inspections
$h_{max}$	maximum number of remaining phase hours an aircraft can have for it to be entered into maintenance
$k$	time periods required to complete maintenance (aircraft unavailable)
$b_i$	flying hours remaining on aircraft $i \in I$ at the beginning of the horizon, $t = 1$ , until it must enter maintenance
$e_q$	end of horizon flying hours target for aircraft assigned to target $q \in Q$

$\overline{M}$  maximum number of aircraft that can be in maintenance at any give time without incurring a “cost” or “penalty”

Next, we present the decisions variables in the formulation. In each time period, several decisions must be made for each aircraft. At the start of each time period,  $t = 1, 2, \dots, T$ , the model must decide whether an aircraft,  $i \in I$ , is entered into phase maintenance. In addition, the model must decide whether the aircraft will be assigned to fly a given sortie type in each time period. These two decisions are represented by the decision variables  $m_i^t$  and  $x_{ij}^t$ . These sets of decisions define the overall flying and maintenance schedule. The decision variables  $v_{iq}$  represent the decision to assign aircraft  $i \in I$  to end of horizon target  $q \in Q$ . Recall that the  $e_q$  parameters define end of horizon targets. If an aircraft is assigned to end of horizon target  $q \in Q$ , then the aircraft must have  $e_q$  phase hours remaining at the end of the planning horizon. Lastly,  $h_i^t$  represents the remaining phase hours for aircraft  $i \in I$  at the start of period  $t$ . The variables  $h_i^t$  represents the phase hours remaining on each aircraft in a given time period. While  $h_i^t$  is considered a variable, its value is determined by the decision variables  $m_i^t$  and  $x_{ij}^t$ .

### Decision Variables.

$$\begin{aligned}
 x_{ij}^t &= \begin{cases} 1, & \text{if aircraft } i \in I \text{ flies sorties type } j \in J \text{ in period } t \\ 0, & \text{otherwise} \end{cases} \\
 m_i^t &= \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases} \\
 v_{iq} &= \begin{cases} 1, & \text{if aircraft } i \in I \text{ is assigned to end of horizon target } e_q \\ 0, & \text{otherwise} \end{cases} \\
 h_i^t &\geq 0 \quad \text{phase hours remaining at the start of period } t \text{ until aircraft } i \in I \text{ must enter} \\
 &\quad \text{maintenance, “life remaining”} \\
 Z &\quad \text{objective function value that is being minimized}
 \end{aligned}$$

### Formulation.

$$(3.2.1) \quad \text{PM-MIP} = \text{minimize } Z$$

subject to

$$(3.2.2) \quad h_i^1 = b_i, \quad , \forall i \in I,$$

$$(3.2.3) \quad h_i^{t+1} \leq h_i^t - \sum_j x_{ij}^t l_j + \bar{h} m_i^t, \quad , \forall t, i \in I,$$

$$(3.2.4) \quad h_i^{t+1} \geq h_i^t - \sum_j x_{ij}^t l_j, \quad , \forall t, i \in I,$$

$$(3.2.5) \quad h_i^{t+1} \leq \bar{h}, \quad , \forall t, i \in I,$$

$$(3.2.6) \quad h_i^{t+1} \geq \bar{h} m_i^t, \quad , \forall t, i \in I,$$

$$(3.2.7) \quad \sum_j x_{ij}^t \leq 1, \quad , \forall t, i \in I,$$

$$(3.2.8) \quad \sum_i x_{ij}^t \geq s_j^t, \quad , \forall t, j \in J,$$

$$(3.2.9) \quad \bar{h} - h_i^t \geq (\bar{h} - \bar{h} h_{\min}^{\cdot}) m_i^t, \quad , \forall t, i \in I,$$

$$(3.2.10) \quad \sum_j x_{ij}^t \geq \sum_j x_{ij}^{t+1}, \quad , \forall \text{odd } t, i \in I,$$

$$(3.2.11) \quad m_i^t + \sum_j x_{ij}^{t+y} \leq 1, \quad , \forall t \in [1, T - k + 1], y \in [0, k - 1], i \in I,$$

$$(3.2.12) \quad m_i^t + m_i^{t+y} \leq 1, \quad , \forall t \in [1, T - k], y \in [1, k], i \in I,$$

$$(3.2.13) \quad 0.9 v_{iq} e_q \leq h_i^T, \quad , \forall i \in I, q \in Q,$$

$$(3.2.14) \quad 1.1 v_{iq} e_q + \bar{h}(1 - v_{iq}) \geq h_i^T, \quad , \forall i \in I, q \in Q,$$

$$(3.2.15) \quad \sum_i v_{iq} = 1, \quad , \forall q \in Q,$$

$$(3.2.16) \quad \sum_q v_{iq} = 1, \quad , \forall i \in I,$$

$$(3.2.17) \quad \sum_i \sum_{\tau=t-k}^t m_i^\tau - \bar{M} \leq Z, \quad , \forall t \in [k + 1, T],$$

$$(3.2.18) \quad Z \geq 0,$$

$$(3.2.19) \quad h_i^t \geq 0, \quad , \forall i \in I, t,$$

$$(3.2.20) \quad x_{ij}^t \in \{0, 1\}, \quad , \forall i \in I, j \in J, t,$$

$$(3.2.21) \quad m_i^t \in \{0, 1\}, \quad , \forall i \in I, t,$$

$$(3.2.22) \quad v_{iq} \in \{0, 1\}, \quad , \forall i \in I, q \in Q.$$

The objective function (3.2.1), in conjunction with constraints (3.2.17) and (3.2.18), minimizes the maximum number of aircraft in maintenance at any given time during the planning

horizon beyond the amount of available fixed maintenance capacity. By minimizing the maximum, we seek to smooth the variability in maintenance demand over time. The parameter  $\bar{M}$  in constraint (3.2.17) leads to a piecewise-linear penalty function that imposes zero cost when  $Z \leq \bar{M}$ . The characteristics of the objective function are discussed in much greater detail in Section (3.2.2).

The first set of constraints, (3.2.3)-(3.2.6), in the formulation ensure that the phase hours remaining on each aircraft are tracked properly depending on whether an aircraft flies or enters maintenance. Constraints (3.2.3)-(3.2.3) can be thought of as “bookkeeping” constraints. If an aircraft enters phase maintenance in period  $t$ , constraints (3.2.5) and (3.2.6) force the phase hours remaining,  $h_i^{t+1}$ , to be set equal to  $\bar{h}$ . In time periods in which an aircraft does not enter maintenance, constraints 3.2.5 and 3.2.6 become irrelevant and instead, constraints (3.2.3) and (3.2.4) ensure that the phase hours remaining on an aircraft in a time period  $t$ , is simply the phase hours remaining at the start of period  $t - 1$  minus any flying hours accrued in period  $t - 1$ . The  $\bar{h}$  in constraint (3.2.3) functions as a “big- $M$ ” since it ensures that when  $m_i^t = 1$ , constraint (3.2.5) restricts  $h_i^{t+1}$  instead of constraint (3.2.3).

Constraint (3.2.8) enforces all of the operational sortie requirements. This constraint ensures that the flying schedule satisfies the specified sortie requirements. Since the sortie requirements are determined to satisfy the FHP it is not necessary to impose an additional set of constraints to ensure that enough total hours are flown. Therefore, constraint (3.2.8) represents all the sortie requirements necessary for aircrew training as well for meeting the FHP.

In practice, maintenance schedulers are not allowed to schedule an aircraft for phase maintenance until the phase hours remaining on the aircraft drop below a given threshold. That is, there is a limit on how early an aircraft can be entered into maintenance. Constraint (3.2.9) models this condition.

Constraint (3.2.10) enforces an earlier discussed policy that is common among operational fighter units. Each time period in the model represents a 1/2 day, either morning or afternoon. In practice, the set of aircraft flown in the afternoon is a subset of the aircraft that flew in the morning. If an aircraft is not flown in the morning (an odd time period), then it cannot be assigned to fly in the afternoon. Of course this constraint assumes that the number of required afternoon sorties is always less than or equal to the number of morning sorties which is always the case in practice. These constraints might or might not be necessary depending on the length of time that each period represents.

Constraint (3.2.11) and (3.2.12) enforces the downtime necessary to complete a phase maintenance inspection. Upon entering maintenance, an aircraft cannot be assigned to fly a sortie for  $k$  periods, the length of time needed to complete phase maintenance. Similarly, the aircraft cannot be assigned to enter maintenance for  $k$  periods.

Constraints (3.2.13)-(3.2.16) are end of the horizon constraints that ensure that all aircraft are relatively evenly spaced in terms of phase hours remaining at the end of the planning period. Constraints (3.2.15) and (3.2.16) require that each aircraft be assigned to a unique end of horizon target. Constraints (3.2.13) and (3.2.14) require that each aircraft's phase hours be within 10% of its assigned target. The targets,  $e_q$ , are evenly distributed in the range  $[0, \bar{h}]$ . This group of constraints is crucial since, in reality, there is no true "end of the horizon". Without these constraints, the solution generated by the model could lead to a situation where all aircraft have minimal phase hours remaining at the end of the model's planning period.

**3.2.2. Improving the PM-MIP using the LP Relaxation.** Recall that  $\bar{M}$  was a part of the constraints associated with the objective function. The objective function and the associated constraints were:

$$Z_{\text{MIP}}^* = \min Z$$

subject to

$$\sum_i \sum_{\tau=t-k}^t m_i^\tau - \bar{M} \leq Z \quad \forall t$$

$$Z \geq 0.$$

Although  $\bar{M}$  was introduced simply as a representation of the fixed maintenance capacity, the value of  $\bar{M}$  also has a significant effect on the computational behavior of the formulation. If we increase  $\bar{M}$ , we relax the formulation because more and more feasible solutions result in the same objective function value. In the extreme case, if  $\bar{M}$  is set to an arbitrarily large value, all feasible solutions will be optimal since they will all result in an objective function value of  $Z_{\text{MIP}}^* = 0$ . While increasing the value  $\bar{M}$  leads to faster solution times, it also degrades the quality of the optimal solutions.

If  $Z_{\text{MIP}}^* = 0$  then it means that the fixed maintenance capacity was never exceeded, but it does not indicate whether the best possible maintenance schedule was generated. As long as the maximum number of aircraft in maintenance at any given time over the planning horizon does not exceed  $\bar{M}$ , the resulting objective function value  $Z_{\text{MIP}}^* = 0$ . Notice that when  $Z_{\text{MIP}}^* = 0$ , the maximum number of aircraft that are in maintenance at any given time over the planning horizon is unknown. An  $Z_{\text{MIP}}^* = 0$  means that the fixed maintenance capacity was never exceeded, but it does not indicate whether the best possible maintenance schedule was generated. For example, suppose  $\bar{M} = 5$  and the PM-MIP generates a maintenance schedule that enters a maximum 4 aircraft into phase maintenance simultaneously. Then the resulting  $Z_{\text{MIP}}^* = 0$ , but it is possible that there exists a feasible solution that would

require fewer than 4 aircraft to be in maintenance at any given time and would result in the same  $Z_{\text{MIP}}^* = 0$ .

On the other hand, if we set  $\bar{M} = 0$ , we are guaranteed an optimal solution that leads to the best possible maintenance schedule, that is, a feasible schedule that results in the fewest aircraft that can be in maintenance at any given time. Let  $Z_{\text{MIP}}^{\bar{M}=0*}$  represent the optimal objective function value when  $\bar{M} = 0$ . Then  $Z_{\text{MIP}}^{\bar{M}=0*}$  is the maximum number of planes in phase maintenance at any given time over the planning horizon and we know that no other feasible maintenance schedule can achieve fewer than  $Z_{\text{MIP}}^{\bar{M}=0*}$  planes in maintenance at any given time. However, if we always set  $\bar{M} = 0$  to ensure we achieve the best possible maintenance schedule, we do not take advantage of the computational benefits of increasing the value of  $\bar{M}$ . Therefore, we seek to find highest value for  $\bar{M}$  that still guarantees the best possible maintenance schedule.

Let  $Z_{\text{LP}}^{\bar{M}=0*}$  represent the optimal objective function value for the linear program (LP) relaxation of the PM-MIP with  $\bar{M} = 0$ . Since we know that the objective function value of the PM-MIP must be integer,  $Z_{\text{LP}}^{\bar{M}=0*}$  leads to a simple lower bound on  $Z_{\text{MIP}}^{\bar{M}=0*}$ ,

$$\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil \leq Z_{\text{MIP}}^{\bar{M}=0*}.$$

Since  $Z_{\text{MIP}}^{\bar{M}=0*}$  represents the maximum number of planes in phase maintenance, we know that at least  $\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil$  aircraft will be in maintenance simultaneously at some point in the planning horizon. Then we can set  $\bar{M} = \lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil$  and solve the PM-MIP knowing that the resulting maintenance schedule will be the best possible maintenance schedule. Previously, we presented an example where if  $Z_{\text{MIP}}^* = 0$  for some arbitrary value of  $\bar{M}$  ( $\bar{M} = 5$ ), we were unsure whether or not the resulting maintenance scheduled was the best possible schedule. However, if we set  $\bar{M} = \lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil$ , we know that

$$Z_{\text{MIP}}^{\bar{M}=\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil*} + \bar{M} = Z_{\text{MIP}}^{\bar{M}=0*},$$

where  $Z_{\text{MIP}}^{\bar{M}=\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil*}$  is the optimal objective function of the PM-MIP with  $Z_{\text{MIP}}^{\bar{M}=\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil*}$ .

Therefore, if  $Z_{\text{MIP}}^{\bar{M}=\lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil*} = 0$ , we precisely know that there is a maximum of  $\bar{M}$  number of aircraft in maintenance over the planning horizon.

In summary, by setting the parameter  $\bar{M} = \lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil$  and then solving the PM-MIP, we take advantage of the computational benefits of increasing the value of  $\bar{M}$  while guaranteeing that the resulting maintenance schedule is the best possible schedule.

3.2.2.1. *Computational Results of setting  $\bar{M} = \lceil Z_{\text{LP}}^{\bar{M}=0*} \rceil$ .* To understand the computational benefits of setting  $\bar{M}$  using the LP relaxation of the PM-MIP formulation, we compare

the solve times for multiple small instances of the phase maintenance model in which  $\overline{M} = 0$  and then  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ . We solve for a planning horizon of approximately six months with a fleet of ten aircraft. There are two sortie types and each aircraft must undergo phase maintenance every 100 hours. The sortie requirements are variable over time but follow a realistic pattern in that sortie requirements are greater during the beginning of the week compared to the end of the week. The end of horizon targets are evenly distributed on the interval  $[0, 100]$ . The number of phase hours remaining for each aircraft at the start of the planning horizon is randomly generated for each run using a discrete uniform distribution from  $[0, 100]$ . Holding all data constant, we solve for the case where  $\overline{M} = 0$  and  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ .

The PM-MIP was written in AMPL and solved using CPLEX 11 on a machine with an Intel Xeon 2.80GHz dual core processor and 8gb of RAM.

In approximately 90% of random instances, the solution time for the formulation in which  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$  is lower than the solution time for the formulation in which  $\overline{M} = 0$  (with the same set of data used in both cases). On average, the solution time for the phase maintenance MIP formulation is reduced by 60% when setting  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ . While the adjustment to the formulation improves computational performance in the vast majority of instances, there remains cases where setting  $\overline{M} = 0$  results in a better solution time due to the branch-and-bound algorithm used by commercial solvers such as CPLEX.

In total, 60 runs<sup>1</sup> were completed and the solve times are shown in Figure 3.2.1. In 52 of 59 cases, by setting  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ , the solve time for a given instance of the PM-MIP is reduced. On average, the solve times are reduced by 60%. Although setting  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$  leads to reduced solve times in most instances, there are instances when the PM-MIP with  $\overline{M} = 0$  actually solves faster. This is due to the branch and bound algorithm used by the commercial solver CPLEX and how it balances breadth versus depth in implementing the algorithm. In practice, the PM-MIP formulations with  $\overline{M} = 0$  and  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$  could be solved in parallel and the optimal solution of whichever solves first could be used as the optimal solutions will be equivalent. Also shown in Figure 3.2.2 is the corresponding value of  $Z_{\text{LP}}^{\overline{M}=0*}$  for each run. In all runs,  $Z_{\text{LP}}^{\overline{M}=0*} < 1$  and the improved formulation was solved with  $\overline{M} = 1$ .

Given the computational benefits of setting  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ , we will assume for here on out that any instance of the phase maintenance MIP is solved after setting  $\overline{M} = \left\lceil Z_{\text{LP}}^{\overline{M}=0*} \right\rceil$ .

<sup>1</sup>The results of only 59 runs are shown. The solver was constrained to a maximum solve time of 1200 seconds and therefore any run in which the solve time for both formulations reached 1200 seconds was discarded.

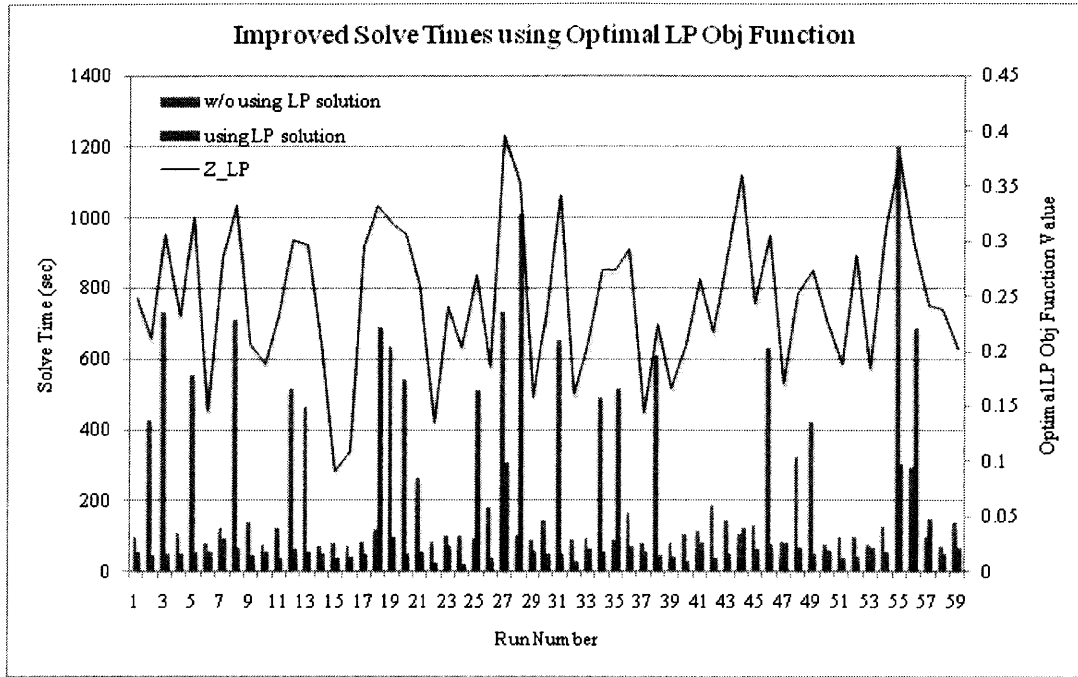


FIGURE 3.2.1. Comparison of solve times for setting  $\bar{M} = 0$  and  $\bar{M} = \left[ Z_{LP}^{\bar{M}=0*} \right]$ , represented by the left and right columns, respectively, for 60 random instances.

### 3.3. Implementation and Practical Considerations

In this section we implement and present the results on the MIP formulation introduced in the previous section for a realistic problem instance drawn from a long-term phase maintenance schedule for a single fighter squadron.

The length of the planning horizon is one year, with a single time period representing half a day. This results in  $T = 520$  since the model does not include weekends. There are two unique sortie types,  $|J| = 2$ . The first sortie type represents an aircraft flying a single sortie in a period with an *average sortie duration (ASD)* of 1.4 hrs. The second sortie type represents an aircraft flying two sorties in a single period (i.e., two sorties in a morning, by “pitting” in between the two) and has an ASD of 2.4 hrs. We assume that an aircraft must undergo a phase maintenance inspection every 300 flight hours and that a phase maintenance inspection requires two weeks, 20 periods, of aircraft downtime. The phase hours remaining on each of the aircraft at the beginning of planning horizon are randomly assigned and are uniformly distributed from 0 to 300. As illustrated in Figure 2.3.3, the end of horizon targets are evenly spaced from 0 to 300 so as to ensure that phase hours are robustly distributed amongst all aircraft. The sortie requirements are input directly from an operations schedule that specifies the number of required sorties in each period. The sortie requirements data



PROBLEM INSTANCE DATA		
Data Description	Symbol	Value
Length of Planning Period	$T$	520 periods
Number of Aircraft	$I$	15
Number of Sortie Types	$J$	2
Flying Hours between Phase Maintenance Inspections	$\bar{h}$	300 hrs
Downtime Associated with Phase Maintenance	$k$	20 periods
Length of Specific Sortie Types	$l_1, l_2$	1.4, 2.4 hrs
Phase Hours Remaining on each Aircraft $i \in I$ at $t = 1$	$b_i$	$U(0, h)$
End of Horizon Phase Hours Remaining Target for $q \in Q$	$e_q$	evenly spaced on $[0, h]$
Number of Sorties of type $j \in J$ required in period $t$	$s_j^t$	varies

TABLE 1. Example Instance Data

for this instance also meet the following scheduling criteria which are standard for most Air Force fighter units:

- no sorties are scheduled on weekends,
- sortie types requiring an aircraft to “pit” are scheduled only for mornings, i.e., an 8 pit 8 would not be allowed in the afternoon,
- the number of afternoon sorties is always less than or equal to the number of morning sorties,
- generally, more sorties are scheduled earlier in the week as opposed to later in the week to accommodate aircraft failures.

Table 1 summarizes all of the data used for this instance.

Using the data show in Table 1, we first solve the LP relaxation of the PM-MIP to find  $Z_{LP}^{\bar{M}=0*}$ . The LP relaxation solves almost instantly (0.01 seconds) and  $Z_{LP}^{\bar{M}=0*} = 0.246$ . We then set  $\bar{M} = \lceil Z_{LP}^{\bar{M}=0*} \rceil = 1$  and solve the original PM-MIP. With approximately 30k variables and 350k constraints, the PM-MIP solves in approximately 3 hours. With  $\bar{M} = 1$ , the optimal objective function value is  $Z_{MIP}^{\bar{M}=1*} = 0$ , meaning that no more than one aircraft is ever in maintenance at any given time.

