TRADEOFFS IN AIR FORCE MAINTENANCE:
Squadron Size, Inventory Policy, and Cannibalization

by

Luis C. Tsuji

B.S. Aeronautics and Astronautics
Massachusetts Institute of Technology, 1996

M.S. Aeronautics and Astronautics
Massachusetts Institute of Technology, 1999

Submitted to the Technology and Policy Program
in partial fulfillment of the requirements for the Degree of

Master of Science
in Technology and Policy

at the

Massachusetts Institute of Technology

September 1999

Signature of Author ____________________________
Technology and Policy Program
August 6, 1999

Certified by ____________________________
Professor Joseph M. Sussman
JR East Professor, Department of Civil and Environmental Engineering
Thesis Supervisor

Accepted by ____________________________
Professor Richard de Neufville
Chairman, Technology and Policy Program

TRADEOFFS IN AIR FORCE MAINTENANCE:
Squadron Size, Inventory Policy, and Cannibalization

by

Luis C. Tsuji

Submitted to the Technology and Policy Program
on August 6, 1999 in partial fulfillment of the requirements for
the Degree of Master of Science in Technology and Policy

ABSTRACT

The Air Force sustainment system, which includes maintenance and logistics, is facing difficult challenges. As the maintenance system is being downsized, operations tempo is increasing and private companies are entering into competition for maintenance workload. The Air Force is under intense pressure to improve maintenance performance. Attempts to change the maintenance system in a piecemeal fashion have often led to unintended consequences and global sub-optimization. High-level simulation models of the maintenance system that could illustrate critical tradeoffs could provide a valuable tool for learning, and help improve system performance in the future.

This thesis uses a simple high-level simulation model to model the sustainment of a unit of C-5 aircraft. It examines high-level tradeoffs in performance and cost due to the number of aircraft, the number of spare parts, and cannibalization practices. The effects of depot repair time and the failure probability of aircraft parts are also considered. In a system like that of the Air Force sustainment system that aims to improve maintenance and logistics performance and reduce cost, yet must deal with large demand variability and must be prepared for wartime surge, cannibalization, a large number of aircraft, and a large inventory of spare parts may be necessary and may even be cost-effective.

Thesis Supervisor:	Prof. Joseph M. Sussman JR East Professor of Civil and Environmental Engineering

ACKNOWLEDGEMENTS

There are many generous people who have helped me along the way in preparing this thesis.

I would first like to thank my advisor, Prof. Joseph Sussman for his guidance, support, and motivation throughout this project. Without a doubt he is one of MIT's greats. I would also like to thank Wesley Harris, Bill Rutley, and Lori McLaughlin for their help in managing the daily issues and events in the Lean Sustainment offices and for their overall concern for the students' well-being. Lean Sustainment faculty members Wesley Harris, Kirkor Bozdogan, Dennis Mathaisel, Donald Rosenfield, Michael Siegel, and Stuart Madnick provided useful feedback and comments on the early research. Neil "Skip" Raymond provided valuable background and must have answered a thousand questions.

Air Force personnel that we visited were very helpful and hospitable. I would like to extend my thanks to the men and women of the Ogden ALC Landing Gear Shop, Air Combat Command headquarters, Air Mobility Command headquarters, the 1st Fighter Wing, and the 436th Air Mobility Wing. Visiting with them and seeing the jobs they perform for this country was an incredible experience.

In the MIT Technology and Policy Program I would like to thank Gail Hickey, Linda Manion, Liz Zotos, and my academic advisor Richard Tabors.

My colleagues Steve Czerwonka, Tom Lee, Spencer Lewis, Jason Van Wey, and and of course Lori, were a great group of friends and helped keep the atmosphere at the office enjoyable. I would also like to thank Melanie Bautista for being a kind heart and for stopping by to say hi when the offices were lonely.

I don't know if it is possible to thank a place, but I will anyway. This school and this town have been good to me for seven years. Well, except for those parking tickets. I would also like to thank my neighbors and friends at Stefani House of Pizza and the Royal East. They were always willing to chat, and were people I could talk to when I wasn't at the office.

Lastly, I would like to thank my family and friends for their ever-present support. I didn't get to see them often while I was here, but I always knew they were thinking of me, praying for me. BB, I miss you. You will always be in my heart.

FORWARD

This project was funded by the United States Air Force Contract Number F33615-96-D-5101 Delivery Order Number 036, issued by the Air Force Research Laboratory to GRC International, Inc. and subcontracted to the Office of Sponsored Programs of the Massachusetts Institute of Technology. The research was performed under the MIT Lean Sustainment Research Program of the Center for Technology Policy and Industrial Development.

TABLE OF CONTENTS

ABSTRACT **3**

ACKNOWLEDGEMENTS **5**

FORWARD **7**

TABLE OF CONTENTS **9**

TABLE OF FIGURES **11**

TABLE OF TABLES **13**

LIST OF ABBREVIATIONS AND ACRONYMS **15**

CHAPTER 1: INTRODUCTION **17**

CHAPTER 2: AIR FORCE SUSTAINMENT SYSTEM OVERVIEW **22**

2.1 SUSTAINMENT SYSTEM OVERVIEW 22
2.2 THE POST-COLD WAR SUSTAINMENT ENVIRONMENT 27
  2.2.1 Decreased Resources *29*
  2.2.2 Increased Demands *32*
  2.2.3 Declining Readiness *34*
  2.2.4 Air Force Adaptations *40*
  2.2.5 Summary *42*

CHAPTER 3: BACKGROUND **43**

3.1 MODELING AND FORECASTING AIR FORCE CAPABILITY AND SPARE PARTS DEMAND 43
3.2 MODELING TO OPTIMIZE SPARES SUPPORT 46
3.3 UNDERSTANDING AND IMPROVING SYSTEMS 47

CHAPTER 4: CANNIBALIZATION **50**

4.1 EXCESSIVE USE OF CANNIBALIZATION 50
4.2 CANNIBALIZATION INEFFICIENCIES 51
4.3 CANNIBALIZATION BENEFITS 53

CHAPTER 5: MODEL DESIGN **58**

5.1 SCOPE AND PURPOSE 58
5.2 MATLAB MODEL DESCRIPTION 59
5.3 INPUTS AND OUTPUTS 62
5.4 SUMMARY 63

CHAPTER 6: SYSTEM TRADEOFFS **65**

6.1 CHARACTERIZING TRADEOFFS IN READINESS 65
  6.1.1 Modeling Assumptions *65*
  6.1.2 Aircraft and Spare Parts *67*
  6.1.3 No Cannibalization *67*
  6.1.4 Using One Cann-Bird *72*
  6.1.5 Unrestricted Cannibalization *75*
  6.1.6 Cannibalization Policy Comparison *79*

6.1.7 Depot Repair Time *85*
6.1.8 Probability of Failure *90*
6.2 CHARACTERIZING TRADEOFFS IN COST DUE TO AIRCRAFT, SPARES, AND CANNIBALIZATION 93
6.2.1 Cost Assumptions *94*
6.2.2 Determining Total Cost *96*
6.2.3 Effect of Relative Costs *98*
6.2.4 Effects of Repair Time *100*
6.2.5 Comparison of Costs for Various Cannibalization Options *101*
6.3 SUMMARY 108
CHAPTER 7: CONCLUSION AND RECOMMENDATIONS FOR FURTHER STUDY 109
7.1 CONCLUSION 109
7.2 RECOMMENDATIONS FOR FURTHER STUDY 113
BIBLIOGRAPHY 115
APPENDIX A 117
APPENDIX B 119

TABLE OF FIGURES

FIGURE 2-1 THE MAINTENANCE SYSTEM 23
FIGURE 2-2 THE FLOW OF COMPONENTS AND SUBCOMPONENTS THROUGH THE AIR FORCE SUSTAINMENT SYSTEM 23
FIGURE 2-3 AIR FORCE BUDGET OUTLAYS, INCLUDING PROJECTIONS. 30
FIGURE 2-4 AIR FORCE OPERATIONS AND MAINTENANCE BUDGET OUTLAYS, INCLUDING PROJECTIONS. ... 30
FIGURE 2-5 PERSONNEL LEVELS 31
FIGURE 2-6 NUMBER OF FIGHTER AND ATTACK AIRCRAFT 31
FIGURE 2-7 MISSION CAPABLE RATE FOR AIR FORCE MAJOR AIRCRAFT. 35
FIGURE 2-8 MISSION CAPABLE RATE FOR USAF AIR COMBAT COMMAND FIGHTERS 36
FIGURE 2-9 TOTAL NOT MISSION CAPABLE DUE TO SUPPLY RATING FOR AIR FORCE MAJOR AIRCRAFT. 36
FIGURE 2-10 CANNIBALIZATION RATE FOR USAF AIR COMBAT COMMAND OPERATIONAL FIGHTERS 38
FIGURE 2-11 CANNIBALIZATION RATE FOR B-1B, C-5, AND F-16. 39
FIGURE 6-1 PERCENT OF REQUESTED FLIGHTS FLOWN, NO CANNIBALIZATION 71
FIGURE 6-2 MISSION CAPABLE RATE, NO CANNIBALIZATION. 71
FIGURE 6-3 NUMBER OF MISSION CAPABLE AIRCRAFT, NO CANNIBALIZATION 72
FIGURE 6-4 PERCENT OF REQUESTED FLIGHTS FLOWN, USING ONE CANN-BIRD. 73
FIGURE 6-5 MISSION CAPABLE RATE, USING ONE CANN-BIRD. 74
FIGURE 6-6 NUMBER OF MISSION CAPABLE AIRCRAFT, USING ONE CANN-BIRD. 74
FIGURE 6-7 NUMBER OF CANNIBALIZATION ACTIONS FOR THE CASE OF ONE CANN-BIRD. 75
FIGURE 6-8 PERCENT OF REQUESTED FLIGHTS FLOWN, USING UNRESTRICTED CANNIBALIZATION. 77
FIGURE 6-9 MISSION CAPABLE RATE, USING UNRESTRICTED CANNIBALIZATION. 77
FIGURE 6-10 NUMBER OF MISSION CAPABLE AIRCRAFT, USING UNRESTRICTED CANNIBALIZATION. 78
FIGURE 6-11 NUMBER OF CANNIBALIZATION ACTIONS FOR THE CASE OF UNRESTRICTED CANNIBALIZATION. 78
FIGURE 6-12 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON PERCENTAGE OF FLIGHTS FLOWN, FOR 15 AIRCRAFT. 80
FIGURE 6-13 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON PERCENTAGE OF FLIGHTS FLOWN, FOR 10 AIRCRAFT. 80
FIGURE 6-14 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON MC RATE AND MC AIRCRAFT, FOR 15 AIRCRAFT. 81
FIGURE 6-15 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON MC RATE AND MC AIRCRAFT, FOR 10 AIRCRAFT. 81
FIGURE 6-16 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON NUMBER OF CANNIBALIZATION ACTIONS, FOR 15 AIRCRAFT. 82
FIGURE 6-17 COMPARISON OF CANNIBALIZATION POLICIES: EFFECT ON NUMBER OF CANNIBALIZATION ACTIONS, FOR 10 AIRCRAFT. 82
FIGURE 6-18 THE MIX OF SPARES, AIRCRAFT, AND LEVEL OF CANNIBALIZATION REQUIRED TO ACHIEVE 90% FOR THE PERCENTAGE OF REQUESTED FLIGHTS FLOWN. 84
FIGURE 6-19 THE MIX OF SPARES, AIRCRAFT, AND LEVEL OF CANNIBALIZATION REQUIRED TO ACHIEVE ON AVERAGE 10 MISSION CAPABLE AIRCRAFT. 84
FIGURE 6-20 PERCENT OF REQUESTED MISSIONS FLOWN RESULTING FROM DIFFERENT REPAIR TIMES, FOR 15 AIRCRAFT, USING ONE CANN-BIRD. 87
FIGURE 6-21 PERCENT OF REQUESTED MISSIONS FLOWN RESULTING FROM DIFFERENT REPAIR TIMES, FOR 10 AIRCRAFT, USING ONE CANN-BIRD. 87
FIGURE 6-22 MISSION CAPABLE RATE AND NUMBER OF MISSION CAPABLE AIRCRAFT RESULTING FROM DIFFERENT REPAIR TIMES, FOR 15 AIRCRAFT, USING ONE CANN-BIRD. 88
FIGURE 6-23 MISSION CAPABLE RATE AND NUMBER OF MISSION CAPABLE AIRCRAFT RESULTING FROM DIFFERENT REPAIR TIMES, FOR 10 AIRCRAFT, USING ONE CANN-BIRD. 88
FIGURE 6-24 NUMBER OF CANNIBALIZATIONS RESULTING FROM DIFFERENT REPAIR TIMES, FOR 15 AIRCRAFT, USING ONE CANN-BIRD. 89

Figure 6-25 Number of cannibalizations resulting from different repair times, for 10 aircraft, using one Cann-bird. 89
Figure 6-26 Percent of requested flights made for varying probabilities of failure, with 15 aircraft, using one Cann-bird. 90
Figure 6-27 MC rate resulting from different probabilities of failure, for 15 aircraft, using one Cann-bird. 91
Figure 6-28 Number of cannibalizations resulting from different probabilities of failure, for 15 aircraft, using one Cann-bird. 91
Figure 6-29 Effects of increasing probability of failure on percent of requested flights flown and on MC rate, for 15 aircraft, 12 spares, and depot repair time of 50 days. 93
Figure 6-30 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50. 98
Figure 6-31 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50, cost of a set of spares equals 0.5 times the cost of an aircraft. 99
Figure 6-32 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50, cost of a set of spares equals the cost of an aircraft. 99
Figure 6-33 Cost for Different Depot Repair Times, with 15 aircraft. 101
Figure 6-34 Comparison of Costs of Different Cannibalization Policies 102
Figure 6-35 Cost for various numbers of aircraft, amount of spares, and cannibalization policies, if the cost of a complete set of spares (one of each) equals 0.5 times the cost of an aircraft. 103
Figure 6-36 Cost for various numbers of aircraft, amount of spares, and cannibalization policies, if the cost of a complete set of spares (one of each) equals the cost of an aircraft. 104
Figure 6-37 Cost Comparison: 15 aircraft with one Cann-bird, 14 aircraft with one Cann-bird, and 14 aircraft with no cannibalization 107
Figure 6-38 Cost Comparison: cost with no cannibalization vs. cost using one Cann-bird, for 15 aircraft and 11 aircraft 107
Figure 6-39 Cost Comparison: cost with no cannibalization vs. cost using one Cann-bird, for 10 aircraft. 108

TABLE OF TABLES

TABLE 1: NUMBER OF AIR FORCE MISSIONS DURING 1998 33

LIST OF ABBREVIATIONS AND ACRONYMS

2LM	Two-level Maintenance
3LM	Three-level Maintenance
ACC	Air Combat Command
AF	Air Force
AFB	Air Force Base
AFMC	Air Force Materiel Command
AFRL	Air Force Research Laboratory
ALC	Air Logistics Center
AMC	Air Mobility Command
ANG	Air National Guard
AREP	Aircraft Repair Enhancement Program
BRAC	Base Realignment and Closure
Cann-bird	Aircraft used as a source of spare parts for other aircraft.
CANN Rate	Cannibalization Rate
FY	Fiscal Year
HQ	Headquarters
IMVP	International Motor Vehicle Program
IPT	Integrated Product Team
LAI	Lean Aircraft Initiative
LG	Logistics
LMA	Logistics Management Agency
LRU	Line Replaceable Unit
LSI	Lean Sustainment Initiative
ManTech	Manufacturing Technology
MC	Mission Capable
MERLIN	Multi-Echelon Resource and Logistics Information Network
MIT	Massachusetts Institute of Technology
MRSP	Mobility Readiness Spares Package
NMC	Not Mission Capable
NMCM	Not Mission Capable Maintenance
NMCS	Not Mission Capable Supply
NRTS	Not Reparable This Station
NSN	National Stock Number
PDM	Programmed Depot Maintenance
REMIS	Reliability and Maintainability Information System
SKU	Stock Keeping Unit
SMAG	Supply Management Activity Group
SRU	Shop Replaceable Unit
TCTO	Time Compliance Technical Order
US	United States
USAF	United States Air Force

CHAPTER 1: INTRODUCTION

The Air Force sustainment system, which includes maintenance and logistics, is facing difficult challenges. In the post-Cold War era of downsizing, the Air Force budget and force structure have decreased dramatically, while at the same time deployments and operations tempo have greatly increased. These factors, combined with the increasing age of the aircraft fleet, have led to declining readiness levels and to decreasing quality of life for personnel. Along with these difficult challenges, there is now increasing competition for maintenance workload from commercial providers. Although the government is embracing a significant level of private sector participation, there is considerable debate over the balance between public and private workload, with resulting institutional friction.

The Air Force is under intense pressure to improve sustainment performance. However, attempts to change the sustainment system in a piecemeal fashion have often led to unintended consequences and system-wide sub-optimization. High-level simulation models of the sustainment system that could illustrate critical tradeoffs could provide a valuable tool for learning, and help improve system performance in the future. The contribution of this thesis is to provide some of these models and the insights derived from them.

The Air Force has been involved in a major effort to bring about improvements by implementing so called "lean" business practices. The term *lean*, which will be used frequently in this thesis, refers to a broad set of practices that are based on the ideas of teamwork, communication, efficient use of resources and elimination of waste, and continuous improvement. The principles of *lean* grew out of the Toyota production system and have brought significant changes to the automobile industry and recently to the aerospace manufacturing industry.[1] *Lean* encompasses many management practices that are rapidly being adopted in the business world, including Just In Time (JIT), quality circles, total quality management, and supply chain integration.

---

[1] James P. Womack, Daniel T. Jones, and Daniel Roos, *The Machine That Changed the World: The Story of Lean Production*, 1st Harper Perennial ed. (New York: Harper Collins, Original Published by Rawson Associates, 1990).

Whereas *lean* practices have been implemented in the automobile manufacturing industry and the aircraft manufacturing industry, it remains to be seen how well and how far the lessons learned in the production environment translate to the military maintenance and logistics environment. There are several fundamental differences between the production environment and the military maintenance environment. Such differences include the amount of variability inherent in the demand for spare parts in the military maintenance system, the variability in the way items fail and consequently the repair processes they require, and the need to be prepared for wartime surge. Such differences may create challenges in the adaptation of lean thinking to the military maintenance environment.

This thesis was completed under the auspices of the Lean Sustainment Initiative (LSI). LSI is a joint effort between Headquarters Air Force Materiel Command (HQ AFMC/LG), the Air Force Research Laboratory (AFRL), and the Massachusetts Institute of Technology. The purpose of LSI is to examine the application of lean practices to the Air Force sustainment system.

LSI emerged from the expertise and background in lean principles and practices found at MIT. The MIT International Motor Vehicle Program (IMVP) with its study of the global automobile industry initiated MIT's study of lean practices. The IMVP coined the term "lean" to refer to the Japanese production practices that they were then studying. In the book *The Machine That Changed the World*, IMVP researchers Womack, Jones, and Roos tell the story of lean production.[2] The MIT Lean Aircraft Initiative (now Lean Aerospace Initiative) carried the study of lean into the arena of aircraft production. LSI is now taking another step by examining the application of lean practices to military aircraft maintenance and logistics.

As part of the initial stages in the Lean Sustainment Initiative, members of the LSI team made a series of familiarization visits to various Air Force facilities. These included Air Force repair depots, command headquarters, and operational bases. LSI

---

[2] James P. Womack, Daniel T. Jones, and Daniel Roos, *The Machine That Changed the World: The Story of Lean Production*, 1st Harper Perennial ed. (New York: Harper Collins, Original Published by Rawson Associates, 1990).

researchers talked with personnel and observed operations, learning about the system from those working within it.

A predominant impression based on observations garnered on these trips was that while the personnel in the Air Force maintenance and logistics system are skilled and knowledgeable in their own areas and quite dedicated to fulfilling their missions, they often lack understanding of the maintenance and logistics system as a whole. Furthermore, as goals, objectives, and metrics for individual personnel or organizations do not necessarily align with those of the Air Force sustainment system as a whole, when personnel try to improve or optimize their own part of the system, they may be taking counterproductive action within the context of the larger Air Force system.

Overall, it seems the system as a whole is not aligned to operate in an optimized manner. For example, the chronic lack of spare parts, the use of parts from one aircraft to repair another aircraft, and the declining percentage of aircraft ready to fly missions were indications that the system may not be optimizing the tradeoffs between the costs of spare parts, labor, aircraft, and missed missions.

From these observations it became clear that in order to improve the Air Force sustainment system, effort must first be made to improve the common understanding and coordination of the entire system. Thus the Lean Sustainment Initiative is focusing initially on characterizing and understanding the Air Force sustainment system.

System models can be valuable tools for characterizing and understanding the Air Force sustainment system. Models that represent the Air Force system can generate insight into system behavior and provide a valuable tool for learning, for both those creating the models and those learning from the results the models produce. The Air Force sustainment system is incredibly complex. High-level simulation models, although only simplified mathematical abstractions of the real world system, can illustrate critical tradeoffs in the complex system.

As an initial step in a larger systems characterization and modeling effort at the MIT Lean Sustainment Initiative, this research will examine high-level tradeoffs between aircraft squadron size, inventory policy, and cannibalization practices (the practice of using parts from one item, in this case an aircraft, to repair another item when spare parts are not readily available). While a small fleet of aircraft, a low spare parts stock level, or limited cannibalization may be desirable from a cost viewpoint when considered individually, the combination of the three may not be lean; in fact, it might not provide the needed level of mission capable aircraft. In a system like that of the Air Force maintenance system that aims to improve maintenance performance and reduce cost, yet must deal with large demand variability and must be prepared for wartime surge, cannibalization or a large inventory in some cases may be required. Using a high-level model and Air Force data, this research aims to explore such issues, and discuss the concept of *lean* as applied to Air Force maintenance.

Previous works have explored the issue of modeling Air Force capability and spare parts demands in depth and have provided some valuable insights. The models that have been created are sophisticated, highly detailed models aimed at predicting Air Force spare parts inventory requirements and Air Force mission capability. These models tend to be highly data-intensive. This thesis in contrast to earlier works, will focus on illustrating high level systems tradeoffs involving issues such as inventory and cannibalization policy based on stylized assumptions about the system and its parameters. The model in this thesis effort is intended to produce insight for learning and negotiation, and not for accurately predicting Air Force requirements and capability.

One area that this thesis will focus on is that of cannibalization. Cannibalization is in large part a means for dealing with the variability in demand for spare parts. Cannibalization often occurs when the maintenance and logistics system for spare parts fails to provide necessary spare parts when they are needed. Cannibalization causes more work for maintenance personnel, and risks breaking the parts being cannibalized. From an operations standpoint, it is highly inefficient. However, due to the high number of parts in an aircraft, many with a very low rate of failure, cannibalization may be a more effective practice than maintaining stock of rarely needed items. Thus while cannibalization is often detrimental, in certain situations it may be the most effective from a system-wide readiness perspective.

To summarize, this thesis applies a high-level simulation model to examine the tradeoffs between aircraft fleet size, inventory policy, and cannibalization practices. Through this examination this thesis will also discuss these tradeoffs in the context of lean practices. This thesis aims to add to the baseline of understanding for the Lean

Sustainment Initiative and contribute to continued improvement of the Air Force sustainment system.

An overview of the sustainment system, discussing the system and its changing environment, will be provided in Chapter 2. Chapter 3 will review relevant literature and discuss previous works. Chapter 4 will examine the issue of cannibalization in detail. Chapter 5 will introduce and describe a simulation model used in this thesis, with Chapter 6 presenting the analysis of the model results. Chapter 7 will conclude the thesis and suggest areas for further research.

CHAPTER 2: AIR FORCE SUSTAINMENT SYSTEM OVERVIEW

In order to create, use, or interpret a model effectively, it is first essential to understand the system that it represents. While the purpose of a model is often to increase understanding of the system, there must be some groundwork upon which to build. This chapter provides an overview of the Air Force sustainment system and the changing environment in which it must operate. The term sustainment in the context of this research refers to the maintenance and logistics support of Air Force aircraft. The Air Force sustainment system is a highly complex collection of organizations that is facing a difficult period of transition. This chapter aims to first provide an understanding of the basic system, and second provide an understanding of the changing times and issues that the system needs to deal with.

2.1 SUSTAINMENT SYSTEM OVERVIEW

The sustainment system is a complex collection of interconnected processes and organizations, stretching from the flight line mechanic back to the manufacturer of spare parts. The system incorporates both the public and private sectors. Figure 2-1 gives a simplified view of the overall sustainment system. Figure 2-2 provides a more detailed illustration of the system, illustrating general flows of material throughout the system. The term field level refers to operations at an Air Force base. Base supply and base level repair (also referred to as intermediate level repair) are at the field level, along with flight line (or operational) activities. Flight line or operational activities refer to activities at the aircraft squadron or wing[3] level, which typically take place on or near the flight line. Depot level refers to the repair system above the field level that includes Air Force and contractor depot repair facilities and such organizations as the Defense Logistics Agency (DLA). Industry level refers to operations in industry that support and supply depot level activities.

---

[3] A wing is an organizational unit that typically consists of several squadrons.

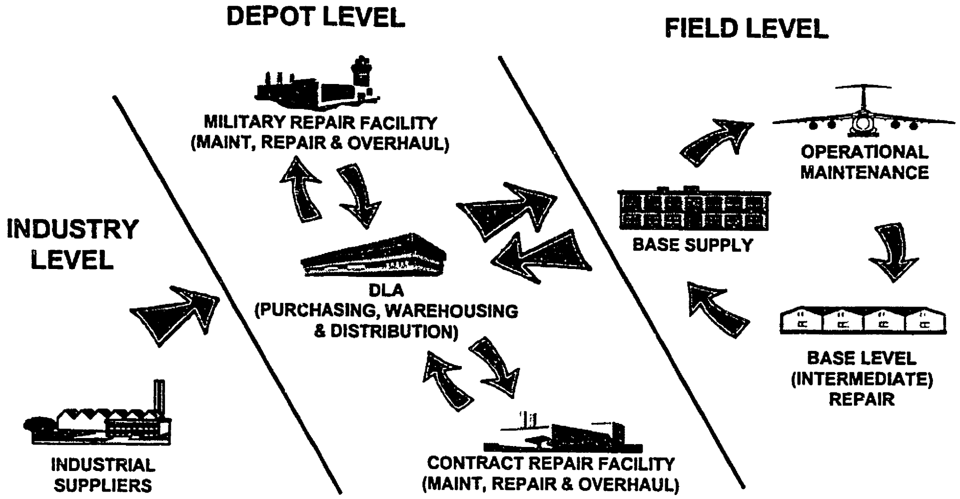

Figure 2-1 The Maintenance System[4]

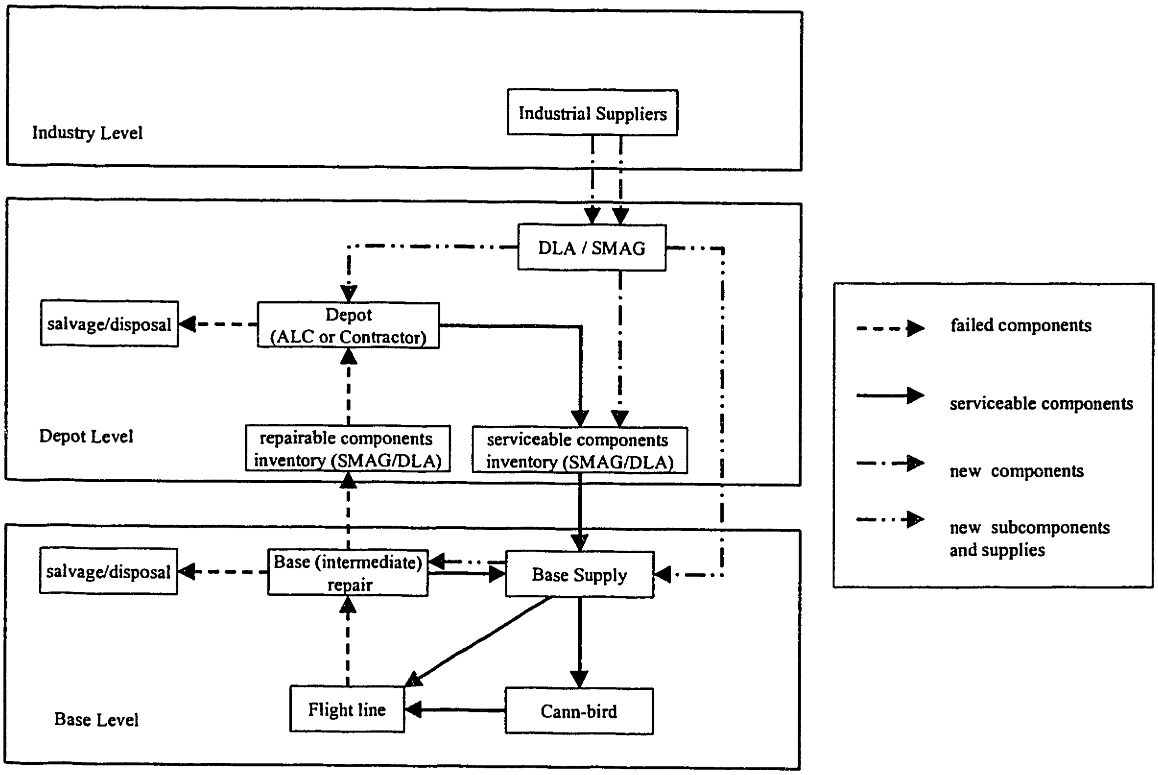

Figure 2-2 The Flow of Components and Subcomponents Through The Air Force Sustainment System

---

[4] Modified from Neil V. Raymond, *Lean Sustainment Initiative Baseline Seminar Logistics Overview*, (GRC International, Inc., 1997).