Figure 3.3.1 illustrates the optimal solution. Each line in the plot represents an aircraft and the phase hours remaining on that aircraft. When an aircraft undergoes phase maintenance, the aircraft’s phase hours remaining is set to  $\bar{h} = 300$  hrs. Therefore, the vertical lines in the plot represent an aircraft entering maintenance and its phase hours being reset to 300 hrs. Recall that the downtime for phase maintenance was 20 periods in this instance. We see that the optimal solution distributes the phase maintenance workload so that an aircraft enters phase maintenance approximately every 20 periods.

**3.3.1. Limitations of the PM-MIP Formulation.** The objective function in the PM-MIP minimizes the maximum number of aircraft in maintenance at any given time. From

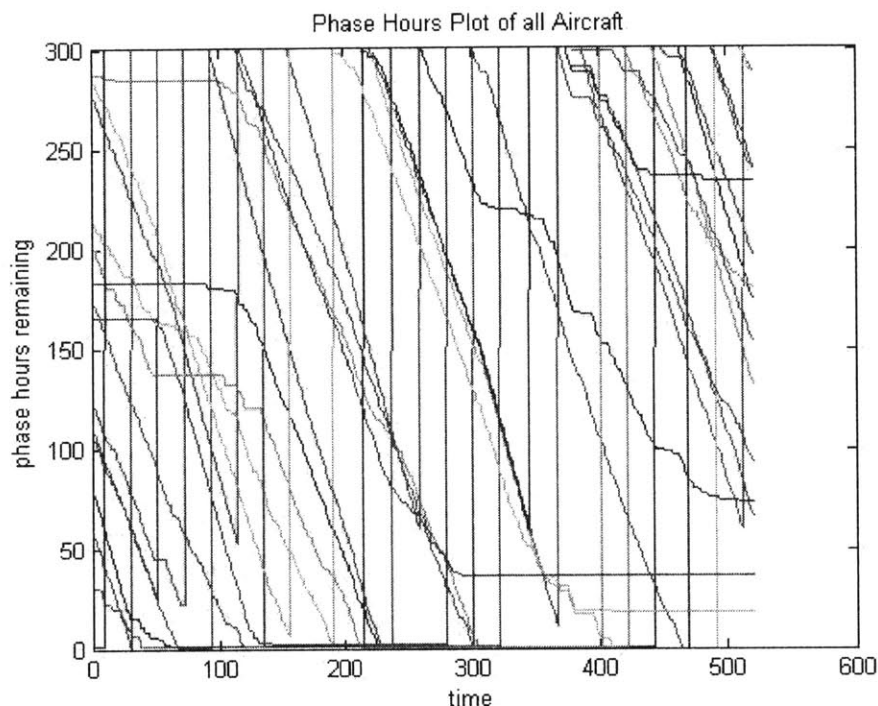


FIGURE 3.3.1. Optimal solution shown as a plot where each line represents an aircraft's phase hours over the planning horizon. The jumps in phase hours correspond to an aircraft entering maintenance.

an operational standpoint, the min-max objective function leads to a flying and maintenance schedule that attempts to smooth the maintenance workload evenly over the time horizon. Reducing the variance in the maintenance workload with respect to time is an important consideration due to the fact that most Air Force maintenance units are undermanned which make sudden increases in demand extremely difficult to manage. Since sortie requirements for a fighter unit can change dramatically over a given time horizon due to large-scale exercises and/or deployments, the utilization rates of aircraft can vary over time. This makes it difficult to schedule phase maintenance and keep phase hours evenly spaced for a fleet of aircraft. In order to accommodate these sudden, if brief, increases in sortie requirements, schedulers must plan well in advance and in many cases they might be unable to predict the consequences of flying and maintenance decisions made early in the planning horizon.

While the objective introduced with the PM-MIP formulation has many benefits, it does have its weaknesses. Since the objective function is minimizing the maximum number of aircraft in maintenance at any given time, the model might produce a solution that enters more aircraft into maintenance than necessary. There is no penalty for entering an aircraft into maintenance as long as the maximum number of aircraft in maintenance at

any give time remains unchanged. With this objective function, a schedule in which a single aircraft enters maintenance over the planning period is equally optimal to a schedule in which aircraft enter maintenance immediately one after another over the entire planning period. Alternative objective functions, such as minimizing the total number of maintenance actions over the planning period, could be implemented to avoid this shortcoming. However, given the important benefits of minimizing the variance in the maintenance workload, the objective function 3.2.1 was determined to be the most reasonable.

In addition to the limitations of the objective function, the PM-MIP has several other limitations. The sortie requirements constraints,

$$\sum_i x_{ij}^t \geq s_j^t, \quad \forall t, j \in J$$

allow assigning more aircraft to fly then absolutely necessary. That is, the sortie requirements are treated as the minimum sorties required and they can be exceeded by any amount. In operational settings, this is not allowed. The sortie requirements in the Air Force are carefully crafted to precisely meet the FHP requirements. By allowing the model to overfly the sortie requirements, the model has the potential to greatly overfly the FHP which is unacceptable. However, the nature of the objective function incentives minimal phase maintenance and therefore the model generally seeks to conserve phase hours. Therefore, although the model does allow for some overflying, the amount of overflying is limited by the objective function.

In terms of computational behavior, the PM-MIP formulation is not ideal. Although the instance used in the previous section solved in a reasonable amount of time, 3 hours, more realistic, detailed data sets result in solution times on the order of days. In formulations where there are three or more sortie types and a highly constrained number of aircraft, the solution times for the PM-MIP increase significantly. Furthermore, other applications of the PM-MIP formulation outside of specifically fighter aircraft phase maintenance might have much tighter data sets. When applied to the fighter aircraft phase maintenance scheduling problem, there is generally a high ratio of active aircraft to aircraft that are in maintenance which makes the problem easier to solve. In other applications, this may not be case. Therefore, we seek to develop an alternate formulation of the PM-MIP in the following chapter that can be applied in more constrained settings.

Lastly, the PM-MIP is limited in that it models only phase maintenance. A typical MDS has tens of preventative maintenance actions that require the plane to be unavailable for a day or longer. If we also consider the maintenance actions that do not necessarily require scheduled downtime, each MDS has hundreds of preventative maintenance requirements. While phase is the most significant maintenance inspection since it requires extended downtime, the minor maintenance actions also play a crucial role in aircraft availability and

sortie generation. Therefore, the MIP solution provides a strong foundation on which to build a more extensive maintenance schedule even if it is limited in that it cannot produce a comprehensive maintenance schedule.

### 3.4. Special Cases

Under a few key assumptions and relaxations, the MIP formulation for phase maintenance scheduling can be solved optimally using a heuristic and a simple network formulation. Specifically, we assume a single sortie type and constant sortie requirements, and we relax the upper bound on the end of horizon target constraint. We first explain the assumptions and relaxations that allow us to simplify the original MIP formulation. We also discuss the justifications and limitations of the simplifying assumptions and relaxations. We then present a heuristic that permits us to quickly determine the optimal values of the maintenance scheduling and end of horizon target variables,  $m_i^t$  and  $v_{iq}$ , respectively, if a feasible solution exists. We then present a framework under which the phase maintenance scheduling problem can be reformulated as a network formulation. The structure of the network formulation is determined using the  $m_i^t$  and  $v_{iq}$  values from the heuristic. The network formulation will determine if a feasible set of values for  $x_{ij}^t$  variables exist. By using the heuristic and network formulation, the optimal solution to the relaxed PM-MIP formulation can be found.

The purpose of the relaxed PM-MIP is very different than that of the original PM-MIP. Recall that the PM-MIP generated a daily flying and maintenance schedule. The relaxed PM-MIP models the phase maintenance scheduling problem from a more aggregate level and it is better used as a strategic planning tool than as a tool for generating a detailed daily schedule. By making several assumptions and relaxations to the PM-MIP, we lose the ability to generate a useful daily schedule. However, we can still capture the impact of major strategic decisions pertaining to items such as personnel, number of aircraft, number of sorties, and average sortie duration. Since the relaxed PM-MIP can be solved for a long planning horizon, it can be used to evaluate the long term effects and trade offs of making various strategic decisions.

**3.4.1. Assumptions and Relaxations to the MIP Formulation.** The assumptions are as follows:

- (1) there is only a single sortie type,  $|J| = 1$ ,
- (2) sortie requirements are constant in each period  $t = 1, \dots, T$ , that is,  $s = s_j^t = s_j^{t+1} = \dots = s_j^T$ .

These assumptions imply a simplified representation of the phase maintenance scheduling problem. The assumptions are reasonable for gaining basic insights into the general problem.

We assume there exists a single sortie type has a duration that is approximately the average sortie duration (ASD) of all sortie types. In operational settings, phase maintenance

for fighter aircraft is required every several hundred flying hours, with the exact requirement differing for each unique MDS. This equates to long intervals between phase maintenance actions that are usually more than a year. During that time period, it is reasonable to assume that each individual fighter aircraft will fly a similar mix of sortie types. There is no reason to believe that one particular aircraft will fly significantly more sorties of a given type than another aircraft. Given that all aircraft fly a similar mix of sortie types, the average duration of all the sorties by each aircraft will be relatively close. Hence, we can make a reasonable assumption that there is only a single sortie type that has a duration that is approximated using the average of all sortie types and thereby simplify the PM-MIP.

The assumption that the sortie requirements remain constant across all time periods is also justified by the relatively infrequent nature of phase maintenance. Fighter aircraft undergo phase maintenance every several hundred flying hours, during which the importance of minor fluctuations in sortie requirements across time periods is minimized. Although sortie requirements might actually vary by time period, we can assume a constant sortie requirement that is estimated as the average of the actual sortie requirements. By assuming a constant sortie requirement that is the average of the actual sortie requirements, the total number of required sorties flown by an aircraft between phase maintenance actions remains largely unchanged. Therefore, even with the constant sortie assumption, the model accurately captures the total number of sorties an aircraft will fly between consecutive phase maintenance inspections. However, the constant sortie assumption does degrade the model in that the flying decisions made in each time period are no longer particularly useful. The purpose of having the  $x_{ij}^t$  variables was to allow the model to assign a specific aircraft to fulfill each individual sortie requirement. With constant sortie requirements, the optimal set of flying decisions do not represent an implementable flying schedule as is the case without the assumption.

Under these two assumptions, the exact flying schedule generated from solving the model is not particularly important, but the relaxed model is still valuable for evaluating policy trade offs. Since these assumptions allow for the implementation of an efficient algorithm, decision makers can quickly evaluate general trade offs using the relaxed model. For example, the relaxed model can be used to understand the impact of increasing the size of the fleet of aircraft, changing the phase maintenance interval, or increasing maintenance resources. Although the relaxed model will not lead to a flying schedule that is directly implementable, it will provide good, general insights regarding phase maintenance scheduling.

In addition to the two assumptions, we slightly relax the constraints associated to the end of horizon targets. Recall that the MIP formulation assigns an end of the horizon flying hours target to each aircraft. The formulation then requires the flying hours remaining on

each aircraft at the end of planning horizon to be within 10% of the assigned target. The following constraints used to enforce these end of horizon targets:

$$\begin{aligned} v_{iq} * 0.9e_q &\leq h_i^T & \forall i, j \\ v_{iq} * 1.1e_q + \bar{h}(1 - v_{iq}) &\geq h_i^T & \forall i, j. \end{aligned}$$

In these constraints  $v_{iq}$  represents the binary decision variable that assigns aircraft  $i$  to end of horizon target  $q$  and the variable  $h_i^T$  represents the flying hours remaining on aircraft  $i$  in the last time period  $T$ . The parameter  $e_q$  represents the the flying hours target associated with end of horizon target  $q$  and the parameter  $\bar{h}$  represents the maximum flying hours between phase maintenance.

The purpose of these constraints in the MIP formulation was to ensure that the flying hours remaining for each aircraft would be well spaced and prevent a scenario in which a large number of aircraft would simultaneously come due for phase maintenance at some point beyond the end of the planning horizon. For this special case, we relax the upper limit on the end of horizon targets and only constrain the flying hours with a lower limit. Each aircraft is still assigned to an end of horizon target, but rather than being within  $\pm 10\%$  of the assigned target, the flying hours remaining on each aircraft must simply exceed the assigned target. The relaxation results in the elimination of the second of the prior constraints (3.2.20) and slightly alters the first constraint (3.2.13) to become

$$v_{iq} * e_q \leq h_i^T \quad \forall i, j$$

The flying hours will no longer be guaranteed to be well spaced at the end of the horizon. However, since we relax only the upper limit on the end of horizon targets, the relaxation of these constraints will, at the extreme, result in a situation in which the flying hours remaining on each aircraft at the end of the planning horizon far exceeds the assigned target. Since the purpose of the original constraints in the MIP formulation is to add a measure of robustness and to reduce the likelihood of a sudden increase in phase maintenance beyond the end of the planning horizon, the optimal solution with the relaxed constraints will also provide sufficient flying hours to easily manage future phase maintenance actions.

Lastly, we eliminate the constraints in the PM-MIP that require the phase hours remaining on an aircraft to be below a certain threshold before it can be entered into maintenance. In operational settings, the constraints

$$\bar{h} - h_i^t \geq (\bar{h} - h_{\max} * \bar{h})m_i^t \quad \forall t, i$$

are enforced to prevent the aircraft from entering phase maintenance excessively early. In reality, entering an aircraft into phase maintenance incurs a cost in terms of manpower,

parts, possibility of error when disassembling/reassembling the aircraft, etc. But the most significant cost comes from the downtime of the aircraft. The relaxed formulation of the PM-MIP still accounts for the cost of downtime of the aircraft and the need to satisfy all sortie requirements, but we simplify the PM-MIP by eliminating this set of constraints.

Notice that with the single sortie type assumption, the phase hours remaining on each aircraft can be used to calculate exactly how many sorties an aircraft can fly before it must enter phase maintenance. Since there is only a single sortie type, keeping track of phase hours is equivalent to keeping track of the number of sorties remaining until phase maintenance. Therefore, each aircraft can be thought of as having a “supply” of sorties in each time period that is used to satisfy the “demand” generated by the sortie requirements in each time period. Similarly, the end of horizon constraints can be thought of as a demand for sorties in the final time period. For the remainder of this section, we will describe the state of an aircraft in terms of sorties remaining rather than phase hours remaining. As a result, the variables  $h_i^t$  and parameters  $e_q$  will be assumed to be measured in terms of sorties remaining rather than flying hours while  $s$  still denotes the number of sorties required in each time period.

**3.4.2. An Optimal Maintenance Scheduling Policy.** With the assumptions and relaxations presented in the previous section, we show that for a fixed objective function value,  $Z_{\text{MIP}}$ , the  $m_i^t$  and  $v_{iq}$  variables can be optimally determined using a simple greedy heuristic leaving only the  $x_i^t$  variables to be solved for. In the next section we use the results of the heuristic values of the  $m_i^t$  and  $v_{iq}$  variables to generate a network flow formulation that is then used to solve for a feasible set of  $x_i^t$  variables. Since we solve the heuristic and network formulation for a fixed value of  $Z_{\text{MIP}}$ , we then search for the minimum  $Z_{\text{MIP}}$  that has a feasible solution.

By entering aircraft into phase maintenance immediately after another aircraft exits phase maintenance, the total number of phase maintenance actions completed over the planning horizon can be maximized without increasing the maximum number of aircraft in maintenance at any given time. In addition, by always entering the aircraft with the fewest sorties remaining into maintenance, we maximize the total supply of sorties available over the planning horizon. Based on these two principles, for a given  $Z_{\text{MIP}}^{\bar{M}=0}$  value, we present a heuristic that determines the optimal values of the  $m_i^t$  and  $v_{iq}$  variables. If a feasible set of values for the  $x_i^t$  variables does not exist with the heuristic determined the values of the  $m_i^t$  and  $v_{iq}$  variables, then a feasible set of values for the  $x_i^t$  variables does not exist for any other values of the  $m_i^t$  and  $v_{iq}$  variables.

It is important to note that we choose the value of the objective function,  $Z_{\text{MIP}}^{\bar{M}=0}$ , a priori and then look for a feasible solution (hence  $Z_{\text{MIP}}^{\bar{M}=0}$  not  $Z_{\text{MIP}}^{\bar{M}=0*}$ ). Although it might seem unreasonable to search for a feasible solution for a fixed objective function value, notice that in almost any realistic instance of the phase maintenance scheduling problem, the objective

function value will be limited to a small set of possible values. The objective function value,  $Z_{\text{MIP}}^{\bar{M}=0}$ , must be integer and non negative. In addition,  $Z_{\text{MIP}}^{\bar{M}=0}$  cannot exceed the total number of aircraft,  $|I|$ . Therefore, if we can efficiently determine whether a feasible solution exists for a given  $Z_{\text{MIP}}^{\bar{M}=0}$ , an optimal solution can be found efficiently by beginning with  $Z_{\text{MIP}}^{\bar{M}=0} = 0$  and solving for each increasing value of  $Z_{\text{MIP}}^{\bar{M}=0}$ . The full implementation of this algorithm is presented in Section 3.4.4, but we mention it here to motivate the need for a heuristic that can be used to determine the values of the  $m_i^t$  and  $v_{iq}$  variables for a fixed  $Z_{\text{MIP}}^{\bar{M}=0}$ .

Notice that  $Z_{\text{MIP}}^{\bar{M}=0}$  is the maximum number of aircraft in maintenance at any given time in the planning. If we can optimally solve the PM-MIP where  $\bar{M} = 0$ , instances where  $\bar{M} > 0$  are trivial since the optimal solution to the PM-MIP with  $\bar{M} = 0$  will also be optimal for any other value of  $\bar{M}$ . To simplify the notation for the remainder of the section, let

$$u = Z_{\text{MIP}}^{\bar{M}=0}.$$

The heuristic for determining the values of  $m_i^t$  and  $v_{iq}$  for a given  $u$  is as follows:

- (1) Index the aircraft so that the lowest indexed aircraft has the fewest sorties remaining at  $t = 1$  and the highest index aircraft has the most sorties remaining.  $h_1^1 \leq h_2^1 \leq \dots \leq h_{|I|}^1$ .
- (2) Index the end of horizon targets so that the lowest index target is the target requiring the fewest sorties remaining,  $e_1 \leq e_2 \leq \dots \leq e_{|Q|}$ .
- (3) If  $u = 0$ , set all  $m_i^t = 0$  and assign aircraft to the end of horizon targets so that  $v_{ii} = 1, \forall i \in I$ . That is, the aircraft with the fewest sorties remaining at the start of the horizon is assigned to the lowest end of horizon target, the aircraft with the second fewest sorties remaining at the start of the horizon is assigned to the second lowest end of horizon target, and so on.
- (4) If  $u > 0$ ,
  - (a) Set  $m_i^t$  the variables by beginning at  $t = 1$ . Assign the  $u$  aircraft with the fewest sorties remaining at the start of the planning horizon,  $i = 1, \dots, u$ , to enter phase maintenance so  $m_1^1, \dots, m_u^1 = 1$ . Continue by assigning the next  $u$  aircraft with the fewest sorties remaining,  $i = u + 1, \dots, 2u$ , to enter phase maintenance in  $t = 1 + k$ . Continue for time periods  $t = 1 + 2k, 1 + 3k, \dots, 1 + \lfloor (T - 1)/k \rfloor k$ , where  $\lfloor (T - 1)/k \rfloor + 1$  represents the number of consecutive phase maintenance periods that can be completed within a planning horizon of length  $T$ . After all aircraft,  $i \in I$ , have been assigned to one phase maintenance period, continue the assignment process with  $i = 1$ . That is, continuously cycle through the aircraft until each  $k^{\text{th}}$  time period has  $u$  aircraft assigned to enter maintenance. This ensures that  $u$  aircraft always enter maintenance as soon as the previous  $u$  aircraft exit maintenance.



- (b) Set the  $v_{iq}$  variables by beginning with the last aircraft to exit phase maintenance,  $i = \{i \in I; m_i^{1+\lfloor(T-1)/k\rfloor k} = 1\}$ . Assign aircraft  $i$  to end of horizon target  $e_Q$  by setting  $v_{iQ} = 1$ . From aircraft  $i$ , we cycle through all the aircraft in reverse order of the indices,  $i = j, j-1, \dots, 1, |I|, |I|-1, \dots, j+1$ , and assign each aircraft to the highest unassigned end of horizon target from  $q = e_Q, e_{Q-1}, \dots, e_1$ . By assigning end of horizon targets in this manner, we are treating each end of horizon target as if it is another maintenance “slot”. The maintenance slots are assigned so as to maximize the total supply of sorties over the planning horizon, just as is the case setting the  $m_i^t$  variables.

Notice that the objective of the heuristic, for a given value of  $u$ , is to complete as much phase maintenance as possible during a fixed planning horizon and thereby maximize the supply of sorties over the planning horizon. We now give an example of the implementation of the heuristic in a small instance.

*Example 3.4.1:*

*Assume we have three aircraft,  $|I| = 3$ , and the sorties remaining on each aircraft at the beginning of the planning horizon are 5, 10, and 15 sorties..*

$$h_1^1 = 5, h_2^1 = 10, h_3^1 = 15$$

*The downtime associated with a phase maintenance action is 3 time periods so  $k = 3$ . The length of the planning horizon is 6 time periods so  $T = 6$ . There are three end of horizon targets with sorties remaining targets of 10, 20, and 30 sorties.*

$$e_1 = 10, e_2 = 20, e_3 = 30$$

*We implement the phase maintenance scheduling heuristic for the case when  $u = 1$ . Therefore, at most one aircraft can be in maintenance at any given time in the planning horizon.*

*We begin by assigning the first  $u$  aircraft with the fewest sorties remaining at the start of the planning horizon to enter maintenance at  $t = 1$ .*

$$m_1^1 = 1$$

*We continue to cycle through the all the aircraft in order of their indices and assign  $u$  aircraft to enter maintenance every  $k^{\text{th}}$  time period until  $t = 1 + \lfloor(T-1)/k\rfloor k = 4$ .*

$$m_2^4 = 1$$

After we have assigned  $u$  aircraft to enter maintenance in every  $k^{\text{th}}$  time period, we assign each aircraft to an end of horizon target. We assign aircraft to end of horizon targets by cycling through all the aircraft in the order in which they exit phase maintenance. Aircraft 2 is the last aircraft to exit phase maintenance in the planning horizon and is therefore assigned to the end of horizon target requiring the most sorties. Aircraft 1 is the second to last aircraft to exit maintenance and is assigned to the second highest end of horizon target. Aircraft 3 never enters maintenance and is therefore assigned to the lowest end of horizon target.

$$v_{12} = 1, v_{23} = 1, v_{31} = 1$$

We set all remaining  $m_i^t$  and  $v_{iq}$  variables that have not yet been fixed equal to zero. This means that the values for all  $m_i^t$  and  $v_{iq}$  variables have been determined and only the  $x_{ij}^t$  variables remain.

We claim that by using this heuristic to determine the values of the  $m_i^t$  and  $v_{iq}$  variables, the feasibility of the problem remains unchanged. That is, for a given objective function value, a feasible set of values for the  $x_{ij}^t$  variables exists if and only if a feasible set of values for the  $x_{ij}^t$  variables exists under the heuristic determined value of the  $m_i^t$  and  $v_{iq}$  variables. We prove this by induction.

LEMMA 1. *For a fixed objective function value to the PM-MIP, a feasible set of values for the  $x_{ij}^t$  variables exists if and only if a feasible set of values for the  $x_{ij}^t$  variables exists under the heuristic determined value of the  $m_i^t$  and  $v_{iq}$  variables.*

PROOF. Case 1:  $u = 0$

When  $u = 0$ , all aircraft must reach the end of the planning horizon without ever entering into maintenance. Hence,

$$m_i^t = 0 \quad \forall i, t$$

The heuristic is also used to assign each aircraft,  $i \in I$ , to an end of horizon target,  $q \in Q$ , by assigning the aircraft with the fewest sorties remaining at the start of the horizon to the lowest end of horizon target, the aircraft with the second fewest sorties remaining at the start of the horizon to the second lowest end of horizon target, and so on. It is obvious that by fixing the values of the  $v_{iq}$  variables in this manner, we maximize the total supply of sorties over the planning horizon.

In order for there to exist a feasible set of values for  $x_i^t$  variables, it must be feasible to satisfy all sortie requirements,  $s$ . All sortie requirements can be satisfied if in each time period there are

1. at least  $s$  aircraft that are not in maintenance and,
2.  $s$  aircraft with at least one sortie remaining. In this case  $u = 0$ , so there will always be  $|I|$  available aircraft in every period.

If  $|I| > s$ , then condition 1 is met. With regards to condition 2, by assigning the end of horizon targets using the heuristic, not only do we maximize the total supply of sorties over the planning horizon, we also maximize the distribution of sorties amongst the fleet of aircraft. That is, if any aircraft were assigned to a feasible end of horizon target different than the one prescribed by the heuristic, the total supply of sorties remains unchanged, but the absolute difference between aircraft's sorties supplies must increase.

*Example 3.4.2:*

Assume  $u = 0$  and we have two aircraft with starting states  $h_1^1 = 8$  and  $h_2^1 = 10$ . There are two end of horizon targets  $e_1 = 3$  and  $e_2 = 6$ . According to the heuristic,  $v_{11} = 1$  and  $v_{22} = 1$ . Then aircraft 1 has a supply of  $h_1^1 - e_1 = 5$  sorties and aircraft 2 has a supply of 4 sorties, a difference in supplies of 1 sortie. It is feasible to change the end of horizon target assignments. Let us now assume we have  $v_{12} = 1$  and  $v_{21} = 1$ . Then aircraft 1 has 2 sorties available and aircraft 2 has 7 sorties available, a difference in supplies on 5 sorties. Regardless of how the end of horizon targets are assigned, there is a total supply of 9 sorties over the planning horizon. However, the distribution of sorties between the aircraft is different. Suppose that there are three time periods with a requirement for 2 sorties in each period, resulting in a total demand of 6 sorties over the planning horizon. These sortie requirements could be met if heuristic values of the  $v_{iq}$  were used, but not the alternate values of  $v_{iq}$  were used.

Since the heuristic maximizes the distribution of sorties amongst aircraft, any other set of values of the  $v_{iq}$  variables would result an increased likelihood that condition 2 would not be met. That is, there will never exist a case where, by setting the  $v_{iq}$  variables in a manner different from the heuristic, feasibility can be gained. Hence, in the case where  $u = 0$ , a feasible set of values for the  $x_i^t$  variables exists if and only if a feasible set of values for the  $x_i^t$  variables exists under the heuristic determined value of the  $m_i^t$  and  $v_{iq}$  variables.

Case 2:  $u = 1$

Just as in the case when  $u = 0$ , we claim when  $u = 1$  a feasible set of values for the  $x_i^t$  variables exists if and only if a feasible set of values for the  $x_i^t$  variables exists under the heuristic determined value of the  $m_i^t$  and  $v_{iq}$  variables. When  $u = 1$ , there is at most one aircraft in maintenance at any given time in the planning horizon. According to the heuristic,

one aircraft will always be in maintenance since they are assigned to enter maintenance one after another.

Unlike the case  $u = 0$  where  $m_i^t = 0 \forall i, t$ , when  $u = 1$  we must consider whether the heuristic values of  $m_i^t$  eliminate any feasible sets of values for  $x_i^t$ . In other words, if it is infeasible to meet the sortie requirements under the heuristic values of  $m_i^t$ , we must consider whether it is possible to meet the sortie requirements under some other values of  $m_i^t$ . There are two possible reasons why the sortie requirements might not be feasible for given set of values for the  $m_i^t$  and  $v_{iq}$  variables. The first possibility is that there are not enough aircraft that are not in maintenance in a given time period. Under the heuristic values of the  $m_i^t$  variables, there are always  $|I| - 1$  aircraft that are not in maintenance. If the sortie requirement,  $s$ , is greater than  $|I| - u$ , the sortie requirements are obviously infeasible in every time period. More generally, for any set of  $m_i^t$  variables, there will at least be one point in time in the planning horizon when one aircraft will be in maintenance. Since sortie requirements are constant, if  $s > |I| - 1$ , the sortie requirements will be infeasible under any values of the  $m_i^t$  variables.

The second case under which the sortie requirements may be infeasible in a given time period is if there are not  $s$  aircraft with a supply of at least one sortie. The heuristic requires that one aircraft be in maintenance at all times. Each time an aircraft is in maintenance, the remaining  $|I| - 1$  aircraft must be able to meet all the sortie requirements. The total demand for sorties during those  $k$  periods is  $k * s$ . However, suppose that the remaining  $|I| - 1$  could not meet the  $s$  sortie requirements in each of the  $k$  periods during which one aircraft was down for maintenance. Then by deviating from the heuristic and not entering an aircraft into maintenance as soon as another one exits, there would temporarily be  $|I|$  aircraft available and an increase in the supply of sorties. In the short term then, it appears that the heuristic values for the  $m_i^t$  variables may not always be optimal. However, even if a maintenance action is delayed, eventually one aircraft will have to enter maintenance. When an aircraft enters delayed maintenance, the remaining  $|I| - 1$  aircraft will have to be able to support the sortie requirements for  $k$  periods. Since sortie requirements are constant, if it was previously infeasible to meet the sortie requirements with  $|I| - 1$  aircraft, it will still be infeasible after delaying the maintenance action. Therefore, ultimately, by altering the heuristic values of the  $m_i^t$  variables and delaying maintenance actions, feasibility cannot be gained.

The arguments presented for the cases when  $u = 0$  and  $u = 1$  hold true for all values of  $u$ . □

### 3.4.3. Network Formulation of the Phase Maintenance Scheduling Problem.

Given the assumptions and relaxations outlined in Section 3.4.1, the phase maintenance scheduling problem can be modeled for a particular value of  $u$  as a network formulation .

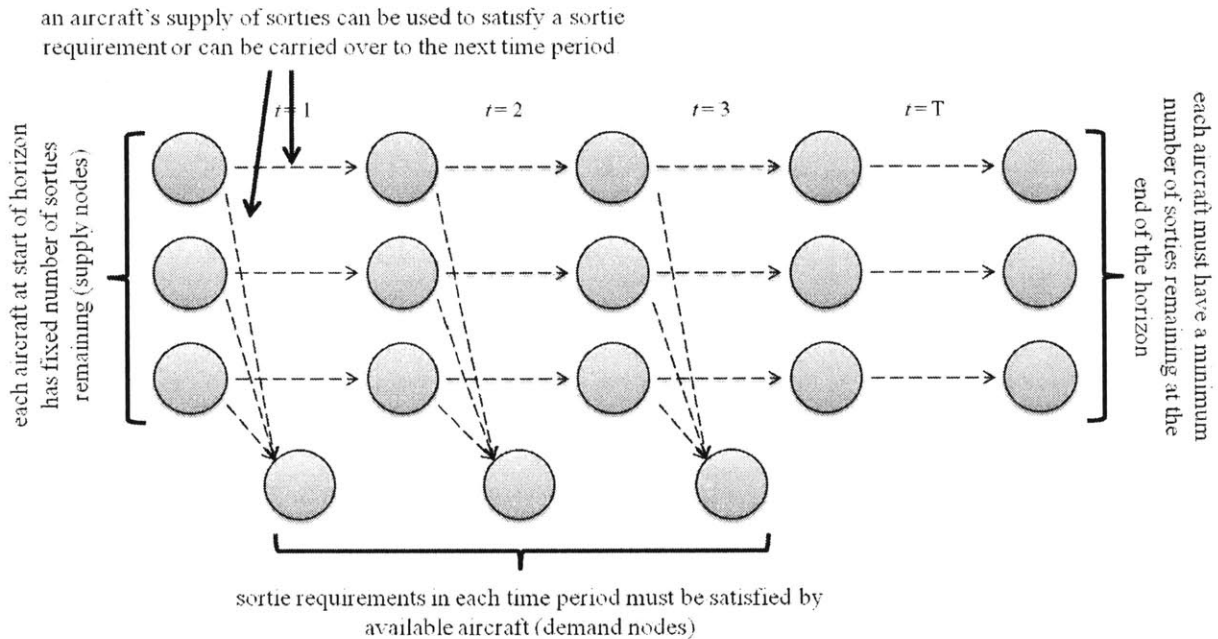


FIGURE 3.4.1. Conceptual framework for network reformulation

Using the heuristic values for the  $m_i^t$  and  $v_{iq}$  variables to define the network structure, we can solve the network formulation to conduct a feasibility check for the  $x_i^t$  variables. The reformulation of the problem hinges on the assumption that there is only a single sortie type. As noted before, under the single sortie type assumption, we only keep track of the sorties remaining on each aircraft rather than the actual phase hours remaining. The sorties remaining on each aircraft can then be thought of as a supply of sorties while the sortie requirements and end of horizon targets are demands for sorties. Figure 3.4.1 illustrates the general framework for the resulting network reformulation.

We can more concisely define the network structure using the heuristic values of the  $m_i^t$  and  $v_{iq}$  variables. We initially build a network in which for every time period,  $t = 1, 2, \dots, T$ , the network has a node representing each aircraft and a node representing the sortie requirements. The values of the  $m_i^t$  and  $v_{iq}$  variables are then used to define the arcs within the network and can also be used to determine the supply or demand at each node. For example, if we know that an aircraft is in maintenance in a given time period, then we can eliminate the arc connecting the aircraft to the sortie requirements in that time period. Also, if we know when an aircraft is entering maintenance, we know when it will receive an additional “supply” of sorties. We know precisely the total supply of sorties within the network since it is simply the sum of the sorties remaining at the start of the horizon and the sorties gained through maintenance over the planning horizon, or

$$\text{total supply of sorties} = \sum_i h_i^1 + \bar{h}u(\lfloor (T-1)/k \rfloor + 1),$$

where  $(\lfloor (T-1)/k \rfloor + 1) \cdot u$  is the total number of phase maintenance actions that are completed over the planning horizon. Similarly, we know the sortie demand over the planning horizon from the sortie requirements and end of horizon targets:

$$\text{total demand for sorties} = sT + \sum_q e_q.$$

Since the total supply of sorties must equal the total demand for sorties within the network, we use a sink node as an additional node to capture the difference between the supply and demand of sorties. Recall that in the PM-MIP, aircraft could be entered into phase maintenance even though it had up to  $h_{\max}$  phase hours remaining on the aircraft. Essentially, the model was allowed to “throw away” phase hours. The sink node in the network formulation serves this purpose; it is a way to capture the extra supply of sorties in the model that are not used to meet sortie requirements or end of horizon targets.

Based on these principles, we present the following phase maintenance network formulation, (PM-N):

#### Sets.

$N$	set of all nodes
$I$	set of all aircraft sortie inventory nodes, $i \in I$
$D$	set of all sortie demand nodes, $d \in D$
$A$	set of all arcs
$A_{ID}$	set of all arcs $(i, j)$ with $i \in I$ and $j \in D$
$A_{IM}$	set of all arcs that lead out of a node in which an aircraft entered maintenance.

For the network formulation we define two primary sets of nodes. The first set of nodes,  $I$ , represent the supply of sorties remaining on each aircraft in each time period. We refer to these nodes as the inventory nodes since they represent the inventory of sorties that are available to meet the sortie requirements demands. The second set of nodes  $D$ , represent the sortie demand nodes in each time period.

The arcs within the network are determined using the heuristic values of the  $m_i^t$  and  $v_{iq}$  variables. We define two primary sets of arcs. The set  $A_{ID}$  contains all arcs that originate at an inventory node  $i \in I$  and terminate at a sortie demand node  $d \in D$ . If an aircraft is in maintenance in a given time period, there will be no arcs from that aircraft’s inventory nodes to the demand nodes in those periods. Otherwise, each aircraft will have an arc from it’s inventory node to the demand node in that time period.

Based on the heuristic values of the  $m_i^t$  variables, we know which nodes in the network formulation are the inventory nodes associated with an aircraft entering phase maintenance.

At these nodes, we know that an aircraft receives a “supply” of  $\bar{h}$  from having undergone phase maintenance. However, when an aircraft returns to service from phase maintenance, it will have  $\bar{h}$  sorties remaining regardless of the number of sorties it had remaining when it entered maintenance. That is, whether an aircraft entered phase with few or many sorties remaining, it will return with  $\bar{h}$  sorties remaining. To capture this condition in the network formulation, it is necessary to distinguish the set of arcs that originate at the inventory nodes associated with an aircraft entering phase maintenance. We let the set  $A_{IM}$  contain all arcs that originate at an inventory node,  $i \in I$ , that is associated with an aircraft entering phase maintenance. Note that each of these maintenance inventory nodes will only have one outgoing arc since an outgoing arc to the demand node is nonexistent. The arcs in the set  $A_{IM}$  have a capacity constraint that requires the flow coming out of the inventory node to equal  $\bar{h}$ . If the aircraft had any sorties remaining when it entered phase maintenance, they will flow to the sink node that captures all supplies of sorties that are not used to meet sortie requirements or end of horizon targets.

#### Parameters.

$b_i$  supply ( $b_i > 0$ ) or demand ( $b_i < 0$ ) at node  $i \in N$

The parameter  $b_i$  defines the supply or demand at every node  $i \in N$ . At the start of the planning horizon, each aircraft has an initial supply of sorties remaining. In addition, each time an aircraft undergoes maintenance, the it receives an additional supply  $\bar{h}$  sorties. For the nodes  $i \in D$ , the demand is  $b_i = -s$ . For the end of the horizon target nodes, the demand is  $b_i = -e_q$ . The demand of the sink node is the difference between the supply and demand for sorties from all other nodes, or

$$b_{\text{sink}} = [s * T + \sum_q e_q] - [\sum_i h_i^1 + \bar{h} * (\lfloor (T - 1)/k \rfloor + 1) * u].$$

#### Variables.

$f_{ij}$  number of sorties that flow from node  $i \in N$  to node  $j \in N$

**Formulation.**  $Z_{PM-N} = \text{minimize } 0$

subject to:

$$(3.4.1) \quad b_i + \sum_{(j,i) \in A} f_{jij} = \sum_{(i,j) \in A} f_{ij} \quad \forall i \in N$$

$$(3.4.2) \quad f_{ij} \leq 1 \quad \forall (i,j) \in A_{ID}$$

$$(3.4.3) \quad f_{ij} = \bar{h} \quad \forall (i,j) \in A_{IM}.$$

The overall objective of this network formulation is to ascertain if there are feasible values for the  $x_i^t$  variables for a given value of  $u$ . Any feasible set of values for the  $f_{ij}$  variables in

the network formulation can easily be translated to equivalent values for the  $x_i^t$  variables. If the network formulation is infeasible, we can conclude that it is infeasible to meet the sortie requirements with at most  $u$  aircraft in maintenance at any given time.

The first set of constraints, (3.4.1), are basic flow balance constraints. They require that the flow out of each node is equal to the flow into the node plus any supply or demand that may exist at that node.

The second set of constraints, (3.4.2), ensure that in each time period, an aircraft can only fly one sortie. These constraints are consistent with those in the PM-MIP formulation where each aircraft was only allowed to fly one sortie per time period.

The third set of constraints, (3.4.3), are necessary to set the number of sorties an aircraft has when coming out of maintenance. If an aircraft enters maintenance with a positive number of sorties remaining, these sorties do not carryover when the aircraft exits maintenance. Rather, regardless of how many sorties remaining an aircraft had prior to entering maintenance, the aircraft will have  $\bar{h}$  sorties remaining when it exits maintenance. In order to enforce this, we require that the flow out of a maintenance node be equal to  $\bar{h}$  while any additional sorties are forced to flow to the sink node.

*Example 3.4.3:*

*We illustrate the network formulation for the same data used in Example 3.4.2 where  $u = 1$ . The data is summarized below*

$$\begin{aligned} h_1^1 &= 5, & h_2^1 &= 10, & h_3^1 &= 15 \\ e_1 &= 10, & e_2 &= 20, & e_3 &= 30 \\ m_1^1 &= 1, & m_2^4 &= 1 \\ v_{12} &= 1, & v_{23} &= 1, & v_{31} &= 1 \end{aligned}$$

*In addition, let  $\bar{h} = 30$  and  $s = 1$ . We can then illustrate the network formulation as shown in Figure 3.4.2. Aircraft 1 enters maintenance at  $t = 1$  and so there is no arc from the inventory of aircraft 1 to the demand at  $t = 1$ . The same is true for aircraft 1 at  $t = 4$ . In addition, at the nodes that represent an aircraft entering maintenance, there is a supply of  $\bar{h}$  sorties. The arcs from the last set of inventory nodes at  $t = 6$  to the end of horizon targets are defined by that values of the  $v_{iq}$  variables.*

*The supply or demand at each node is shown in the figure. In this case, we know that total supply of sorties in the network is 90, 30 from the initial states of the aircraft and an additional 60 from maintenance actions. The total demand is 66, 6 from the 6 periods of sortie requirements and 60 from*



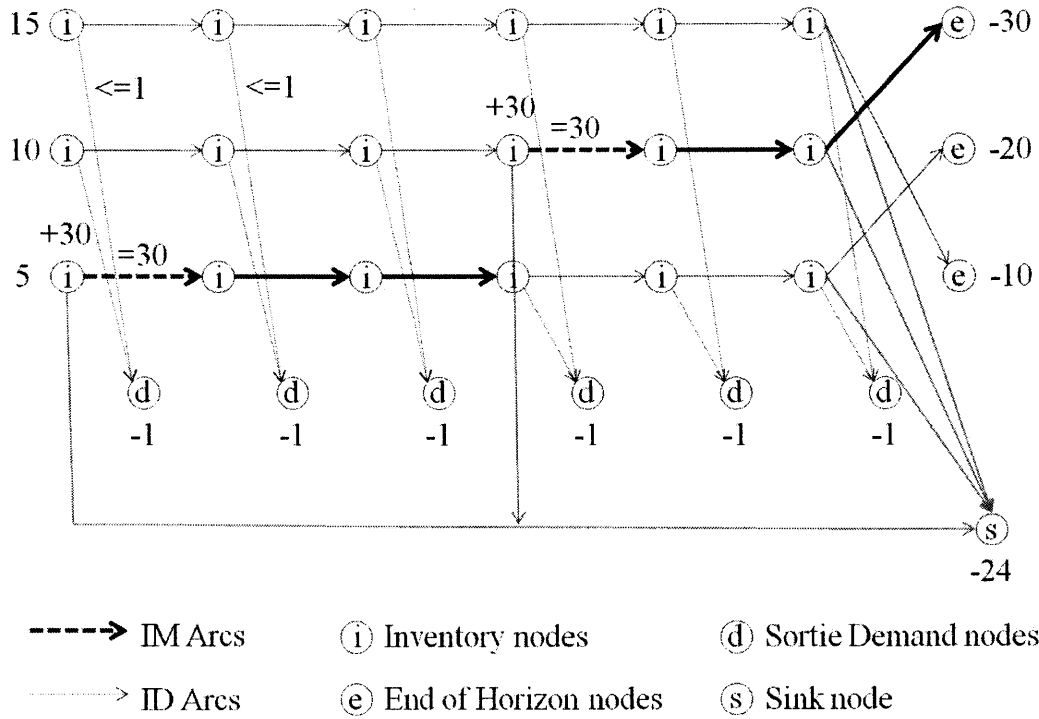


FIGURE 3.4.2. Example Network Formulation Instance

*the end of horizon targets. Therefore, the demand at the sink node must account for the difference between the supply and demand in the network and is therefore  $66 - 90 = -24$ .*

From a computational perspective, the network formulation performs extremely well. Even for realistic instances of the problem that have close to 20,000 arcs and 10,000 nodes, the formulation solves instantly. This is a significant improvement from the solves times of the PM-MIP that was on the order of several hours.

**3.4.4. Algorithm for Optimally Solving the Special Case of the Phase Maintenance Scheduling Problem.** In Section 3.4.1 we outlined several simplifying assumptions and relaxations to the original PM-MIP formulation. Then in Section 3.4.2, we presented a heuristic that would lead to the optimal values of the  $m_i^t$  and  $v_{iq}$  variables for a given value of  $u$ . Then in Section 3.4.3 we presented a network reformulation, PM-N, of the special case of the PM-MIP that incorporates the heuristic values of the  $m_i^t$  and  $v_{iq}$  variables and can be used to efficiently check the feasibility of the special case PM-MIP for a given value of  $u$ .