To provide an introduction and explanation to the Air Force maintenance and logistics system, it helps to trace the flow of a hypothetical broken aircraft part through the system. Imagine that an aircraft just returned from a mission and it is discovered either by the pilot or by the aircraft ground crew that a part has failed. The flight line maintenance personnel (also referred to as operational maintenance personnel) remove the failed part and replace it with a spare part from the inventory in base supply. In situations when spares are not available, cannibalization may be used to provide a replacement part. Cannibalization will be discussed more extensively in Chapter 4. If the failed part is categorized as consumable (such as bolts, fuses, seals, etc.), the failed part is discarded. If the failed part is categorized as reparable (avionics components, landing gear, etc.), it is sent to its designated source of repair such as a shop or a depot. In this thesis, aircraft parts that are categorized as repairable will be referred to generically as components. A component that is removed from the aircraft for repair and replaced with a spare by maintenance personnel on the flight line is referred to as an LRU (line replaceable unit). LRUs often contain subcomponents that can be removed and replaced in a repair shop. Such a subcomponent is referred to as an SRU (shop replaceable unit).

Currently there are five government-operated depots for the Air Force throughout the US. These Air Force repair depots are referred to as Air Logistics Centers (ALCs), and are under the command of the Air Force Materiel Command (AFMC), which is located at Wright-Patterson Air Force Base, OH. The five depots are Ogden Air Logistics Center in Ogden, UT, Oklahoma City ALC, Oklahoma City, OK, Warner Robins ALC, Warner Robins, GA, San Antonio ALC, San Antonio, TX, and Sacramento ALC, Sacramento, CA. Two of the ALCs, San Antonio and Sacramento, were scheduled to be closed as part of the Base Realignment and Closure (BRAC) Commission, in an attempt to reduce excess infrastructure and capacity. However, the plans for closure were changed to a plan for privatization in place. Some of the repair activity will remain at the two ALCs, performed under private contractors, while some of the workload will be transferred to the three remaining ALCs.

In addition to the government depots, private industry also provides depot level repair under government contract. A significant portion of the depot maintenance, approaching 50%, is now contracted out to industry.[5] Industry involvement in depot level repair is increasing, as maintenance is becoming a more significant portion of the available workload for aerospace defense contractors, and as the push to greater efficiency is causing politicians and decision makers to consider privatizing more of the maintenance workload.

Repair takes place at several different levels, also referred to as echelons. Items that are designated for three-level maintenance (3LM) are removed from aircraft on the flight line, then sent to an intermediate repair facility. The intermediate repair facility may be located on the operational base, or may be located in a more central location to serve several bases. The intermediate repair facility has limited repair capabilities, but is able to troubleshoot and diagnose problems, remove and replace subcomponents, and make simple repairs. Items that cannot be repaired at the intermediate repair facility are declared NRTS (Not Repairable This Station, pronounced as "nerts") then sent on to a central depot for repair.

Items that are designated for what is called two-level maintenance (2LM) are removed from aircraft on the flight line and sent to a central repair depot without repair first being attempted at an intermediate level repair facility. The central depot is often far from the operational base.

The advantages of the three-level system with the intermediate facilities are shorter average pipeline flow times and thus fewer required spare components in the inventory system, while the advantages of the two-level system are reduced infrastructure and personnel. In the past repair was primarily 3LM, but recently the Air Force switched many repair processes from 3LM to 2LM, expecting the reductions in necessary manpower and equipment to save money. One of the assumptions that motivated this change was that anticipated simultaneous improvements in transportation and repair processes would allow repair pipeline flow times to decrease, and thus not require massive investments in additional spares that a longer repair pipeline would need. The change from 3LM to 2LM was made largely before these pipeline improvements were a

---

[5] 10 U.S.C. 2466. As federal law, the "50-50 rule" stipulates that no more than 50% of the funds allocated to a military department or defense agency for depot-level maintenance and repair work may be used to contract the workload to private industry. This was recently increased from the "60-40 rule" that limited private workshare to 40%.

reality, and the performance of the maintenance system suffered. Increased spare parts shortages caused more aircraft to be not mission capable (NMC).

Due to constraints on financial and repair resources and the need to focus on repairing those parts that are most needed, depot repairs are prioritized before they are scheduled. When a repairable aircraft part fails and is sent from its air base to be repaired, it first goes into a queue of repairable components inventory before it is repaired at a depot. The Air Force Supply Management Activity Group (SMAG) manages this repairable components inventory, and is responsible for the overall management of repairable items. This inventory is generally maintained and handled by the Defense Logistics Agency (DLA), which is responsible for shipping, warehousing, and inventory control of spare parts and the many other items used by the military. The DLA provides services for all branches of the military as well as other federal agencies. The repairable inventory may be located in a central warehouse but is often located at the depots. Repair needs are prioritized Air Force-wide, through complex algorithms that take into account the effect of the needed part on Air Force readiness, the availability of repair resources, and many other factors. The highest ranked items are to be repaired first, as funds are not typically available to repair all broken parts. Once the need for the broken part becomes critical enough that it is ranked high enough to be funded, the part is brought into repair at the depot facility. In 1997, when this prioritization process was still new, repair budgets were so constrained that typically only those parts that were resulting in aircraft being grounded were critical enough to be repaired.

Once components are repaired they are referred to as serviceable and are placed into the inventory system for serviceable components and stored in DLA warehouses. From there SMAG directs DLA to distribute the serviceable components to replenish base inventories as needed.

Although repairable parts can often be repaired many times, they eventually reach a point when they can no longer be returned to service. For example, in the case of a wheel, the hole for the axle can only be bored or honed to remove damage so many times before the hole becomes too large for tolerances. Larger and larger bushings could be used, but eventually a limit is reached. When repairables reach the limit of repair, they are condemned. New parts may need to be ordered from suppliers to replace those that were condemned.

There are two different modes of maintenance for Air Force aircraft: scheduled and unscheduled maintenance. Up to this point primarily unscheduled maintenance has been discussed. Unscheduled maintenance responds to parts that fail during use. Scheduled maintenance on the other hand is intended to be preventative. Scheduled maintenance includes time dependent maintenance actions such as isochronal inspections, time compliance technical orders (TCTO), and programmed depot maintenance (PDM). Scheduled maintenance actions are planned and funded in advance and are not dependent on the prioritization process that occurs for unscheduled maintenance.

The repair and replacement of components requires the support of industrial suppliers. These suppliers provide the materials and parts necessary to complete repairs. This includes everything from nuts and bolts to complex subcomponents or even replacement components for those parts that are too old or damaged to repair and need to be condemned. In the past, contracts were often awarded to the lowest cost bidder, but a major effort is under way to improve the contracting process and to award bids to the best value bidder.[6] Improving contractor performance and speeding the supply of parts is a major challenge the Air Force is facing.

This brief introduction and overview of the sustainment system provides only a glimpse into the whole system. There are many other aspects that have not been discussed, such as the financial system and the information system, which are beyond the scope of this thesis. The improvement of these systems as well is vital to the future of the Air Force.

2.2 THE POST-COLD WAR SUSTAINMENT ENVIRONMENT

In addition to understanding the basics of the sustainment system operations, it is important to have an understanding of the changing environment in which the sustainment system must operate. The environment surrounding the sustainment system

---

[6] This effort is part of a program called CREP, for Contract Repair Enhancement Program, which will be introduced later in this chapter.

has much to do with the constraints imposed upon the system, and impacts the sustainment system's performance and behavior.

The United States Air Force is undergoing a historic shift to adapt to the post-Cold War world. In the 1980's during the Cold War, the Air Force prepared primarily for a massive conflict with the Soviet Union. The adversary and battlefield were known and the threats were largely understood. In the post-Cold War world of the 1990's, the adversaries and threats are not as clearly known. The U.S. Air Force must be prepared to face multiple regional conflicts and humanitarian missions simultaneously. Furthermore, as the Air Force reduces its fixed forward infrastructure, such as bases in Europe that were designed to counter the Soviets, it must increasingly deploy from bases in the US. This adds to the logistics burden.

The Air Force must be prepared for both peacetime and war, a challenge that requires many tradeoffs. These two modes have often contradictory requirements and expectations and drive towards different behavior. The Air Force is expected to operate efficiently during peacetime and not waste taxpayers' money. Yet it is also expected to be ready to fight at any moment. The need to operate cost-effectively during peacetime suggests fewer resources. The need to be ready for war suggests more resources. The Post-Cold War environment of the 1990's has emphasized the divide between being prepared for peacetime and prepared for war. Greater fiscal pressure has increased the need to be cost conscious and efficiently use resources. A more unpredictable world where conflicts require deployments and do not allow long periods for mobilization stresses the need to be prepared with little or no advanced notice.

Although it has been nearly a decade since the end of the Cold War, the US Air Force is still in a time of transition and is facing difficult challenges. Overall, defense budgets have been significantly reduced while force deployments and operations have increased, requiring the Air Force to maintain high levels of activity with fewer forces and fewer resources. These factors have placed strains on the Air Force system and its combat effectiveness. Because of the changing global environment, decreased forces, and increased activity, the emerging Air Force strategy emphasizes readiness, flexibility, mobility, and responsiveness. Due to budgetary constraints, costs are becoming a much greater consideration. Improved logistics and maintenance operations plays a major role in enabling the Air Force to implement its strategy effectively.

2.2.1 Decreased Resources

To provide a more detailed understanding of the changes in the Air Force environment it helps to look at some historical data that illustrates the decrease in Air Force resources. Since the end of the Cold War, the Air Force has been steadily downsizing, much as it has after previous major wars. After World War II and after Vietnam, the Air Force budget was cut dramatically, to the point that capability was severely affected and readiness for subsequent conflicts was rather low. Similarly, defense budgets have decreased dramatically throughout the 1990's, as shown in Figure 2-3. This decrease occurred despite the Gulf War and other ongoing regional actions. The budget is now 57% of it Cold War peak (reached at the end of the Reagan-era buildup), on par with the budget level corresponding with the "hollow Air Force" of the 1970's.[7]

Reflecting the overall defense budget, the budget level for operations and maintenance activities, shown in Figure 2-4, has seen a similarly dramatic decrease. Only very recently have alarms been given as to the severity of the decrease in readiness that in part is resulting from budget reductions. In response, plans have recently been made to slightly increase future year budgets. Coinciding with the reduced budgets over the past decade, force structure, personnel levels, and the number of bases have been sharply reduced. Figure 2-5 illustrates the decrease in personnel levels, and Figure 2-6 illustrates the trend in force structure levels by illustrating the decrease in fighter and attack aircraft.

---

[7] During the 1970's, Air Force readiness declined dramatically. For example, overall mission capable rates (the percentage of aircraft capable of flying missions)for fighter aircraft fell to a low of 54% in 1978.

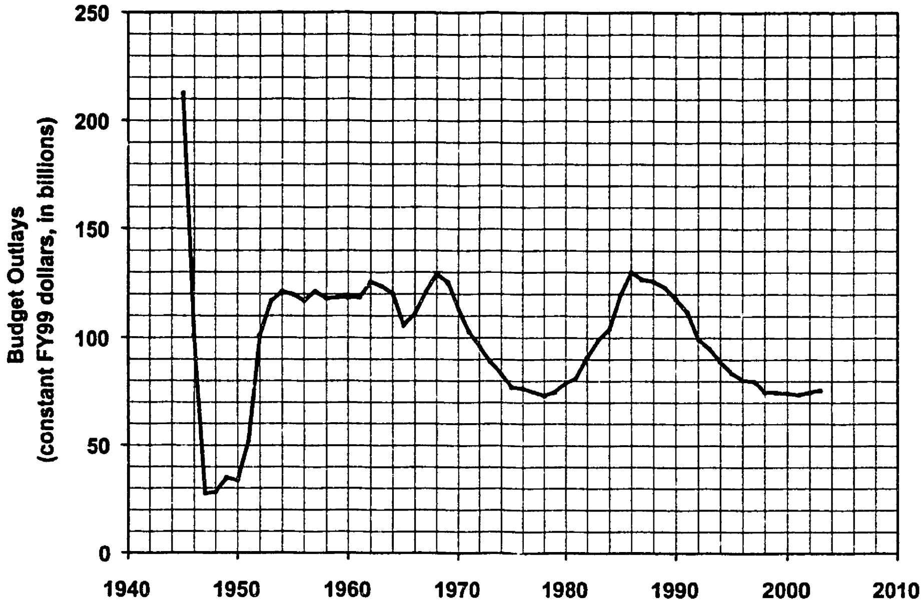

Figure 2-3 Air Force Budget Outlays[8], including projections.

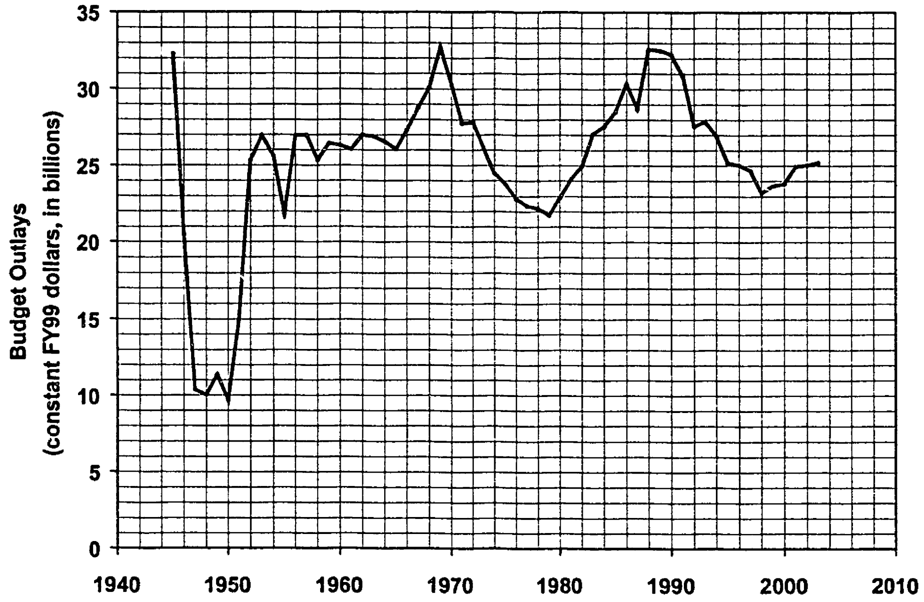

Figure 2-4 Air Force Operations and Maintenance Budget Outlays[9], including projections.

---

[8] Data from Office of the Under Secretary of Defense, *National Defense Budget Estimates for FY 1999*, (U.S. Department of Defense, 1998).

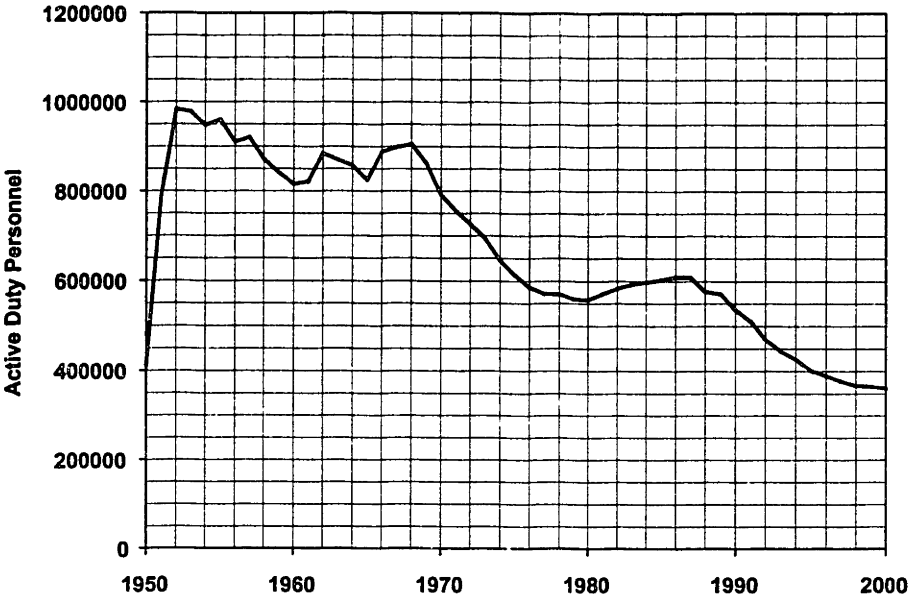

Figure 2-5 Personnel Levels

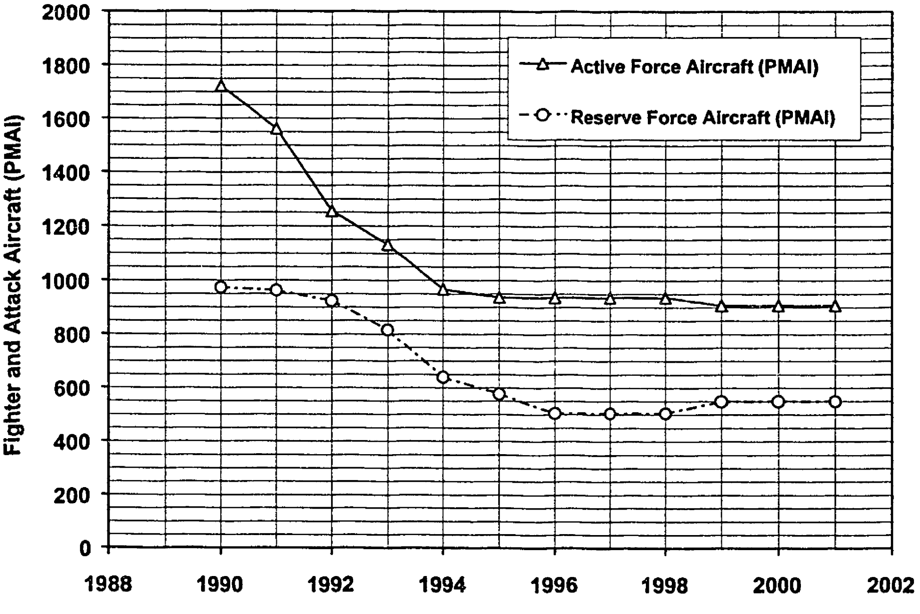

Figure 2-6 Number of Fighter and Attack Aircraft

---

[9] Data from Office of the Under Secretary of Defense, *National Defense Budget Estimates for FY 1999*, (U.S. Department of Defense, 1998).

2.2.2 Increased Demands

While the Air Force has been downsizing, the workload has remained high. This is one of the main factors behind the challenges the Air Force is facing. There are several factors that have contributed to the increased demands on the now smaller Air Force system.

2.2.2.1 Deployments and Operations Tempo

Deployments and high operations tempo strongly impact the maintenance and logistics system as well as the Air Force as a whole. Increases in operations tempo generally lead to increased demands for maintenance.[10] Deployments also strain the maintenance system, as forces deployed into areas of conflict must be maintained at high levels of readiness, often at the expense of the rest of the system.

With the end of the Cold War, the world has become more politically unstable and the Air Force has been called upon to serve in a greater number of actions around the world. While the Air Force has been downsizing since the end of the Cold War, the number of forces deployed from their home bases has actually increased 400% from Cold War levels. This increased level of activity has been fairly consistent throughout the 90's, from the Gulf War to the Balkans. The USAF has been flying a large number of missions in support of the Bosnian peace agreement and has also been flying missions over the northern and southern no-fly zones in Iraq. Aircraft flying in Iraq have been regularly involved in combat. Aircraft are also currently seeing action in Kosovo and Serbia. Table 1 gives an idea of the number of missions that the Air Force undertook in the past year, before operations in Serbia and Kosovo. In a sense, with ongoing missions in Iraq and Bosnia, the Air Force was operating on a constant war footing with a high operations tempo, even before NATO operations in the Kosovo crisis. This places greater demands on the Air Force maintenance and logistics system as well as the Air Force as a whole.

---

[10] Raymond A. Pyles and Hyman L. Shulman, *United States Air Force Fighter Support in Operation Desert Storm*, MR-468-AF, (Santa Monica, CA: RAND, 1995). pp.17-18,24,45. Many factors influence maintenance demands, but generally demands on maintenance increase as the rate of flights increases.

Table 1: Number of Air Force missions during 1998[11]

Number of sorties in the Balkans	2,200
Number of sorties in southwest Asia	27,000
Airlift missions	30,000 to 90 countries
Exercises	1,600 in 35 countries
Military-to-military contact visits	300
Humanitarian relief missions	100 to 30 countries

2.2.2.2 Aging Aircraft

The issue of aging aircraft is also significant for the Air Force maintenance and logistics system. While the issue of aging aircraft may seem to be a supply-side issue in terms of the Air Force meeting the increased demand for missions, the issue of aging aircraft itself increases demands on the Air Force sustainment system. As aircraft age, they become more difficult and more expensive to maintain. Material degradation caused by fatigue and corrosion becomes a greater concern for the airframe. This requires more intensive and more frequent inspection and maintenance. As components near the end of their service lives they fail more frequently. Parts that rarely failed may begin to fail more frequently. Obsolescence becomes a concern with many high technology components, as manufacturers may not wish to or may not be able to produce spares for components that are outdated. Another problem that the Air Force faces with aging systems is that some of the original contractors are no longer in business.

Many of the aircraft platforms used by the Air Force are reaching or have exceeded their planned service lives. However, many will remain in the active inventory for decades. For many aircraft, there are no current plans for future replacement. For example, the C-5A has been in service since 1969, and is planned to remain a core part of the airlift fleet until at least 2045.[12] The B-52G entered service in 1962, and is also projected to remain in service if necessary until around 2045.[13] The current average age of USAF aircraft is 20 years, and will increase to 30 years in the year 2015 despite

---

[11] United States Air Force, *Posture Statement 1999*, http://www.af.mil/lib/afissues/1999/posture/index.html, 1999. p.7. This is before the action in Kosovo.

[12] USAF Air Mobility Command, *1998 Air Mobility Master Plan*, (HQAMC, 1998).

[13] U.S. Department of the Air Force, *U.S. Air Force White Paper on Long Range Bombers*, (1999)., p.21.

current modernization plans.[14] Maintaining such an aged fleet will continue to place increasing stresses on the maintenance and logistics system. Overall, the Air Force fleet is smaller and older, creating a challenge to readiness.

2.2.3 Declining Readiness

The Air Force has been struggling to meet the challenges of supporting more deployments and operations with a smaller defense budget, fewer and older aircraft, fewer spare parts, and fewer personnel. The difficulty of the challenge has become apparent, as readiness rates have plummeted. According to the 1999 Air Force Posture Statement, major unit readiness decreased by 18% in the past two and a half years, and stateside combat readiness declined by 56%.[15]

Stateside units suffer greater decreases in combat readiness than forward-deployed forces because their resources are taken to enhance the readiness of forward-deployed units. Stateside forces in a sense are a source of supply in times of need for units that are deployed. Non-deployed aircraft are cannibalized as necessary to support deployed aircraft. Pyles and Shulman note that with the decrease in the size of the Air Force, the effectiveness of stateside non-deployed units as a resource for deployed units is diminished because there are fewer resources to draw upon.[16] This suggests that combat readiness will decrease as the number of aircraft decreases, and may explain at least a portion of the dramatic decrease in readiness in recent years.

2.2.3.1 Declining Mission Capability

One measure of aircraft readiness is the mission capability (MC) rate. MC rate is the percentage of aircraft that a unit possesses that are capable of flying missions. As shown in Figure 2-7, the average percentage of Air Force aircraft that are mission capable has declined in recent years, falling from 84.6% in 1990 to 74.3% in 1998. The most

---

[14] United States Air Force, *Posture Statement 1999*, http://www.af.mil/lib/afissues/1999/posture/index.html, 1999. p.26.

[15] United States Air Force, *Posture Statement 1999*, http://www.af.mil/lib/afissues/1999/posture/index.html, 1999. p.19.

[16] Raymond A. Pyles and Hyman L. Shulman, *United States Air Force Fighter Support in Operation Desert Storm*, MR-468-AF, (Santa Monica, CA: RAND, 1995)., p. 98. Pyles and Shulman specifically refer to the decrease in the number of aircraft of a given type effecting their combat readiness. As the Air Force is roughly keeping the same number of types of aircraft as it reduces the number of aircraft, they are generally reducing the number of aircraft of a given type.

pronounced drops have occurred in 1998, when MC rates for the Air Force fell from 76.6% to 74.3%.[17] The MC rate for fighters is shown in Figure 2-8 to further illustrate the trend. With both the number of fighters decreasing (as seen in Figure 2-6) and the MC rate falling, the total number of mission capable fighters is falling.

Much of the decrease in mission capable rates is due to problems with supply of spare parts. The percentage of aircraft that are not mission capable due to lack of spare parts has been steadily increasing over the past decade. Figure 2-9 illustrates this increase by showing the average Total Not Mission Capable Supply (TNMCS) rating for Air Force aircraft since 1990. The TNMCS rating measures the percentage of aircraft that are not mission capable due to lack of spare parts supply. The TNMCS rating has more than doubled, increasing from approximately 6% in 1990 to 14% in 1998. This accounts for the majority of the decrease in overall MC rates.[18]

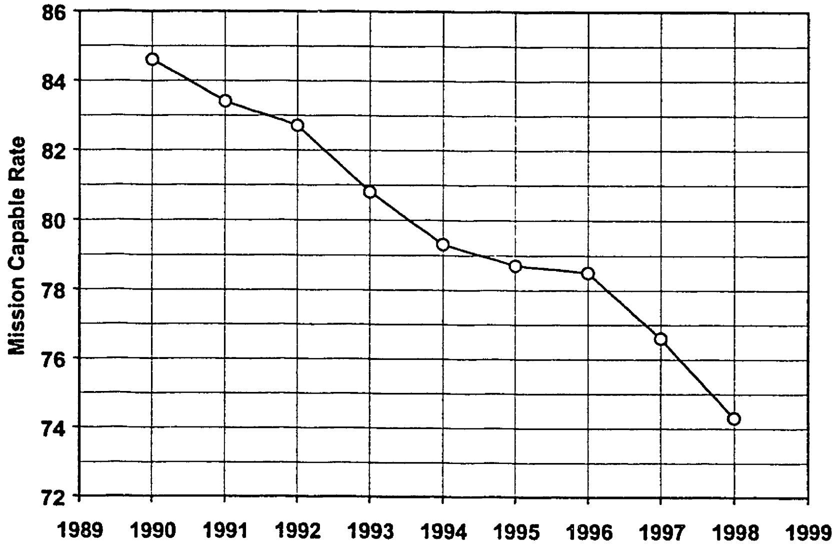

Figure 2-7 Mission Capable Rate for Air Force major aircraft.[19]

---

[17] Data from United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).

[18] United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).

[19] Data originally from the Air Force's Multi-Echelon Resource and Logistics Information Network (MERLIN) and the Reliability and Maintainability Information System (REMIS), as reported by United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).p.21. MC rates of major aircraft do not include helicopters, training aircraft, or aircraft with very low quantities.