By using the heuristic and network reformulation, we can optimally solve the special case of the PM-MIP. The algorithm is as follows.

- (1) Set  $u = 0$ .

- (2) Implement the heuristic to fix the values of the  $m_i^t$  and  $v_{iq}$  variables.
- (3) Using the heuristic values of the  $m_i^t$  and  $v_{iq}$  variables, reformulate the special case PM-MIP as a PM-N network formulation. Solve the PM-N formulation. If a feasible solution exists, the optimal objective function for the special case PM-MIP is  $u$  and the optimal solution consists of the  $m_i^t$  and  $v_{iq}$  variables as determined by the heuristic and the  $x_i^t$  values that correspond to the feasible solution of the network formulation. If a feasible solution to the network formulation does not exist, proceed with the algorithm.
- (4) Set  $u = u + 1$ . If  $u = |I|$ , the problem is infeasible. If  $u < |I|$ , return to step 2.

Recall that  $u$  represents the maximum number of aircraft in maintenance at any given time in the planning horizon. Therefore, it is lower bounded by zero and upper bounded by the total number of aircraft. By starting the algorithm at  $u = 0$  and increasing in one unit increments, we check every possible value of the objective function of the special case PM-MIP. As soon as we find a  $u$  for which a feasible solution exists, we know the optimal solution to the special case of the PM-MIP.

### 3.5. Summary

In this chapter we presented a mathematical formulation for the scheduling of phase maintenance. The PM-MIP presented in this chapter uses a min-max objective function to minimize the variance in maintenance demand while also minimizing the overall amount of phase maintenance. While the formulation does not provide a comprehensive preventative maintenance schedule since it only considers phase maintenance, it provides a strong foundation on which to build more complex maintenance scheduling models. From a computational standpoint, PM-MIP is not ideal. In realistic instances that include multiple sortie types and limited numbers of available aircraft, the solve times can be on the order of a day or longer. Hence, we seek to develop an alternate formulation of the PM-MIP that lends better computational behavior and provides an implementable flying schedule.

In addition to presenting the primary PM-MIP formulation, we presented a special case of the PM-MIP that can be solved optimally using a heuristic and simple network formulation. Under the primary assumptions that there is only a single sortie type and constant sortie requirements, we show that a greedy heuristic can optimally determine the values of the  $m_i^t$  and  $v_{iq}$  variables. We then show that for a given objective function value, the heuristic results can be used to reformulate the PM-MIP as a network formulation, PM-N. The PM-N formulation efficiently solves for a feasible set of  $x_i^t$  variables based on a predetermined set of maintenance actions and end of horizon target assignments. By enumerating over possible values of the PM-MIP objective function, we can easily solve the special of the PM-MIP to optimality.

## Two-Stage MIP for Phase Maintenance Scheduling

In this chapter, we present alternative methods of modeling the phase maintenance scheduling process. In Chapter 3, we introduced the PM-MIP formulation and discussed the less than ideal computational behavior of the formulation. Motivated by the need for a model with better computational behavior, we disaggregate the PM-MIP into two interrelated subproblems that are solved sequentially. The disaggregation leads to far better computational behavior and makes it tractable to solve for expanded data sets that were previously intractable. Although the disaggregation is suboptimal relative to the PM-MIP, it lends a good solution that justifies the trade off between optimality and tractability.

We begin by presenting the framework under which we disaggregate the PM-MIP. The subproblems are then developed, discussed, and compared to the PM-MIP. Lastly, we implement the subproblems and compare the solutions with those from the PM-MIP.

### 4.1. Framework for Disaggregation of Phase Maintenance Scheduling

The overall purpose of the PM-MIP was to optimally balance sortie requirements and phase maintenance requirements. The PM-MIP formulation made flying and maintenance decisions for each time period over the entire planning horizon. Therefore, the optimal solution to the PM-MIP provided a highly detailed schedule that specified precisely when each aircraft would enter phase maintenance and on what days each aircraft would fly. However, since the model simultaneously considered every time period over the planning horizon, the PM-MIP was intractable for expanded data sets.

In order to improve the computational tractability of the phase maintenance scheduling problem, we seek to disaggregate the problem into subproblems that are more tractable. Recall that in the PM-MIP, the flying and maintenance decisions were made simultaneously and were represented by the  $x_{ij}^t$  and  $m_i^t$ , variables, respectively. We now consider the possibility of making the flying and maintenance decisions sequentially rather than simultaneously. That is, we first determine the maintenance schedule and then solve for a flying schedule assuming that the maintenance schedule is fixed, or vice versa. By considering these sets of decisions separately in sequence, we substantially reduce the number of decision variables that need to be simultaneously considered and therefore improve the tractability of the formulation.

In the disaggregation of the PM-MIP, we first solve for the maintenance schedule and then solve for the daily flying schedule. This disaggregation hinges on the granularity of the

time periods in the planning horizon for each of the two subproblems. When solving the first subproblem for the maintenance schedule, we let each time period represent the duration of a single phase maintenance inspection. In practice, maintenance schedulers almost always enter aircraft into phase maintenance as soon as the previous aircraft exits phase maintenance. This leads to the concept of non overlapping, sequential phase maintenance time slots. By treating each of these slots as a single time period when solving for a maintenance schedule, each time an aircraft is assigned to enter maintenance it has a downtime of one period. This reduction in the time granularity of the planning horizon results in far fewer variables and constraints.

After the maintenance schedule has been fixed after solving the first subproblem, the second subproblem is solved to determine the daily flying schedule. The second subproblem has a planning horizon that is the duration of a single phase maintenance inspection and each a time period represents a 1/2 day. Therefore, the second subproblem must be solved for every time period of the first subproblem. For each instance of the second problem, the solution from the first subproblem provides the data regarding the state of each aircraft which is then used to assign individual aircraft to sortie requirements. Due to the short planning horizon of the second subproblem, expanded data sets that previously made the PM-MIP intractable have marginal impact on the tractability of the model.

Ultimately, the disaggregation of PM-MIP uses two subproblems that have shortened planning horizons in order to generate a long-term daily flying and maintenance schedule. In the next section, we present and develop the formulation for the two subproblems in detail and discuss further the interactions between the them.

#### 4.2. Formulations of Subproblems

In this section we present two subproblems, one that generates a maintenance schedule and one that generates a flying schedule, which we will refer to as the M-Sub and F-Sub. The M-Sub is solved first. Each time period in the M-Sub represents the downtime associated with one phase maintenance inspection. Therefore, when an aircraft is assigned to phase maintenance in the M-Sub, it is unavailable to fly for one time period. Unlike the PM-MIP, the M-Sub does not assign aircraft to individual sortie requirements. Rather, the M-Sub decides how many hours each aircraft will be flown in each time period. By assigning flying hours to aircraft in each time period rather than assigning individual sorties, the M-Sub can still keep track of the phase hours remaining for each aircraft and decide when to enter each aircraft into phase maintenance. Ultimately, a solution to the M-Sub provides a maintenance schedule and it specifies how many hours each aircraft can fly in each time period.

After the M-Sub has been solved, the F-Sub is solved for each time period of the M-Sub. That is, the planning horizon for the F-Sub is equivalent to the duration of the downtime

associated with one phase maintenance inspection. Since the maintenance schedule is determined using the M-Sub, the F-Sub does not consider any decisions pertaining to maintenance scheduling. Rather, the F-Sub focuses on assigning individual aircraft to sortie requirements while ensuring that each aircraft does not fly more hours than was assigned by the M-Sub. The F-Sub generates a daily flying schedule for each period of the M-Sub. By joining the optimal solutions of the M-Sub and F-Sub, a long-term daily schedule, similar to the one generated by the PM-MIP, can be produced.

**4.2.1. M-Sub Ingredients and Formulation.** As discussed earlier, maintenance slots are generally non overlapping in practice. This allows us to segment the planning horizon into time periods that are each the length of one phase maintenance inspection since we are primarily concerned with generating a maintenance schedule with the M-Sub. In the PM-MIP, each time period was  $1/2$  a day and a phase maintenance inspection required a downtime of  $k$   $1/2$  day periods. In the M-Sub, each time period is equivalent to the downtime associated with one maintenance inspection, or  $k/2$  days. Also, in the PM-MIP, the model precisely tracked the phase hours remaining for each aircraft since it could assign aircraft to individual sorties of known duration. In the M-Sub, the model cannot consider daily sortie requirements so it must consider the sortie requirements over every  $k/2$  days in aggregate. Using the known daily sortie requirements, we can determine the total number of flying hours that must be completed over every  $k/2$  days. The M-Sub is then required to assign flying hours to each aircraft for each time period so that there are enough cumulative assigned flying hours to meet the aggregate sortie requirements. By assigning flying hours in each time period, the M-Sub can keep track of the phase hours remaining on each aircraft which is necessary to generate a maintenance schedule since phase maintenance is solely based on accrued flying hours. Ultimately, the M-Sub generates a phase maintenance schedule and flying hours assignments for each time period.

**M-Sub Formulation.** We begin by defining all the relevant sets to the M-Sub. Just as in the PM-MIP, we let  $I$  denote the set of all aircraft and  $|I|$  is the total number of aircraft. We let  $Q$  denote the set of all end of horizon targets. Also as in the PM-MIP, we assume that each aircraft is assigned to a unique end of horizon target so  $|Q| = |I|$ .

The sets in the M-Sub formulation are summarized below.

**Sets.**

- $I$             set of all aircraft  $i \in I$
- $Q$             set of all end of horizon targets  $q \in Q$

Next, we present the data associated with the M-Sub. Much of the data is identical to the data used in the PM-MIP. In the formulation,  $T$  is the length of the planning horizon, with each time period  $t = 1, 2, \dots, T$ , representing the length of a phase maintenance inspection.

Since flying operations and phase maintenance are generally restricted to weekdays, we assume that weekends are not included in the planning horizon.

The parameter  $\bar{h}$  represents the phase maintenance interval, that is, the maximum flying hours an aircraft can accrue between phase maintenance actions. In operational settings, to prevent aircraft from entering phase maintenance excessively early, maintenance schedulers can only schedule an aircraft for phase maintenance if its phase hours remaining is below a certain limit,  $h_{max}$ . This parameter  $h_{max}$  represents the maximum number of phase hours an aircraft can have when entered into phase maintenance.

The parameters  $b_i$  and  $e_q$  describe the states of aircraft at the beginning and end of the planning horizon. At the start of the planning horizon, each aircraft,  $i \in I$ , has a given number of phase hours,  $b_i$ , remaining. At the end of the planning horizon, it is necessary to ensure a distribution of phase hours among the aircraft so that a large number of aircraft will not later come due for phase maintenance at the same time. The parameters  $e_q$  are used to define the end of horizon targets. For each target  $q \in Q$ ,  $e_q$  is the target number of flying hours for any plane assigned to that target at the end of the horizon.

The parameter  $\bar{M}$  is a fixed maintenance capacity that can be used without cost or penalty. It has the same properties as the  $\bar{M}$  in the PM-MIP which were discussed in Section 3.2.2.

The parameter  $F_t$  is unique to the M-Sub and was not present in the PM-MIP.  $F_t$  represents the aggregation of the sortie requirements in period  $t$ . Specifically, it is the sum of the total flying hours necessary to complete all the sortie requirements over a given  $k/2$  days, where  $k$  is the number of  $1/2$  day periods necessary to complete a single phase maintenance inspection. Since the daily sortie requirements provide the number of sorties required in each  $1/2$  period as well as the duration of those sortie,  $F_t$  can easily be calculated.

The parameter  $f_{max}^t$  is also unique to the M-Sub and is crucial to the disaggregation of the PM-MIP. The parameter  $f_{max}^t$  is the maximum number of flying hours each aircraft can fly in a single period  $t$ . In each time period, the M-Sub decides how many flying hours each aircraft will fly in order to meet the sortie requirements. However, since the parameter  $F_t$  only provides information regarding the total flying hours that are required, the M-Sub has no information regarding how many aircraft are needed on each day. In an extreme case, the M-Sub could assign a single aircraft to fly all  $F_t$  hours in a time period even though there must actually be multiple aircraft available per day. In order to force the M-Sub to distribute the  $F_t$  required flying hours amongst multiple aircraft, we use the parameter  $f_{max}^t$ . We determine the value of  $f_{max}^t$  in the following manner,

$$f_{max}^t = \frac{F_t}{s_t},$$

where  $s_t$  is the maximum number of aircraft required in period  $t$  of the M-Sub. For example, suppose each period of the M-Sub is equivalent to 2 weeks. During the first two weeks of the planning horizon, on one day, 8 sorties are required in the morning and that is the most sorties required in any 1/2 day of the first two weeks. Then  $s_t = 8$ . By setting  $f_{\max}^t$  in this manner, we are guaranteed to have at least  $s_t$  aircraft that are assigned to fly in a given time period.

It may also be desirable to increase  $s_t$  by multiplying by some factor  $\alpha > 1$  so that the  $F_t$  hours are guaranteed to be distributed across more than  $s_t$  aircraft. This could be used to add a measure of robustness and protect against the possibility of random aircraft failures. For the implementation of the M-Sub in this section, we do not use an  $\alpha$  multiplier.

### Data.

$T$	length of time horizon where $t = 1, 2, \dots, T$
$\bar{h}$	maximum accrued flying hours between phase maintenance inspections
$h_{\max}$	maximum number of phase hours an aircraft can have remaining for it to be entered into maintenance
$b_i$	phase hours remaining on aircraft $i \in I$ at the beginning of the horizon, $t = 1$
$e_q$	end of horizon flying hours target for aircraft assigned to target $q \in Q$
$\bar{M}$	maximum number of aircraft that can be in maintenance at any given time without incurring a "cost" or "penalty"
$F_t$	flying hours required in period $t$ to meet all sortie requirements
$f_{\max}^t$	maximum flying hours that an aircraft can fly in a single period

The decision variables in the M-Sub are largely the same as those in the PM-MIP. The variable  $m_i^t$  represents the binary decision to enter aircraft  $i$  into phase maintenance at time  $t$ . The variable  $v_{iq}$  represents the binary decision to assign aircraft  $i$  to end of horizon target  $q$ . The variable  $h_i^t$  is used to keep track of the phase hours remaining for each aircraft  $i$  in period  $t$ .

The only variable that is unique to the M-Sub, is  $f_i^t$ . The variable  $f_i^t$  represents the number of flying hours assigned to be flown by aircraft  $i$  in period  $t$ . Unlike the PM-MIP formulation that contained a binary variable that represented the decision to fly an aircraft or not, the M-Sub uses a continuous variable that represents the flying hours assigned to each aircraft in each period.

### Decision Variables.

$$m_i^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases}$$

$v_{iq} =$	$\begin{cases} 1, & \text{if aircraft } i \in I \text{ is assigned to end of horizon target } e_q \\ 0, & \text{otherwise} \end{cases}$
$h_i^t \geq 0$	phase hours remaining at the start of period $t$ until aircraft $i \in I$ must enter maintenance, "life remaining"
$f_i^t \geq 0$	the number of hours flown by aircraft $i \in I$ in time period $t$
$Z$	objective function value that is being minimized

**Formulation.**

$$(4.2.1) \quad \text{M-Sub} = \text{minimize } Z$$

s. t.

$$(4.2.2) \quad h_i^1 = b_i, \quad , \forall i \in I,$$

$$(4.2.3) \quad h_i^{t+1} \leq h_i^t - f_i^t + \bar{h}m_i^t, \quad , \forall t, i \in I,$$

$$(4.2.4) \quad h_i^{t+1} \geq h_i^t - f_i^t, \quad , \forall t, i \in I,$$

$$(4.2.5) \quad h_i^{t+1} \leq \bar{h}, \quad , \forall t, i \in I,$$

$$(4.2.6) \quad h_i^{t+1} \geq \bar{h}m_i^t, \quad , \forall t, i \in I,$$

$$(4.2.7) \quad \bar{h} - h_i^t \geq (\bar{h} - \bar{h}h_{\max})m_i^t, \quad , \forall t, i \in I,$$

$$(4.2.8) \quad \sum_i f_i^t \geq F_t, \quad , \forall t,$$

$$(4.2.9) \quad f_i^t \leq (1 - m_i^t)f_i^{\max}, \quad , \forall t, i \in I,$$

$$(4.2.10) \quad 0.9v_{iq}e_q \leq h_i^T, \quad , \forall i \in I, q \in Q,$$

$$(4.2.11) \quad 1.1v_{iq}e_q + \bar{h}(1 - v_{iq}) \geq h_i^T, \quad , \forall i \in I, q \in Q,$$

$$(4.2.12) \quad \sum_i v_{iq} = 1, \quad , \forall q \in Q,$$

$$(4.2.13) \quad \sum_q v_{iq} = 1, \quad , \forall i \in I,$$

$$(4.2.14) \quad \sum_i m_i^T - \bar{M} \leq Z, \quad , \forall t,$$

$$(4.2.15) \quad Z \geq 0,$$

$$(4.2.16) \quad h_i^t \geq 0, \quad , \forall i \in I, t,$$

$$(4.2.17) \quad f_i^t \geq 0, \quad , \forall i \in I, t,$$

$$(4.2.18) \quad m_i^t \in \{0, 1\}, \quad , \forall i \in I, t,$$

$$(4.2.19) \quad v_{iq} \in \{0, 1\}, \quad , \forall i \in I, q \in Q.$$



The formulation for the M-Sub shares many constraints with the PM-MIP. The constraints that are unique to the M-Sub formulation are constraints (4.2.8) and (4.2.9). In the M-Sub, we replace the sortie requirements constraints from the PM-MIP with constraints on the flying hours assigned to each aircraft in each time period, represented by  $f_i^t$ .

Constraint (4.2.8) requires that the sum of flying hours assigned to the entire fleet of aircraft in each period must exceed the total duration of all sorties required over the  $k/2$  days that are aggregated into each period of the M-Sub, which is  $F_t$ . Constraint (4.2.9) ensures that the  $F_t$  required flying hours are allocated amongst multiple aircraft. Based on the value of  $f_{\max}^t$ , the M-Sub is required to assign the required flying hours across some minimum number of aircraft. This ensures that in the F-Sub, there are enough aircraft available to meet the daily sortie requirements.

The detailed explanations for the remaining constraints can be found in Section (3.2.1).

**4.2.2. F-Sub Ingredients and Formulation.** After the M-Sub has been solved, we then solve the F-Sub. Each instance of the F-Sub considers a planning horizon that is equivalent to one period in the M-Sub. Recall that each period in the M-Sub was the duration of a single phase maintenance action or  $k/2$  days, where  $k$  is the number of  $1/2$  days periods required to complete a phase maintenance action. The F-Sub has a planning horizon of  $k$   $1/2$  day periods.

Over the planning horizon of  $k$  periods, the F-Sub seeks to assign aircraft to daily sortie requirements using the solution from the M-Sub as inputs. From the M-Sub solution, we know the number of flying hours that have been assigned to each aircraft for the planning horizon of a single instance of the F-Sub. The goal of the F-Sub is to assign aircraft to sortie requirements so that the total flying hours flown by each aircraft over the  $k$  periods in the F-Sub closely matches the flying hours assignments that were handed down by the M-Sub. In other words, the flying hours assigned in the M-Sub are treated as targets in the F-Sub. The F-Sub ensures that all sortie requirements are met while also seeking to meet the flying hours targets assigned from the M-Sub. Ultimately, the F-Sub generates a daily flying schedule for each time period in the M-Sub.

**F-Sub Formulation.** We begin with the sets relevant to the F-Sub formulation. As with the PM-MIP and M-Sub, we let  $I$  represent the set of all aircraft and  $|I|$  is the total number of aircraft. Since the M-Sub determines when each aircraft is assigned to phase maintenance, the set  $I$  in the F-Sub excludes all aircraft that are assigned to phase maintenance. We let  $J$  denote the set of all sortie types, just as in the PM-MIP.

**Sets.**

- $I$  set of all aircraft  $i \in I$
- $J$  set of all sortie types  $j \in J$

Next we present the data in the F-Sub. The length of the planning horizon is  $T$ , where each period represents a 1/2 day. The parameter  $f_i$  comes directly from the solution of the M-Sub and it represents the number of flying hours assigned to aircraft  $i$ . Since we are solving the F-Sub for each time period of the M-Sub, we drop the superscript the, variable  $f_i^t$  in the M-Sub becomes the parameter  $f_i$  in the F-Sub. The parameter  $l_j$  and  $s_j^t$  are the same in the F-Sub as they were in the PM-MIP.

**Data.**

$T$	length of time horizon where $t = 1, 2, \dots, T$
$f_i$	number of flying hours assigned to aircraft $i \in I$ (from the M-Sub solution)
$l_j$	length of sortie type $j \in J$
$s_j^t$	minimum number of sorties of type $j \in J$ required in period $t$

There are only two decision variables in the F-Sub. The variables  $x_{ij}^t$  is the binary decision to assign aircraft  $i$  to sortie type  $j$  in period  $t$ . It is the same variable that was found in the PM-MIP formulation. The second variable is  $z_i$ . It represents the absolute difference between the flying hours assigned by the M-Sub and the number of flying that were assigned to aircraft  $i$  in the F-Sub. The variable  $z_i$  is minimized in the objective function of the F-Sub.

**Decision Variables.**

$x_{ij}^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases}$	
$z_i$	difference between the F-Sub assigned flying hours and the M-Sub assigned flying hours for aircraft $i \in I$

**Formulation.**

$$(4.2.20) \quad \text{F-Sub} = \text{minimize } \sum_i z_i$$

s.t.

$$(4.2.21) \quad \sum_i x_{ij}^t \geq s_j^t, \quad \forall t, j \in J,$$

$$(4.2.22) \quad \sum_j x_{ij}^t \leq 1, \quad \forall t, i \in I,$$

$$(4.2.23) \quad \sum_j x_{ij}^t \geq \sum_j x_{ij}^{t+1}, \quad \forall \text{odd } t, i \in I,$$

$$(4.2.24) \quad \left| f_i - \sum_t \sum_j x_{ij}^t l_j \right| \leq z_i, \quad \forall i,$$

$$(4.2.25) \quad z_i \geq 0, \quad \forall i.$$

The objective function 4.2.20, in conjunction with constraints 4.2.24 and 4.2.15, minimizes the sum of the differences between the number of flying hours assigned to aircraft  $i$  in the F-Sub and the flying hours targets from the M-Sub. In constraint 4.2.24, the parameter  $f_i$  represents the flying hours that were assigned to aircraft  $i$  from the M-Sub. The F-Sub seeks to assign aircraft to the sortie requirements so that total flying hours assigned to each aircraft over the planning horizon is as close as possible to the targets from the M-Sub.

Constraints 4.2.21-4.2.23 were first presented with the PM-MIP and they serve the same purposes in the F-Sub. Constraint 4.2.21 ensures that the sortie requirements in each time period are satisfied. Constraint 4.2.22 ensures that each aircraft is only assigned to a single sortie in each time period. Constraint 4.2.23 requires that sorties in the afternoon (even numbered time periods) are only flown by aircraft that were also assigned to fly in the morning (odd number time periods).

Notice that the F-Sub does not track the number of flying hours in each time period and it does not make any maintenance decisions. The set  $I$  already excludes any aircraft that are assigned to phase maintenance and the maintenance schedule is fully determined by the M-Sub. This leads to a very simple MIP that needs only to consider one set of decision variables, the  $x_{ij}^t$  variables.

### 4.3. Discussion of the Disaggregation Models

We begin by summarizing the basic interaction between the two subproblems. The M-Sub is solved first and it generates a maintenance schedule and flying hours targets for each aircraft in each period. As discussed earlier, the M-Sub considers the sortie requirements in aggregate since each time period of the M-Sub is equivalent multiple periods in the PM-MIP. The M-Sub then assigns flying hours to aircraft in each period to meet the aggregated sortie requirements. This allows the M-Sub to keep track of the phase hours remaining for all the aircraft and it uses this information to schedule aircraft for maintenance.

An instance of F-Sub is solved for each period of the M-Sub and it uses the maintenance schedule and periodic flying hours targets from the M-Sub as data. Using the same sortie requirements that were used in the PM-MIP, the F-Sub assigns aircraft to individual sorties so that each aircraft's total assigned flying hours matches the flying hours targets from the M-Sub as closely as possible. Combined with the maintenance schedule generated by the M-Sub, the F-Sub solution leads to a long-term daily flying and maintenance schedule.

Notice that the maintenance schedule generated by the M-Sub relies on each aircraft flying the number of hours it was assigned in each period of the M-Sub. However, the F-Sub is not required to have every aircraft precisely fly the number of hours that were assigned to it by the M-Sub. Therefore, if the flying schedule from the F-Sub results in an aircraft deviating greatly from the M-Sub target, the entire maintenance schedule from the M-Sub

could be compromised. For example, suppose the M-Sub assigns an aircraft to fly 30 hours in a given period. Then an instance of the F-Sub is solved for that same period and the aircraft is assigned to fly 26 hrs. This could happen in any number of periods of the M-Sub which would mean the maintenance schedule is entering aircraft into phase maintenance long before necessary. In the context of the fighter phase maintenance scheduling problem, however, the relatively long interval between phase maintenance inspections ( $\sim 300$  hrs) and the much smaller sortie durations ( $\sim 1-2$  hrs) makes it uncommon for the F-Sub assignments to deviate much from the M-Sub targets (see Section 5.5 for empirical results). The potential for deviations in the two subproblems may be more of a concern if the two stage approach is applied to other problems.

Another important characteristic of the two stage approach is that the F-Sub will always generate a feasible flying schedule. Notice that in the F-Sub formulation, there are no constraints that involve the phase hours remaining for each aircraft. Therefore, an aircraft that was assigned zero flying hours by the M-Sub can still be assigned to a sortie. While this gives the model flexibility as it seeks to generate a flying schedule, it also leads to unnecessarily long solution times. Since any aircraft can be assigned to a sortie regardless of whether it was assigned flying hours by the M-Sub, the F-Sub must evaluate a large number of feasible solutions. Although the solution times can be very long, a near optimal solution is found within seconds as demonstrated in Section 5.5. The vast majority of the solution time leads to no change in the objective function value or very little improvement.

Just as in the PM-MIP, the two-stage model is still limited in its scope since it only applies to the scheduling of phase maintenance. However, given that the two stage model has far better computational behavior than the PM-MIP, it is possible to consider expanding the model to include various usage based preventative maintenance requirements. Previously, the PM-MIP was intractable for large data sets and when only considering phase maintenance. With the two stage approach, we could expand the model to consider multiple usage based preventative maintenance requirements that share maintenance resources.

#### 4.4. Implementation of the Two-Stage Model

In this section we implement the two-stage model for phase maintenance scheduling and present the results. We use the same sortie requirements and data that we used in the implementation of the PM-MIP in Section 3.3. Table 1 summarized the data that was previously presented in the implementation of the PM-MIP

To incorporate the same data into the M-Sub and F-Sub formulations, we must adapt the data. Table 1 shows the data used for the M-Sub and F-Sub. Notice that the planning horizons for the M-Sub is only 26 periods where as in the PM-MIP the horizon is 520 periods since one period of the M-Sub is equivalent to 20 periods in the PM-MIP. In addition, the

(A) Data for the M-Sub

M-SUB INSTANCE DATA		
Data Description	Symbol	Value
Length of planning horizon	$T$	26 periods
Number of AC	$I$	15
Flying hours required in period $t$	$F_t$	varies
Max flying hours assignable to each aircraft in period $t$	$f_{max}^t$	$F_t/s_t$
Flying hours between Phase MX Inspections	$\bar{h}$	300 hrs
Max phase hours remaining when AC is entered into MX	$h_{max}$	60 hrs
Phase hours remaining on each AC $i \in I$ at $t = 1$	$b_i$	$U(0, h)$
End of horizon phase hours remaining target for $q \in Q$	$e_q$	$[0, h]$

(B) Data for the F-Sub

F-SUB INSTANCE DATA		
Data Description	Symbol	Value
Length of planning horizon	$T$	20 periods
Number of aircraft	$I$	all not in phase ( $\leq 15$ )
Flying hours assigned to aircraft $i$ (from M-Sub)	$f_i$	varies
Length of Specific Sortie Types	$l_1, l_2$	1.4, 2.4 hrs
Number of sorties of type $j \in J$ required in period $t$	$s_j^t$	varies
Number of sortie types	$J$	2

TABLE 1. Data from PM-MIP implementation adapted for the two subproblems

F-Sub only considers a planning horizon of 20 1/2 day periods rather than the 530 1/2 day periods considered in the PM-MIP.

The two subproblems was written in AMPL and solved using CPLEX 11 on a machine with an Intel Xeon 2.80GHz dual core processor and 8gb of RAM.

The M-Sub solves instantly (0.01s) and generates a maintenance schedule and flying hours assignments. Figure 4.4.1 plots the phase hours remaining on each aircraft over the planning horizon. Compared to the plot generated by the PM-MIP (Figure 3.3.1), the M-Sub solution has noticeably better spacing between the fleet of aircraft throughout the planning horizon. This is in large part due to the parameter  $f_{max}^t$  which restricts the rate at which an aircraft can be flown and requires multiple aircraft to be flown. In addition, the M-Sub solution enters fewer aircraft into phase maintenance over the course of the planning horizon. This suggests that the M-Sub has far fewer issues with overflying the sortie requirements than the PM-MIP.

After solving the M-Sub, we then solve an instance of the F-Sub for each period in the M-Sub. For each of the 26 instances of the F-Sub, we use the flying hours assignments as input to the M-Sub. In addition, for each instance of the F-Sub we only consider the aircraft

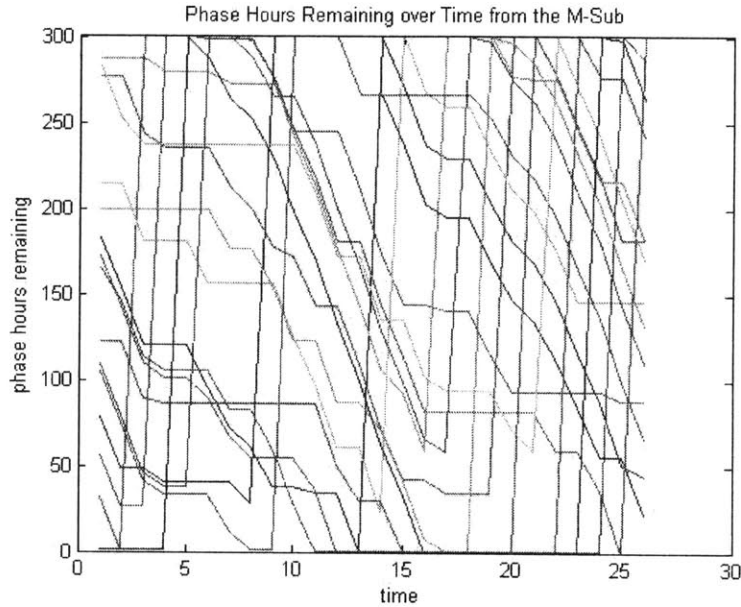


FIGURE 4.4.1. Phase hours plot from the optimal solution to the M-Sub

that were not entered into phase maintenance by the M-Sub. The sortie requirements that are used in the F-Sub are the same as those used in the PM-MIP.

Figure 4.4.2 shows an example of an optimal solution from one instance of the F-Sub. An 'x' denotes that an aircraft has been assigned to fly on a given day. The second and third row show the sortie requirements for the planning horizon. Recall that each period represents a 1/2 day period so the odd numbered periods are mornings and the even numbered periods are afternoons. There are a total of 132 sorties that are required over the planning horizon. The F-Sub slightly overflies the sortie requirements by assigning aircraft to 149 sorties. In practice, these additional sorties could without much consequence. The sum of the flying hours assigned to the fleet of aircraft by the M-Sub was 245.8 hours and the total duration of sorties assigned by the F-Sub is only 0.4 hours from that target. That is, the objective function value for this instance is 0.4 hours. Therefore, if the 17 additional sorties were ignored, the aircraft would have more phase hours remaining at the end of the 20 periods than the M-Sub expects, but it would not result in a significant deviation from the maintenance schedule when considering that the interval between phase maintenance actions is 300 hrs.

Figure 2 summarizes the results from each run of the F-Sub. The objective function represents the difference between the flying hours assigned by the M-Sub and the actual duration of sortie assignments from the F-Sub. As shown in the table, the F-Sub does lend very good solutions that come close to flying hours targets assigned by the M-Sub. Given that the interval between phase maintenance inspections is 300 hrs and the F-Sub is only

SORTIE REQUIREMENTS																				
period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
sortie type 1	8	6	0	8	0	0	0	0	7	6	8	6	0	8	0	0	0	0	8	6
sortie type 2	0	0	8	0	8	8	6	8	0	0	0	0	8	0	0	0	8	7	0	0
DAILY FLYING SCHEDULE																				
AC 1	x	x	x	x	x	x	x	x	x		x	x	x				x	x	x	
AC 2	x	x	x	x	x		x	x	x	x	x	x	x	x			x	x	x	x
AC 3	x	x	x	x	x	x	x	x					x	x	x		x	x	x	x
AC 4			x		x				x	x	x		x	x					x	
AC 5	x		x	x	x	x	x	x	x		x	x	x	x	x		x	x	x	x
AC 6	undergoing phase maintenance																			
AC 7																				
AC 8			x		x	x	x	x	x	x	x	x	x		x		x	x	x	x
AC 9																				
AC 10	x	x	x	x	x	x	x	x	x	x			x	x			x	x	x	
AC 11	x	x	x	x	x	x	x	x	x	x	x	x	x	x			x		x	x
AC 12																				
AC 13																				
AC 14	x	x	x	x		x	x		x	x	x	x	x	x			x	x	x	x
AC 15	x		x	x	x	x	x	x	x		x		x	x	x	x	x		x	

FIGURE 4.4.2. Optimal flying schedule from one instance of the F-Sub

deviating from the M-Sub targets by a few hours, the solutions are highly reasonable and the resulting flying schedule is implementable.

In terms of computational behavior, each instance of the F-Sub was given a time limit of 10 seconds to solve. Approximately a 1/3 of the instances solved in less than 10 seconds. The remaining instances all reached near optimal solutions in seconds and terminated after 10 seconds. In total, the two stage model produced a long-term flying and maintenance within minutes while the PM-MIP required several hours.

#### 4.5. Summary

In this chapter, we presented a disaggregation of the PM-MIP that leads to two subproblems that are solved sequentially to generate a long-term flying schedule. The disaggregation results in significant improvements in computational behavior. This allows for solving instances of the phase maintenance scheduling problem with expanded data set that were previously intractable with the PM-MIP. Although the gains in computational behavior require a trade off with optimality, the results from the disaggregation of the PM-MIP are still highly implementable. Therefore, it is justifiable to use the two-stage model in generating a long-term flying schedule.

Period 1/2 days	Obj. Fn. $\sum_i z_i$	Total Hours $\sum_i f_i$	% Diff
1	0.4	245.8	0.16%
2	1.6	265.6	0.60%
3	3.4	58.8	5.78%
4	0	0	0.00%
5	1	75.8	1.32%
6	1.4	165.6	0.85%
7	3	85.4	3.51%
8	1	166.6	0.60%
9	1.2	214.2	0.56%
10	0	207.2	0.00%
11	5.4	252	2.14%
12	1.4	243.6	0.57%
13	1	260.4	0.38%
14	0.4	245.8	0.16%
15	0.6	265.6	0.23%
16	0.6	58.8	1.02%
17	0.8	0	0.00%
18	2.6	75.8	3.43%
19	1.4	165.6	0.85%
20	3	85.4	3.51%
21	2.2	166.6	1.32%
22	1.6	214.2	0.75%
23	0	207.2	0.00%
24	5.4	252	2.14%
25	1.6	243.6	0.66%
26	1.2	260.4	0.46%

TABLE 2. Results for each instance of the F-Sub; one for each period of the M-Sub



## Low-Observable Maintenance Scheduling

In this chapter, we shift our focus to maintenance scheduling issues that are unique to low-observable (LO) aircraft. The Air Force's newest generation of fighter aircraft with LO capabilities have preventative maintenance requirements that did not exist for previous generations of fighter aircraft. In particular, an aircraft's LO capabilities degrade randomly over time which makes it difficult to make maintenance scheduling decisions. We model the LO maintenance scheduling problem as a variant of the restless multi-armed bandit (MAB) problem [19] and apply a variant of the heuristic developed by Whittle to generate a simple index policy [8]. The index policy allows maintenance schedulers to quickly rank a fleet of aircraft based on the state of each aircraft's LO capability and decide which aircraft to enter into LO maintenance and for how long [19].

We begin this chapter by presenting a detailed explanation of the LO maintenance process and the dynamics associated with it. We then characterize the process as a restless MAB problem and outline Whittle's algorithm. Next, we develop and discuss the dynamic programming (DP) formulation used to generate the desired index policy [4]. Lastly, we simulate the index policy under a range of conditions to quantify the effectiveness of the policy.

It is important to note, the analysis in this chapter is focused solely on the data and processes pertaining to a specific Air Force MDS, but the results are relevant to all aircraft with similar LO capabilities and maintenance requirements.

### 5.1. Low-Observable Maintenance Process

The newest generation of fighter aircraft in the Air Force, have low-observable technologies that present unique maintenance issues that did not exist for previous generations of fighter aircraft. In particular, the outer surfaces of the LO aircraft are coated with a metallic paint that is designed to minimize the radar signature of the aircraft. While LO aircraft have many design features that contribute to the LO capability of the aircraft, such as the shape and angles of the aircraft, the metallic coating is the primary contributor to the increased maintenance requirements for LO aircraft. If the coating is damaged in any way, the radar signature of the aircraft can be affected. Since LO aircraft are not considered to be fully mission capable (FMC) unless their radar signature is below a certain level, maintenance personnel must continuously repair the metallic coating on LO aircraft in order to sustain

an acceptable FMC rate for a fleet of aircraft. Due to limited maintenance capacity and the downtime required to complete LO maintenance, the scheduling of LO maintenance is crucial to fleet health.

In this section, we begin by discussing how LO damages occur, major operational factors that influence LO maintenance decisions, and the basic maintenance process for repairing damages. We then present and discuss some data pertaining to LO maintenance.

**5.1.1. Operational LO Maintenance.** There are two primary categories of LO maintenance. The first category pertains to the case when an internal component of an aircraft needs to be maintained and therefore requires that an outer panel be removed to access the inside of the aircraft. Since the coating on the external surfaces of the aircraft is continuous and smooth across all panels, LO maintenance personnel must break the metallic coating around the edge of the panel to remove the panel and then they must restore the coating when the panel is reinstalled. This action of “breaking the shell” is often simply referred to as LO restore.

The second category of LO maintenance deals with damages to the metallic coating due to regular flying activities. Simple scrapes and dings can have a significant impact on the overall radar signature of an aircraft. Therefore, LO maintenance personnel must carefully track the damage on each aircraft and decide when to repair them. This type of maintenance is solely aimed at reducing the radar signature of the aircraft and improving its LO capability. We will refer to these maintenance actions as LO reduction, or LO redux, actions.

In this chapter, we limit the scope of our analysis to the scheduling of LO redux maintenance actions. The schedule for LO restore actions is driven by the maintenance requirements to internal components of the aircraft and is independent of the state of an aircraft’s LO capability. Therefore, it would require far more data and knowledge to fully model the LO restore scheduling process. In contrast, LO redux actions are based solely on the state of an aircraft’s radar cross-section and the distribution of damages on an aircraft. Therefore, we limit ourselves to the scheduling of LO redux actions.

To properly plan LO redux maintenance, maintenance personnel must be able to characterize the state of each aircraft’s LO capability. As an aircraft collects minor damages, the LO capability of the aircraft deteriorates. Depending on the size, location, and shape of each specific damage, the overall impact of a single damage can range from being negligible to causing the aircraft to no longer be FMC. Examples of damages include scratches to the metallic coating and chipping of the coating. Each time an aircraft flies, maintenance personnel record all damages in a database that stores the LO damage information for each aircraft. To estimate the cumulative impact of all the damages and to characterize the state of a aircraft’s LO capability, maintenance personnel use a system called the Signature Assessment System (SAS). The SAS takes into account all the recorded damages in the database

and estimates the LO capability of an aircraft with a SAS number. Since SAS number is a measure of LO capability, we use the terms LO redux and SAS redux interchangeably.

This SAS number is crucial in making LO maintenance decisions. In general, a higher SAS number corresponds to a larger radar signature and a lower SAS number corresponds to a smaller radar signature. The minimum possible SAS number is 0 which represents an aircraft with maximum stealth capability. If an aircraft's SAS number exceeds a certain threshold, then the aircraft is no longer considered to be FMC since its stealth characteristics have been sufficiently degraded. It is important to note however, that even if an aircraft is no longer FMC due to its high SAS number, it can still be flown in support of the sortie requirements. This is due to the fact that the minor damages that lead to increases in the SAS number do not have a significant impact on the flight characteristics of the aircraft. The overall state of a fleet of aircraft is often measured in terms of the FMC rate. Although a high SAS number can cause an aircraft to no longer be FMC, there are a multitude of other maintenance issues unrelated to SAS number that can also cause an aircraft to no longer be FMC. Still, maintenance schedulers and personnel must carefully monitor the cumulative damage for each aircraft and try to keep each aircraft's SAS number below the FMC threshold.

When deciding which aircraft to enter into SAS redux, schedulers must consider not only the SAS number for each aircraft but also the nature of the damages that contribute to the SAS number. In particular, two aircraft may have the same SAS number, but the distribution of damages on each aircraft could be drastically different. One aircraft may have only a few damages that each cause a significant increase in the SAS number while the second aircraft may have a large number of damages that each contribute a small amount to the overall SAS. Individual damages are classified as either a "heavy hitter" (HH) or a non heavy hitter, where a heavy hitter is a damage that singlehandedly increases the SAS number of an aircraft by a significant amount. The number of man hours necessary to complete a SAS redux action with a desired reduction in SAS number depends largely on the distribution, both physically and severity, of damages on a particular aircraft. Therefore, when scheduling aircraft for SAS redux, schedulers evaluate the total SAS of each aircraft as well as the distribution of damages.

Based on the distribution of damages on an aircraft, the amount by which the aircraft's SAS number will decrease after a fixed time in SAS redux maintenance will vary drastically. The reduction in SAS number due to a SAS redux action is referred to as the "buy back". Maintenance personnel generally repair damages in order of how much they contribute to the overall SAS number, repairing the damages that contribute the most to the SAS number first. Furthermore, most damages take approximately the same amount of time to repair regardless of how much they contribute to the overall SAS number. Therefore, if an aircraft has a large

number of heavy hitters, a single day of SAS redux will result in a large reduction in the SAS number. If the same aircraft undergoes SAS redux for multiple days, each subsequent day will have decreasing marginal returns in terms of total SAS buy back.

It may then seem reasonable to only complete SAS redux on aircraft with heavy hitters since they offer the most buy back per day of SAS redux. Notice however, that if only heavy hitters are repaired, the FMC rate of the fleet may increase in the short term, but the number of smaller damages on each aircraft will slowly increase over time and will push the SAS number closer to the FMC threshold. All aircraft will then have to undergo extensive SAS redux in order to repair the smaller damages that had long been ignored. On the other hand, if schedulers try to maintain every damage as soon as it occurs, the limited maintenance capacity will be backlogged and the FMC rate will likely suffer in the short term. Therefore, schedulers must carefully choose when to maintain each aircraft and to what extent so as to sustain a high FMC rate for the entire fleet. Notice that unlike with phase maintenance, schedulers must choose not only the aircraft to assign to maintenance but also how long to leave the aircraft in maintenance. In general, schedulers can assign an aircraft to undergo a SAS redux action for durations that range from one to five days, or a complete overhaul which takes weeks.

Finally, another important consideration in scheduling aircraft for LO maintenance is the maintenance capacity. In practice, there is a fixed capacity for LO maintenance since it must be completed in a devoted bay that is specially designed for LO maintenance. These LO maintenance bays are used for both LO restore and redux. Therefore, the feasibility of assigning an aircraft to SAS redux is obviously dependent on the available maintenance bays.

Currently, the decision process regarding SAS redux is largely dependent on the personalities of each maintenance unit. Since there are no published policies regarding LO maintenance, each maintenance unit has the flexibility to make LO maintenance decisions however they see fit. Therefore, the LO maintenance policies used by flying units throughout the Air Force can vary. In conversations with several experienced maintenance personnel, the LO maintenance decision process was described as being "somewhat haphazard". The focus of this chapter is to mathematically model the LO maintenance decision process and develop a policy that can be easily implemented.