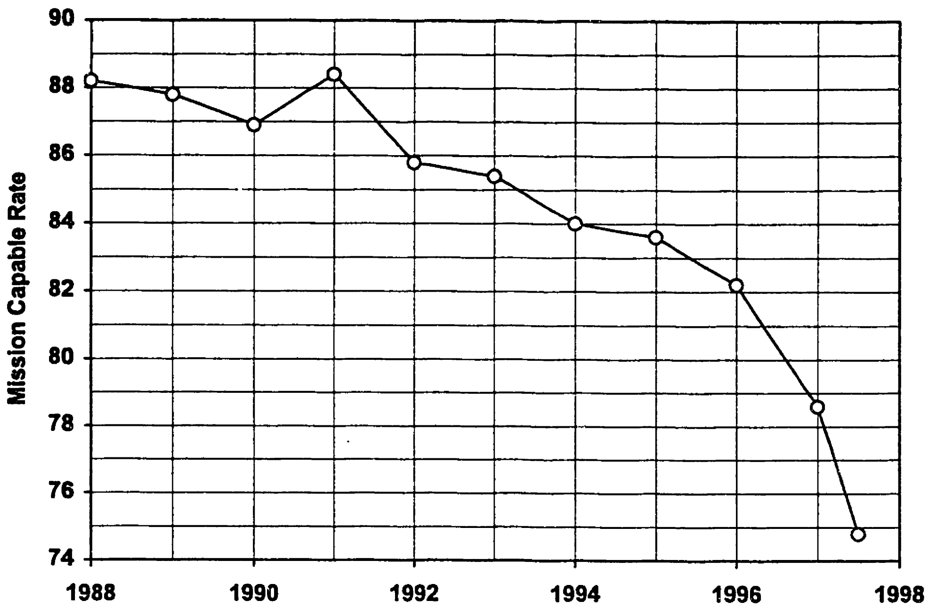

Figure 2-8 Mission Capable Rate for USAF Air Combat Command Fighters[20]

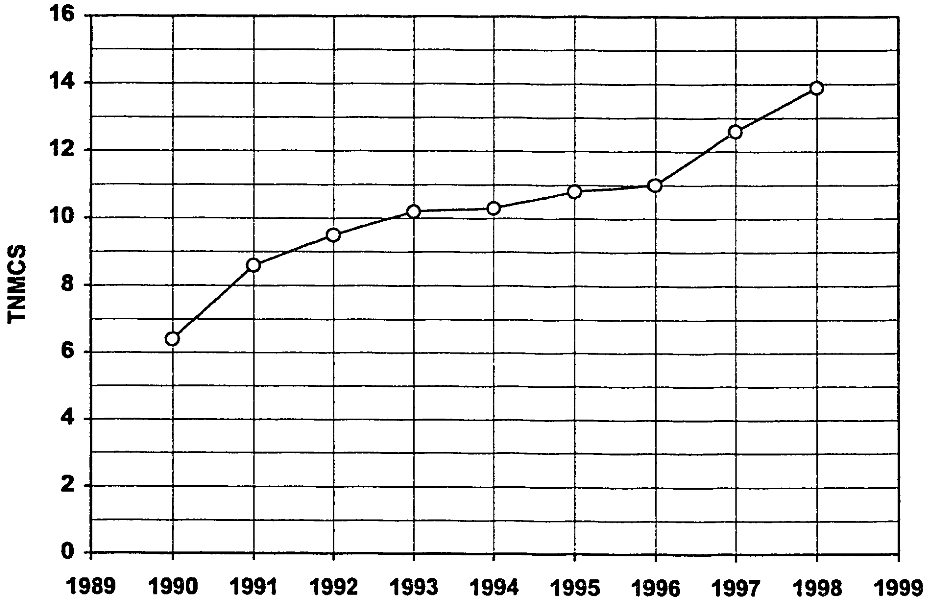

Figure 2-9 Total Not Mission Capable due to Supply rating for Air Force major aircraft.[21]

---

[20] Data from USAF Air Combat Command, *Air Combat Command Director of Logistics Ten Year Lookback Standards and Performance FY88-FY97*, (USAF HQ ACC/LGPA, 1997).

2.2.3.2 Personnel Retention Difficulties

Due to several factors, including the stresses on personnel due to high operations tempo and frequent deployments, and a strong economy, personnel retention rates Air Force-wide have dropped in recent years. Critical fields for the sustainment system such as avionics and aircraft maintenance are among those having the most difficulty retaining skilled personnel. Retention problems have been especially serious for mid-level non-commissioned officers (NCOs), a vital core of experience, skill, and leadership.[22] Retention problems can form a vicious cycle, or a reinforcing process to use the system dynamics terms of Senge[23]. The loss of skilled mid-level NCOs places more strain on personnel that remain, creating even more retention problems.

On the civilian side of the maintenance process, ALCs are reportedly also having some difficulty attracting and retaining skilled technicians. A strong economy and a disparity in pay and benefits are significant factors. The relocation of workloads from one ALC to another has also caused difficulties for retention in some cases, as a significant percentage of experienced laborers and technicians are unwilling to relocate. Overall, retention problems in both the military and civilian sides of Air Force sustainment system are having a significant effect on the maintenance and logistics system's readiness and effectiveness.

2.2.3.3 Increased Cannibalization and Use of Readiness Spares Packages

Due to tight budgets and increased demands, funding for spare parts and for repairs has not been able to meet needs, and thus cannibalization is becoming more and more necessary. Cannibalization is the taking of parts from an aircraft for use in the repair of another aircraft. The issue of cannibalization will be discussed in more depth in Chapter 4. Cannibalization typically occurs when spare parts are not in supply and allows a

---

[21] Data originally from the Air Force's Multi-Echelon Resource and Logistics Information Network (MERLIN) and the Reliability and Maintainability Information System (REMIS), reported by United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).p.22. TNMCS rates of major aircraft do not include helicopters, training aircraft, or aircraft with low quantities, such as the SR-71.

[22] United States Air Force, *Posture Statement 1999*, http://www.af.mil/lib/afissues/1999/posture/index.html, 1999. p.20.

[23] Peter M. Senge, *The Fifth Discipline: The Art & Practice of The Learning Organization*, 1st ed. (New York: Doubleday, 1990).

squadron to consolidate failed parts that lack spares into the fewest number of aircraft. Cannibalization also sometimes occurs when supply cannot deliver the required spare in a timely manner. Squadrons are rated by such metrics as their on-time performance and their mission capable rate. However, they have little control over the supply of spare parts. Cannibalization gives them the means to deal with supply problems and helps the squadrons' performance ratings. Over the past several years aircraft cannibalization has been increasing steadily, and in the past two years, sharply. The cannibalization rate, the number of cannibalizations per 100 sorties, is used as an indicator of the health of the maintenance and logistics system, as excessive cannibalization is seen as undesirable. According to the Air Force Posture Statement 1999, Air Force-wide, the cannibalization rate has increased 78 percent since 1995.[24] Figure 2-10 illustrates the increasing trend in the cannibalization rate for fighter aircraft from 1988 to 1997. Figure 2-11 illustrates the cannibalization rate for three representative aircraft, the B-1B bomber, C-5 transport, and the F-16 fighter. All show significant increases in cannibalization.

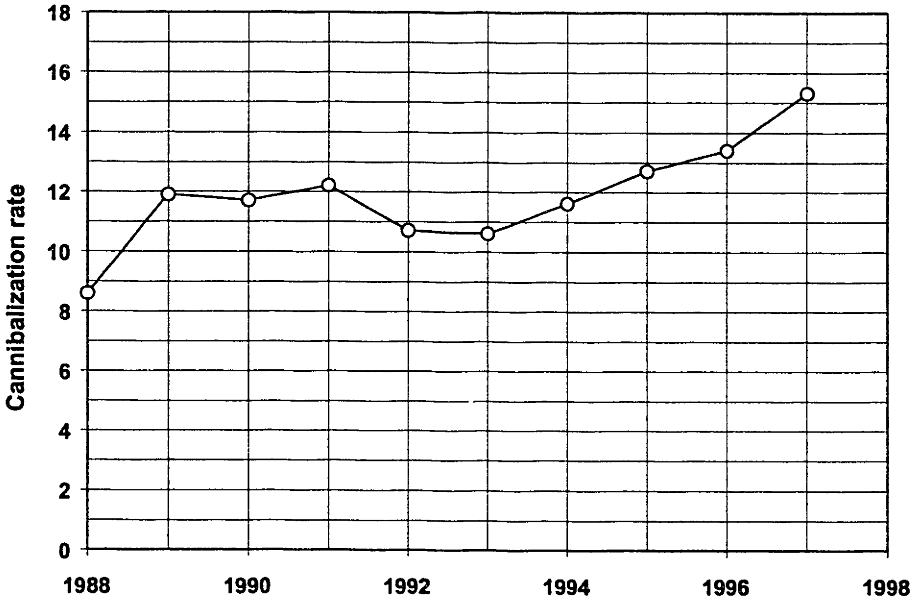

Figure 2-10 Cannibalization Rate for USAF Air Combat Command operational fighters[25]

---

[24] United States Air Force, *Posture Statement 1999*, http://www.af.mil/lib/afissues/1999/posture/index.html, 1999. p. 26.

[25] USAF Air Combat Command, *Air Combat Command Director of Logistics Ten Year Lookback Standards and Performance FY88-FY97*, (USAF HQ ACC/LGPA, 1997).

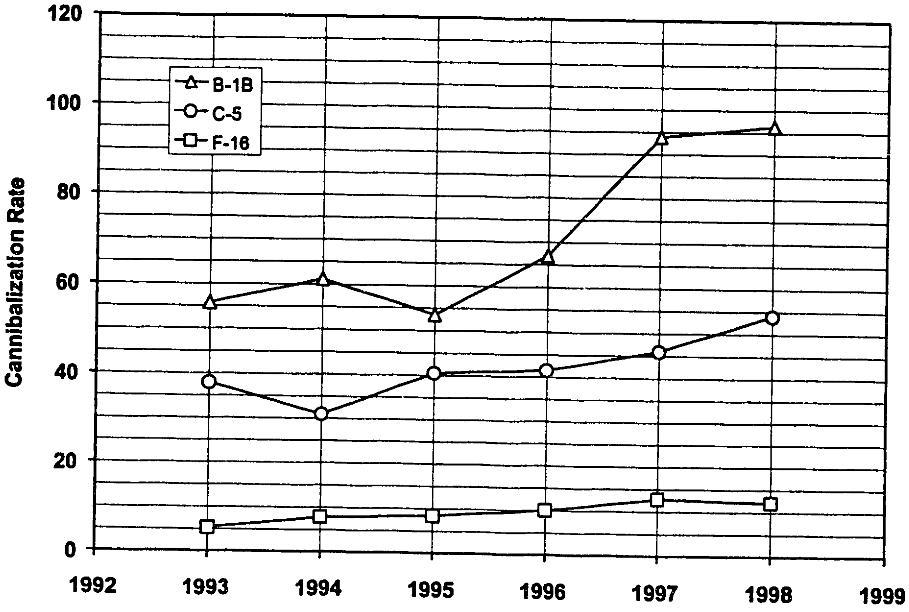

Figure 2-11 Cannibalization Rate for B-1B, C-5, and F-16.[26]

The use of parts from Mobility Readiness Spares Packages (MRSP) to fill spare parts needs has also been increasing. Mobility Readiness Spares Packages are kits of spare parts and supplies that are prepared for potential use in squadron deployments. The packages are air transportable and are designed to be deployed with a squadron of aircraft to provide a source of spare parts in the early stages of a deployment when supply lines may not yet be established. These kits of supplies are critical in supporting the early stages of wartime actions, and are meant to be kept prepared so that rapid deployments can be supported. If these kits are used to support peacetime operations because of parts shortages, wartime preparedness is sacrificed. When these kits are then needed for deployments, such as in the case of Desert Storm, aircraft that are not being deployed will need to be cannibalized to fill up the spares kits so that deployed war-fighting aircraft can be maintained.

---

[26] This data for this chart is documented and discussed in United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999). p.23. The data is originally from Air Combat Command and Air Mobility Command.

Maintenance personnel at a C-5 wing and an F-15 wing that LSI researchers visited reported that the use of MRSPs to support peacetime operations was becoming excessive. The soldiers reported that spares kits were typically only about 70% complete, due to parts being taken out when spare parts were not available, and that in certain cases it was worse than 70%. In a recent study, the General Accounting Office (GAO) made similar observations, noting, "This practice, in effect, trades off the Air Force's ability to sustain its forces during wartime for the ability to prepare its peacetime forces for war."[27]

The use of MRSPs allows greater stockage effectiveness, which is a measure of the percent of time that base supply will have a part available when one is requested. Air Mobility Command reported an average stockage effectiveness of about 85% for July 1997 to June 1998. However, this relied heavily upon using MRSPs to provide parts. Excluding the parts taken from MRSPs, the average stockage effectiveness would have been about 45%.[28]

2.2.4 Air Force Adaptations

The Air Force has implemented many programs and changes to adapt to changing times. The Air Force has been rather active in attempting process improvements and is attempting to become more lean. Elements of lean thinking can be found in several programs and ideas implemented in the past several years. One example is the expanded use of rapid transportation for the movement of parts and supplies. This can help reduce the overall time required for repairs by improving the responsiveness of the logistics system. Another is the restructuring of the Air Force with the Expeditionary Aerospace Force (EAF) concept, which is designed to improve force flexibility and responsiveness in deployments and improve the quality of life for personnel. A third example is the Lean Logistics or Agile Logistics program.

Under the EAF concept, the Air Force will be organized into 10 Aerospace Expeditionary Forces (AEFs). Although the units in an AEF will be drawn from

---

[27] United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).

[28] USAF Air Mobility Command, *HQ AMC Orientation*, (Scott AFB, IL: HQ AMC/LG, 1998).

geographically separated bases and will include active duty, Air Force Reserve, and Air National Guard units, the AEF will deploy as one combined force. The AEFs will be a total force, consisting of combat aircraft and all necessary supporting aircraft, spacecraft, and ground systems and units.

A major advantage of the EAF concept is the stability of deployment scheduling. The cyclical 15-month schedule consists of 3-month "vulnerability periods" in which 2 AEFs must be ready for deployment. Units and personnel can prepare in advance for their potential deployment and know that the potential deployment is limited to 3 months. This improves preparedness and eases stresses on the quality of life for personnel.

Under the Lean Logistics program office (now replaced by Agile Logistics), the Air Force Materiel Command implemented a series of initiatives to improve the maintenance process. These initiatives were called DREP, AREP, and CREP and each applied lean concepts to different aspects of Air Force maintenance. Each program worked on improving communication, information flow, and teamwork, and each emphasized placing focus on the war-fighters' needs. DREP, for Depot Repair Enhancement Program, sought to reduce the time required to perform depot repair of broken parts by shifting from quarterly negotiated batch processing of repairs to a repair on demand system. AREP, for Aircraft Repair Enhancement Program, focused on reducing the time required to overhaul or upgrade an aircraft by applying lean principles to the Programmed Depot Maintenance (PDM) of aircraft. CREP, or Contract Repair Enhancement Program, aimed to improve contractor performance by changing contract policies to include lean practices such as performance incentives and long term contracts.

Although DREP and AREP made significant improvements in the maintenance process, the benefits have been hampered by parts supply problems. Depots report that their number one problem is the supply of parts that they need to make repairs. Untimely supply of subcomponents and parts necessary for the depots to repair broken components in turn delays the repair of the components necessary to repair aircraft. Thus CREP, or at least the improvement in contractor performance that it aims to bring about, is a vital necessity. Such improvements in contracting and in contractor performance are important to improve the parts supply and logistics support of the depot repair process.

2.2.5 Summary

The Air Force has been facing a period of difficult transition. With the end of the Cold War and the new threats that must be faced, there have been many changes. Declining budgets and increasing activity have strained the system. In response, the Air Force has been implementing changes in an effort to adapt and become a more lean fighting force. Improvements have been made, but improvements in one area continue to reveal troubles in others, such as the shortage of subcomponents and parts. There is still much to be done before the Air Force can be considered lean.

CHAPTER 3: BACKGROUND

Much valuable work on understanding and modeling the Air Force sustainment system has already been performed. This chapter provides a brief discussion of earlier work. RAND has published a series of works on spare parts demand and on capability assessment that are relevant to this thesis and provided important background. Similarly the Logistics Management Institute has produced some important work on modeling Air Force capability and determining optimal spares kits. The General Accounting Office has produced many reports quite germane to this thesis that provide a broad background covering almost all aspects of the sustainment system.

3.1 MODELING AND FORECASTING AIR FORCE CAPABILITY AND SPARE PARTS DEMAND

Project AIR FORCE, a division of RAND Corporation, has been deeply involved in research on Air Force logistics support since the early years of the Air Force. RAND is an independent, nonprofit institution that provides research and analysis to policymakers. Through a multitude of projects and reports, their research has helped to deliver insights into the complex system behavior of the Air Force maintenance and logistics system. From research in the 1950's on spare parts demand forecasting to current research on Lean Logistics, there is much to be learned from their work.

RAND addressed two issues of particular interest to this thesis. The first is modeling Air Force capability and the second is identifying and characterizing the high variability in spares demands. Recent RAND work on capability assessment revolves around RAND's Dyna-METRIC model.[29] In *Dyna-METRIC Version 6: An Advanced Capability Assessment Model*, Isaacson and Boren introduce version 6 and illustrate its potential uses. The Dyna-METRIC model, originally developed by R.J. Hillestad[30], has been under development for nearly 20 years, evolving from an analytic model of theater

---

[29] Karen E. Isaacson and Patricia M. Boren, *Dyna-METRIC Version 6: An Advanced Capability Assessment Model*, R-4214-AF, (Santa Monica, CA: RAND, 1993).. See also Karen E. Isaacson and Patricia M. Boren, *Dyna-METRIC Version 5: A Capability Assessment Model Including Constrained Repair and Management Adaptations*, R-3612-AF, (Santa Monica: RAND, 1988).

[30] R.J. Hillestad, *Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control*, R-2785-AF, (Santa Monica, CA: RAND, 1982).

tactical Air Force capability to a simulation version with global perspective. Dyna-METRIC Version 4 has been incorporated into the Air Force's Weapon System Management Information System (WSMIS) as its capability assessment model.

Dyna-METRIC models the Air Force maintenance system and evaluates wartime readiness and sustainability based on logistics resources and pipelines. The model is very flexible and can accommodate most scenarios. Variables such as the number of aircraft, the flying schedule, and the turn rate, can vary with time and with the base of operation. The number of bases and the deployment of aircraft to certain bases can change with time.

Dyna-METRIC models three echelons of repair activity (air base, intermediate repair, and depot) and considers both components and subcomponents. It can consider cannibalization on three levels: full cannibalization, designated cannibalization, or no cannibalization. Full cannibalization assumes all parts can be cannibalized, while designated cannibalization allows input that designates which components can be cannibalized and which cannot. It allows the option of considering lateral supply and repair among bases, which is the practice of one base using its resources to support another in times of need. It allows priority scheduling of repairs or random scheduling of repairs. Under priority scheduling, repairs are prioritized by the number of grounded aircraft that a particular type of part is causing. Random scheduling selects the next part to be repaired at random from the total group of parts awaiting repair. Dyna-METRIC also accounts for battle damage to repair resources and repair queues, as well as delay of transportation. It can measure daily aircraft availability and sortie generation capability. It can also measure the status of components in the pipeline, such as backorders, number in transit, and number in work. It generates a report of problem parts, and can be used to illustrate the causes of poor performance. While the current simulation based Version 6 of Dyna-METRIC does not have the capability to directly calculate spares requirements, previous analytical (as opposed to simulation) versions of the model, such as Version 4, do have this capability.

Dyna-METRIC has system wide scope and significant use in evaluating larger issues. It has been used to evaluate two-level vs. three-level maintenance as well as several other management practices and structures. In *Evaluations of Alternative*

Maintenance Structures, Abell and Shulman utilized Dyna-METRIC to perform a study on two-level vs. three-level maintenance.[31] They determined that since the two-level system needs either greater investment in spare parts or a more responsive repair system, two-level maintenance could be more cost effective only if the depot repair system significantly improved in delivering "relevant, timely, and robust... repair".[32] In making this determination they modeled the system under its then current practices, then evaluated scenarios with expedited handling, improved processes, and improved demand determination.

Another important topic of research that has relevance to any Lean Sustainment modeling effort is the issue of uncertainty. RAND has examined this issue in depth, providing useful insight on the validity of any model that predicts Air Force system behavior. Crawford initiated much of this work, identifying and exploring two major sources of uncertainty:

1. the extreme magnitude of the variability in demands for aircraft spare parts
2. the variability in the repair pipeline

and exploring implications of these variabilities.[33]

In the report *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts*, Adams, Abell, and Isaacson provide a good review of the decades of research on uncertainty in Air Force spare parts demand.[34] It is interesting to note that many of the goals of Lean Logistics, such as reducing resupply time, procurement lead time, and repair-cycle time, were espoused over 40 years ago as a means to counter demand uncertainty, in a report by Brown in 1956[35]. Although Adams et al. propose an

---

[31] John B. Abell and H. L. Shulman, *Evaluations of Alternative Maintenance Structures*, R-4205-AF, (Santa Monica, CA: RAND, 1992).

[32] In the report summary, the emphasized recommendations were that two-level maintenance was more cost effective and more beneficial than three-level. The main caveats were not emphasized in the summary, but were strongly stated in the results section of the report.

[33] Gordon B. Crawford, *Variability in the Demands for Aircraft Spare Parts: Its Magnitude and Implications*, R-3318-AF, (Santa Monica, CA: RAND, 1988).

[34] John L. Adams, John B. Abell, and Karen E. Isaacson, *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts*, R-4211-AF/OSD, (Santa Monica, CA: RAND, 1993).

[35] B. B. Brown, *Characteristics of Demand for Aircraft Spare Parts*, R-292, (Santa Monica, CA: RAND, 1956).

improved methods for modeling and forecasting demands, the overall lesson is that demands for spare parts even in peacetime are highly variable and difficult to predict. Thus forecasting demands is very difficult, which makes effective stockage of spares rather challenging.

Predicting aircraft spare parts needs for wartime is an even more daunting task than for peacetime, as shown by Pyles and Shulman in their study of the Gulf War.[36] The Gulf War demonstrated that the specific operational environment for a given aircraft dramatically affects the type and number of part failures an aircraft experiences. Such operational environments are often vastly different then those expected or planned for, as the Air Force must adapt to rapidly changing battlefield situations. Pyles and Shulman identified six broad sets of factors that affected spare parts demand variations: changing tactics, new technologies, changing campaign plans and missions, changing performance criteria or tolerances, changing demand processes, and unexpected support constraints. Because of these factors, peacetime predictions for spare parts needs prove to be inaccurate in times of conflict. The traditional system of "pushing" inventory of spare parts into the battle theater has inefficiencies, as unexpected needs for spare parts cannot be filled, and unneeded parts consume logistics resources. Responsive solutions in the Gulf War such as Desert Express, a daily express cargo flight of critical parts and supplies, were found to be valuable for meeting changing demands. This reinforces the argument for agile or lean logistics, which focus on logistics system responsiveness rather than stockpiling anticipated inventory needs.

3.2 MODELING TO OPTIMIZE SPARES SUPPORT

The Logistics Management Institute (LMI) has also contributed heavily to the modeling of Air Force issues. LMI is a private, nonprofit corporation that provides consulting, research, and analysis to the government and other nonprofit organizations. Of particular interest to this thesis is LMI's Aircraft Sustainability Model (ASM). The Air Force currently uses the ASM for determining the composition of spares kits for

---

[36] Raymond A. Pyles and Hyman L. Shulman, *United States Air Force Fighter Support in Operation Desert Storm*, MR-468-AF, (Santa Monica, CA: RAND, 1995).

wartime deployment. As noted earlier, the spares kits are packages of spare parts that are deployed along with aircraft squadrons in time of conflict. The kits provide spares support to the squadrons during the initial stages of a conflict when the logistics pipeline may not yet be established. The ASM has also been used to determine initial spare provisioning for new types of aircraft.

The Aircraft Sustainability Model is a mathematical statistical model that computes optimal mixes of spare parts.[37] The optimal mix is determined by a marginal analysis approach, where each spare part is ranked according to its effect on weapon system readiness divided by its procurement cost. Using this approach, the model determines the spares mix that allows the greatest readiness for a given budget, or similarly, the least cost for a target readiness level.

The ASM, like Dyna-METRIC, is quite detailed and able to consider a wide range of scenarios. It takes into account levels of indenture (LRU vs. SRU), a multi-echelon supply structure, and cannibalization. However, it only looks at the effects of cannibalization on aircraft readiness. It does not quantify the amount or cost of cannibalization.

The Dyna-METRIC model and the ASM are both rather effective models at determining spares mixes and readiness. They have the potential to be used by decision-makers in a more widespread manner to analyze much larger issues.

3.3 UNDERSTANDING AND IMPROVING SYSTEMS

The United States General Accounting Office (GAO) has published many reports that provided background for this thesis. The GAO, the investigative arm of the US Congress and watchdog over the usage of public funds, has wide –ranging authority to conduct audits and evaluations of any program or activity that uses public funds. The GAO generally investigates programs at the request of Congress, but may initiate investigations on its own if they are in the public interest. More than a simple auditor, the GAO evaluates government programs for efficiency and effectiveness and makes

---

[37] F. Michael Slay et al., *Optimizing Spares Support: The Aircraft Sustainability Model*, AF501MR1, (McLean, VA: Logistics Management Institute, 1996).

recommendations on improving the programs' operations. The GAO is quite prolific in performing these investigations and writing reports, and has provided extensive material on such issues as the performance of defense inventory management, depot maintenance, logistics support, and many other wide ranging topics. These reports provide a wealth of general background material, as well as providing relevant data and insights on the Air Force sustainment system.

One recent report of particular interest was a report evaluating the effectiveness of the Supply Management Activity Group (SMAG) in providing inventory such as spare parts to its warfighter customers.[38] The GAO report discussed many of the impacts of the Air Force's supply-related readiness problems, such as decreased mission capability rates, increased usage of cannibalization, the peacetime use of spares from Mobility Readiness Spares Packages, and the loss of revenue due to missed missions. The GAO report on SMAG reports three main reasons that Air Force readiness decreased over the past few years. These three are:

1. problems with inventory requirements forecasting and budgeting;
2. the depots inability to provide timely repair due to subcomponent parts shortages;
3. the Air Force reduced SMAGs budget based on overly optimistic Agile Logistics goals that were not all achieved.

This report provided valuable background material for this thesis, much of which corroborated the observations in Chapter 2 of this thesis. It contained data to add support to many of the observations contained in Chapter 2, such as declining MC rates and increasing cannibalization. Whereas the observations of Chapter 2 were commonly reported by Air Force personnel on LSI site visits, specific data to support the observations were limited to Air Combat Command data. The GAO report provided Air Force-wide data to more adequately support the observations.

---

[38] United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999).

The GAO will soon be publishing a report on the implementation of the Air Force's Agile Logistics program. This report should provide another source of valuable information on a topic highly relevant to this thesis.

An important reference for this thesis was *The Machine That Changed the World*, by Womack, Jones, and Roos.[39] This book provided an interesting introduction to *lean*. Through examining and comparing the world's automobile industries, the authors identified and described principles and practices that they coined "lean". Lean principles include responsiveness to change, minimizing waste, having effective relationships within the value stream, getting the right thing to the right place at the right time in the right amount, a commitment to quality, and continuous improvement. Lean practices are means to achieve lean principles, such as a Just In Time inventory system, seamless information flow, an integrated supply chain, and integrated product teams. This book was the primary source in providing a core understanding of what it means to be lean. In addition to defining what was lean, they illustrated the benefits of being lean. In the automobile industry, lean manufacturers have lower costs, shorter product development cycles, shorter production times, and higher quality.

In this chapter we have discussed earlier works that are relevant to this thesis. In the next chapter, Chapter 4, we will discuss the issue of cannibalization.

---

[39] James P. Womack, Daniel T. Jones, and Daniel Roos, *The Machine That Changed the World: The Story of Lean Production*, 1st Harper Perennial ed. (New York: Harper Collins, Original Published by Rawson Associates, 1990).

CHAPTER 4: CANNIBALIZATION

Cannibalization is an issue of particular interest in this thesis. Cannibalization is the practice of taking usable parts from one item to complete or repair another similar item. The practice is used in a wide variety of environments and industries, from truck fleets to industrial machinery. Cannibalization of aircraft parts is a practice that is used widely in the Air Force today. The practice is controversial because while it is often considered an operational necessity, it is also seen as inefficient and wasteful.