**5.1.2. SAS Evolution Data.** In this section we present data pertaining to the LO maintenance process. The data in this section shows the daily SAS increase for each aircraft that was flown on a given day. That is, for each aircraft that was flown in a given day, the SAS number was recorded before the first flight of the day and recorded again after the last flight of the day. The difference between the preflight SAS and post SAS flight is the data we

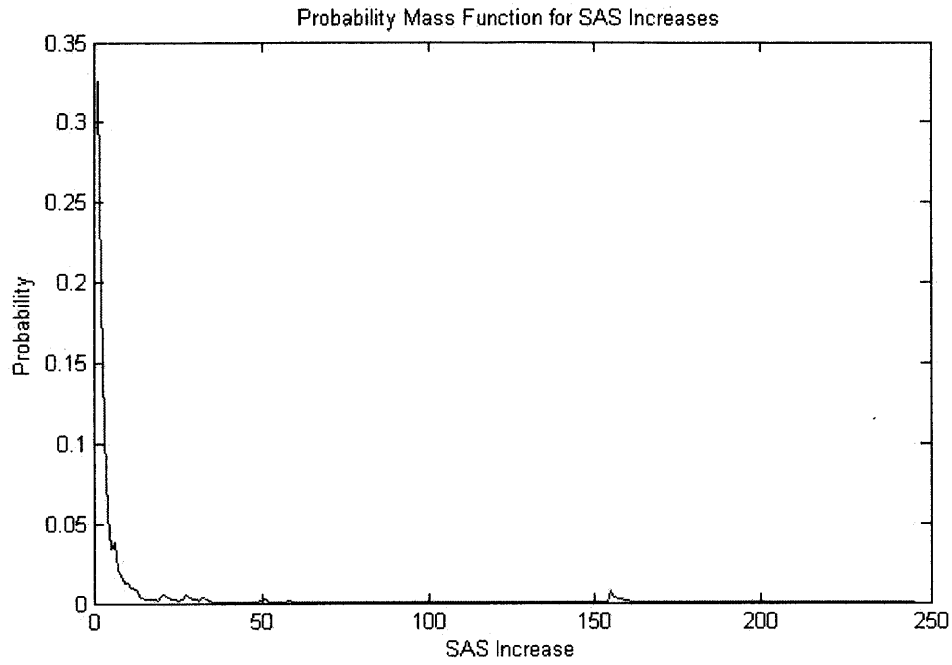
SAS Increase	Probability	Cumulative Prob.
0	0.325724151	0.325724151
1	0.193938231	0.519662383
2	0.107040092	0.626702475
3	0.055821984	0.682524458
4	0.034145406	0.716669864
5	0.038749281	0.755419144
6	0.020909265	0.77632841
7	0.017648187	0.793976597
8	0.012660656	0.806637253
9	0.013044312	0.819681565
10	0.009399578	0.829081143
11	0.009015922	0.838097065
12	0.007864953	0.845962018
13	0.003644734	0.849606752
14	0.00306925	0.852676002
15	0.002110109	0.854786112

TABLE 1. SAS Increase Probabilities

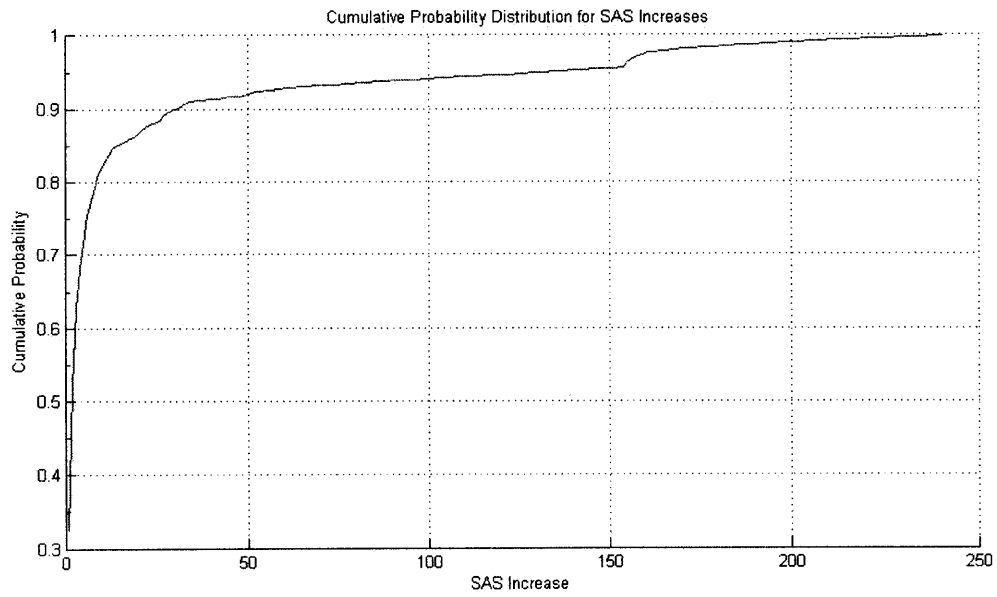
present in this section. The data is from a period of 3.5 years which includes approximately 10,000 flights and 65 aircraft.

We begin by presenting the probability distribution of the SAS increase. Figure (5.1.1) shows the empirical probability distribution of the daily SAS increases. Notice that the vast majority of SAS increase are relatively small, 10 or lower. Approximately half the times an aircraft flies, the SAS increases by less than 1 unit, and approximately 83% of the time the SAS increase is less than 10 units. Near a SAS increase of 150, the SAS increase CDF has a slight jump. This jump can be attributed to a specific type of damage, canopy damage, that immediately increases the aircraft's SAS number by more than 150 units. A canopy damage is a damage to the metallic coating the covers the glass covering of the aircraft cockpit. Table 1 shows specific probabilities of SAS increases between 0 and 15 which make up 85% of increases.

Although the data provides us with a very good understanding of how the total SAS number for each aircraft evolves over time, the data does not provide any information regarding the distribution of damages on each aircraft. For example, the data may show that an aircraft had a preflight SAS number of 70 and a postflight SAS number of 90. If the SAS increase was all due to one damage, then we would know that there is a significant HH, but if the 20 unit increase were due to several damages, then the aircraft may not have any HH. However, since the data does not specify the number of damages that caused the SAS increase, we must make assumptions regarding the distribution of damages. We assume that if an aircraft experiences an increase of 20 units or more from a single day, it has a



(A) PMF of daily SAS increases



(B) CDF of daily SAS increases

FIGURE 5.1.1. Probability distribution of daily SAS increases

	AC Number	Num Flights	Flight Hours	Start SAS	End SAS	ΔSAS Increase
SAS Increase	-0.071	0.016	0.045	0.102	0.444	1.000

TABLE 2. SAS Increase Correlation Coefficients

heavy hitter. Under that assumption, the probability of a HH on any given day of flying is approximately 13%.

To fully understand the stochastic nature of the evolution of an aircraft's SAS number, we examine whether there are significant factors that influence the SAS increases experienced by an aircraft. The SAS increase data also included the number of flights, total hours flown, and aircraft tail number associated with each SAS increase. Table 2 shows the correlation coefficients between the daily SAS increase and a number of different variables. Intuition would suggest that if an aircraft accrues more flights and flight hours in a day, the daily SAS increase would be higher. However, the correlation coefficients for these two variables, 0.016 and 0.045, respectively, are statistically insignificant. The only variable with a significant correlation coefficient is the postflight SAS number. This is to be expected since aircraft that experience a large SAS increase will immediately have a high postflight SAS number. Further analysis using regression models similarly showed that factors such as aircraft tail number, number of flights, total flight hours, days since last redux action, and flights since last redux action are all insignificant. In addition, we found that the age of the aircraft was also a statistically insignificant independent variable. Since there are no apparent significant factors that influence the distribution of SAS increases, for the remainder of the chapter we assume that SAS increases are random and follow the empirical distribution shown in Figure 5.1.1.

While the upward transitions of the SAS number are well characterized by the data, the downward transitions that result from a SAS redux maintenance action are not as well defined. It is obvious that if the duration of the SAS redux maintenance action is longer, the resulting SAS number will be lower. In addition, the distribution of damages on an aircraft will also greatly affect the resulting SAS number of an aircraft return from SAS redux maintenance. While both of these factors are associated with the amount by which an aircraft's SAS number decreases, the decreases are still stochastic since it is impossible to precisely forecast the SAS number of an aircraft coming out of SAS redux. Information regarding the length of time in maintenance and the distribution of the damages helps to predict the resulting SAS number, but it does not provide all the information necessary. To better quantify the reduction in SAS number that results from a SAS redux action, we used

the limited historical data available and spoke in-depth with maintenance personnel who currently work on the flightline.

It is important to note that although the work presented in the remainder of this chapter is based on the limited SAS redux data that was available, the modeling approach could easily be extended and improved to incorporate better data. In fact, if more precise data regarding SAS redux actions were available, the modeling approach would generate a maintenance policy that would give better operational results.

Ultimately, the manner in which the SAS number of an aircraft evolves over time is stochastic. In the case of upward transitions of the SAS number, there are no apparent factors that influence that nature of the SAS increases. Therefore, a single empirical distribution can be used to accurately model the stochastic evolution. In the case of downward transitions, the probability of decreasing the SAS number by a given amount highly depends on the length of maintenance an aircraft undergoes as well as its distribution of damages, but it is still a stochastic process.

## 5.2. Modeling Framework for the LO Scheduling Problem

The LO scheduling problem is a dynamic system in which decisions must be made to sustain a high FMC rate. In this section we identify and discuss the following three components to the dynamic system: 1) the stochastic evolution, 2) the maintenance decisions, and 3) the cost or objective. Given these components, the LO scheduling problems can be modeled as a dynamic program.

**5.2.1. Stochastic Evolution of the System.** The evolution of the SAS number and the overall fleet's LO capabilities are characterized by stochastic behavior. When an aircraft is not in maintenance, its SAS number will change by a non-negative amount. Although it is likely that the change in SAS number will be relatively small, there is always a possibility of sudden, drastic increase in an aircraft's SAS number. In Section 5.1.2, we found that the upward transitions in the SAS number were not significantly related to several factors, including flight hours and number of flights. Therefore, the probability distribution of the SAS increases remains constant independent of the state of a given aircraft. That is, whether an aircraft has a SAS number of 20 and no heavy hitters, or a SAS of 150 and multiple heavy hitters, the probability of the SAS number of either aircraft increasing, assuming that they both fly on a given day, by some amount  $x$  is equal. Figure 5.1.1 shows the entire cumulative probability distribution of the upward transitions of the SAS number.

The downward transition of the SAS number results from a SAS redux maintenance action. The reduction in SAS number due to a SAS redux actions depends on two factors: the duration of the SAS maintenance and the distribution of damages. In general, when an aircraft undergoes SAS redux, maintenance personnel will repair the damages that contribute



the most to the total SAS number of the aircraft first and complete the damages that contribute the least last. Therefore, the first day of maintenance has a high marginal return while each subsequent day of maintenance has decreasing marginal returns. However, the benefit of completing a multi-day SAS redux action is that there is only a single set up and tear down. Therefore, entering an aircraft into two one day redux actions, one shortly after the other and assuming the aircraft did not experience a significant increase in SAS between the two days, will always be less efficient in terms SAS number reduction than entering an aircraft into a two day long SAS redux action. In addition to the duration of the maintenance action, the distribution of damages on an aircraft will impact the buy back gained from a SAS redux action. If an aircraft has heavy hitter damages, a single repair will result in a significant drop in the SAS number. Conversely, if an aircraft only has minor damages that each contribute a small amount to the aircraft's total SAS number, several repairs will have to be completed before any significant reduction in SAS number is achieved.

Although information regarding the distribution of damages on an aircraft and the duration of SAS redux maintenance allows us to more accurately predict the outcoming SAS number, the reduction in SAS number is still randomly distributed across a given range. This is due to several reasons. First, there are multiple ways of repairing a damage and depending on the method that is employed, the amount of SAS reduction can vary. The different methods trade off effectiveness and speediness. Some methods are quick but do not completely eliminate the SAS increase caused by a damage while other methods completely eliminate a damage but require much more time. In addition to varying methods, another factor that affects the amount of buy back gained from a SAS redux action is the skill set of the personnel. Depending on the skill set of the personnel assigned to maintain an aircraft, the SAS reduction can vary between a given range. Therefore, information regarding the distribution of damages on an aircraft and the duration of SAS redux maintenance allows us to narrow the range of possible post redux SAS values, but we can not precisely know the decrease in SAS that will result.

**5.2.2. Decisions.** As the SAS numbers of a fleet of aircraft evolve over time, decisions must be made regarding when to enter each aircraft into SAS redux and for how many days. In each time period, if there is available maintenance capacity, an aircraft can be assigned to a SAS redux action of a given duration. There are a limited number of SAS redux "packages" and each package is a defined by its duration. For example, there can be SAS redux packages with durations of one, two, three, or four days. In general, if an aircraft is not already in maintenance, it is feasible to assign the aircraft to any SAS redux package. That is, the state of an aircraft does not limit the set of feasible maintenance actions. For aircraft that are already in maintenance, the only feasible decision in a given time period is to finish the current maintenance action. We assume that once an aircraft is assigned a

specific maintenance package, it can not be made available for flight until the maintenance action is complete.

Another key factor in determining the set of feasible decisions is the maintenance capacity. If the LO maintenance capacity is full, additional aircraft can not be entered into maintenance until another aircraft exits maintenance. Therefore, as long as maintenance capacity is full, the only feasible decision for all other aircraft is to simply leave them available for flight and not enter them into maintenance.

In summary, assuming there is capacity available, the set of feasible decisions for any aircraft that is not in maintenance includes all SAS redux actions and leaving the aircraft available for flight. The only feasible decision for an aircraft that is already in maintenance is to complete the current maintenance action. If there is no available maintenance capacity, then the only feasible decision is to leave the aircraft available for flight.

**5.2.3. Cost/Objective of the System.** Ultimately, the desire to maintain a high FMC rate is the central motivator for all LO maintenance scheduling decisions. Over time, the SAS numbers of aircraft will increase according to the stochastic process described earlier. Without any SAS redux, the entire fleet of aircraft will eventually become non fully mission capable (NFMC) as each aircraft's SAS number would exceed the prescribed FMC threshold. We can quantify the state of the fleet by assigning a cost or reward to each aircraft based on its SAS number. By assigning a state dependent cost for each aircraft, we can then incentive decisions that induce a high FMC rate. To sustain a high FMC rate, maintenance scheduling decisions must be carefully made in each time period of the planning horizon. Without the proper balance of short term and long term effects, the SAS redux decisions made in any given period can have a significant impact on the sustainable FMC rate over the planning horizon.

### 5.3. Multi-Armed Bandit Problems

The LO maintenance scheduling problem lends itself to be solved using dynamic programming (DP). However, due to the maintenance capacity constraint that limits that number of aircraft that can be in maintenance, the size of the state space would quickly make the DP intractable. In order to capture the maintenance capacity constraint within a DP, the state of the system at any given time period would have to include information regarding every aircraft in the fleet. In addition, maintenance actions can last multiple periods which means the capacity constraint cannot be considered separately in each time period. Given that, in practice, there can be approximately 40 aircraft in a fleet of aircraft and SAS numbers can take on a wide range of values, the size of the state space necessary to capture the capacity constraint would easily make the DP intractable.

Due to the intractability of solving the capacitated LO maintenance scheduling problem, we are motivated to find an alternative method for generating LO maintenance schedule. In this section we present the multi-armed bandit (MAB) problem and classify the LO maintenance scheduling problem as such. This leads to a DP that incorporates a relaxation of the maintenance capacity constraint, thereby making possible to consider each aircraft independently and greatly reduce the state space[19]. We then present Whittle's Index which has been shown to give near optimal performance in a wide range numerical examples of the MAB problem [12].

**5.3.1. Multi-Armed Bandit Problems.** A multi-armed-bandit (MAB) problem deals with the problem of optimal allocation of resources amongst a number of competing projects. In general, the MAB problem models the trade off between high short term rewards with the prospect of better future rewards[8]. In a classic MAB problem, there are  $n$  projects and the state of each project  $i = 1, 2, \dots, n$  is known to be  $s_1, s_2, \dots, s_n$ . In each time period, a decision must be made to operate one of the  $n$  projects. If project  $i$  is operated, then an immediate reward of  $g(i)$  is gained. Furthermore, project  $i$  transitions from state  $s_i(t)$  to  $s_i(t + 1)$  according to a Markov rule that is project and state dependent. The unoperated projects in a given time period do not yield a reward and they do not transition. The objective is to maximize the expected discounted reward over an infinite planning horizon.

There are many variations of the classic MAB, but we will focus specifically on the restless MAB. In the classic MAB, the unoperated projects remained static and did not undergo a state transition. In the restless case, unoperated and operated projects both undergo a state transition in each time period. In another extension of the classic MAB problem, we no longer restrict ourselves to operating one project at a time. Rather, we allow up to  $m$  of the  $n$  projects to be operated at any given time [19].

For the restless MAB, Whittle proposed a heuristic that leads to a simple index policy which is used to decide what projects to operate at any given time. The heuristic hinges on relaxing the capacity constraint on the number of projects that can be operated at one time. Recall that  $m$  of  $n$  projects could be operated in each period. By requiring that this constraint be met in expectation rather than in each period and then moving this constraint to the objective with a suitable Lagrange multiplier, Whittle showed that the Lagrange multiplier could be used as an index value[19]. By relaxing the capacity constraint, the problem becomes much easier to solve as each state can be considered independent of the other states. The Lagrange multiplier serves as a surrogate for the capacity constraint and indicates the value of using one unit of capacity in a given state.

In general, the heuristic seeks a subsidy  $\lambda_{s_i}$  that makes it equally attractive to operate or not operate a project  $i$  in state  $x_i$ . This subsidy value can be calculated for every project since the state of each project is known. The projects are then ranked according to their

subsidy values. In each period, the  $m$  highest ranked projects are operated. Notice that this leads to a simple index policy that is easily implemented and followed. We will refer to this index policy as Whittle's index. In practice, Whittle's policy has been shown to be near optimal in a wide range of cases [12].

**5.3.2. LO Maintenance Scheduling as an MAB.** The LO maintenance scheduling problem very closely matches the construct of the restless MAB problem. Each aircraft can be thought of as a project and deciding to maintain an aircraft is analogous to operating a project. For aircraft that are not assigned to maintenance, the state of the aircraft will transition based on the known probabilities of upward SAS number transitions. For aircraft that are maintained, the transition probabilities are also known based on the state of the aircraft. The maintenance capacity constraint is analogous to only being allowed to operate  $m$  out of the  $n$  projects in each period. Furthermore, the reward for operating an aircraft is a function of the aircraft's SAS number and where it is relative to the FMC threshold. By providing a greater reward to aircraft that are below the FMC threshold and then maximizing the expected reward over an infinite horizon, we maximize the sustainable FMC rate for a fleet of aircraft.

A primary difference between the LO maintenance scheduling problem and the restless MAB problem described in the previous section lies in that there are multiple feasible SAS redux packages. In the MAB context we must not only decide which projects to operate, but must also decide how to operate the project. Even with multiple ways to operate a project, Whittle's index policy can still be implemented. Another difference between the LO maintenance scheduling problem and the restless MAB problem is that not all aircraft transition when not in maintenance. Of the aircraft not in maintenance, only a subset may y on a given day. We assume that the SAS number of aircraft that are not in maintenance and are not own does not change. We refer to this as the partially restless problem.

#### 5.4. Application of Whittle's Index to the SAS Redux Scheduling Problem

In this section, we apply Whittle's heuristic to the LO maintenance scheduling problem and produce an index policy that decides when and how long to enter aircraft into SAS redux. Recall that the heuristic requires the calculation of a subsidy value,  $\lambda_{s_i}$ , that makes it equally attractive to maintain or not maintain an aircraft  $i$  in state  $s_i$ . These subsidy values are then the index values associated with being in state  $s_i$ . Based on the previous regression analysis, we make the assumption that all aircraft are identical in their dynamics and thus the index values depend only on the state and not the aircraft. That is  $\lambda_s = \lambda_{s_i}, \lambda_{s_{i+1}}, \dots$ . Once the index values have been calculated for all possible states, aircraft are ranked according to their state dependent index values. Then the highest ranking aircraft are entered into maintenance to fill the available maintenance capacity. However, since there

are multiple types of maintenance, we must also know what type of maintenance to enter an aircraft into. The maintenance decision associated with each index value is the decision,  $u_s = \arg \max_{u \in M} J_u^*(s)$ , where  $M$  is the set of all maintenance decisions. Ultimately the SAS index policy produces an index value and associate maintenance decision for each state  $s \in S_f$ , where  $S_f$  is the set of states for aircraft that are not currently in maintenance.

To implement Whittle's heuristic for the LO maintenance scheduling problem we consider every state  $s \in S_f$ , calculate the associated the  $\lambda_s$  and determine the associated decision. To do this, we use the following algorithm:

- (1) Begin with a state  $s \in S_f$  and a fixed  $\lambda$  value.
- (2) Solve the infinite horizon DP that models the LO maintenance scheduling problem using policy iteration
- (3) Use bisection search to find the  $\lambda$  for which  $J_{u=\text{no}}^* \text{MX}(s) = \max_{u \in M} J_u^*(s)$  and set  $\lambda_s = \lambda$
- (4) once  $J_{u=\text{no}}^* \text{MX}(s) = \max_{u \in M} J_u^*(s)$
- (5) Set  $u_s = \arg \max_{u \in M} J_u^*(s)$
- (6) Repeat for all states  $s \in S_f$

**5.4.1. DP Formulation.** In this section, we begin by presenting the DP formulation that is solved for each state  $s \in S$  and a fixed  $\lambda$ .

For a given state  $s$  and a fixed  $\lambda$ , we solve a discounted infinite horizon DP. First, we will present the state space used in our implementation of Whittle's heuristic. Next we will define the decisions and the sets of state dependent feasible decisions. Lastly we will introduce the transition probabilities that will be used to capture the stochastic nature of the process.

After explaining the formulation we then explain the bisection search used to find the  $\lambda_s$  value that results in  $J_{u=\text{no}}^* \text{MX}(s) = \max_{u \in M} J_u^*(s)$ .

*State Space.* Our objective is to form a state space that captures every possible state that an aircraft could be in over a planning horizon. Recall in Section 5.2.1, we stated that the upward transitions in the SAS number are captured by the empirical distribution from the historical data presented in Section 5.1.2 and that correlation and regression analyses indicate that factors such as number of flights in a day, number of hours own, days since last redux, flights since last redux, etc. are not significant factors that influence the upward transition probabilities. For downward transitions in the SAS number due to SAS redux, the transition probabilities are dependent upon the SAS number of the aircraft when it entered maintenance as well as the distribution of damages on the aircraft. Therefore, the state space of the model must include information regarding the SAS number of each aircraft as well as some information regarding the distribution of damages on each aircraft.

Along with the SAS number and distribution of damages, the state space must capture the state of an aircraft that is currently in maintenance. Recall that there are multiple

types of SAS redux actions with different durations. We must be able to distinguish when an aircraft is in each of the different types of maintenance. In addition, since maintenance actions last multiple days, the state space must track how many days are remaining in the maintenance action.

Given these considerations we define state,  $s_i$ , of each aircraft  $i$  by four pieces of information:

- The current SAS number of the aircraft (0-300) if not in redux or the incoming SAS number if currently in redux.
- Whether or not the aircraft has a HH (0 or 1).
- What SAS redux action the aircraft is in (0-5).
- The number of days remaining in maintenance (0-11).

Therefore,  $s_i$  is a vector that contains four pieces of information. The first value will be an integer value on the range  $[0, 300]$  and will represent the SAS number of the aircraft. Although the SAS number can actually take on values greater than 300, we limit it due to the fact that aircraft with SAS numbers exceeding 300 are rare and are mostly likely a special case. The second value of the state vector is a binary indicator that represents whether an aircraft has a heavy hitter or not. Although this is a highly generalized representation of the distribution of damages on an aircraft, the seemingly simple distinction between an aircraft with a HH and one without a HH allows us to much more accurately represent the state transitions without greatly increasing the size of the state space. If historical data were available regarding the specific distribution of damages on aircraft, the state space could be expanded further and the LO maintenance scheduling process could be more accurately modeled.

The third and fourth values of the state vector are used to capture the state of an aircraft that is in maintenance. If an aircraft is not in maintenance, both values will be 0. However, if an aircraft is in SAS redux maintenance, the third value can take on one of several values, each value representing a SAS redux action of a different duration. Based on our discussions with maintenance personnel and SAS redux data, we define five possible types of maintenance:

- 1 day maintenance,
- 2 day maintenance,
- 3 day maintenance,
- 4 day maintenance,
- 11 day maintenance,

If an aircraft is in maintenance, the first two values in the state vector represents the SAS state of the aircraft when it entered maintenance. It is necessary to track the incoming SAS number and heavy hitter status since it obviously influences the feasible range of SAS

numbers an aircraft can take on when it exits SAS redux (i.e., an aircraft can not exit maintenance with a higher SAS number than it entered with). The heavy hitter indicator is important because short maintenance actions have a larger reduction in the SAS number when an aircraft has a heavy hitter than when an aircraft does not have a heavy hitter. An 11 day maintenance action, however, represents a complete overhaul of the aircraft's external coating and so the outcoming SAS number will be independent of the incoming SAS number and heavy hitter status.

Given a state vector with the four aforementioned components, there are 6611 states in the LO maintenance scheduling DP.

*Decisions.* For an aircraft that is not in maintenance and is in a given SAS and heavy hitter state, a decision must be made to either enter the aircraft into maintenance or leave it available to fly. For aircraft currently in maintenance, the only feasible decision is to continue maintenance until the action is complete. In total, there are six possible decisions: one to not enter an aircraft into maintenance and five for each of the five types of SAS redux maintenance actions. The six possible decisions are as follow:

- Do not enter into maintenance
- Enter into 1 day maintenance
- Enter into 2 day maintenance
- Enter into 3 day maintenance
- Enter into 4 day maintenance
- Enter into 11 day maintenance

Although there are a total of six possible decisions, the set of feasible decisions depends on the state of the aircraft. For an aircraft that is not in maintenance, all decisions are feasible. For aircraft that are in a maintenance, the only feasible decision is to do nothing. We use  $U(s)$  to denote the feasible set of decisions for given state  $s$ .

*State Transition Probabilities.* The transition probabilities used in the DP algorithm are state and decision dependent. In Section 5.1.2, we presented the empirical distribution of SAS increase. The data used to build the empirical distribution was collected from aircraft that had flown on a given day. If an aircraft did not fly on a given day, then there was no SAS increase data recorded for the aircraft. This is due to the fact that aircraft that do not fly can be assumed to have a SAS increase of zero. In the DP, we do not have any information regarding which aircraft are flown and which are not. Therefore, we capture the fact that only a subset of aircraft fly in each period by increasing the probability of a SAS increase of zero. In practice, approximately 40% of a fleet will fly on a given day. This means that there is a 60% chance of no SAS increase and a 40% chance that an aircraft will be subject to the empirical distribution of SAS increase previously mentioned. This is captured in the

	<b>MX Duration</b>				
	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>11</b>
<b>non HH</b>	<b>0.1</b>	<b>0.2</b>	<b>0.29</b>	<b>0.37</b>	<b>n/a</b>
<b>HH w/SAS &lt;175</b>	<b>0.4</b>	<b>0.47</b>	<b>0.53</b>	<b>0.58</b>	<b>n/a</b>
<b>HH w/SAS &gt;175</b>	<b>0.6</b>	<b>0.65</b>	<b>0.69</b>	<b>0.72</b>	<b>n/a</b>

TABLE 3. SAS redux downward transition percentages

DP by using an augmented version of the empirical distribution in which the probability of a zero SAS increase is increased appropriately.

For aircraft that undergo SAS redux, the transition probabilities depend on the length of the SAS redux maintenance, the incoming SAS number of the aircraft and whether it not it has a HH. We do not have a complete empirical distribution for the downward transition probabilities due to the large number of possible incoming and outgoing SAS and heavy hitter states. Accordingly, we assume the expected reduction in SAS number from a maintenance action is a percentage of the incoming SAS number. Furthermore, we assume the outgoing SAS number is uniformly distributed around the expected reduction in SAS number. The expected reduction and uniform distribution were determined by SAS redux data which lists incoming SAS number, outgoing SAS number, and length of maintenance as well as in depth discussions with maintenance personnel. Due to the maintenance practice of fixing SAS damages in decreasing contribution, we must also ensure that the expected SAS reduction from a multi-day maintenance action is more than multiple shorter maintenance actions (i.e., the SAS reduction from a 2 day maintenance action is more than two back to back 1 day maintenance actions).

Table 3 shows the percentages used to determine the expected SAS reduction from each of the five different maintenance durations. Note the increase in SAS reduction percentage for aircraft with an incoming SAS number greater than 175 which have a heavy hitter. This is due to the fact that canopy damages lead to an immediate 150 unit increase in the SAS number are the most common reason for an aircraft to have a SAS number in the high 100s or 200s. Since the canopy damage can be repaired quickly, the reduction resulting from a SAS redux maintenance in this region is higher than for aircraft with HHs with lower SAS numbers or no HHs. After the expected SAS reduction has been determined using the reduction percentages, the outgoing SAS number is uniformly distributed around  $\pm 10\%$  of the expected SAS reduction. For example, if an aircraft entered a 1 day maintenance action with HHs and a SAS of 100, the aircraft will return from maintenance with an integral SAS number in the interval [54, 66].



After the upward and downward transitions in the SAS number, all other transition probabilities represent an aircraft transitioning from a non maintenance state to a maintenance state or an aircraft moving through a multi day maintenance action. If an aircraft is in a non maintenance state and the decision is made to enter the aircraft into maintenance, it will transition with probability 1 from its current state to the corresponding maintenance state. Similarly, if an aircraft is already in the middle of a multi day maintenance action, it will transition with probability 1 to the state that corresponds to the next day of maintenance.

We let  $p(u_k(s))$  denote the transition probabilities associated with a state  $s$  and the decision  $u_k(s)$ , where  $u_k$  is the current policy associated with state  $s$ .

*Reward Function.* The overall objective of the LO maintenance scheduling problem is to enter aircraft into SAS redux so as to maintain a high FMC rate. Therefore, we provide a positive reward for aircraft that have a SAS number below the FMC threshold. In current practice, the FMC threshold is 100. Any aircraft with a SAS number less than 100 that is not in maintenance is considered to be FMC, regardless of whether or not it has a heavy hitter. If an aircraft has a SAS number greater than 100, it is no longer considered FMC, but it can still be flown to meet the sortie requirements. Therefore, since an aircraft with a SAS number greater than 100 still provides some benefit, it is also given a positive reward, albeit a much smaller reward than that given to an FMC aircraft.

Aircraft that are in maintenance are given a reward of  $\lambda$ . Recall that Whittle's index relies on determining the subsidy value that makes operating a project equally as attractive as not operating a project. In our implementation, maintaining an aircraft is analogous to operating a project so  $\lambda$  is the reward given to an aircraft that is in maintenance. When  $\lambda = \lambda_s$ , the expected reward for doing nothing in state  $s$  is equal to the expected reward for doing maintenance in state  $s$ . Note that  $\lambda_s$  may be positive or negative. If  $\lambda_s$  is negative, there is a cost associated with maintenance while  $\lambda_s$  positive implies a subsidy for entering maintenance.

We let  $g(s)$  denote the reward gained from being in state  $s$ . We define  $g(s)$  as follows:

$$g(s) = \begin{cases} 1 & s \in \text{set of states with SAS number} \leq 100 \text{ and not in maintenance} \\ 0.2 & s \in \text{set of states with SAS number} > 100 \text{ and not in maintenance} \\ \lambda & s \in \text{set of states that represent being in maintenance} \end{cases}$$

We provide a reward based solely on the state of the aircraft and not the action chosen. This is a result of the fact that once an aircraft is entered into maintenance, it will receive  $\lambda_s$  for each day it is in maintenance.

*Objective Function.* The objective function used in the DP seeks to maximize the expected reward. The expected reward for a given state  $s$  is,

$$J^*(s) = \max_{u \in U(s)} \left[ g(s) + \alpha \sum_{s'=1}^{|S|} p_{ss'}(u) J^*(s') \right] \quad \forall s \in S.$$

In our DP formulation, we use an  $\alpha = 0.999$ . Since sudden increases in SAS number occur at a relatively low rate, but have a significant impact on the objective function, we use a high  $\alpha$  value to “look” farther into the future.

**5.4.2. Bisection Search for  $\lambda_s$ .** As mentioned before, we are looking for a value of  $\lambda$  for each state  $s$  such that  $J_{u=\text{no}}^* \text{MX}(s) = \max_{u \in M} J_u^*(s)$ . This value of  $\lambda$  will serve as the index value for state  $s$ . Furthermore, the maintenance decision associate with  $\lambda_s$  is  $u_s = \arg \max_{u \in M} J_u^*(s)$  which is the maintenance action we will use if an aircraft is in state  $s$  and we decide to place it into SAS redux.

For a given state  $s$ , and the state specific subsidy value  $\lambda_s$  that results in  $J_{u=\text{no}}^* \text{MX}(s) = \max_{u \in M} J_u^*(s)$  is found using bisection search. After solving the DP for a given value of  $\lambda$ , if  $J_{u=\text{no}}^* \text{MX}(s) \neq \max_{u \in M} J_u^*(s)$ , the value of  $\lambda$  is updated. If  $J_{u=\text{no}}^* \text{MX}(s) < \max_{u \in M} J_u^*(s)$ , it means that the expected reward for deciding to enter an aircraft in state  $s$  into maintenance is higher than the expected reward of not entering the aircraft into maintenance. In this case, the reward for being in maintenance is too high and  $\lambda$  must be decreased. Conversely, if  $J_{u=\text{no}}^* \text{MX}(s) > \max_{u \in M} J_u^*(s)$ , the reward for being in maintenance must be decreased.

The initial upper and lower bounds for  $\lambda_s$ , denoted as  $\bar{\lambda}$  and  $\underline{\lambda}$  respectively, are unknown a priori. This is due to the fact that the value of  $\lambda_s$  is a function of the several factors including the reward function, transition probabilities, and discount factor. However, an upper bound,  $\bar{\lambda}$ , for  $\lambda_s$  can easily be found. If the DP is solved using some value of  $\lambda$  and it results in  $J_{u=\text{no}}^* \text{MX}(s) > \max_{u \in M} J_u^*(s)$ , that is,  $\lambda$  is too small. We set  $\underline{\lambda} = \lambda$  and resolve the DP with a new value of  $2\lambda$ . Continue this process until  $J_{u=\text{no}}^* \text{MX}(s) < \max_{u \in M} J_u^*(s)$  at which point a suitable  $\bar{\lambda}$  has been found. A similar process can be used to find a lower bound for  $\lambda_s$  if  $J_{u=\text{no}}^* \text{MX}(s) < \max_{u \in M} J_u^*(s)$  for the initial value of  $\lambda$ .

## 5.5. Implementation of the DP

The dynamic program described in the previous section was implemented in MATLAB R2010b. Based on computational experience, we chose policy iteration as our solution method.

Figure 5.5.1 shows the index values associated with all non maintenance aircraft states. The two lines differentiate between aircraft with heavy hitter and aircraft without heavy hitters. As expected, for a given SAS number, an aircraft with a HH will always be ranked above an aircraft without a HH. For both HH states and non HH states, the index values peak right above a SAS number of 100. This is in line with intuition since aircraft that




FIGURE 5.5.1. Index values associated with each possible SAS state

are slightly above the FMC threshold of 100 are almost guaranteed to become FMC after undergoing SAS redux. The large drop, and subsequent spike, in index values right before 175 for states with HHs is a result of the SAS redux percentages used to model canopy damages. Recall that aircraft with a SAS number greater than 175 were considered likely to have a canopy damage and a redux action was likely to eliminate the 150 unit contribution from the single damage. Therefore, it is expected that the index values for states with a  $SAS > 175$  to be high than those for states with a  $SAS < 175$ .

According to the index values, there are cases when an aircraft with no HHs should be given maintenance priority over aircraft with HHs. For example, an aircraft with no HHs and a SAS number of 101 should be entered into maintenance before an aircraft with a HH and a SAS number of 80. Although this may seem intuitive, historical data from actual flying units show that, often times, aircraft with HH are given priority regardless of other factors. The index values make it possible to easily determine when aircraft with HHs should be given priority and when they should not.

Figure 5.5.2 shows the maintenance decisions,  $u_s$ , associated with each non maintenance state. The only feasible decision for maintenance states is decision 0, to continue the maintenance action until complete, so they are not shown in the figure. For all HH states, the index policy decision is to enter an aircraft into a one day maintenance action. This is expected due to the high marginal returns from the first day of SAS redux on an aircraft with a HH. The index policy decisions for non HH states are far more varied. In general, as the

SAS number increases the associated maintenance action increases in duration. This follows intuition since an aircraft with a higher SAS number needs to undergo a longer SAS redux maintenance action to return to an FMC state.

Although the general trend of longer maintenance actions for higher SAS numbers meets intuition, the variations in the decisions associated with SAS numbers from 179 to 200 are not intuitive. Between SAS numbers of 179 and 200, the decision associated with the non HH states oscillates between a 2, 3, and 4 days maintenance action. Further analysis of the expected reward values associated with the various feasible decisions for a given state show that the oscillations are likely due to numerical rounding errors. For example, the state with a SAS number of 179 had a corresponding maintenance action that enters an aircraft in that given state into a 2 day maintenance action. However, the difference between the expected reward of a 2 day maintenance action and a 4 day maintenance action is actually less than 0.01%. Therefore, it is reasonable to ignore the oscillations in the decisions and assume that the maintenance durations increase as SAS number increases. This is consistent with intuition.

In summary, the SAS redux index policy assigns an index value and maintenance decision to each non maintenance state. Given a fleet of aircraft, each aircraft can be assigned an index value based on the state of the aircraft. After each aircraft has been assigned an index value, the fleet is rank ordered from highest to lowest using the index values. The highest ranking aircraft has the highest priority for maintenance and is entered into SAS redux as soon as there is open maintenance capacity. When an aircraft is entered into SAS redux, the duration of the SAS redux action is the maintenance decision that is associated with the state dependent index value.

## 5.6. Simulations

In this section, we simulate a long planning horizon and evaluate the effectiveness of the MAB index policy generated in the previous section. The simulation data is based on data gathered from an active Air Force flying unit. Each simulation run has a planning horizon of 1000 days. On each day 16 aircraft are required to fly, 8 from each of two squadrons. This represents an aggressive flying schedule since each squadron usually flies fewer than 8 aircraft per day. Aircraft that fly are subject to SAS increases that follow the empirical distribution presented in Section 5.1.2. Aircraft that do not fly are assumed to have no change in SAS number. For each trial, we specify the flying rule used to select what aircraft fly in each period. The total fleet size is determined by historical data and is approximately 40 aircraft, although it varies slightly over the course of the planning horizon. Lastly, we assume a LO maintenance capacity of five aircraft.

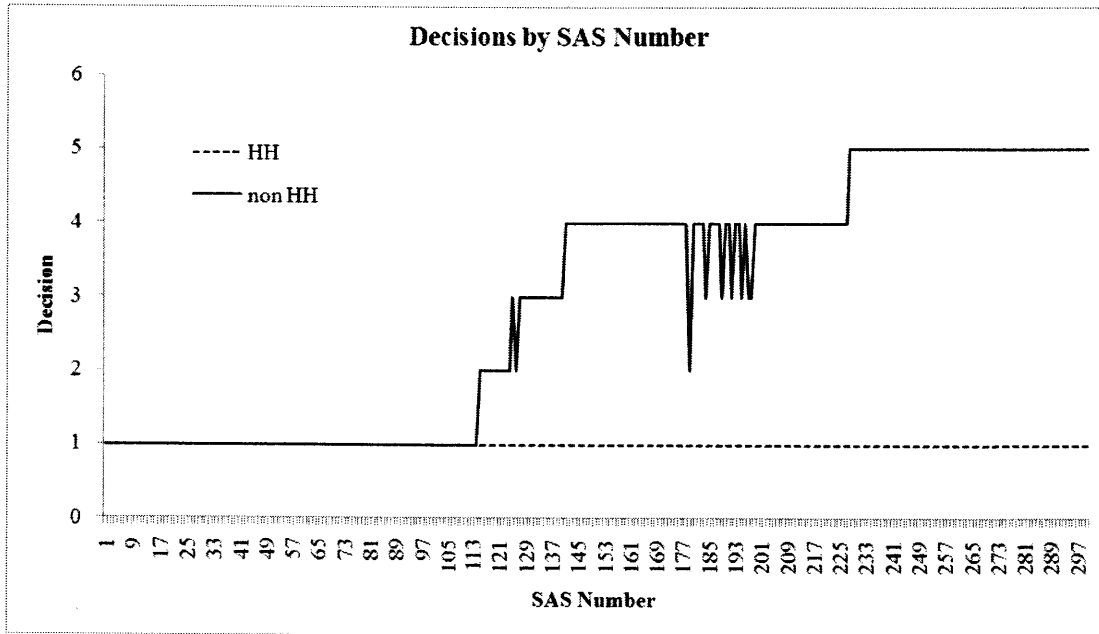


FIGURE 5.5.2. Decisions associated with possible SAS state

In each simulation run, compare maintenance policies that are used to determine which aircraft to maintain and what type of maintenance to perform. In the first set of trials, we compare the performance of the SAS index policy and the optimal uncapacitated policy. The uncapacitated policy is the policy that is generated if the DP presented in Section 5.4.1 is solved using a  $\lambda = 0$ . If  $\lambda = 0$ , it means there is no maintenance capacity constraint. The optimal DP solution will find the policy that maximizes the expected reward for each state.

To compare the effectiveness of each of the policies, we will consider their average daily FMC rate over the planning horizon. The daily FMC rate is equal to the number of aircraft not in maintenance with a SAS number less than or equal to 100 divided by the total number of aircraft. Figure 5.6.1 shows the average FMC over the planning horizon for both policies in five separate trials in which the aircraft to fly are chosen at random among all aircraft not in maintenance. Also shown, for reference as an upper bound, is the FMC rate of using the optimal uncapacitated policy in an uncapacitated scenario. Notice that, in the capacitated scenario, the SAS index policy results in a higher average FMC rate than the uncapacitated policy in each trial. This is to be expected since the SAS index policy implicitly takes into account the maintenance capacity constraint while the uncapacitated policy does not. The SAS index policy results in FMC rates around 60%. Based on the results of the simulation, we show that the SAS index policy outperforms the uncapacitated policy which is a simpler alternative to the SAS index policy.

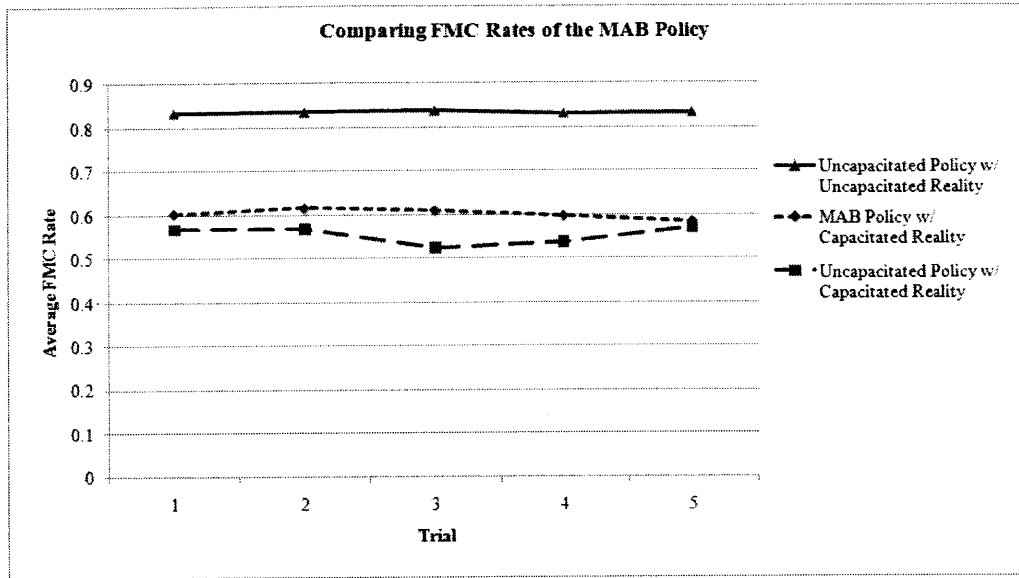


FIGURE 5.6.1. FMC rates for 5 trials. Each trial had a planning horizon of 1000 periods.