The issue of cannibalization became a topic of discussion for members of the MIT Lean Sustainment Initiative after a series of visits to Air Force bases. The MIT researchers were struck by the apparent inefficiency of utilizing cannibalization and of dedicating one aircraft or more for the purpose of cannibalization.

Currently, each wing of aircraft typically operates with one Cann-bird. A Cann-bird is an aircraft that is designated as the primary source for cannibalization. When mission critical parts break on other aircraft on the base and there are not spares readily available in supply, the parts are removed from the Cann-bird and placed on the other aircraft so that the other aircraft can be mission capable. Thus, the Cann-bird consolidates "holes" or parts shortages into one aircraft. When spares eventually arrive, they then are used to repair the Cann-bird.

4.1 EXCESSIVE USE OF CANNIBALIZATION

Cannibalization is intended to be used primarily as a last resort for supply, in cases where the logistics system is unable to provide necessary spare parts. However, there is a growing perception by many stakeholders in the Air Force system that cannibalization is now relied upon too heavily. As shown earlier in Chapter 2, in Figure 2-10 and Figure 2-11, cannibalization has been increasing in the past several years.

To get an idea of the amount of cannibalization that is performed, consider the cases of Dover AFB and Travis AFB, which operate C-5 transports. From July 1997 to June 1998, Dover AFB averaged about 980 man-hours for cannibalization per month, and Travis averaged about 710 man-hours for cannibalization per month. Dover's rate of 980 man-hours each month equates to more than 6 workers working full time all month on cannibalization, and Travis's rate of 710 man-hours equates to about 4.5 full time workers. The cannibalization rate, the number of cannibalizations per 100 flights, averaged to 51 for Dover AFB, or just over one cannibalization for every two flights. This cannibalization rate is high, but not unusually high, given that the standard established by AMC is a cannibalization rate of 46.6. The standard itself seems to acknowledge that the system is designed to require significant dependence on cannibalization.

In some cases certain parts are canned so frequently that when the spares arrive, they are not installed on the Cann-bird. One squadron spoke of the practice of placing a trailer next to the Cann-bird, so that frequently cannibalized parts could be placed in the trailer rather than be reinstalled on the Cann-bird. This saved the labor of placing the parts on the Cann-bird and removing them the next time they needed to be cannibalized.

In cases of heavy cannibalization, one Cann-bird may not be sufficient. In the case of a deep stock out, where many parts of the same type are needed and there are no more in supply, squadrons may need to begin cannibalizing more than one aircraft. Maintenance crews then cannibalize other aircraft that are for various reasons not in flying condition. Typically squadrons have an aircraft undergoing isochronal inspection (referred to as the "ISO-bird") or an aircraft being refurbished (the "refurb-bird") or aircraft out of service for other reasons, so these aircraft are then used to back-up the Cann-bird. One squadron reported that in desperate cases they have used up to five aircraft at a time (out of a total of about 30 aircraft) as sources for cannibalization.

4.2 CANNIBALIZATION INEFFICIENCIES

Cannibalization causes a high level of maintenance labor inefficiency. The amount of labor to fix a plane using cannibalization is double the labor required to make a normal repair. There is also a chance that the part being cannibalized will be damaged in the process. If the part is damaged, the situation is even worse than before, with two parts that must be repaired. Another factor that adds to the labor inefficiency of cannibalization is that as a matter of policy an aircraft should not be kept as a Cann-bird for more than 45 days. If an aircraft does not fly for more than 45 days, it must be reported up the chain of command as a "hangar-queen", which reflects poorly on the wing's ratings. This 45-day limit is established because experience has shown that aircraft that remain grounded for long periods of time tend to develop additional maintenance problems due to their inactivity. The longer an aircraft sits, the more difficult it becomes to bring it back into service. This means that Cann-birds must be rotated, so roughly every 45 days, a new Cann-bird is cannibalized to replace all the missing parts in the old Cann-bird to bring the old Cann-bird back into normal flight operations. This process of returning a Cann-bird to flyable status, called Cann-recovery, takes a significant amount of time and labor. At Dover AFB Cann-recovery typically requires 10 to 15 maintenance personnel working 12-hour shifts to work 10 to 14 days.[40] During this time there are effectively two aircraft not mission capable as Cann-birds, one as the new Cann-bird, and one as the old Cann-bird undergoing Cann-recovery.

The increase in workload due to cannibalization has impacts on the maintenance workforce. Maintenance crews are already stressed due to the high operations tempo that has resulted from the increased activity of the Air Force. There is a limit to how far maintenance crews can be pushed before they reach their limits and performance begins to suffer. By performing cannibalizations, maintenance crews do not have as much time to perform their other work. Currently, maintenance crews make up for this by working longer hours in order to finish necessary tasks. As military personnel do not earn overtime pay, the immediate cost to the Air Force is not affected. However, it does affect the morale and quality of life for the personnel, which in turn affects their willingness to reenlist. Retaining skilled maintenance technicians has been more difficult in recent years, as reenlistment rates have dropped. Longer hours only add to this problem.

Keeping an aircraft on the ground as a Cann-bird also seems to be inefficient when considering the missions the plane could accomplish if it were in flying condition. There is an opportunity cost to having an aircraft act as a parts bin and not fly missions. This assumes that the aircraft would be needed and could be utilized if it were in flying condition. In the case of today's Air Force where airlift capacity is strained, such an assumption of need is generally valid. Air Mobility Command charges its customers for

---

[40] United States General Accounting Office, *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems*, GAO/NSIAD/AIMD-99-77, (GAO, 1999). p. 24.

flying missions and uses that revenue for operations. AMC thus loses significant revenue if it is unable to fly requested missions. The value of a mission to the Air Force or the US say, in times of armed conflict, may be much higher than the lost revenue to AMC.

4.3 CANNIBALIZATION BENEFITS

Although cannibalization places additional stresses and costs on the maintenance system and causes an aircraft to be grounded, there are benefits in terms of mission success and readiness.

Cannibalizations are performed for two reasons: needed spare parts are not in supply, or it would take too long to retrieve the necessary part from supply. Cannibalization allows greater readiness when spare parts are not readily available by allowing parts needs to be consolidated into one aircraft. This provides a squadron or a wing[41] a means to deal with stock-outs of needed parts. Cannibalizations are also sometimes used to provide a quick source of supply when spares are not conveniently available. For example, when a failure occurs just before launch, it is sometimes beneficial to quickly cannibalize a replacement part from a nearby aircraft, in order to achieve better on-time ratings.

Cannibalization may be a useful management tool in a lean inventory system, due to a combination of the characteristics of parts failures and the large number of parts in an aircraft. Part failures for most components are quite unpredictable. Failure rates do not correlate well to either hours flown or sorties flown.[42] Demands for spares over time are often "lumpy", with long periods of low demand interspersed with sudden spikes or increases in demand. Many different aircraft operational factors cause these fluctuations, with changes in demand levels influenced by the environment, training schedules, deployments, new tactics and modes of use, discovery of unexpected failure modes, and many other factors. For example, when an unusual failure is discovered on one aircraft,

---

[41] A wing is an organizational unit that typically consists of several squadrons. Typically there is one aircraft used for cannibalization for the entire wing, although it is not unheard of to have one cannibalization aircraft per squadron in extreme situations.

[42] Gordon B. Crawford, *Variability in the Demands for Aircraft Spare Parts: Its Magnitude and Implications*, R-3318-AF, (Santa Monica, CA: RAND, 1988).

other aircraft are then checked more rigorously for that problem, leading to a spike in demands for spares for the damaged part. High variability in spares demands makes it difficult to predict which spares will be needed and how many will be needed. Determining accurate stock needs is rather difficult even in peacetime. In wartime with greater operational uncertainties thrown in, it is practically impossible. The variability in demands leads to stock-outs. Cannibalization can provide a means for dealing with variation in the demand for spare parts by consolidating the stock-outs into one aircraft.

Furthermore, most parts typically fail rather infrequently. This adds to support for cannibalization because there are many thousands of parts that fail very infrequently. The Air Force opts to not stock the majority of these very low demand parts, because maintaining spares stock levels for so many low demand parts would be costly. However, due to the sheer number of these parts, the probability that at least one of these parts will have failed somewhere in the squadron resulting in a grounded aircraft is reasonably significant. With one plane thus grounded due to a lack of spares, it then becomes a reasonably inexpensive source of inventory through cannibalization for the remaining aircraft.

As an example of the numbers of parts involved in maintaining one aircraft, consider the C-5 Galaxy, a vital part of the U.S. airlift capability. There are currently about 81 C-5s in the active Air Force fleet. Another 13 are in the Air National Guard, and 32 are in the Reserves. A C-5 transport has 90,000 parts; 38,000 of them are numbered SKU's (stock-keeping units). Of these, 33,000 are consumable and 5,000 are reparable. Of the 5,000 reparable, only 3,000 are typically stocked. Of the 3,000 stocked, only 900 have demands fleet-wide of over one per year. Only about 26 parts have more than 100 failures fleet-wide per year. About half of all MICAPs (a status given to aircraft that are not mission capable due to parts supply problems) result from the stock out of parts that have not had a failure in the entire Air Force system in the past 2 years.[43]

For a simple illustration of the probability of stock out resulting from many low-failure parts, we will make some simple assumptions and calculations. Consider that the fleet of C-5s flies on average a total of approximately 40 aircraft missions a day. Assume

---

[43] USAF Air Mobility Command, *HQ AMC Orientation Briefing*, (AMC/LG, HQ, 1998).

that each of the 5,000 reparable part types fails on average only once in the entire fleet every 20 years. This is a very low assumption for failure rates, but helps illustrate the point. In 20 years, with 40 missions per day, there will be a total of 292,000 missions. For a given type of part, if on average only one part fails during 292,000 missions (20 years worth), the probability of that part failing on a given flight, we will call this P_{part_fail}, is roughly 1/292000. The probability that the part does not fail, $P_{part\_not\_fail}$, is 291999/292000. The probability that a given part fails is very low. However, there are many parts. Assume the number of parts n, where n equals 5000. The probability that at least one of the many parts fails on a given flight is:

$$P_{flight\_fail} = 1 - \left(P_{part\_not\_fail}\right)^{n} \cong 0.01698$$

Thus the probability that no parts fail on a given flight is:

$$P_{flight\_no\_fail} = 1 - P_{flight\_fail} \cong 0.9830$$

There are multiple flights per day. We will call the number of flights m, where in this case m equals 40. The probability that a part fails on a given day is thus:

$$P_{day\_fail} = 1 - (P_{flight\_no\_fail})^{m} \cong 0.4959$$

Thus the probability of any part failing during the 40 daily flights on a given day is about 50%. If we assume that only 3,000 of the 5,000 parts is in stock (which is typical Air Force practice for the C-5), the probability of a random single failure resulting in a stock-out, we will call this $P_{stockout}$, is 2/5. Thus the probability of a failure resulting in a stock-out on a given day in this simple illustration is:

$$P_{fail\_stockout} = P_{day\_fail} P_{stockout} \cong 0.20$$

Thus, even with very low probability of failure for each of the parts, the probability of a failure with a stock-out somewhere in the fleet on a given day is still high, in this example approximately 20%. Not all part failures will result in a plane being grounded, as many parts are not critical for mission capability. Nonetheless, it is clear that even with a very low probability of failure for each part, with the high number of parts there is a significant chance of a part resulting in a stock-out on any given day. Since it may take many days, perhaps many months, to acquire a part that has not been stocked, it is likely that on any given day, at least one plane will be down due to a failed part that had no spare in supply. This then lends support to having a Cann-bird, as it is likely that one plane would be down anyway.

This then raises the question: why not stock all of the parts? In the previous example 40% of the low failure parts were not stocked. There are costs to stocking all of the parts that are rarely used. However, assume that all parts were stocked. Even parts that are normally stocked have an occasional stock-out. Even if each part has only a very small probability of a stock out, the cumulative probability of a mission critical part stocking out is reasonably high. Typical rates for stockage effectiveness for the C-5 are about 88%[44], so the probability of stock-out of a needed part is about 12%. Using the same probabilities of failure as before, the probability of a failure resulting in a stock-out on a given day would be about 6%. With long lead times to receive those parts that stock out, the chance that a plane would be down on a given day due to a stock-out is still rather high. Eliminating cannibalization in such an environment may not be a realistic possibility. Cannibalization in situations such as this may be an efficient means of maximizing aircraft readiness and minimizing cost.

Cannibalization in some instances occurs even when there are spares in supply. Many cannibalization actions are what are referred to unofficially as "convenience canns". Squadrons are pressed to meet strict time deadlines in preparing for flights. When an aircraft is warming up for flight and a failed component is discovered during the preflight checks, there may not be time to retrieve a replacement from spares inventory, which may be located away from the flight line. In such situations, cannibalization from nearby aircraft, not necessarily even the designated Cann-bird, provides the squadron with the ability to launch the aircraft on time.

Cannibalization is often seen as an indication of a failure in the inventory management system. It is only a partial indicator. Failures in the inventory system will lead to greater cannibalization, but cannibalization itself does not necessarily mean the logistics system failed. In some ways cannibalization may be the most effective means to

---

[44] From data in *Logistics Executive Summary*, (436th AW/512th AW, Dover AFB, 1998). Stockage effectiveness is a measure of the percentage of spares demands that supply fills without backordering.. Stockage effectiveness excludes from consideration those parts that are a first time request, so it essentially only considers those parts that normally would be stocked.

increase readiness and reduce cost. The question that should be asked is not if the logistics system failed to do what it was designed to do, but whether the system was designed to take into account the tradeoffs due to cannibalization and produce a cost effective solution. In Chapter 6 we will explore this issue further.

CHAPTER 5: MODEL DESIGN

It is a very fundamental principle indeed that knowledge is always gained by the orderly loss of information, that is, by condensing and abstracting and indexing the great buzzing confusion of information that comes from the world around us into a form which we appreciate and understand.

K. E. Boulding from "Economics as a Science"[45]

5.1 SCOPE AND PURPOSE

To explore some of the fundamental tradeoffs in the sustainment system, a model was created using Matlab[46]. The model can be used to illustrate the tradeoffs between aircraft squadron size, spares inventory policy, aircraft component reliability, repair system responsiveness, and cannibalization policy. System models can be valuable tools for characterizing and understanding the Air Force sustainment system. Models that represent the Air Force system can generate insight into system behavior and provide a valuable tool for learning, for both those creating the models and those learning from the results the models produce.

The Air Force sustainment system is incredibly complex. High-level simulation models, although only simplified mathematical abstractions of the real world system, can illustrate critical tradeoffs in the complex system. The model used in this thesis can provide illustrations of general tradeoffs and characteristics of the system, and generate insights that may be masked in a more complex model. While there are many simplifying assumptions and approximations in the model, the goal of the model is to generate some insights into system behavior in order to educate and train decision makers about system characteristics. The quantitative results from the model should be used to simply guide and educate decision-makers about system tradeoffs and may not represent absolute estimates of system performance. Rather relative performance of various options can be judged from model outputs.

---

[45] Kenneth E. Boulding, *Economics as a Science* (New York: McGraw-Hill, 1970).

[46] Matlab, by The MathWorks, Inc., is a computer language for technical computing.

5.2 MATLAB MODEL DESCRIPTION

This section will provide a brief description of the model used in this thesis, as well as an overview of the assumptions and limitations in the Matlab model. A users guide for the model is contained in Appendix B.

The model represents one squadron of a fixed number of aircraft operating out of one airbase. In the model, each day the number of aircraft in the squadron at the base is the same. This can represent a fighter squadron operating out of one base. It can also be used to represent a transport squadron, where aircraft leave the base for extended missions, as long as the number of aircraft leaving each day equals the number that are returning. It is assumed that there is no attrition and no purchase of additional aircraft.

The squadron is assigned a fixed number of missions each day. It is assumed that each aircraft in the squadron may only make one flight per day, so if there are not enough mission capable aircraft to meet the day's requirements some missions are not flown. If there are more aircraft that are mission capable than there are assigned missions on a given day, only the assigned missions are flown and the extra aircraft do not fly.

This approximation that aircraft fly a maximum of one flight a day is generally reasonable for transport aircraft that fly long missions. For fighter aircraft, it is possible for an aircraft to fly multiple sorties in one day. It is reasonable to assume a maximum of one flight per fighter per day during peacetime operations, but in times of need it is possible for fighters to fly as many as 3 sorties or more per day. For this analysis multiple flights per day were not considered, and the maximum number of flights per aircraft per day is set at one.

A simplifying assumption is that the number of requested missions per day remains constant. In the real world this is not the case, as there is variability in the demand for missions. This is especially true in wartime, where the number of missions would surge to greater levels. However, in just considering peacetime, day to day schedules are somewhat regular. There is greater justification for this assumption when considering transport aircraft, in that transport missions are also generally more consistent. This is an acceptable assumption that provides a baseline for comparison.

An "aircraft" is in essence a collection of components, each of which must be in working order for the aircraft to be mission capable. In reality some parts do not need to be in working order for an aircraft to be mission capable, but in the model each part is assumed to be mission critical, or necessary for mission capability.

For those aircraft that fly, there are probabilities of generating failures in their components. Each component that flies has a probabilistic chance of failure, which is modeled as a fixed value for each type of component. The many thousands of components in reality are modeled by a much smaller number of "bundled" components. Each component in the model represents a collection of many components in reality. It is assumed that each component in the model has the same probability of failure. The "bundles" of components can be chosen such that this can be approximately true.

Components that fail are removed and immediately submitted for depot repair. Depot repair times are modeled as deterministic and constant for each component type. Reality is quite different, with each individual component taking a different amount of time, and perhaps requiring different repair parts and repair processes. Repair times vary quite a bit depending on many factors, including the number of failures vs. repair capacity, the difficulty of the repair, the availability of necessary repair parts and repair funds, and the priority of the part.

In the discussion of the model, we describe the repair time as "depot" repair time, but the repair time in the model is actually the total time between when a component fails and when the component returns to supply. This represents the time it takes to ship a broken component to its source of repair, the time it waits for repair, the time it takes in repair, and the time it takes to ship the repaired component back to the base supply. It could even include the additional steps involved in three-level maintenance where there is an intermediate repair facility. For simplicity, this total repair time is simply referred to as depot repair time. Depot repair times are modeled as fixed and are specified for each type of component. It is assumed that all components can be repaired, and that there is no condemnation of components.

When components are eventually returned from repair, they enter the base supply spares inventory. Those components that fail and are removed from the aircraft are replaced the same day by a spare from base inventory, if there are spares available. If a spare is not available and cannibalization is allowed, the part will be cannibalized. Thus aircraft only become not mission capable if there is a stock out for a broken component and no parts are available to cannibalize.

The model calculates mission capable ratings based on whether an aircraft is able to fly on a given day or not due to the availability of spare parts. This is a simplification of the real world, where mission capability is determined not once a day per aircraft, but continuously. In the real world, the time that it takes to repair an aircraft after a flight contributes heavily to Not Mission Capable (NMC) ratings. Not Mission Capable due to Maintenance (NMCM) is the real world Air Force measure of the rate that aircraft are down due to maintenance. NMCM makes up a significant portion of the overall NMC rating. For the C-5, aircraft are typically down for maintenance about 28 percent of the time and down because of supply about 14 percent of the time. Some of this time overlaps, as aircraft are down for both causes approximately 7 percent of the time, so that the total NMC rate is about 35 percent. A separate model that calculates mission capability on a nearly continuous basis was created and the results were compared to the model used in this thesis.[47] The output of the model used in this thesis compared favorably to results from the model that computed mission capability in a much more realistic manner.

One assumption concerning maintenance labor availability is that flight line aircraft repairs can be finished in time for the next day's missions, as long as spares are in stock. This assumes that repairs on the flight line take less than one day and that maintenance personnel have adequate time and resources to complete the repairs. Air Force data shows that this is not always the case. For example, a C-5 that lands and is determined to be NMC takes on average about 12 hours to fix.[48] For the C-5, roughly 35 percent of aircraft are fixed within 4 hours, 55 percent within 8 hours, 65 percent within 12 hours, 90 percent within 24 hours, and 95 percent in 48 hours.[49] Currently the maintenance system is strained and flight line personnel often must work long hours to complete their

---

[47] This "other" model was created using a commercial simulation package called iThink, and calculated MC rates on a nearly continuous basis and considered both NMCM time and NMCS time. The iThink model was in many ways more realistic than the Matlab model used in this thesis, but had the major drawback of a very long run time.

[48] *Logistics Executive Summary.*, (436th AW/512th AW, Dover AFB, 1998). p. 21.

[49] *Logistics Executive Summary.*, (436th AW/512th AW, Dover AFB, 1998). pp. 21-23.

workload. It is quite probable that repairs are not always able to be completed by the next day, and that some repairs may be delayed when workload overwhelms maintenance personnel.

Three different options for cannibalization are considered in the model: unrestricted cannibalization, cannibalization limited to one Cann-bird, and no cannibalization. Under the unrestricted cannibalization scenario, aircraft that have the most broken parts are used to repair the aircraft with the fewest broken parts. Cannibalization only occurs when there are not enough mission capable aircraft on a given day to fly the requested missions. If there are not enough aircraft that are mission capable, the unrestricted cannibalization model sorts the aircraft according to the number of repairs it needs. The least damaged aircraft are first to be repaired and the most damaged aircraft are the first to be cannibalized. Cannibalizations are only performed if the plane can be made ready (i.e. if a plane requires two parts to be repaired, but only one is available for cannibalization, then the available part is not cannibalized.)

The "one Cann-bird" scenario differs from unrestricted cannibalization in that cannibalization is limited to using one designated aircraft as a source of spare parts. Also, in the one Cann-bird model, the Cann-bird does not fly, even when there are enough spares to make it mission capable. The aircraft is held in case it is needed. Of course, in the no cannibalization scenario, there is no Cann-bird.

5.3 INPUTS AND OUTPUTS

Model inputs are:

- Number of aircraft
- Number of components that compose an aircraft
- Number of initial spares for components
- Number of requested flights per day
- Depot repair time
- Probability of failure of a component
- Cannibalization policy: no cannibalization, use of one Cann-bird, or unrestricted cannibalization.
- Length of simulation in days
- Number of trials (the number of times a model runs for each set of inputs)

The model runs multiple trials of the simulation model, and outputs results that are the average over the trials. A standard deviation is also calculated for the average over the trials. Outputs include:

- Percent of requested missions flown: Computed for a simulation as the sum of missions completed during the simulation, divided by the total number requested. Results from the simulation trials are then averaged.
- Percent of aircraft that are mission capable: Computed for a simulation as the mean of daily mission capable rates. Results from the simulation trials are then averaged.
- Number of cannibalizations: Computed for a simulation as the sum of cannibalizations during the simulation. Results from the simulation trials are then averaged.
- Number of cannibalizations due to Cann-recovery: Computed for a simulation as the sum of cannibalizations required for Cann-recovery during the simulation. Results from the simulation trials are then averaged.
- Number of part failures: Computed for a simulation as the sum of failures during the simulation. Results from the simulation trials are then averaged.

The outputs from the Matlab model can then be used to calculate the number of mission capable aircraft, and costs. These calculations are performed separately from the Matlab model, using a spreadsheet.

5.4 SUMMARY

The model presented in this thesis is a high-level model that has limitations when it comes to accurately representing the real world of Air Force maintenance. Anyone with experience in the system would clearly recognize this after reading the list of assumptions. In the real world, the number of aircraft may not remain constant, the daily flying regimen is not constant, and aircraft repairs may not be finished by the next day. Parts failure rates are very unpredictable and highly variable and not easy to model or forecast. Repair times are also rather variable.

With all of these caveats it may seem that the model is of little use in representing the real world of Air Force maintenance. If one wanted to accurately predict daily operations or use the results from the model in a quantitative way, then the model would not be very helpful. However, the model can provide rather useful illustrations of general tradeoffs and characteristics of the system, and generate insights that may not be easily seen in a more complex model.

In Chapter 6, the model will be implemented to illustrate tradeoffs in the Air Force sustainment system. Appendix B contains the model code and users instructions.

CHAPTER 6: SYSTEM TRADEOFFS

In this chapter we will begin by characterizing the basic tradeoffs in readiness and in cost due to the number of aircraft, the number of spares, the amount of cannibalization, the depot repair time, and the probability of part failure. We then will examine several scenarios to illustrate how simple models can be used to examine system behavior. To begin the discussion of tradeoffs, section 6.1 will first examine the impacts on readiness. Tradeoffs in cost will be addressed in the section 6.2.

6.1 CHARACTERIZING TRADEOFFS IN READINESS

Upon visiting Air Force bases and witnessing the use of Cann-birds, MIT researchers speculated about the trade-off between cannibalization and simply having more spare parts. A C-5 wing[50] such as the one at Dover AFB was chosen as the type of wing to be modeled, as the MIT researchers through their site visits were more familiar with the C-5 and had collected some basic data on the operational performance of C-5 squadrons.

6.1.1 Modeling Assumptions

The model described in Chapter 5 was used to begin the exploration of tradeoffs.

In these runs, an aircraft consisted of 25 "components", for modeling practicality. These "components" can be thought of as representing a collection of several actual parts. Each component was given a per-flight failure probability of 0.015. This value was chosen so that the number of failures per year in the model approximated the number of failures per year a similar number of aircraft in the real world would have for the 25 parts that failed most on the C-5, as reported by Air Mobility Command.[51] The number of days in the simulation was set at 360 days. Each part was given a repair time of 50 days, which is roughly equal to the Air Force's 1998 goal for the overall average for depot

---

[50] A wing is an organizational unit that typically consists of several squadrons. In the case of the 436th Air Wing at Dover AFB, a wing consists of two subgroups called sortie generation flights, each of which consists of approximately 15 aircraft.

[51] USAF Air Mobility Command, *HQ AMC Orientation*, (Scott AFB, IL: HQ AMC/LG, 1998). The data from AMC was in the form of a single chart, approximating the number of parts that had numbers of failures in a certain range. This was rather imprecise in nature, but sufficient to generate rough approximations. The approximation omits the top failing part, tires.

repair time.[52] The model was initialized with all spares starting in the base supply and none in the depot repair process.[53]

The model only represents operations at one Air Force base. In reality, although C-5's have a home base they operate worldwide. In translating C-5 operations to the model, some approximations were made. In reality for a wing of about 30 aircraft, there are on average approximately 200 departures from home base per month, and 320 departures from worldwide locations.[54] Typically at the start of any given day there are only about 15 aircraft at home base, with the roughly 15 remaining aircraft scattered throughout the world. The one base in the model can be thought of as representing one end, or approximately one-half, of the transport system the aircraft from the base serve. Thus to approximate the C-5 wing's operations at home base, the wing was modeled as 15 aircraft operating at one base. For each day in the model, the number of aircraft that depart equals the number of aircraft that arrive. The daily requested flight schedule was set at 10 flights per day, a reasonable approximation for the daily number of flights for half of a C-5 wing.

The percentage of aircraft that are scheduled for use is called the commitment rate. The Air Force standard established for the maximum commitment rate is 65%, which equates to about 20 aircraft for a full wing, and 10 for half a wing. The standard for the maximum commitment rate is established to allow aircraft time for inspections and programmed depot maintenance. Due to need for airlift missions, the Air Force goal is to be as close as possible to the 65% commitment rate standard without going over. Data from Dover shows an average of about 17 aircraft scheduled for flight per day for a complete wing. Overall, the approximation of 15 aircraft operating from one base with a goal of 10 flights per day seems to be a reasonable abstraction of reality.

---

[52] Maj. Craig Romero, *Quarterly Agile Logistics Pipeline Analysis Update*, http://38.178.7.130/alconference/briefings.htm, August 1998. Fifty days approximates the established budgeted standard for 1998. The actual pipelines have been much longer: about 75 days was typical for the average during 1998.

[53] The validity of this initial condition was tested by running the model for two years and comparing the results from the second year to the results from the first. This allowed the depot repair pipeline to contain a typical amount of spares. The results comparison showed that there was not significant difference due to the initialization conditions.

[54] Approximations based on data from *Logistics Executive Summary.*, (436th AW/512th AW, Dover AFB, 1998).