Next, we evaluate the effect of different flying rules while using the SAS index policy. That is, we simulate different methods of selecting the 16 aircraft that fly in each period and compare the resulting FMC rates. Specifically, we simulate the following three flying rules:

- (1) randomly select 16 aircraft,
- (2) always fly the 16 aircraft with the highest SAS numbers,
- (3) always fly the 16 aircraft with the lowest SAS number.

It is important to note that we select from the aircraft that are available *after* maintenance decisions have already been made. In all trials, we assume a maintenance capacity of five aircraft. We run five trials for each of the three flying rules again with a planning horizon of 1000 periods.

Figure 5.6.2 shows the average FMC rates that result from each of the three flying rules. The FMC rates generated by the random selection rule are the same as the FMC rates shown in Figure 5.6.1. Interestingly, choosing to always fly the aircraft with the highest SAS numbers results in a significant jump in FMC rate. The average FMC rate is approximately 73% when employing flying rule 2. This is due to the fact that aircraft with low SAS numbers are essentially saved for the future. Conversely, flying rule 3 leads to a drastic drop and an average FMC rate below 30%. The results suggest that aircraft with the highest SAS numbers should be assigned to the sortie requirements as much as possible.

Ultimately, the simulations show the effectiveness of the SAS index policy relative to other policies in realistic scenarios. In addition, we find that the flying rule used to select the aircraft that will fly in each time period can have a significant impact on the FMC rate.

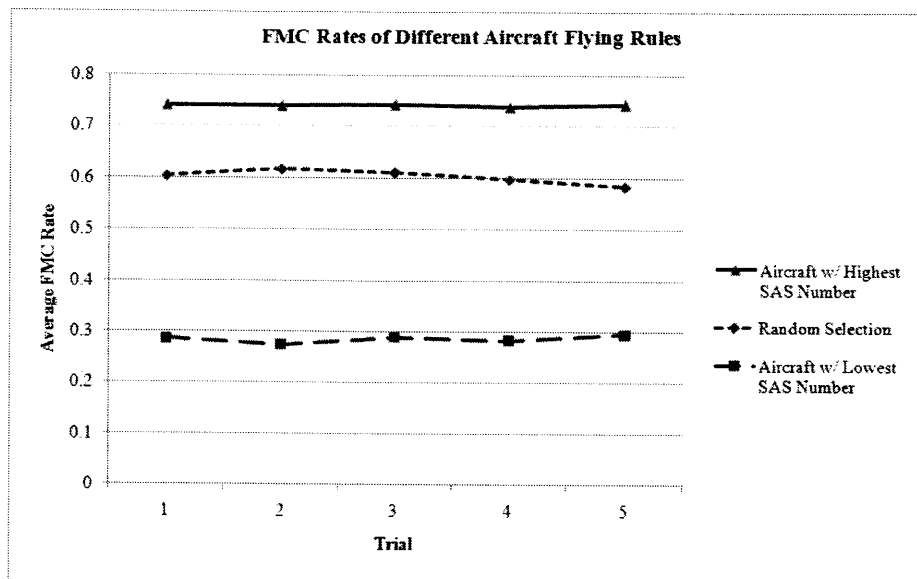


FIGURE 5.6.2. FMC rates for three different methods of selecting which aircraft will fly on a given day.

### 5.7. Expanding the DP Formulation

The DP formulation presented in Section 5.4.1 has been shown to generate a good SAS index policy. However, if more detailed data were available, the DP formulation could be expanded to more accurately model the LO maintenance scheduling problem. In particular, with more data regarding the distribution of damages, we could expand the state space to more accurately capture the relationship between the distribution of damages and the transition probabilities.

In the current DP, the binary indicator for whether or not a HH exists is the only information we store regarding the distribution of damages on an aircraft. Since the downward transition probabilities resulting from a SAS redux action are highly dependent on the distribution of damages, an expanded state space that captures more information about the distribution of damages would greatly enhance the resulting SAS index policy.

The expanded state space would require a longer solution time to generate the SAS index policy, but the SAS index policy is an offline calculation that only needs to be completed once. The resulting SAS index policy indicating the state dependent index values and corresponding maintenance actions can quickly be used to determine the daily maintenance schedule. Therefore, even if the algorithm above has a solution time of weeks, this would be acceptable since the resulting SAS index policy does not need to be regenerated for the foreseeable future. Also of note from a computational perspective is the fact that the algorithm can be implemented in a parallel computing environment. Specifically, the index

values, and resulting maintenance actions, for different states are independent and can be run in parallel. This leads to significant reductions in solution times.

### 5.8. Summary

In this chapter, we presented the unique maintenance scheduling challenges pertaining to aircraft with LO capabilities. Over time, the LO capabilities of an aircraft deteriorate in a stochastic manner. Therefore, maintenance personnel must ensure that the external coating of each aircraft is periodically repaired so as to sustain a high FMC rate for a fleet of aircraft. We developed a simple index policy that is generated offline and can be used to quickly decide what aircraft to maintain and what type of maintenance to perform. The SAS index policy for a realistic set of data is generated using a DP formulation and the effectiveness of the policy is evaluated using a simulation. We find that the SAS index policy results in good average FMC rates over a planning horizon when compared to reasonable alternative policies. Lastly, we propose possible expansions to the DP formulation that would more accurately model the LO maintenance scheduling problem.



## CHAPTER 6

### Conclusion and Recommendations for Future Works

The focus of this research was to model aspects of the maintenance scheduling process for fighter aircraft using mathematical models. Given that the current maintenance scheduling process is both time consuming and potentially inefficient, the models have the potential to significantly improve maintenance operations within the Air Force. We specifically addressed the scheduling of phase maintenance by modeling the process using a mixed integer program (MIP). Given the poor computational performance of the MIP, we disaggregate the formulation into two subproblems that can be solved sequentially. We show that the disaggregated formulation results in significantly better computational performance while generating a long-term daily flying and maintenance schedule.

We then shifted our attention to maintenance scheduling issues that are unique to aircraft with low-observable (LO) capabilities. We show that the characteristics of the LO maintenance process allow it to be modeled as a variant of the multi-armed bandit (MAB) problem. We then present a variant of the heuristic proposed by Whittle that has been shown to provide near optimal solutions for MAB problems. Applying Whittle's heuristic to the LO maintenance scheduling problem, we generate a simple index policy that can be used to schedule aircraft for LO maintenance. We then compare the index policy to alternate policies and show by simulation that the index policy leads to relatively better fully mission capable (FMC) rates, a common measure of overall fleet health.

#### 6.1. Future Work

There are significant opportunities for future work in fighter aircraft maintenance scheduling as well as maintenance scheduling in general. The research presented in this thesis represents a small portion of the potential research in this area. We begin by presenting opportunities for future work regarding the scheduling of traditional preventative maintenance actions, such as phase maintenance. Next we present opportunities for future work specifically related to LO maintenance.

**6.1.1. Traditional Preventative Maintenance.** In Section 2.2.1, we outlined the major categories of preventative maintenance. In this thesis, we only addressed the scheduling of a single type of maintenance, phase maintenance. Although phase maintenance is arguably the most significant preventative maintenance action for a fighter aircraft, the

multitude of other preventative maintenance requirements also require careful planning and scheduling. Further work should focus on incorporating all other preventative maintenance requirements in a comprehensive model.

The PM-MIP and the associated disaggregated model were specifically for phase maintenance scheduling. Although both models could easily be expanded to incorporate other maintenance requirements that are based on flying hours, they cannot be easily adapted to account for calendar based maintenance actions. Including calendar based maintenance actions in a MIP similar to the PM-MIP would likely result in severe issues with tractability. Therefore, alternative methods of simultaneously modeling calendar and usage based maintenance requirements must be developed. One possibility would be to fix the calendar based maintenance actions using a rounding algorithm similar to the one proposed by Zarybnisky et al., and then solve for the usage based maintenance requirements [19].

**6.1.2. Low-Observable Maintenance.** As discussed in Section (5.1.2), the simple index policy for LO maintenance generated by using Whittle's heuristic could very easily be improved with the availability of more data. The data used to develop the index policy in Chapter (5) was highly limited in that we did not have accurate data regarding the transition probabilities associated with SAS redux actions. Understandably, the detailed data necessary to accurately project the effectiveness of a SAS redux action has been deemed classified. Given the strong empirical results from the index policy that was generated using limited data, we strongly support the application of Whittle's heuristic using the full range of available data.

In addition to expanding the data incorporated into the index policy, another area of continued work regarding LO maintenance pertains to LO restore. Recall that LO restore and redux are the two major categories of LO maintenance. While LO redux deals with reducing aircraft SAS numbers, LO restore is the process of repairing panels that were removed in order to access internal components of an aircraft. Since LO redux and restore actions share maintenance resources, schedulers must balance the need for both types of maintenance. The index policy provides a strong policy for determining the maintenance schedule for LO redux actions, but it does not consider the scheduling of LO restore actions. Future work should focus on simultaneously considering and scheduling both types of maintenance.

Given that aircraft with LO capabilities are the future of the Air Force's fighter fleet, it is crucial to establish an effective and consistent maintenance policy. The current process of haphazard LO maintenance scheduling is likely to lead to further issues with aircraft availability and combat readiness as more and more aircraft with LO capabilities enter service. Therefore, we feel that work in the area of LO maintenance is of utmost importance.

## Appendix

### PM-MIP Formulation

*Sets.*

$I =$	set of all aircraft $i \in I$
$J =$	set of all sortie types $j \in J$
$Q =$	set of all end of horizon targets $q \in Q$
$ I  =$	number of aircraft available in the model
$ J  =$	number of unique sortie types
$ Q  =$	number of end of horizon targets (we assume $ Q  =  I $ unless otherwise noted)

*Data.*

$T$	length of time horizon, $t = 1, 2, \dots, T$
$l_j$	length of sortie type $j \in J$
$s_j^t$	minimum number of sorties of type $j \in J$ required in period $t$
$\bar{h}$	maximum accrued flying hours between phase maintenance inspections
$h_{\max}$	maximum number of remaining phase hours an aircraft can have for it to be entered into maintenance
$k$	time periods required to complete maintenance (aircraft unavailable)
$b_i$	flying hours remaining on aircraft $i \in I$ at the beginning of the horizon, $t = 1$ , until it must enter maintenance
$e_q$	end of horizon flying hours target for aircraft assigned to target $q \in Q$
$\overline{M}$	maximum number of aircraft that can be in maintenance at any give time without incurring a “cost” or “penalty”

*Decision Variables.*

$$x_{ij}^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ flies sorties type } j \in J \text{ in period } t \\ 0, & \text{otherwise} \end{cases}$$

$$m_i^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases}$$

$$v_{iq} = \begin{cases} 1, & \text{if aircraft } i \in I \text{ is assigned to end of horizon target } e_q \\ 0, & \text{otherwise} \end{cases}$$

- $h_i^t \geq 0$  phase hours remaining at the start of period  $t$  until aircraft  $i \in I$  must enter maintenance, "life remaining"
- $Z$  objective function value that is being minimized

*Formulation.*

PM-MIP = minimize  $Z$

subject to

$$\begin{aligned}
& h_i^1 = b_i, & , \forall i \in I, \\
& h_i^{t+1} \leq h_i^t - \sum_j x_{ij}^t l_j + \bar{h} m_i^t, & , \forall t, i \in I, \\
& h_i^{t+1} \geq h_i^t - \sum_j x_{ij}^t l_j, & , \forall t, i \in I, \\
& h_i^{t+1} \leq \bar{h}, & , \forall t, i \in I, \\
& h_i^{t+1} \geq \bar{h} m_i^t, & , \forall t, i \in I, \\
& \sum_j x_{ij}^t \leq 1, & , \forall t, i \in I, \\
& \sum_i x_{ij}^t \geq s_j^t, & , \forall t, j \in J, \\
& \bar{h} - h_i^t \geq (\bar{h} - \bar{h} h_{\min} \cdot) m_i^t, & , \forall t, i \in I, \\
& \sum_j x_{ij}^t \geq \sum_j x_{ij}^{t+1}, & , \forall \text{odd } t, i \in I, \\
& m_i^t + \sum_j x_{ij}^{t+y} \leq 1, & , \forall t \in [1, T - k + 1], y \in [0, k - 1], i \in I, \\
& m_i^t + m_i^{t+y} \leq 1, & , \forall t \in [1, T - k], y \in [1, k], i \in I, \\
& 0.9 v_{iq} e_q \leq h_i^T, & , \forall i \in I, q \in Q, \\
& 1.1 v_{iq} e_q + \bar{h} (1 - v_{iq}) \geq h_i^T, & , \forall i \in I, q \in Q, \\
& \sum_i v_{iq} = 1, & , \forall q \in Q, \\
& \sum_q v_{iq} = 1, & , \forall i \in I, \\
& \sum_i \sum_{\tau=t-k}^t m_i^\tau - \bar{M} \leq Z, & , \forall t \in [k + 1, T], \\
& Z \geq 0,
\end{aligned}$$

$$\begin{aligned}h_i^t &\geq 0, && \forall i \in I, t, \\x_{ij}^t &\in \{0, 1\}, && \forall i \in I, j \in J, t, \\m_i^t &\in \{0, 1\}, && \forall i \in I, t, \\v_{iq} &\in \{0, 1\}, && \forall i \in I, q \in Q.\end{aligned}$$

**Network Formulation for Special Case PM-MIP***Sets.*

$N$	set of all nodes
$I$	set of all aircraft sortie inventory nodes, $i \in I$
$D$	set of all sortie demand nodes, $d \in D$
$A$	set of all arcs
$A_{ID}$	set of all arcs $(i, j)$ with $i \in I$ and $j \in D$
$A_{IM}$	set of all arcs that lead out of a node in which an aircraft entered maintenance.

*Parameters.*

$b_i$	supply ( $b_i > 0$ ) or demand ( $b_i < 0$ ) at node $i \in N$
-------	--

*Variables.*

$f_{ij}$	number of sorties that flow from node $i \in N$ to node $j \in N$
----------	---

*Formulation.*  $Z_{PM-N} = \text{minimize } 0$   
 subject to

$$\begin{aligned}
 b_i + \sum_{(j,i) \in A} f_{ji} &= \sum_{(i,j) \in A} f_{ij} & \forall i \in N \\
 f_{ij} &\leq 1 & \forall (i,j) \in A_{ID} \\
 f_{ij} &= \bar{h} & \forall (i,j) \in A_{IM}
 \end{aligned}$$

**M-Sub Formulation***Sets.*

$I$	set of all aircraft $i \in I$
$Q$	set of all end of horizon targets $q \in Q$

*Data.*

$T$	length of time horizon where $t = 1, 2, \dots, T$
$\bar{h}$	maximum accrued flying hours between phase maintenance inspections
$h_{\max}$	maximum number of phase hours an aircraft can have remaining for it to be entered into maintenance
$b_i$	phase hours remaining on aircraft $i \in I$ at the beginning of the horizon, $t = 1$
$e_q$	end of horizon flying hours target for aircraft assigned to target $q \in Q$
$\bar{M}$	maximum number of aircraft that can be in maintenance at any given time without incurring a “cost” or “penalty”
$F_t$	flying hours required in period $t$ to meet all sortie requirements
$f_{\max}^t$	maximum flying hours that an aircraft can fly in a single period

*Decision Variables.*

$m_i^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases}$	
$v_{iq} = \begin{cases} 1, & \text{if aircraft } i \in I \text{ is assigned to end of horizon target } e_q \\ 0, & \text{otherwise} \end{cases}$	
$h_i^t \geq 0$	phase hours remaining at the start of period $t$ until aircraft $i \in I$ must enter maintenance, “life remaining”
$f_i^t \geq 0$	the number of hours flown by aircraft $i \in I$ in time period $t$
$Z$	objective function value that is being minimized

*Formulation.*

$$(6.1.1) \quad \text{M-Sub} = \text{minimize } Z$$

s.t.

$$\begin{aligned}
& h_i^1 = b_i, & , \forall i \in I, \\
& h_i^{t+1} \leq h_i^t - f_i^t + \bar{h}m_i^t, & , \forall t, i \in I, \\
& h_i^{t+1} \geq h_i^t - f_i^t, & , \forall t, i \in I, \\
& h_i^{t+1} \leq \bar{h}, & , \forall t, i \in I, \\
& h_i^{t+1} \geq \bar{h}m_i^t, & , \forall t, i \in I, \\
& \bar{h} - h_i^t \geq (\bar{h} - \bar{h}h_{\max})m_i^t, & , \forall t, i \in I, \\
& \sum_i f_i^t \geq F_t, & , \forall t, \\
& f_i^t \leq (1 - m_i^t)f_{\max}^t, & , \forall t, i \in I, \\
& 0.9v_{iq}e_q \leq h_i^T, & , \forall i \in I, q \in Q, \\
& 1.1v_{iq}e_q + \bar{h}(1 - v_{iq}) \geq h_i^T, & , \forall i \in I, q \in Q, \\
& \sum_i v_{iq} = 1, & , \forall q \in Q, \\
& \sum_q v_{iq} = 1, & , \forall i \in I, \\
& \sum_i m_i^t - \bar{M} \leq Z, & , \forall t, \\
& Z \geq 0, \\
& h_i^t \geq 0, & , \forall i \in I, t, \\
& f_i^t \geq 0, & , \forall i \in I, t, \\
& m_i^t \in \{0, 1\}, & , \forall i \in I, t, \\
& v_{iq} \in \{0, 1\}, & , \forall i \in I, q \in Q.
\end{aligned}$$



**F-Sub Formulation***Sets.*

- $I$  set of all aircraft  $i \in I$   
 $J$  set of all sortie types  $j \in J$

*Data.*

- $T$  length of time horizon where  $t = 1, 2, \dots, T$   
 $f_i$  number of flying hours assigned to aircraft  $i \in I$  (from the M-Sub solution)  
 $l_j$  length of sortie type  $j \in J$   
 $s_j^t$  minimum number of sorties of type  $j \in J$  required in period  $t$

*Decision Variables.*

- $x_{ij}^t = \begin{cases} 1, & \text{if aircraft } i \in I \text{ enters maintenance in period } t \\ 0, & \text{otherwise} \end{cases}$   
 $z_i$  difference between the F-Sub assigned flying hours and the M-Sub assigned flying hours for aircraft  $i \in I$

*Formulation.*

$$\text{F-Sub} = \text{minimize } \sum_i z_i$$

s.t.

$$\sum_i x_{ij}^t \geq s_j^t, \quad , \forall t, j \in J,$$

$$\sum_j x_{ij}^t \leq 1, \quad , \forall t, i \in I,$$

$$\sum_j x_{ij}^t \geq \sum_j x_{ij}^{t+1}, \quad , \forall \text{odd } t, i \in I,$$

$$\left| f_i - \sum_t \sum_j x_{ij}^t l_j \right| \leq z_i, \quad , \forall i,$$

$$z_i \geq 0, \quad , \forall i,$$

## Bibliography

- [1] Aircraft Flying and Maintenance Scheduling Procedures. *ACC Instructions 21-165*. 22 April 2008.
- [2] Aircraft Flying and Maintenance Scheduling Procedures. *Air Force Materiel Command Instruction 21-165*. 12 October 2005.
- [3] Aircraft and Equipment Maintenance Management. *Air Force Instruction 21-101*. 26 July 2010.
- [4] Cohn, A. C. Barnhart. "Improving Crew Scheduling by Incorporating Key Maintenance Routing Decisions," *Operations Research*, Vol. 51, No. 3, pp. 387-396, May/June 2003
- [5] Dahlman CJ, Kerchner R, Thaler D. "Setting Requirements for Maintenance Manpower in the U.S. Air Force." RAND Project Air Force.
- [6] Dewitte L.. "US Air Force further extends phase maintenance interval on F-16." *F-16.net*. 5 December 2003. 16 March 2011. [http://www.f-16.net/news\\_article917.html](http://www.f-16.net/news_article917.html).
- [7] Farias V.F., R. Madan . "The Irrevocable Mult-Armed Bandit Problem." 22 June 2008.
- [8] Fundamentals Aircraft Maintenance. Air Force Tactics, Techniques, and Procedures. 13 October 2010.
- [9] Gittins, J.C., D.M. Jones. "A dynamic allocation index for the sequential design of experiments." 241-266. 1974.
- [10] Levin, A.. "Scheduling and fleet routing models for transportation system." *Transportation Science* 5, 232-255. 1971.
- [11] Liu, K., Q. Zhao, "Indexability of Restless Bandit Problems and Optimality of Whittle Index for Dynamic Multichannel Access." *IEEE Transactions on Information Theory*, vol. 56, no. 11, pp. 5547-5567, November, 2010.
- [12] Lohatepanont, C. Barnhart, "Airline Schedule Planning: Integrated Models and Algorithms for Schedule Design and Fleet Assignment," *Transportation Science*, Vol. 38, No. 1, pp. 19-32, February 2004.
- [13] Mahajan, A., D. Teneketzis. "Multi-armed bandit problems." in A. Hero, D. Castanon, D. Cochran & K. Kastella, eds, *Foundations and Applications of Sensor Management*, Springer, pp. 121-152, 2008.
- [14] "Observations on the Air Force Flying Hour Program." United States General Accounting Office. July 1999.
- [15] Schogol J.. "Air Force wants to trim 5,750 people by 2012." *Stars and Stripes*. 26 March 2010. 19 May 2011. <http://www.stripes.com/news/air-force-wants-to-trim-5-750-people-by-2012-1.100306>
- [16] Siram C., A. Haghani. "An optimization model for aircraft maintenance scheduling and re-assignment." *Transportation Research Part A: Policy and Practice*, vol 37, issue 1, pp 29-48, 2003.
- [17] Tripp RS, et al.. "A Repair Network Concept for Air Force Maintenance." RAND Project Air Force. 2010.
- [18] Whittle, P.. "Restless bandits: Activity allocation in a changing world." *Journal of Applied Probability*, 25A, pp. 287-298, 1988.
- [19] Zarybnisky, E.. "Maintenance Scheduling for Modular Systems – Models and Algorithms" *Operations Research*. Forthcoming.