There is one difference due to this abstraction that must be noted. In the model, mission capability ratings and the percentage of requested missions flown are both calculated based on a total of 15 aircraft (or fewer when the number of aircraft is decreased in the model). When one Cann-bird is used, there is one Cann-bird for 15 aircraft rather than one Cann-bird for 30 aircraft. This causes a slight underestimation in the calculation of MC rate and the percentage of flights flown.

6.1.2 Aircraft and Spare Parts

To begin examining the tradeoffs in readiness due to aircraft, spares, and cannibalization, the model was used to examine how varying numbers of aircraft and spares could affect the percentage of requested flights that could be flown, the mission capable (MC) rate, the number of MC aircraft, and the amount of cannibalization.

Through a series of runs of the model, varying numbers of aircraft and spares were input. While 15 aircraft was chosen as the baseline to represent a current actual C-5 wing, fewer aircraft in a wing were also modeled to determine the tradeoffs in reducing the number of aircraft in a wing. Since the number of requested missions is 10 per day, the maximum commitment rate will be 100% for the 10 aircraft case. Different levels of spare parts inventory were also modeled. In the charts that follow "spares" refers to the total initial number of spare parts for each of the 25 parts that make up the plane in the model. In this set of runs, the number of spares varied from 0 to 16.

We define the expected lead-time demand as the expected number of spare parts that will be needed to fix aircraft during the time period that is required for the depot repair of a broken part. The lead time demand in the model is simply calculated as the probability of failure per flight multiplied by the number of flights per day multiplied by the number of days that it takes for a part to be repaired. The expected depot repair lead-time demand for the parts in this set of runs is 7.5 spares.

In addition to varying the number of aircraft and spares, three different cases of cannibalization were examined: no cannibalization, cannibalization limited to one Cann-bird, and unrestricted cannibalization.

6.1.3 No Cannibalization

To begin our discussion of model results, Figure 6-1 shows the percentage of requested missions that a wing can fly with a given number of aircraft and spare parts. In this case, no cannibalization is performed. As expected the percentage of requested flights that a wing is able to fly increases with both the number of spares and the number of aircraft. The percentage of flights increases steadily as the amount of spares increases, until the amount of spares reaches roughly the expected lead time demand of 7.5 spares. As the amount of spares rises to higher levels, the percentage of flights continues to increase but begins to approach 100%. Increases in the percentage of requested flights flown become more costly in terms of spares as the percentage approaches 100%.

It is interesting to note that in this case of no cannibalization, the number of aircraft does not greatly affect the percentage of requested flights that can be flown. Consider the case of 10 aircraft and 8 spares. If a large investment in aircraft is made to increase the number of aircraft by 50% to 15 aircraft, the percentage of flights flown only increases from about 70% to 79%, which is an increase of only about 13%. The number of spares has greater impact than the number of aircraft. The impact on missions from one more spare of each part is greater than the impact of one more aircraft. For the case of 10 aircraft and 8 spares, an investment in spares to increase the number of spares by 50% to a total of 12 would increase the percentage of flights by about 34%, raising the percentage of flights from about 70% to 93%. This in a sense suggests (if one doesn't consider such issues as the labor involved and the initial cost of an assembled plane) that if the squadron had more than 10 aircraft, it would hypothetically be better in terms of the percentage of requested flights flown to disassemble the "extra" aircraft and put all the parts on the supply shelf.

The fact that additional spares have a greater impact than additional aircraft makes sense in the case of no cannibalization. Aircraft in excess of 10 will only be needed if there is already a stock-out for at least one part. An additional aircraft will allow more flights to be flown, but the additional aircraft will eventually fail as well, add to the total number of failed parts, and tend to deepen stock-outs. An additional set of spares on the other hand can be used to fix several aircraft, and will decrease the probability of stock-out for each of the parts.

Figure 6-2 illustrates mission capability rate as a function of number of spares and number of aircraft. MC rate increases with the amount of spares held for each part, as would be expected. As with percentage of flights flown, MC rate increases rapidly until roughly the amount of spares equals the lead-time demand (7.5 spares), then begins to more slowly approach the limit of 100%. Increases in performance become more costly in terms of spares as MC rate approaches the limit. Again, like percentage of flights flown, the number of spares has a greater impact than the number of aircraft.

There is one major difference between the characteristics of MC rate and the percentage of requested missions flown. Having fewer aircraft in a squadron does not decrease the MC rate in the model as it decreases the percentage of missions flown. Rather, the MC rate increases as the number of aircraft decreases from 15 to 10. To explain this behavior, consider a group of 11 aircraft. With the given assumption that aircraft maintenance crews can replace broken parts in time for the next needed mission as long as spares are in supply, the 11th aircraft isn't needed unless there is a stock-out of a needed part and a plane is not able to fly. When the 11th aircraft flies, it thus prevents a mission from being missed, but it also generates more parts failures and can cause a greater shortage of parts or a deeper stock-out. Over time this leads to a greater percentage of aircraft being unable to fly than if there were only 10 aircraft.

Perhaps a better indicator of mission capability is the number of mission capable aircraft rather than MC rate. Figure 6-3 shows the number of mission capable aircraft as a function of the number of spares and the number of aircraft. Whereas the MC rate decreases when the number of aircraft increases, the number of MC aircraft increases with more aircraft. The number of MC aircraft increases more significantly with greater number of spares.

The number of MC aircraft is important because it is an indicator in the model for the surge capacity of the wing. Surge refers to a period of heightened activity such as war. While only 10 flights are requested per day in peacetime, a higher number of flights may be required during times of war or other need. For example, from Figure 6-1 it can be shown that a squadron of 15 aircraft with 12 sets of spares can fly about 99% of its requested missions during peacetime, or on average 9.9 flights per day. However, Figure 6-3 shows that if needed, the squadron of 15 aircraft would be ready to fly on average 13.5 flights per day at the beginning of a surge. A smaller squadron, while being able to generate nearly as many flights during peacetime, can not generate as many flights during times of surge. For example, a squadron of 11 aircraft with 12 sets of spares would only be ready to fly about 10 flights at the beginning of a surge.

This illustrates one of the main benefits of having more aircraft. The Air Force must operate under dichotomous expectations; while the Air Force is expected to be efficient and lean during peacetime so as to not appear to be wasting taxpayers' money, it is also expected to be able to surge and provide necessary capacity in times of conflict.

One of the general principles of *lean* is to have the right resources at the right place at the right time in the right quantity. This poses a challenge to the Air Force, because the different requirements of peacetime and wartime have different "right" quantities of aircraft and other resources. Military aircraft, spare parts, supplies, and skilled labor are not resources that can be increased and decreased quickly to adjust to the competing requirements of peacetime and wartime. The demands of wartime (for missions, for spare parts, for labor, etc.) are much higher than the demands for peacetime. And predicting when the demands are going to change is very difficult. Lean systems work very well in environments where demand is stable. Where there is large variation in demand, lean systems do not perform as well.

Running a lean "inventory" of resources such as aircraft, spare parts, or labor in peacetime would make meeting wartime demands very challenging. Thus in order to have surge capacity, the Air Force may have excess capacity during peacetime which is not lean. For example, it may be acceptable and efficient to have 10 aircraft and 12 spares during peacetime. As Figure 6-1 shows, almost all of the requested missions can be met. However, with only 10 aircraft, if demands increase over the peacetime level, there is no additional capacity available. For the sake of having additional capacity for wartime surge, the Air Force may be required to retain excess capacity, such as 15 aircraft in the case of the model, that has little benefit during peacetime.

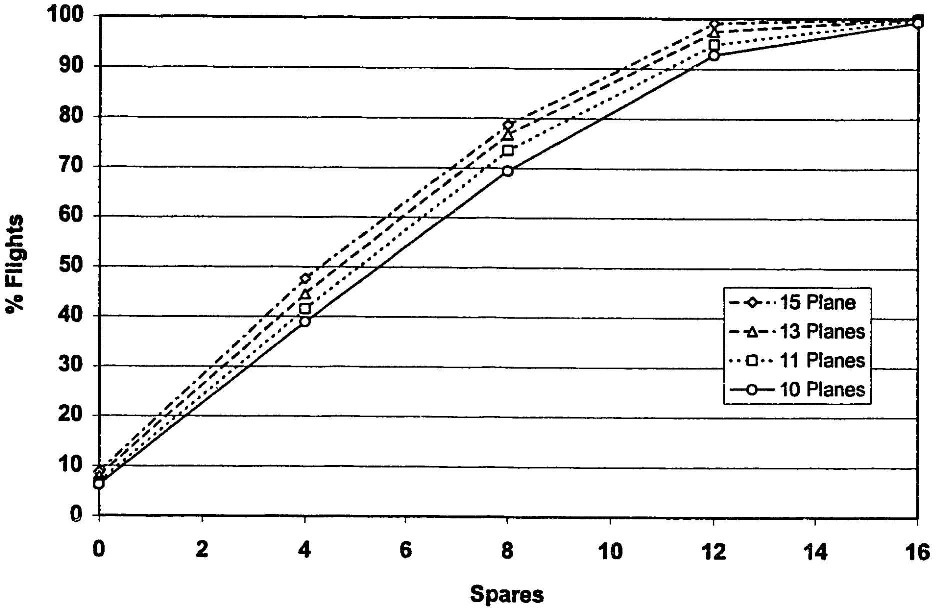

Figure 6-1 Percent of requested flights flown, no cannibalization

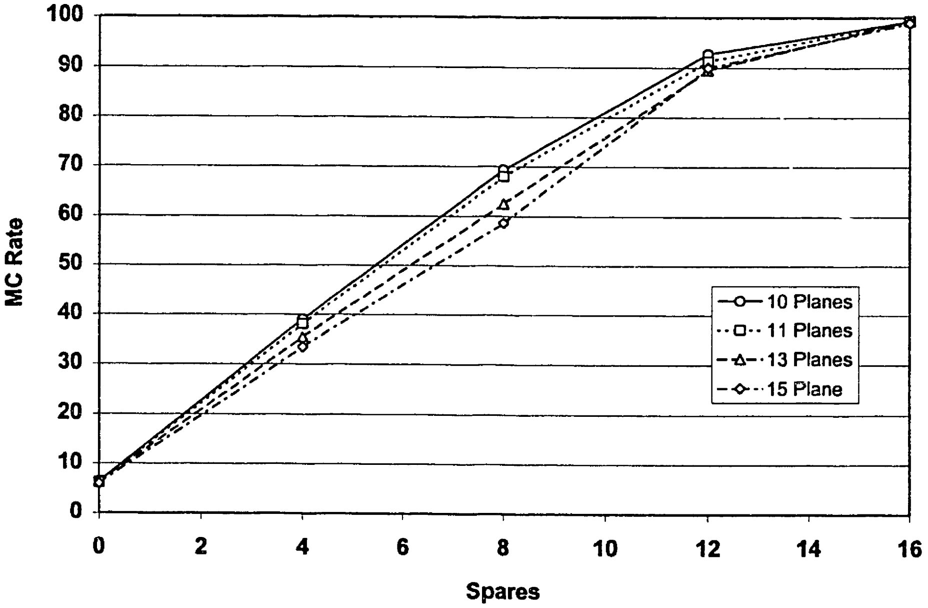

Figure 6-2 Mission Capable rate, no cannibalization.

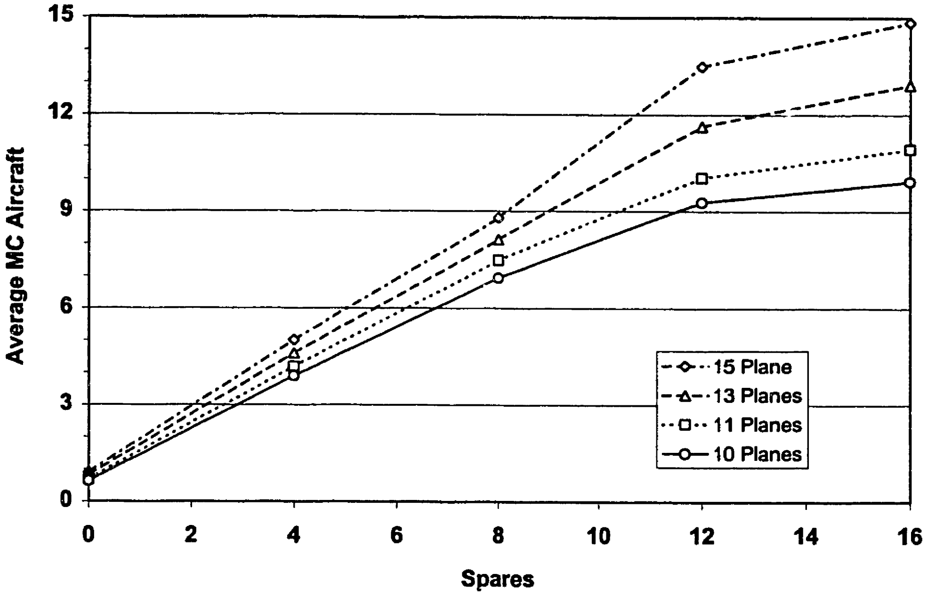

Figure 6-3 Number of Mission Capable aircraft, no cannibalization.

6.1.4 Using One Cann-Bird

Next, let us consider the case of limited cannibalization. The Air Force utilizes cannibalization in normal peacetime operations and would depend significantly on cannibalization in times of conflict. A real C-5 wing typically keeps one aircraft designated for cannibalization, and cannibalizes other aircraft only in extreme cases. To approximate this case in the model, cannibalization was limited to one aircraft. In the model it is assumed the Cann-bird does not fly, even if there are sufficient spares to allow it to be ready. It is maintained as a Cann-bird so that it is available if needed.

Figure 6-4, Figure 6-5, and Figure 6-6 illustrate the percentage of missions flown, the MC rate, and the number of MC aircraft for different numbers of aircraft and numbers of spare parts, for the case of cannibalization limited to one Cann-bird. The tradeoffs are similar to the case of no cannibalization, with a few exceptions. For the percentage of flights flown, shown in Figure 6-4, the percentage is somewhat higher with the use of a Cann-bird, except in the case of 10 aircraft with greater than 8 spares. The exception is due to the Cann-bird not flying even in the case where there are enough spares to make it ready to fly. The mission capable rate and the number of MC aircraft are slightly higher in the case of cannibalization than in the case of no cannibalization when the number of spares is 8 or less. This is due to the ability to consolidate parts shortages to the Cann-bird. However, for a number of spares greater than 8, the MC rate and number of MC aircraft are lower than in the case of no cannibalization. This is because the Cann-bird is assumed to remain grounded even when it would otherwise be ready to fly. For example, in the case of one Cann-bird, a squadron of 13 aircraft averages at most 12 MC aircraft even if there are plenty of spares, because the Cann-bird is assumed by the model to be classified as not mission capable. The curves for percentage of flights made, MC rate, and MC aircraft are asymptotic to different maximums, depending on the number of aircraft. This explains why in Figure 6-5 the curves for different numbers of aircraft cross as they near their maximum.

Cannibalization requires greater labor due to the added maintenance of performing cannibalizations. The number of cannibalizations for the case of one Cann-bird is shown in Figure 6-7. Due to the abstractions of the model, the magnitude of the number of cannibalizations will not directly relate to reality. There are many more parts in reality, so there would also be many more cannibalizations. However, the relative changes in the model will be indicative of the relative changes in reality. At high spares levels, little cannibalization is needed. As the spares level decreases, the amount of cannibalization increases.

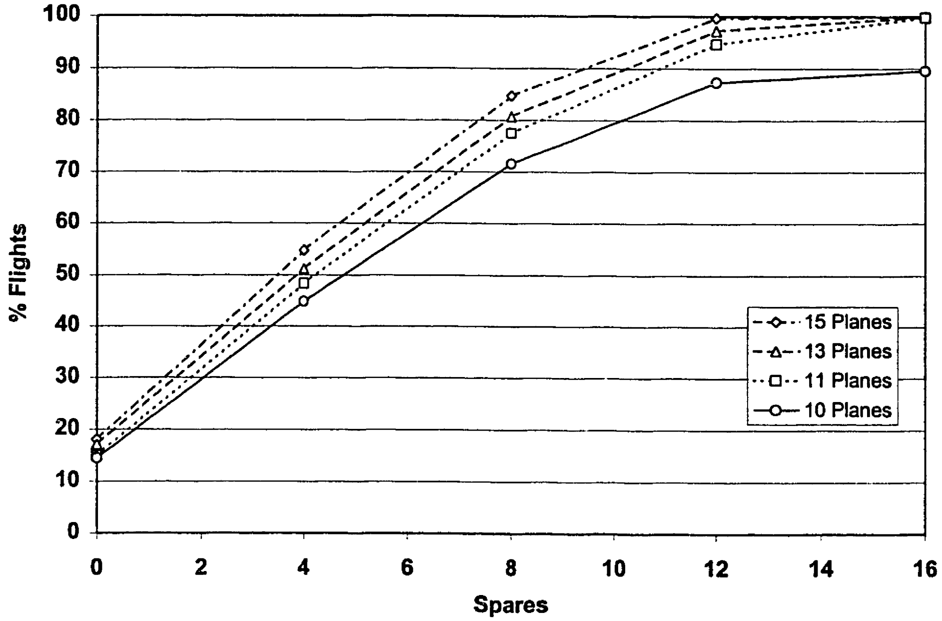

Figure 6-4 Percent of requested flights flown, using one Cann-bird.

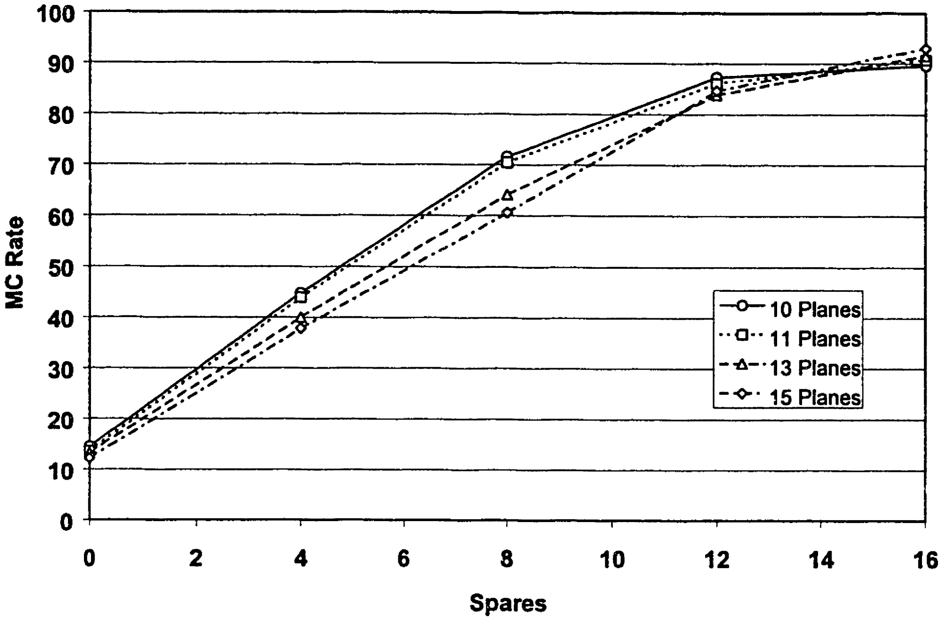

Figure 6-5 Mission Capable rate, using one Cann-bird.

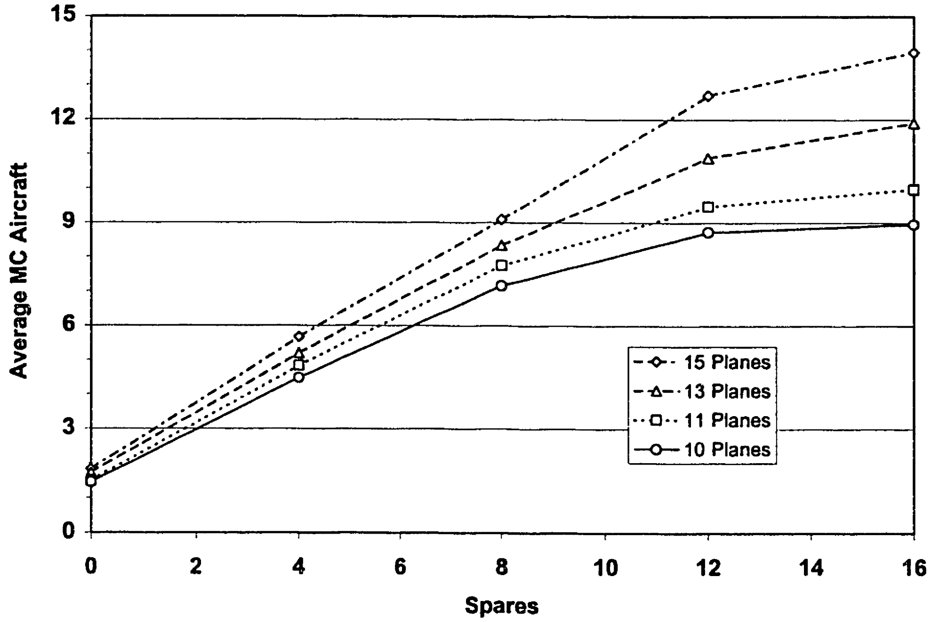

Figure 6-6 Number of Mission Capable aircraft, using one Cann-bird.

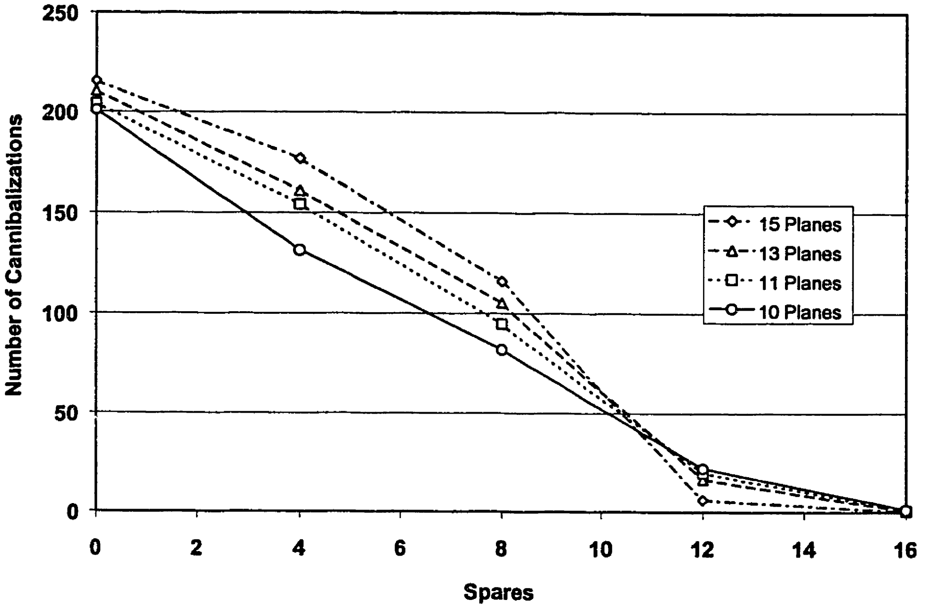

Figure 6-7 Number of cannibalization actions for the case of one Cann-bird.

6.1.5 Unrestricted Cannibalization

Now that the cases of no cannibalization and limited cannibalization have been introduced, the next step to take is to consider unrestricted cannibalization. The term "unrestricted" for this model refers to the number of aircraft that can serve as a parts source for other aircraft. In the modeling of unrestricted cannibalization, the aircraft with the most failed parts are used to fix the aircraft with the fewest failed parts. In the model there is no dedicated Cann-bird, as the aircraft used for spare parts can change from day to day. Thus no Cann-bird must remain grounded if all of the aircraft are mission capable. Essentially this case maximizes the percent of requested missions flown, without concern for the amount of cannibalization required.

This case reflects a level of cannibalization that in the real world is only approached during times of extreme need, such as war or an extreme parts shortage. Figure 6-8 shows the percentage of requested flights flown for the case of unrestricted cannibalization, Figure 6-9 shows the MC rate, Figure 6-10 shows the number of MC aircraft, and Figure 6-11 shows the number of cannibalizations. Each of the four charts shows dramatic increases in their respective measures over the cases of no cannibalization and cannibalization limited to one Cann-bird. Large improvements in readiness must be traded off against the large increase in the cannibalization required and its inherent labor requirements.

Figure 6-8 and Figure 6-10 show a more significant impact due to the number of aircraft than in the cases of no cannibalization or limited cannibalization, especially for lower spares levels. Because unrestricted cannibalization is allowed when needed, adding one more aircraft has roughly the same impact on the percentage of flights as adding one more of each part to the spares inventory. For example, for a starting point of 10 aircraft and 8 spares, the percentage of flights made is about 76%. Adding either one aircraft or one set of spares raises the percentage of flights made to about 80%. Adding three aircraft raises the percentage of flights to about 89%, roughly the same increase that adding three sets of spares would generate. For a smaller initial number of spares, adding an aircraft has about the same impact on the number of MC aircraft as adding a set of spares. For higher initial numbers of spares, the impact of adding aircraft is greater than the impact of adding spares. These results again show the benefit of having more aircraft as well as more spares. The benefit from more aircraft is greatest during times of surge when mission requirements will be high, spares may not be readily available, and widespread cannibalization will be utilized.

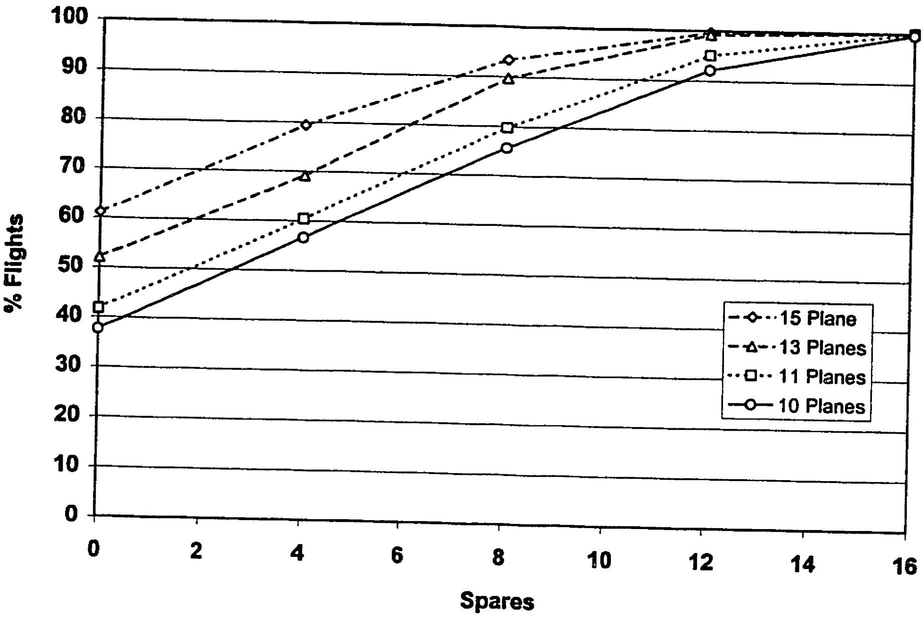

Figure 6-8 Percent of requested flights flown, using unrestricted cannibalization.

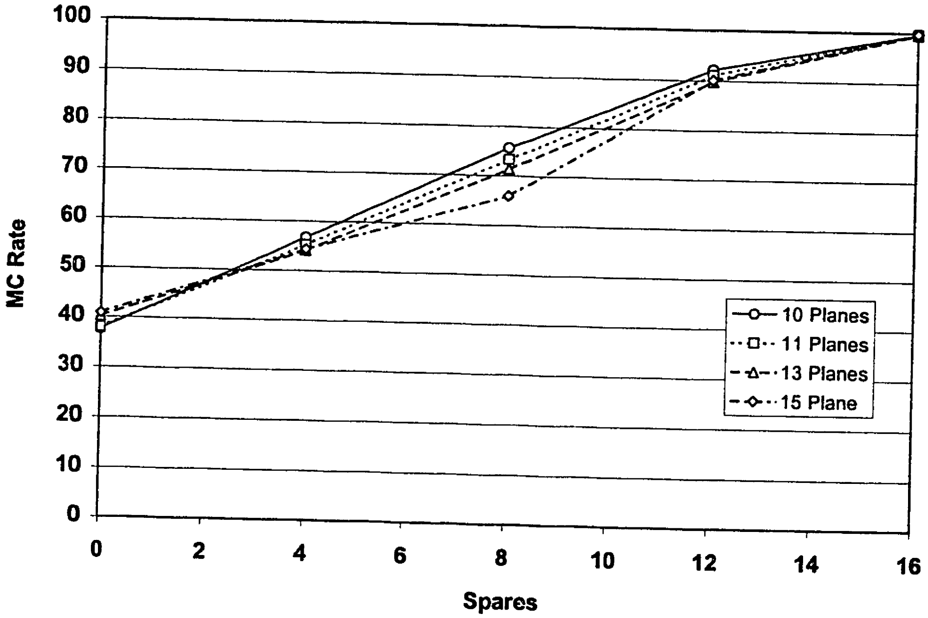

Figure 6-9 Mission Capable rate, using unrestricted cannibalization.

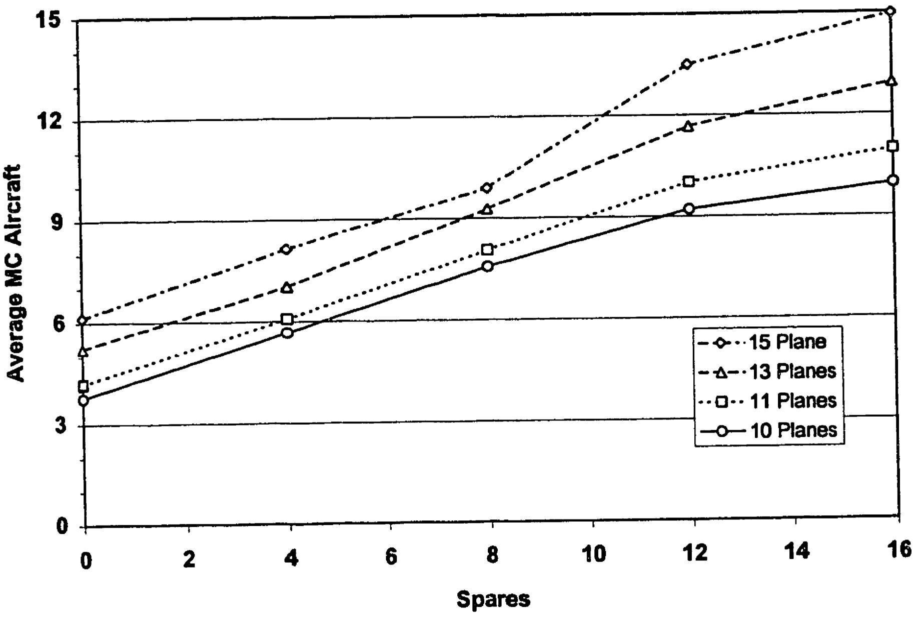

Figure 6-10 Number of Mission Capable aircraft, using unrestricted cannibalization.

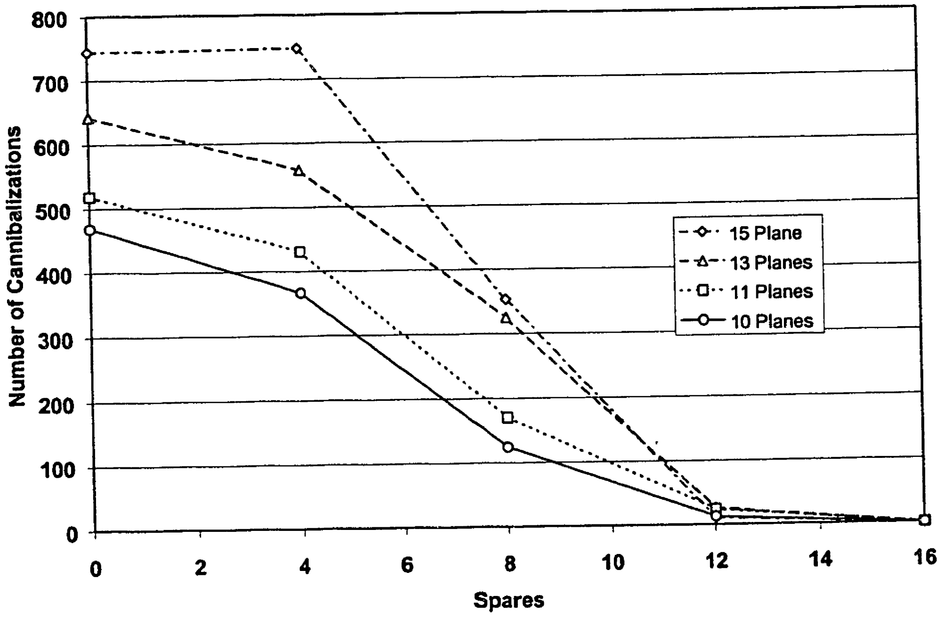

Figure 6-11 Number of cannibalization actions for the case of unrestricted cannibalization.

6.1.6 Cannibalization Policy Comparison

For a clearer comparison between the cases of no cannibalization, cannibalization with only one Cann-bird, and unrestricted cannibalization, see Figure 6-12 through Figure 6-17. These figures show the results from 15 aircraft and 10 aircraft, respectively, for a varying number of spares, comparing the three different cannibalization policies. For a different number of aircraft, the curves would translate up or down, and the magnitudes of the change due to cannibalization would be different, but the order and shape of the curves would be roughly the same. Overall, cannibalization has greater impact the lower the number of spares and the higher the number of aircraft, but also requires more labor and thus has higher labor costs. The figures highlight the large gains in the percentage of missions flown and the MC rate due to cannibalization, as well as the increased amount of cannibalization required to make those gains.

In some cases, retaining a Cann-bird does not show much benefit. One is shown in Figure 6-13. For 10 aircraft with a large number of spares, retaining a Cann-bird actually results in a lower percentage of flights flown. This is due to the fact that 10 missions are requested a day, and because the Cann-bird doesn't fly, there are only at most 9 aircraft available. Similarly for MC rate and the number of MC aircraft, there is no benefit to retaining a Cann-bird when there are ample spares. Figure 6-14 and Figure 6-15 show that for spares levels greater than about 8 the mission capability lost due to keeping a Cann-bird grounded outweighs the mission capability gained by cannibalizing it..

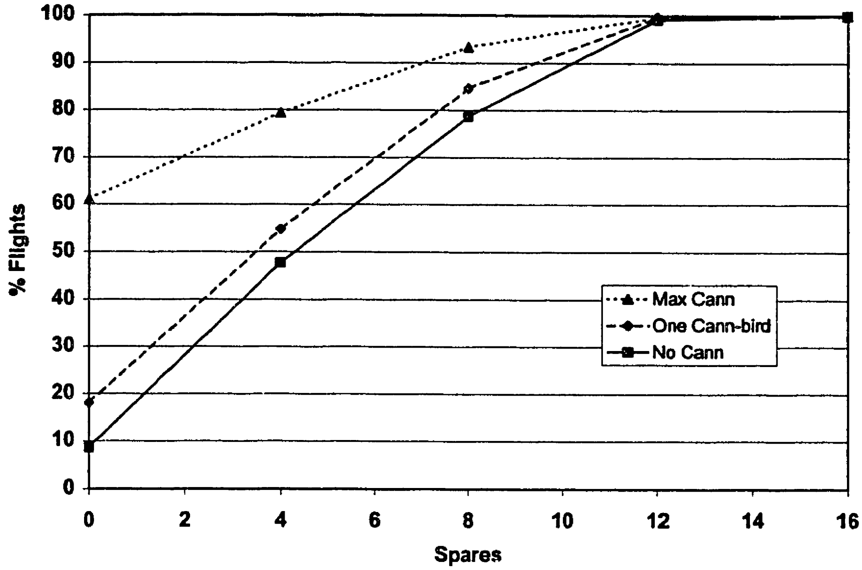

Figure 6-12 Comparison of Cannibalization Policies: effect on percentage of flights flown, for 15 aircraft.

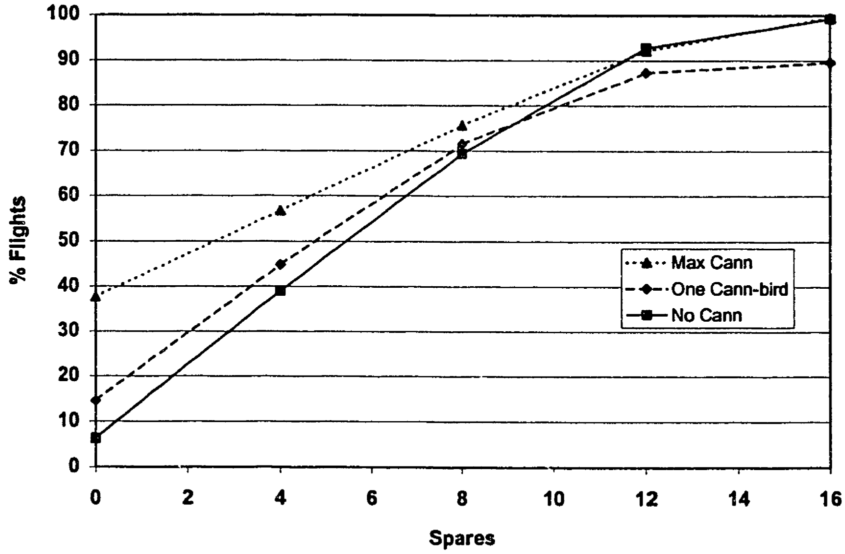

Figure 6-13 Comparison of Cannibalization Policies: effect on percentage of flights flown, for 10 aircraft.

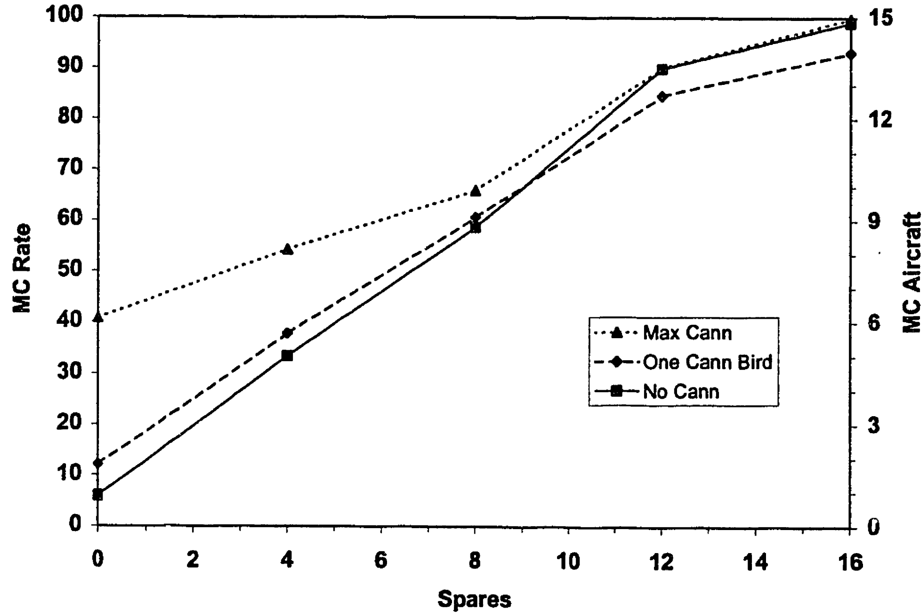

Figure 6-14 Comparison of Cannibalization Policies: effect on MC rate and MC aircraft, for 15 aircraft.

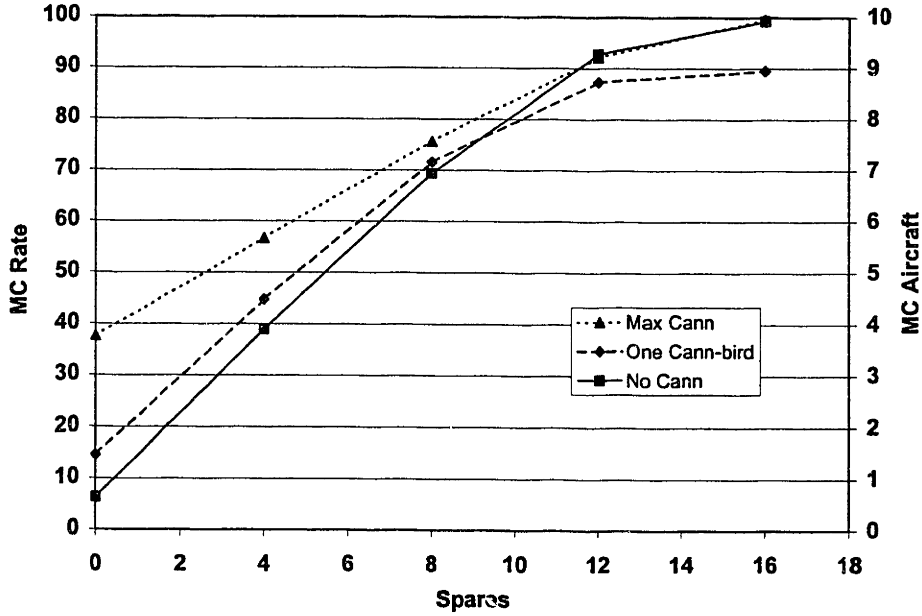

Figure 6-15 Comparison of Cannibalization Policies: effect on MC rate and MC aircraft, for 10 aircraft.

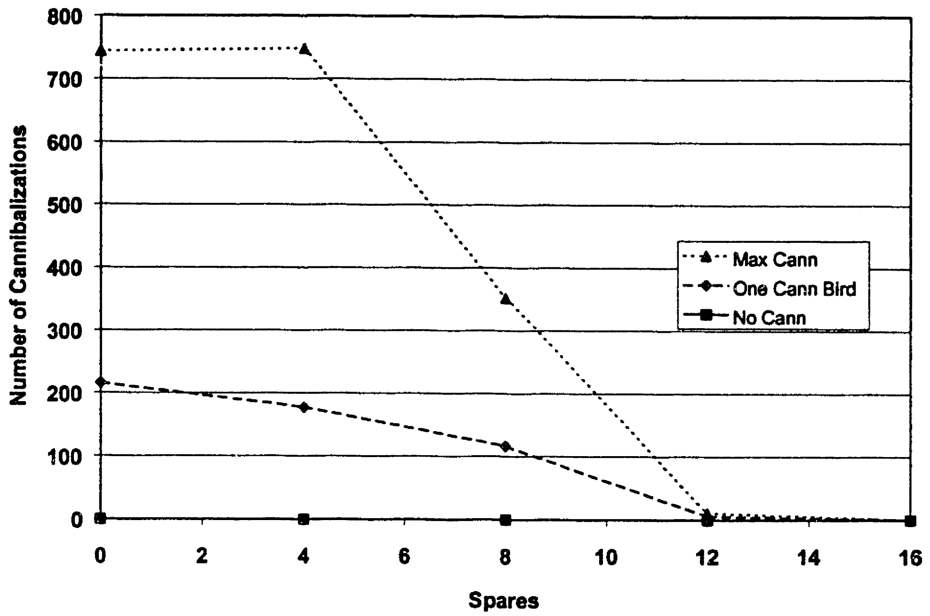

Figure 6-16 Comparison of Cannibalization Policies: effect on number of cannibalization actions, for 15 aircraft.

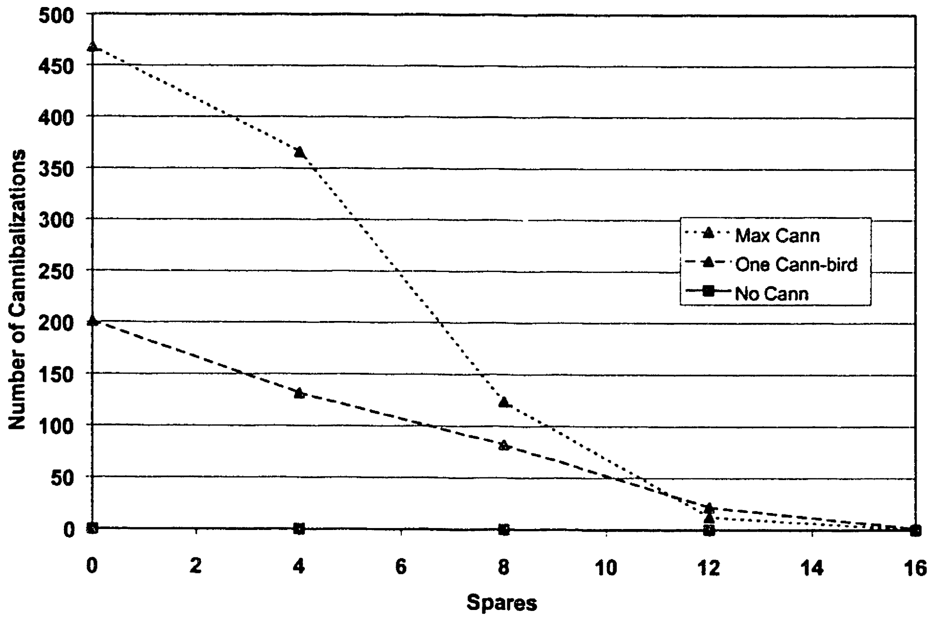

Figure 6-17 Comparison of Cannibalization Policies: effect on number of cannibalization actions, for 10 aircraft.

Another way to consider the comparison between cannibalization policies is to examine the various ways that a certain level of performance can be reached. For example, say that a target of 90% is set for the percentage of requested missions flown. There are many ways to achieve this target, with different numbers of aircraft or spares or different levels of cannibalization. The appropriate mixture to reach the goal may depend on many factors, such as cost, labor constraints, the number of aircraft and spares already on hand, etc. Note that these mixture s of aircraft, spares, and cannibalization to reach the 90% goal are not equal cost options. Figure 6-18 plots lines showing the mixes of aircraft, spares, and cannibalization policy that can achieve a 90% level of requested flights flown. For example, to achieve 90% of requested flights while allowing one Cann-bird, 15 aircraft would require 9 spares, while 12 aircraft would require 10 spares. The case of 10 aircraft with one Cann-bird appears anomalous because only 9 aircraft are available to fly (in the case of one Cann-bird, the Cann-bird never flies). In order to reach the 90% goal the 9 aircraft must always be ready, requiring a higher level of spares than the other cannibalization cases where 10 aircraft are potentially available for flight.

Figure 6-19 shows the mix of aircraft, spares, and level of cannibalization to achieve a different target: an average of 10 mission capable aircraft. As shown in the chart, the case of one Cann-bird typically requires a larger mix of aircraft and spares to reach the target of 10 MC aircraft, because the Cann-bird is not MC by definition in this model. The case of unrestricted cannibalization requires the fewest spares to achieve the target of 10 MC aircraft, but it is interesting to note that with few aircraft there is very little savings in spares resulting from cannibalization.

Overall, the results comparing the cannibalization policies verify that cannibalization has the most benefit (in terms of percentage of flights flown, MC rate, and the number of MC aircraft) in cases where the number of aircraft is high and the number of spares is low. Cannibalization has the least benefit in cases where the number of aircraft is low and the number of spares is high. Retaining one aircraft as a dedicated Cann-bird is unnecessary and detrimental in cases where the number of spares is high..

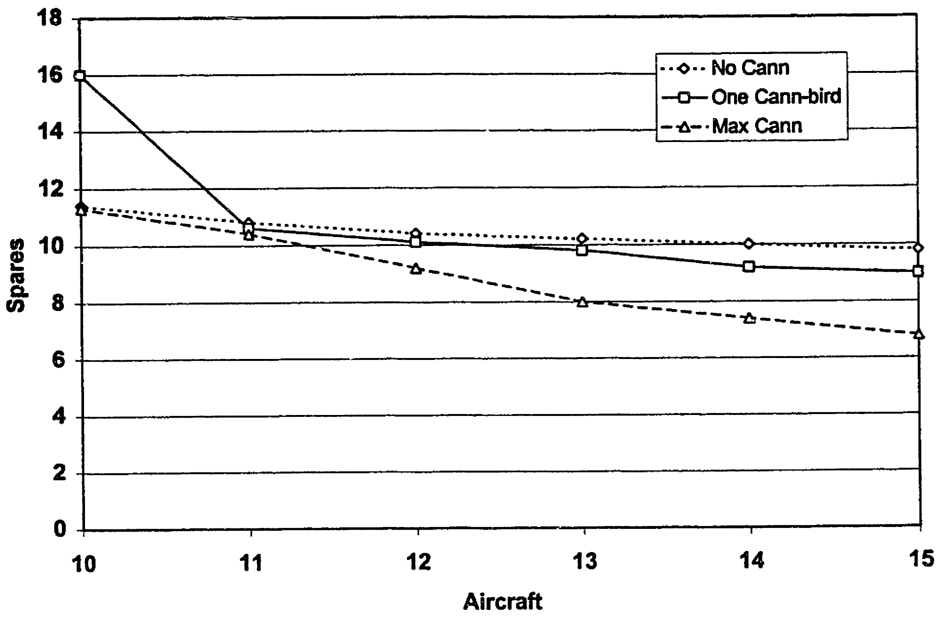

Figure 6-18 The Mix of Spares, Aircraft, and Level of Cannibalization required to achieve 90% for the percentage of requested flights flown.

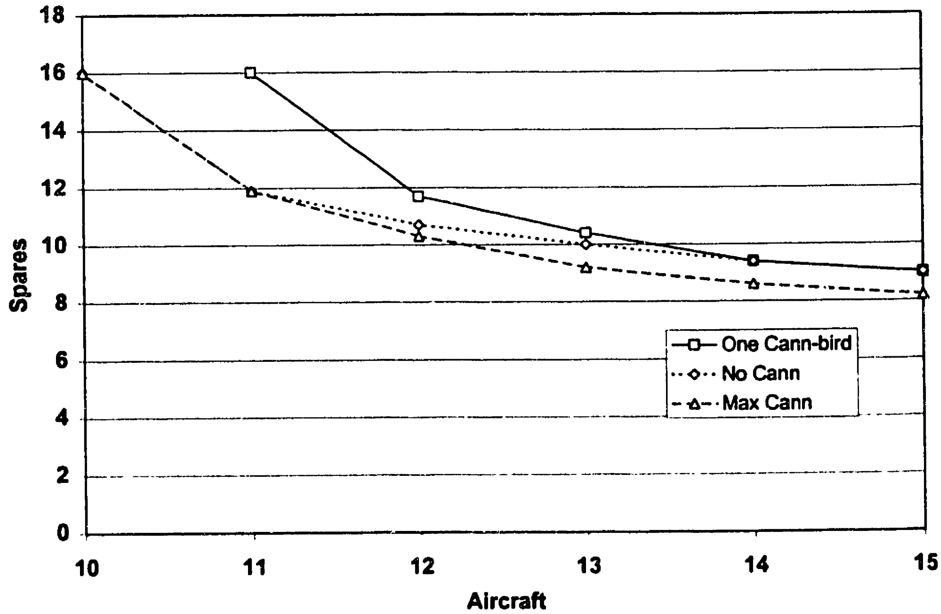

Figure 6-19 The Mix of Spares, Aircraft, and Level of Cannibalization required to achieve on average 10 Mission Capable aircraft.

6.1.7 Depot Repair Time

In our examination of tradeoffs to this point we have not considered the impact of depot repair time. In the previous section, the depot repair time was held constant at 50 days. Depot repair time, along with the parts' probability of failure, determines what an effective stock level for spares must be.

Depot repair time can be thought of as an inventory reorder time. The stock level for spares should be sized so that the inventory can cover the demands that occur during the time that it takes to repair or "reorder" the spares. A shorter repair time requires a smaller buffer inventory of spares to cover it. Thus, for a given readiness level fewer spares are needed as repair time goes down. Or, for a given spares level, if the repair time is shorter, the squadron readiness is greater. The model results illustrating the effects of different repair times are shown in Figure 6-20 through Figure 6-25. These results all assume one Cann-bird is used. Figure 6-20 and Figure 6-21 show the percentage of requested missions flown, for the cases of 15 aircraft and 10 aircraft respectively. Similarly Figure 6-22 and Figure 6-23 each show the MC rate and the number of MC aircraft. Figure 6-24 and Figure 6-25 show the number of cannibalizations.

The results show a strong impact due to the repair time. In Figure 6-20, which show the results for 15 aircraft, a depot repair time of 70 days and a stock level of 8 generates a percentage of requested flights flown of 68%. Keeping the same stock level of 8 and reducing the repair time to 50 days, a reduction of about 29%, increases the percentage of flights from 68% to 85%, an increase of 25%. Looking at the results in a different way, for a depot repair time of 70 days and a stock level of 8, if the percentage of flights is maintained at 68% but the repair time is reduced by 29% to 50 days, the number of spares can be reduced by about 30%, from 8 to roughly 5.5 (of course if we must consider a whole number of 6, the reduction is 25%). In the case where the spares level is 4, the results are even stronger. A reduction in repair time of 29% (from 70 days to 50) results in an increase in percentage of flights flown of about 35% (from 41% to 55%) or a decrease in the required stock of about 40% (from 4 to about 2.25).

For comparison to earlier calculations of the impact of more spares or more aircraft, consider the case of 10 aircraft and 8 spares with a repair time of 50 days, as shown in

Figure 6-21. A reduction in repair time of 40%, from 50 days to 30 days, could increase the percentage of flights flown from 71.5% to 87.4%, an increase of about 22%. As shown earlier in Figure 6-4, an increase in spares of about 50% or an increase in aircraft of over 50% would be required to produce a similar increase in percentage of flights.

Reduced repair time has a similarly strong effect on MC rate, the number of MC aircraft, and the amount of cannibalization. The behavior of MC rate and MC aircraft for different repair times resembles the behavior of the percentage of requested missions flown. The amount of cannibalization changes in an interesting way with the repair time. As repair time decreases, the amount of cannibalization decreases for higher spares levels. This occurs because as repair time decreases, the existing spares level is able to more adequately cover demands, resulting in fewer cannibalizations. However, for lower spares levels, the amount of cannibalization tends to increase as repair time decreases. This increase in cannibalization for low spares levels as repair time decreases occurs because the lower repair time is leading to more flights, and thus more failures and more cannibalizations.

The impact of repair time illustrates the potential of such programs as Agile Logistics and other such initiatives that aim to reduce depot repair times. The cost of reducing depot repair time was not captured in this model. If such programs can cost-effectively produce reductions in repair time, they will be of great value to the Air Force.

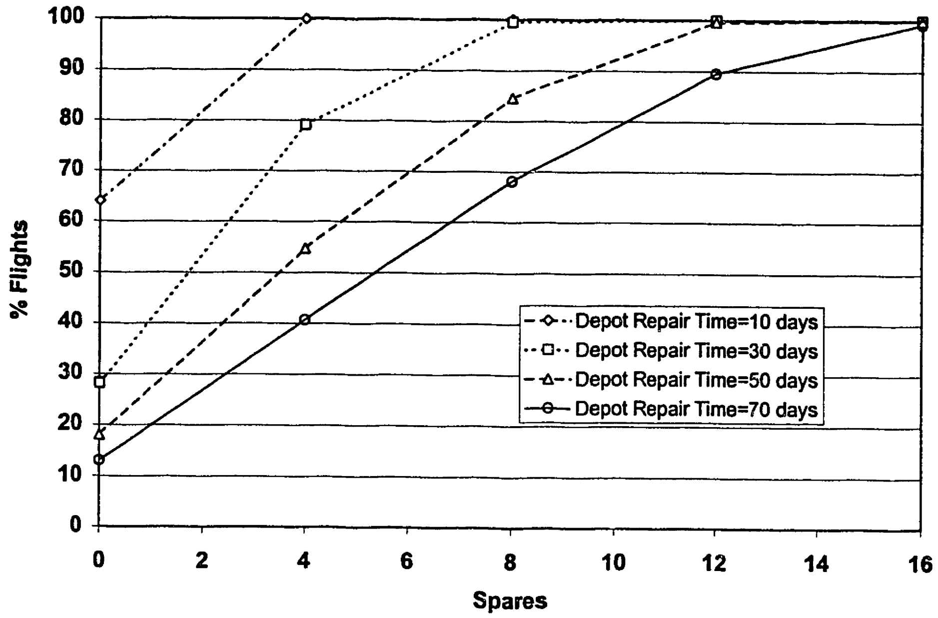

Figure 6-20 Percent of requested missions flown resulting from different repair times, for 15 aircraft, using one Cann-bird.

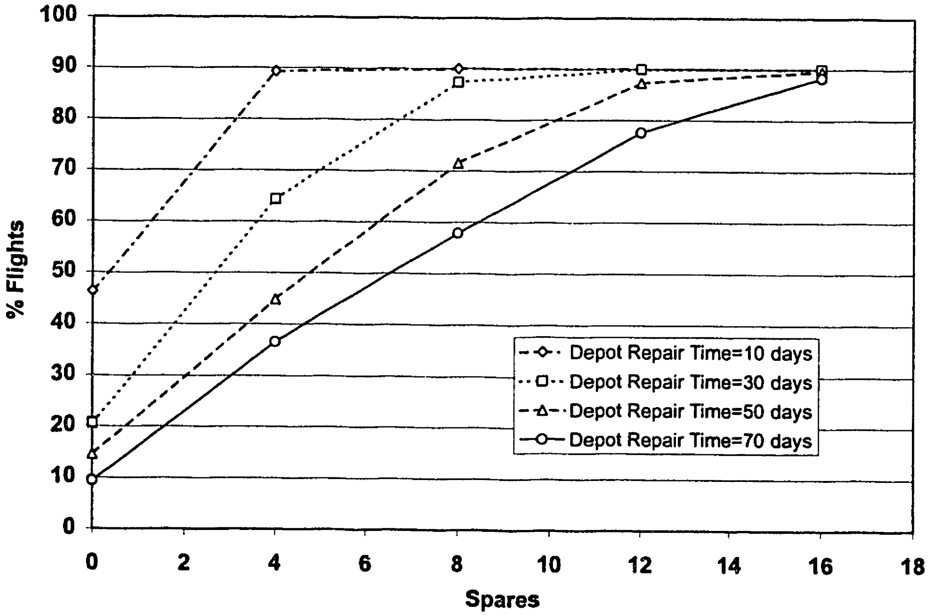

Figure 6-21 Percent of requested missions flown resulting from different repair times, for 10 aircraft, using one Cann-bird.

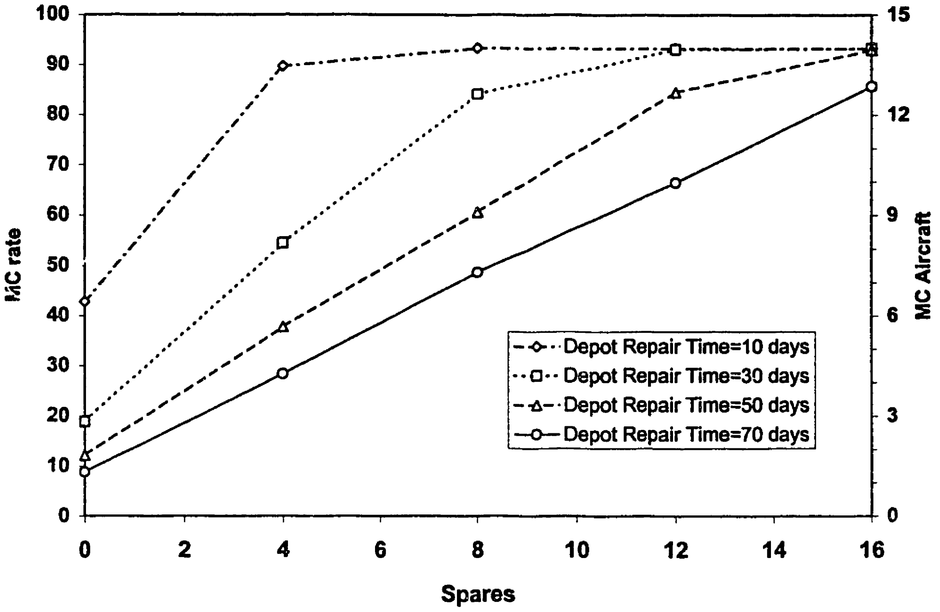

Figure 6-22 Mission Capable rate and number of Mission Capable aircraft resulting from different repair times, for 15 aircraft, using one Cann-bird.

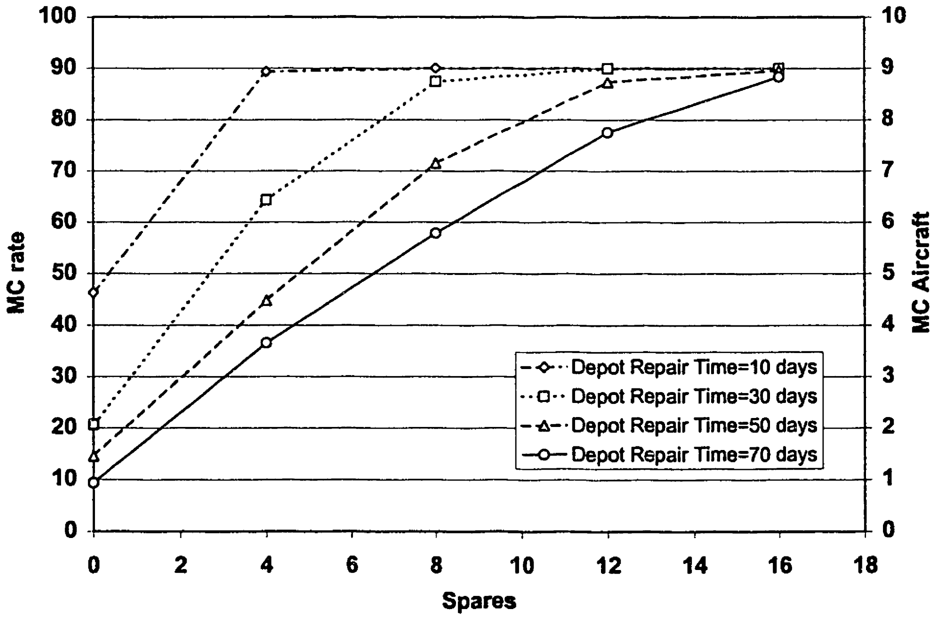

Figure 6-23 Mission Capable rate and number of Mission Capable aircraft resulting from different repair times, for 10 aircraft, using one Cann-bird.

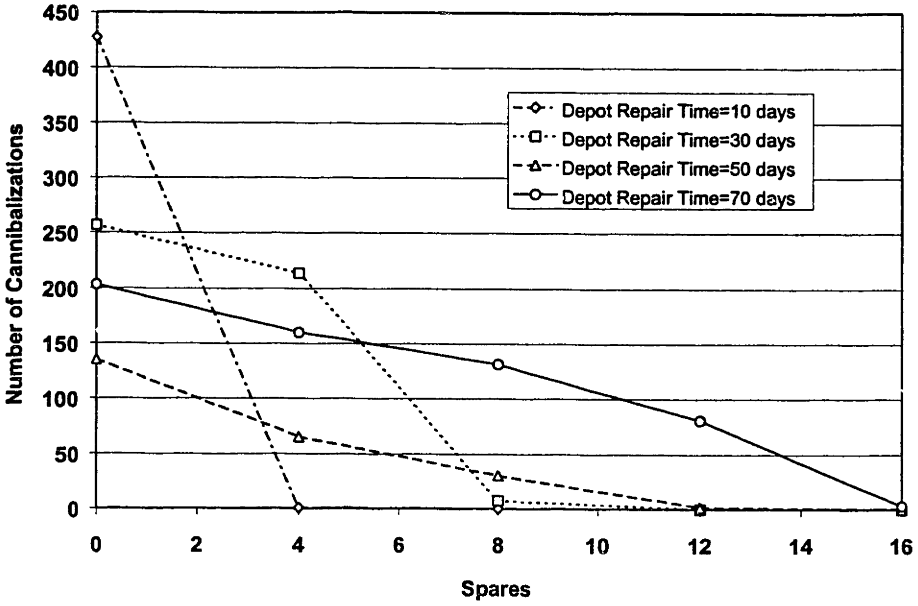

Figure 6-24 Number of cannibalizations resulting from different repair times, for 15 aircraft, using one Cann-bird.

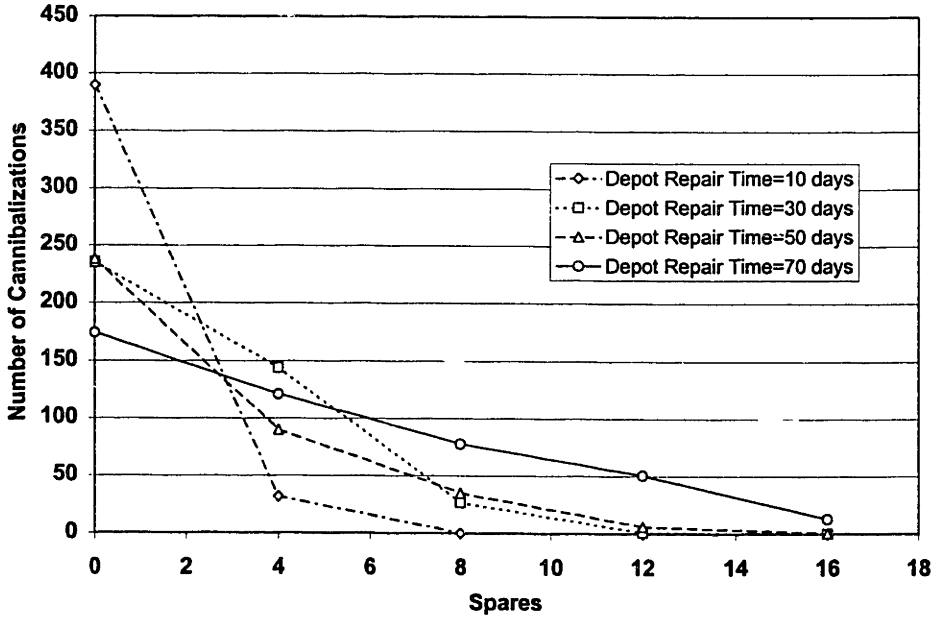

Figure 6-25 Number of cannibalizations resulting from different repair times, for 10 aircraft, using one Cann-bird.

6.1.8 Probability of Failure

Probability of failure is another parameter that determines the appropriate stock level. So far the model has assumed that each part has a probability of failure of 0.015. This thesis, with the exception of this section, assumes that this probability remains constant and equal for all parts. This section will address varying the probability of failure while all the parts still maintain equal probabilities. The issue of aging will be explored, expressed as an increasing probability of failure.

In reality the probability of failure fluctuates based on many factors including the aircraft flight environment, the amount and intensity of use, aging, etc. The magnitude of the fluctuation is not known, particularly for the aggregated "parts" that the model uses, so this section will briefly examine the impacts of different probabilities of failure.

Figure 6-26, Figure 6-27, and Figure 6-28 illustrate the effects of different probabilities of failure, for the case of 15 aircraft using one Cann-bird. As probability of failure increases, the percentage of flights flown, the MC rate, and the number of MC aircraft decrease. The number of cannibalizations tends to increase, until a limit is reached where the number of flights declines to the point that cannibalizations also decline. This behavior is shown by the probability of failure of 0.10.

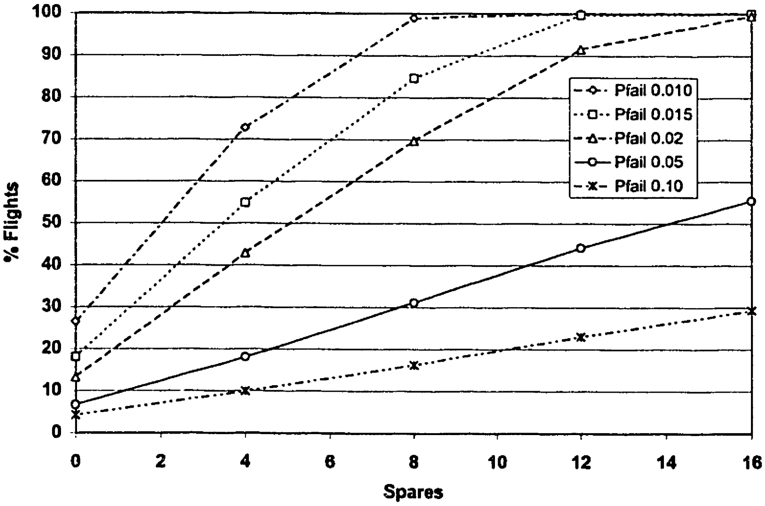

Figure 6-26 Percent of requested flights made for varying probabilities of failure, with 15 aircraft, using one Cann-bird.

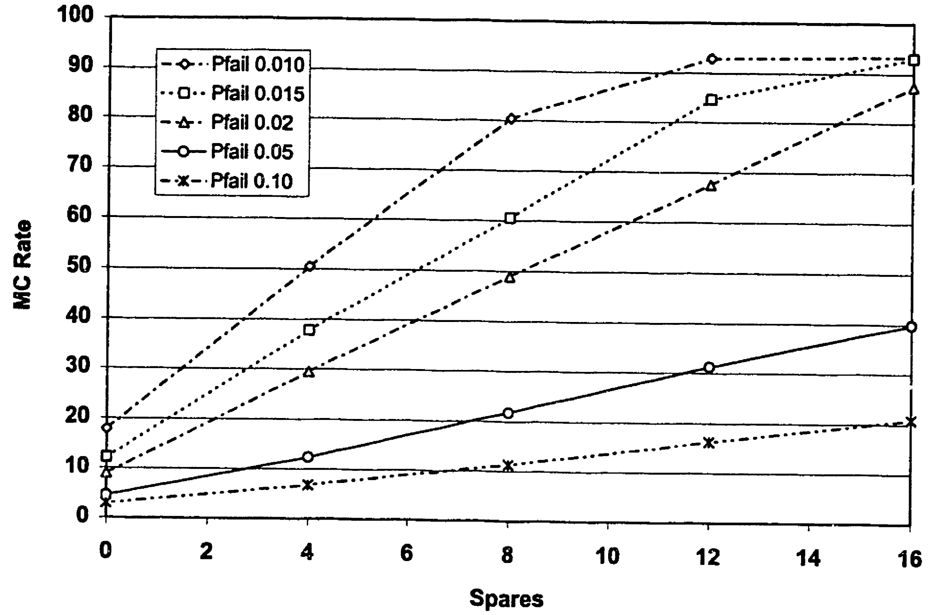

Figure 6-27 MC rate resulting from different probabilities of failure, for 15 aircraft, using one Cann-bird.

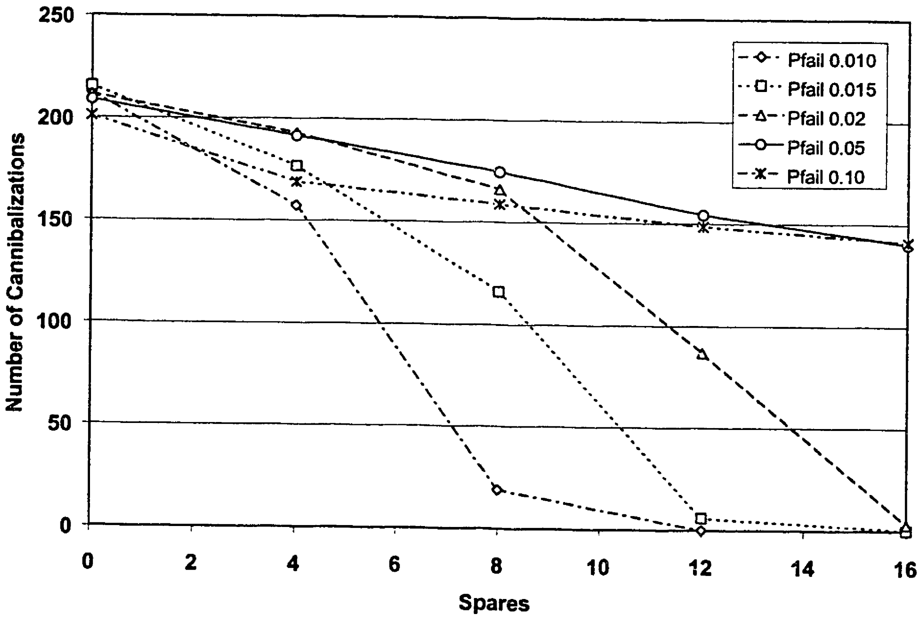

Figure 6-28 Number of cannibalizations resulting from different probabilities of failure, for 15 aircraft, using one Cann-bird.

To get a sense of the impact of reducing the probability of failure, consider the case where the probability of failure is 0.015 and the number of spares is 8. Reducing the probability of failure to 0.010, a decrease of about 33%, increases the percentage of flights flown from about 85% to roughly 99%, an increase of about 16%. A similar reduction in probability of failure for the case where spares equals 4 would increase the percentage of flights flown from 55% to 73%, an increase of roughly 33%. Designing components with lower probability of failure could significantly impact system performance.

The probability of failure can be used to represent aging parts. As parts age, their probability of failure tends to increase. The failure probability increases more rapidly as parts reach the end of their useful lifetimes, a behavior described as reaching the end of a "bathtub" (after the shape of the curve of failure probability over time). This increase in failure probability tends to be exponential.

Figure 6-29 shows the behavior of the percentage of flights made and the MC rate as the probability of failure increases, for the given case of 15 aircraft with 12 spares and a depot repair time of 50 days. It suggests that performance is strong until a threshold of probability is reached where the maintenance and logistics system can no longer fix all failures in a timely fashion. Once this probability threshold is passed, MC rate and percentage of missions flown decrease rapidly as the probability of failure increases. The decrease would be made even more dramatic if the probability of failure increased exponentially with time rather than linearly. The rapid decrease that could result from the combined behavior would suggest that performance problems could arise and worsen rather quickly once a certain threshold of supportability for probability of failure was passed.

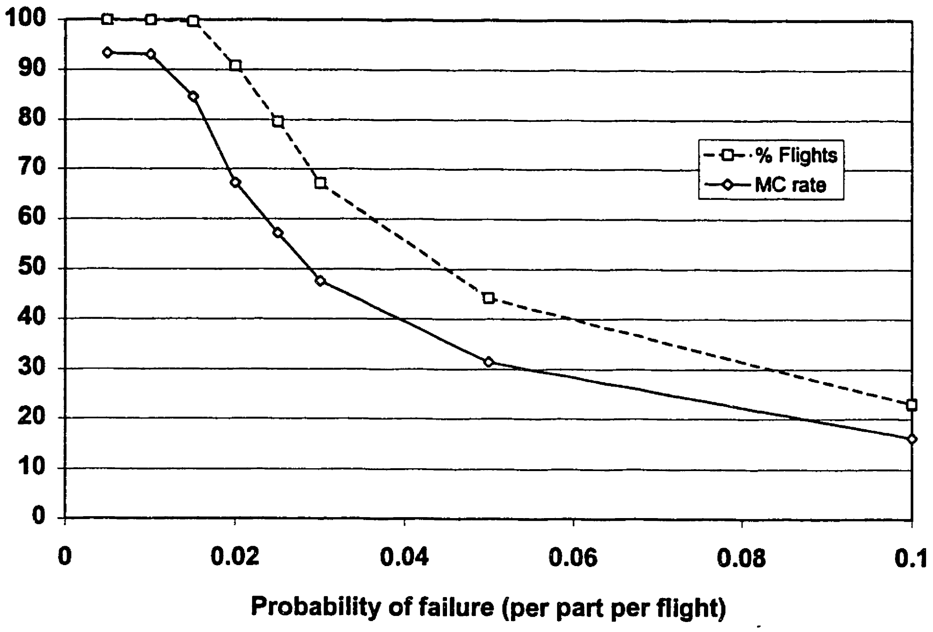

Figure 6-29 Effects of increasing probability of failure on percent of requested flights flown and on MC rate, for 15 aircraft, 12 spares, and depot repair time of 50 days.

6.2 CHARACTERIZING TRADEOFFS IN COST DUE TO AIRCRAFT, SPARES, AND CANNIBALIZATION

Thus far we have talked about performance, in terms of percentage of missions flown, MC rate, number of MC aircraft, and cannibalization. However, there are important costs that must be considered in looking at these tradeoffs between the number of aircraft, the amount of spares, and the amount of cannibalization. For example, more flights can be flown if the number of spares are increased, or the number of aircraft increased, or the level of cannibalization increased. But these additional measures have additional costs. Reducing the number of spares or the number of aircraft would reduce holding costs and capital costs, and reducing the amount of cannibalization would reduce operational costs, especially labor costs. But there are also costs due to not flying missions. The cost of not flying a flight is in effect a stock-out cost, resulting from a "stock out" of mission capable aircraft. These penalty costs, or stock-out costs depending on how one thinks of them, must be weighed against the capital and operating costs.

6.2.1 Cost Assumptions

For this model we will first consider the tradeoffs between the following costs: the cost of the aircraft, the cost of the spares, the cost of missing a flight, and the cost of cannibalization. As getting data to determine these costs proved to be difficult and as the model aggregates the real parts in a plane into 25 "parts", we must rely on making reasonable approximations consistent with the level of detail and purposes of this model.

6.2.1.1 Aircraft and Spares Costs

We approximate the cost of an aircraft to be $185 million.[55] This is based on the cost per C-5 aircraft in 1996 dollars as published by the Air Force. Data to determine the capital cost of a complete set of spares (one of each) is not readily available so values for the cost of spares will be varied to study the effect. The holding cost per year for the spares is assumed to be about 10% of the assumed capital cost of the spares. This holding cost is a best guess approximation for a collection of costs including warehousing, insurance, and breakage and spoilage.

6.2.1.2 Missed Flight Cost

The cost of missing a flight on a given day is very difficult to determine, as it can vary greatly from flight to flight. For purposes of this thesis, we will approximate this cost as a fraction of the cost of lost revenue for the Air Mobility Command. The Air Mobility Command is required to charge its customers to cover costs for operations. Lost revenue from not flying a mission is approximately $250,000 per mission. Since a flight missed on one day can generally be rescheduled and not missed altogether, this may be a high approximation for missing a typical flight on a given day. However, if we consider the value of the missed or delayed flight to the Air Force or to the US, the value may be higher in certain circumstances. For example, the value of delivering critical supplies to the frontlines to enable troops to fight effectively may be extraordinarily high, and the cost of a delay may be very high. For this model, assume that the Air Force determines the cost to be on average half of the lost revenue for completely losing a mission. Perhaps sometimes when a flight is missed missions are lost but other times missions can be rescheduled. Thus the cost of a missed flight is approximated as $125,000

---

[55] Based on the cost per aircraft as stated in *USAF Fact Sheet: C-5 Galaxy*, http://www.af.mil/news/factsheets/C_5_Galaxy.html

6.2.1.3 Cannibalization Cost

Another cost that is difficult to determine is the cost of cannibalization. In this model it will be approximated as the cost of its associated labor. Some Air Force maintenance personnel believe the Air Force considers the cost of the additional labor due to cannibalization to be zero, since Air Force personnel must work overtime to finish their work but are not paid based on the number of hours they work. For this model, however, we will consider the cost of cannibalization to be based on the composite rate of pay for a mid-level (E-5) Air Force maintenance technician. The composite pay includes basic pay, retirement pay, incentives, housing allowances, etc., and represents the cost to the Air Force of military personnel. The annual standard composite pay for an E-5 is approximately $40,000.[56] This does not include such costs as government furnished housing costs, medical costs, or overhead. Assume the total cost, including standard composite pay, housing, medical, and overhead, adds to $65,000. This equates to an hourly cost of approximately $32. Based on a years total number of cannibalizations and the total number of man-hours required to perform those cannibalizations, as reported by the 364th Air Mobility Wing for July 1997 to June 1998, we can approximate that a cannibalization on average required approximately 4.25 man-hours. Thus the average cannibalization costs about $136 in labor. Translating the cost of a cannibalization in the real world to the cost of a cannibalization in the model requires a conversion, since the number of cannibalizations in the model greatly understates the true number of cannibalizations in reality. For this model, $136 per cannibalization in the real world

---

[56] Air Force Instruction (AFI) 65-503, *US Air Force Cost and Planning Factors, Table A19-1,* (1999).

translates to about $2100 per cannibalization in the model, as shown in the footnote below.[57]

It is important to keep in mind that in the larger picture, this cost is an understatement. This cost does not consider the costs of low personnel retention rates, which may be caused to some extent by long hours, which may be in part caused by cannibalization. It also does not consider the cost of part breakage during cannibalization. The cost of part breakage may have significant effect, not only due to the cost for another repair, but due to the cost of more missed flights.

6.2.2 Determining Total Cost

With our assumptions for costs, we can then plot total cost over some time period as a function of the number of spares and the number of aircraft, to explore the response surfaces. Imagine in our model world that the Air Force was planning to begin operating a new wing of "C-5 like" aircraft, and wanted to determine an approximate mix of aircraft and spares to buy to minimize costs over 20 years. The costs to be considered are the initial purchase cost of the aircraft and spares, and the discounted cost over 20 years of cannibalization, holding costs, and missed missions (assuming as we have been that there is a fixed mission requirement of ten flights a day). All charts of cost presented in this thesis will be of the 20 year discounted cost. The initial cost of the aircraft and spares will be at year 0 and not discounted, while the cannibalization costs, holding costs, and missed mission costs will be calculated on a yearly basis for twenty years and discounted to year zero.

---

[57] The number of cannibalizations in the model is understated because the model represents an aircraft as only 25 parts. In order to scale the cost we need a point of comparison between reality and the model. Air Force data was acquired for 1998. To reflect the actual 1998 data as closely as possible, we use the base case model assumptions for the inputs, except we use a depot repair time of 75 days instead of 50. The repair cycle time of 75 more accurately reflects the actual times in 1998, the year we have data for, and we are trying to calibrate the model to reality. So we have a wing of 15 aircraft, each composed of 25 parts, trying to fly 10 missions a day. For the calibration, the parts each have a probability of failure of 0.015 and a depot repair cycle time of 75 days. Because we don't know what level of spares stock corresponds most to reality, once we run the model we plot the percentage of missions completed and the mission capability rate as a function of the spares stock level. Data from the Air Force shows C-5 MC rates averaged about 65% for the wing being modeled, and the number of flights flown was about 89% of the target based on the standard for commitment rate. The number of spares that produces results most similar to the real data is 12, giving an MC rate of 63% and a percentage of flights made of 86%. For a spares level of 12 the model indicates there will be approximately 90 cannibalizations in a year. In reality there were 2778 cannibalizations in 1998 for a full wing of 30 aircraft. Considering only half of the wing, there would be 1389 cannibalizations. Scaling the 90 cannibalizations in the model to 1389 cannibalizations in reality, we approximate that one cannibalization in the model represents approximately 15.4 cannibalizations in reality. Thus a cannibalization in the model should cost about 15.4 times a cannibalization in reality. This gives a cost for a cannibalization in the model as approximately $2100.

Assume the Air Force approximates the initial cost of a plane to be $185 million, and the initial cost of a complete set of spares (one of each) to be $37 million, or about 20% of the cost of an airplane. The ratio of 20% was chosen as a reasonable assumption. The discount rate is assumed to be 10%, and the annual holding cost for spares is assumed to be 10% of the initial capital cost of the spares. Assume that the Air Force approximates the cost of a missed flight on a given day as $125,000. The cost of a cannibalization in the model is set at $2100, and the Air Force plans to operate only one Cann-bird. The depot repair cycle time is assumed to be 50 days. We will call this set of assumptions the "base case".

Using the results from the model, the cost analysis produces the results shown in Figure 6-30 that indicate that it is most cost-effective with the given model assumptions to have about 12 spares for each part. With the cost numbers selected, the cost of missing flights over 20 years dominates the cost of initial spares and the holding costs, so enough spares must be purchased to ensure that few requested flights are missed. The curves for each number of aircraft reach a minimum between 12 and 16 spares. In this model, the minimum cost is achieved by having 11 aircraft and about 12 spares (the case of 10 aircraft is not as cost effective because when one plane sits as the Cann-bird, there are not enough aircraft to fly all the missions).

Let us say that after looking at the model's output, the Air Force chooses to operate 15 aircraft at an Air Force base (representing a 30 aircraft wing in the real world) and 12 sets of spares over other options of similar cost. Assume the Air Force prefers to have more aircraft and less spares, as having more aircraft enables greater surge capacity and allows easier scheduling of aircraft inspections, overhauls, etc., and having less spares seems to appear more "lean". The model does not take these factors into consideration when determining cost. The model predicts the choice of 15 aircraft and 12 spares to be a relatively cost effective option, although not the lowest cost option.

As shown in previous charts (Figure 6-4, Figure 6-5, Figure 6-6, and Figure 6-7), 15 aircraft and 12 sets of spares would generate a percent of requested missions flown rating of over 99%, an MC rate of about 85%, and on average nearly 13 MC aircraft. Cannibalizations would be almost unnecessary, with only a few performed all year.

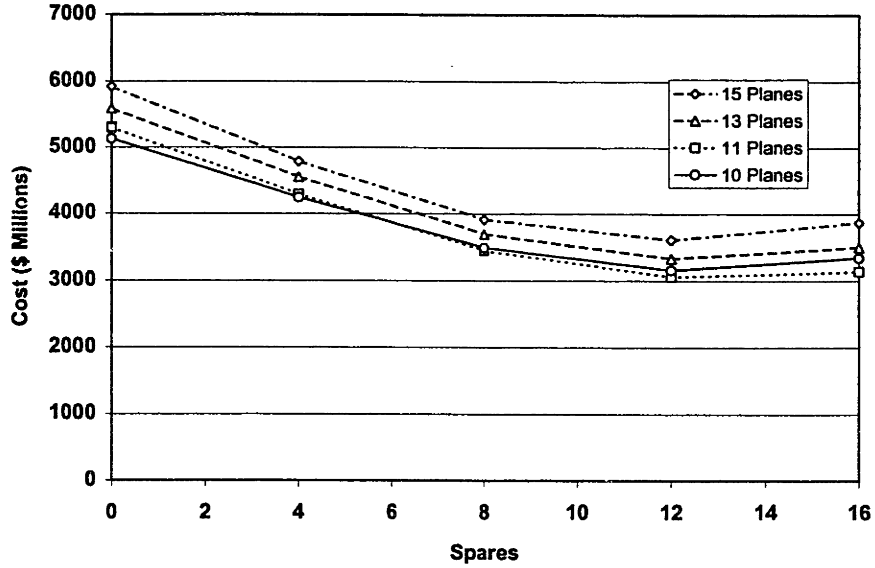

Figure 6-30 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50.

6.2.3 Effect of Relative Costs

The optimum number of aircraft and spares depends on the relative costs of aircraft, spares, and missed flights. Higher parts costs, higher holding costs, or lower costs for missing a flight would push the optimum number of spares lower. Higher costs for missing a flight, lower parts costs, or lower holding costs would push the optimum number of spares higher.

One of the parameters that the results are sensitive to is the cost of the spare parts. Given that it was not possible to make approximations based on data, the variations in total cost that result from changes in the cost of parts should be examined. If the cost of spare parts is higher, naturally the optimum number of spare parts decreases. To illustrate, Figure 6-31 and Figure 6-32 show the cost curves with higher assumed spare parts costs. In Figure 6-31 the cost of a complete set of spares (one of each) is assumed to be half the cost of an aircraft, or $92.5 million. In Figure 6-32 the cost of a set of spares is assumed to be equal to the cost of an aircraft, $185 million. Setting the cost of a complete set of spares equal to the cost of an assembled aircraft is an extreme case. The location of the lowest point in the cost curves shifts to the left as the cost of spares increases. For the case where the cost of a set of spares equals the cost of an aircraft, the low cost solution is to not have any spares at all.

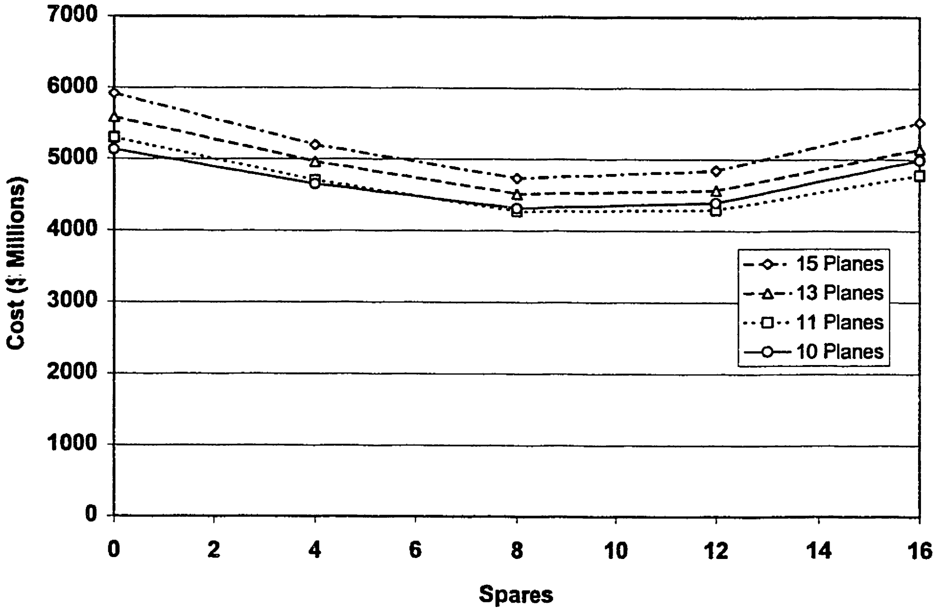

Figure 6-31 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50, cost of a set of spares equals 0.5 times the cost of an aircraft.

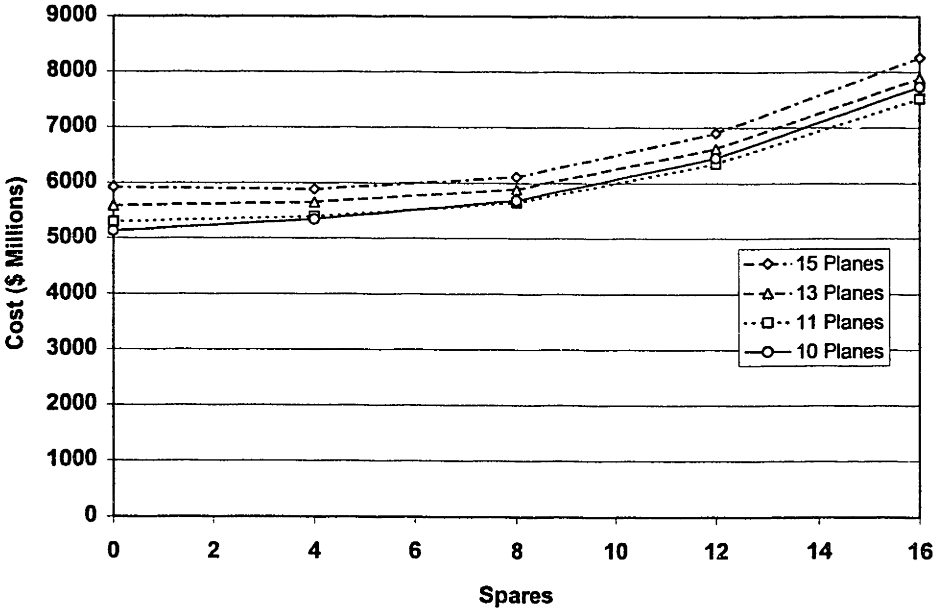

Figure 6-32 Costs for different numbers of spares and aircraft, with one Cann-bird, depot repair time of 50, cost of a set of spares equals the cost of an aircraft.

6.2.4 Effects of Repair Time

The results are also strongly dependent on other assumptions. For example, the model assumed a depot repair time of 50 days for the parts. If the actual repair time were higher or lower, the system would operate differently. The effects of different repair times were previously illustrated for 15 aircraft in Section 6.1.6 and Figure 6-20, Figure 6-22, and Figure 6-24. We now add consideration of the costs, shown in Figure 6-33. Results are plotted for 10, 30, 50, and 70 day depot repair times. The cost does not include the cost required to make the process improvements necessary for the depot repair time reductions. The differences in the curves are only due to the different repair times and the resulting costs of cannibalization and missed missions. Figure 6-33 shows that reducing depot repair time reduces cost and the number of spares needed. The difference in cost due to reduced repair time gives an indicator of the value of programs that will reduce the repair times. For example, if reducing the repair time from 70 to 10 days saves nearly $1 billion, then the value of a program that could achieve this repair time reduction would be nearly $1 billion.

To examine the sensitivities of our model results to the assumption for repair time, consider what could happen if a mistake were made in the assumptions. In our scenario, if Air Force planners assumed and budgeted for a 50-day depot repair cycle time and the actual time turned out to be 70 days, the results of the error would be quite troublesome. The percentage of missions flown would drop from the expected 99% to about 89% and the MC rate would drop from about 85% to about 67%. Cannibalization increases dramatically, from virtually none, to a large amount. As Figure 6-33 shows, the optimum number of spares to minimize costs is higher in the case where depot repair cycle time equals 70 days than it is for 50 days. For 12 spares, the cost is higher for 70 days than 50 days. If the actual depot repair time was 70 days, making a mistake in the planning by assuming 50 days would cause lower mission capable rates, lower percentage of requested missions flown, higher cannibalization, and higher cost.

This scenario showing the effects of making errors in the assumptions echoes of troubles that the Air Force has faced in the implementation of two-level maintenance and the implementation of Agile Logistics. The successful implementation of two-level maintenance depended on the assumption that lean practices would reduce the depot repair time. When two-level maintenance was put in place largely before reductions in depot repair times were made, the performance of the system suffered. In the case of the implementation of Agile Logistics, the budget for spare parts and repairs was cut by nearly $1 billion dollars between 1997 and 1999 in anticipation of expected improvements in depot repair times. The "savings" from the Agile Logistics program were calculated into future budgets, before the actual improvements occurred. When the depots were not able to meet performance expectations, largely due to supply problems of their own, system performance suffered and aircraft MC rates fell.

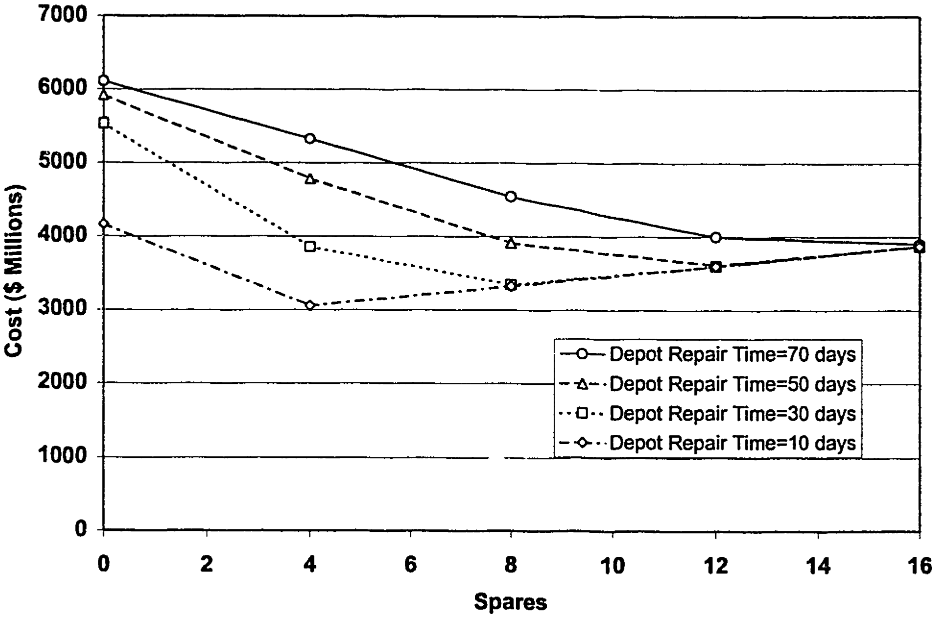

Figure 6-33 Cost for Different Depot Repair Times, with 15 aircraft.

6.2.5 Comparison of Costs for Various Cannibalization Options

The scenario we have discussed so far has only examined the case of cannibalization limited to one Cann-bird. This most closely reflects Air Force peacetime operations. Here, we consider the costs for no cannibalization or unrestricted cannibalization. To compare the cannibalization cases, Figure 6-34 plots the cost for 10 aircraft and 15 aircraft (as well as 11 aircraft in the case of one Cann-bird), for each of the three different cannibalization cases. These plots provide the cost "boundaries" of each of the cannibalization policies for the number of aircraft we have been examining. By boundaries we mean that the costs for 11, 12, 13, or 14 aircraft will fall between the curves for 10 and 15 aircraft (or 10, 11, and 15 aircraft for the case of one Cann-bird). The results show that the most cost-effective option, given the input parameters, is to have few aircraft and many spares. For example, for the case of unrestricted cannibalization or no cannibalization the minimum cost is for 10 aircraft and between 12 and 16 spares. For the case of one dedicated Cann-bird the minimum cost is for 11 aircraft and between 12 and 16 spares. In this case it is 11 aircraft and not 10 because one aircraft, the Cann-bird, is always not mission capable, and because the number of daily missions requested is 10. For a total of 10 aircraft, only nine can fly, causing more flights to be missed resulting in higher costs.

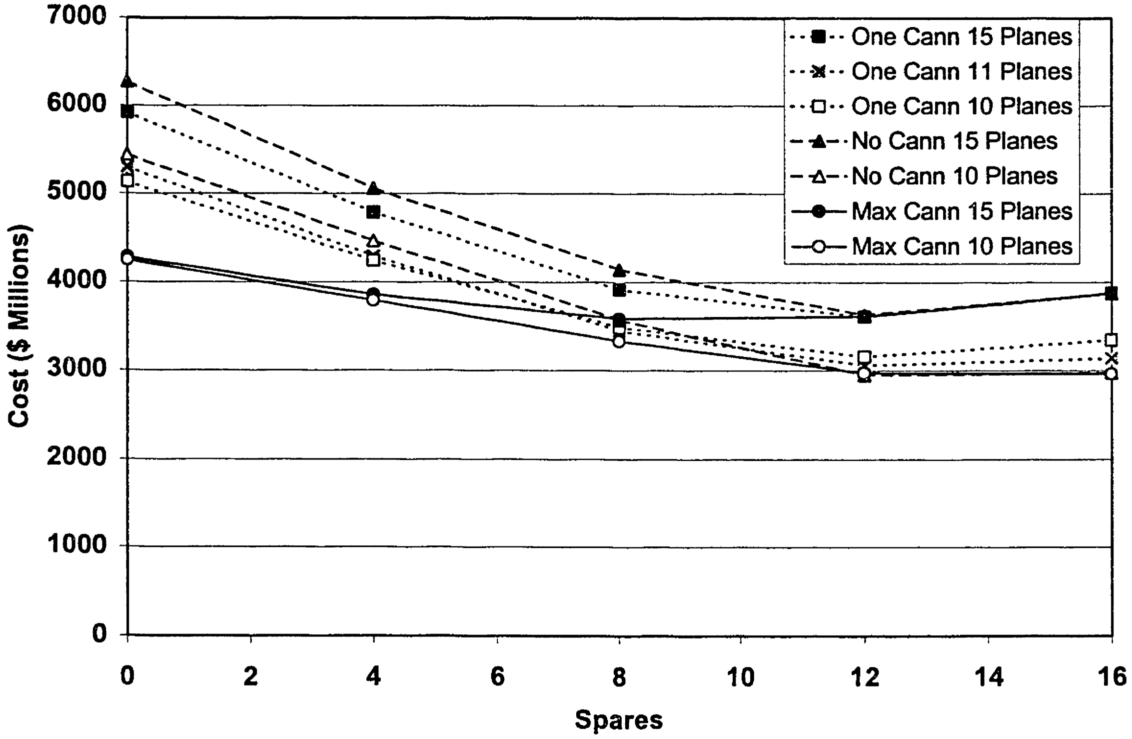

Figure 6-34 Comparison of Costs of Different Cannibalization Policies

To examine how the results are dependent on the cost of spares, Figure 6-35 plots the cost for the various cannibalization cases for the case where a set of spares costs half as much as an aircraft. Figure 6-36 plots the cost of various cannibalization cases where a set of spares costs the same as an aircraft. As the cost of spares increases, the lowest cost solution shifts to the left. Again, if the cost of a set of a set of spares is equal to the cost of an aircraft, then the cost is high enough that it is most cost effective to not have any spares at all.

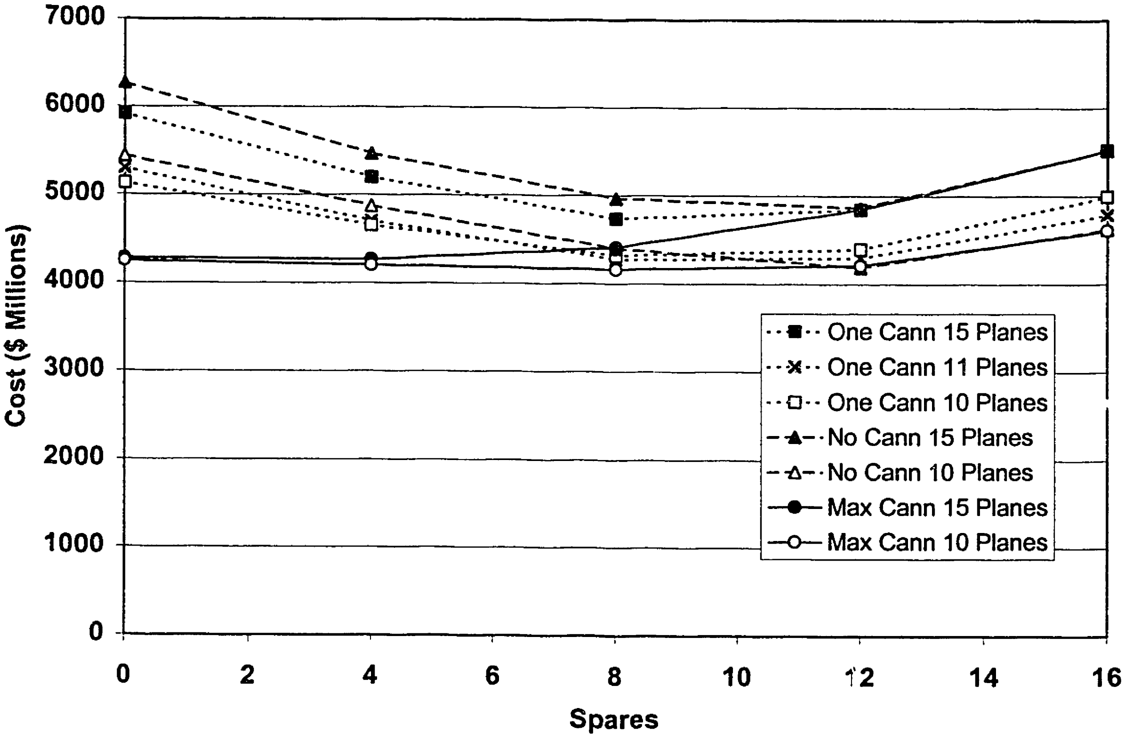

Figure 6-35 Cost for various numbers of aircraft, amount of spares, and cannibalization policies, if the cost of a complete set of spares (one of each) equals 0.5 times the cost of an aircraft.

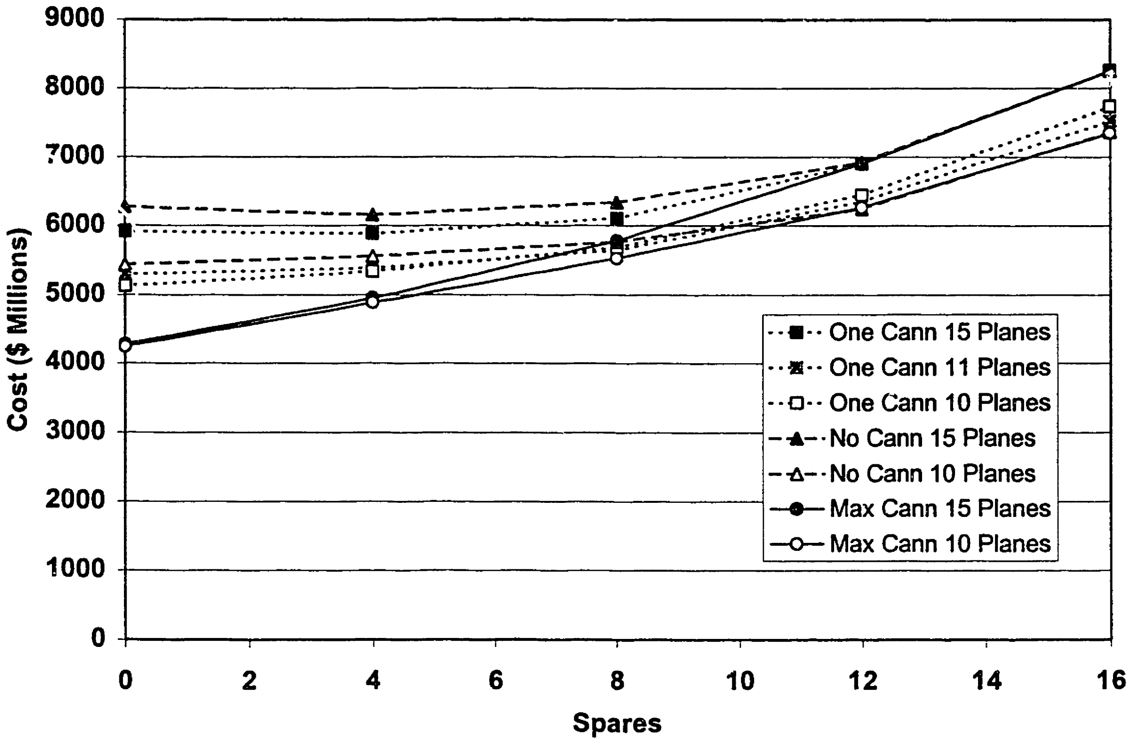

Figure 6-36 Cost for various numbers of aircraft, amount of spares, and cannibalization policies, if the cost of a complete set of spares (one of each) equals the cost of an aircraft.

Figure 6-34 shows some interesting implications. It shows that if the Air Force did not buy as many aircraft, but instead only bought enough to fly the 10 requested daily missions and bought a large number of spare parts (in the model more than 12 spares of each part), the cost to the system would be minimized. However, this option may not be realistically feasible. The model assumes that the daily demand for missions is fixed. In reality it is rather variable. In addition, in reality some aircraft need to be scheduled for inspections, programmed depot maintenance, or other such actions, and are not available for flight. Furthermore, the Air Force would like to have the ability to surge, to increase its capacity in times of war. Thus for these reasons, the Air Force may wish to have a higher number of aircraft, like 15. The model says that with 15 aircraft, the most cost-effective option is to have fewer spares, about 8, and cannibalize as much as needed. This however, assumes that labor is not constrained. In reality labor is a resource and is constrained, to the point that even using only one Cann-bird may cause some stresses. Thus having a low number of spares and cannibalizing to the maximum extent possible may not in reality be a cost-effective option. It depends on how much cannibalization is required and how much labor capacity is available. This leads us to the next best solution of having 12 spares and 15 aircraft. At this level of spares, cannibalization is not much of a factor.

When LSI researchers first observed the use of a Cann-bird, the question was asked, "Instead of having a Cann-bird, why not just have one more part of each kind on the shelf?" This question can be explored with the cost curves calculated from the model results. Figure 6-37 shows the cost curves for 14 aircraft with no cannibalization, 14 aircraft using one Cann-bird, and 15 aircraft using one Cann-bird. The base case costs are used for these curves (an aircraft costs \$185 million, set of spares costs \$37 million, cannibalization costs \$2100, etc.) Consider a wing that has 15 aircraft, 6 sets of spares, and is using one of the aircraft as a Cann-bird. The curve indicates that the total 20 year discounted cost is about \$4.3 billion. Would it be more cost effective if the wing had instead 14 aircraft and 7 sets of spares and did not cannibalize? Indeed, the cost would be slightly less, a little under \$4.2 billion. However, if the wing did have 14 aircraft and 7 spares, it would be even more cost effective to use one of the 14 aircraft as a Cann-bird. For the number of aircraft and spares considered in this model and the cost assumptions that we have made, it is generally more cost effective to cannibalize one of the aircraft than it is to not cannibalize at all.

For a given set of planes, the result that it would be more cost-effective to have an extra set of spares, one less aircraft, and not cannibalize than it is to use a Cann-bird depends on the relative cost of aircraft and spares. The cost of a set of spares must be low compared to the cost of an aircraft. In the model, if the cost of a set of spares is less than about 0.5 times the cost of an aircraft, it is more cost effective to have an additional set of spares, one less aircraft, and not cannibalize.

The result that it is more cost effective for a given set of aircraft to cannibalize one of the aircraft rather than not cannibalize at all depends on a few conditions. For the cost assumptions we have used, the result is true if the number of aircraft exceeds the number of daily requested missions (in this case 11 aircraft or more), for any level of spares up to the level required to fly all of the requested flights (which in this case is about 12 spares). To illustrate these points, Figure 6-38 compares the cost curve for 15 aircraft with no cannibalization to the curve for 15 aircraft with one Cann-bird, and compares the curve for 11 aircraft with no cannibalization to the curve for 11 aircraft with one Cann-bird. For the case of 10 aircraft in this model, cannibalization is more cost effective as long as the number of spares is less than approximately 9 spares. Figure 6-39 compares the curves for 10 aircraft.

The result that it is generally more cost effective to use one aircraft as a Cann-bird rather than not cannibalize at all also depends on the relative costs of missed missions and cannibalization. For the model inputs we have used, the cost of missed missions outweighs the cost of cannibalizations. If this were not the case, cannibalization would not be cost effective. If the cost of cannibalization were high enough relative to the cost of a missed mission, it would not be more cost effective to cannibalize. With the given cost of $125,000 for a missed mission, in order for the use of one aircraft as a Cann-bird to be less cost effective than not cannibalizing at all, the cost of a cannibalization in the model would have to be at least $200,000.[58] This translates to a cost of a cannibalization in real life of around $13,000.[59] This cost is very high if we are only considering labor costs. If we consider the potential cost of failed cannibalizations or the impact of overworked maintenance personnel, this might not seem quite so high.

While it is generally more cost effective to cannibalize one aircraft, this does not mean that there are no tradeoffs. The cost calculation does not take into account the number of mission capable aircraft. When there are adequate spares the use of a Cann-bird can reduce the number of mission capable aircraft. In the model, this reduces surge capacity, as the Cann-bird is assumed to not fly. Of course, in the real world if there was a need for wartime surge and there were adequate spares to make the Cann-bird ready, the cannibalization policy could always be changed so that the aircraft that was the Cann-bird could be allowed to fly.

---

[58] Determined by iterating the cost calculation with increasing cannibalization costs until the overall cost for the cases of no cannibalization and one Cann-bird were approximately equal.

[59] Recall that a conversion was made because the model, due to its abstraction of an aircraft as only 25 parts, understates the real number of cannibalizations. See footnote 57.

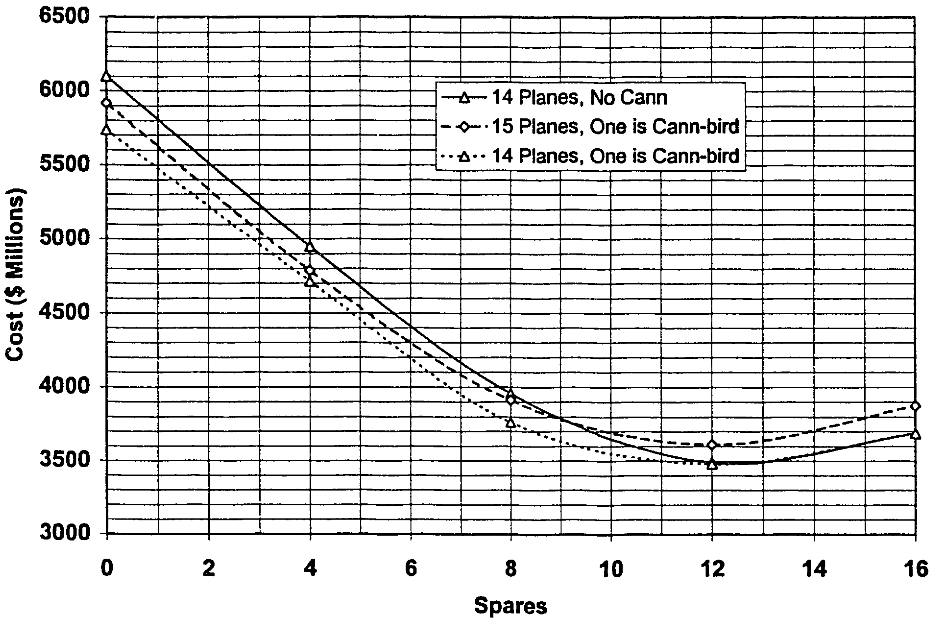

Figure 6-37 Cost Comparison: 15 aircraft with one Cann-bird, 14 aircraft with one Cann-bird, and 14 aircraft with no cannibalization.

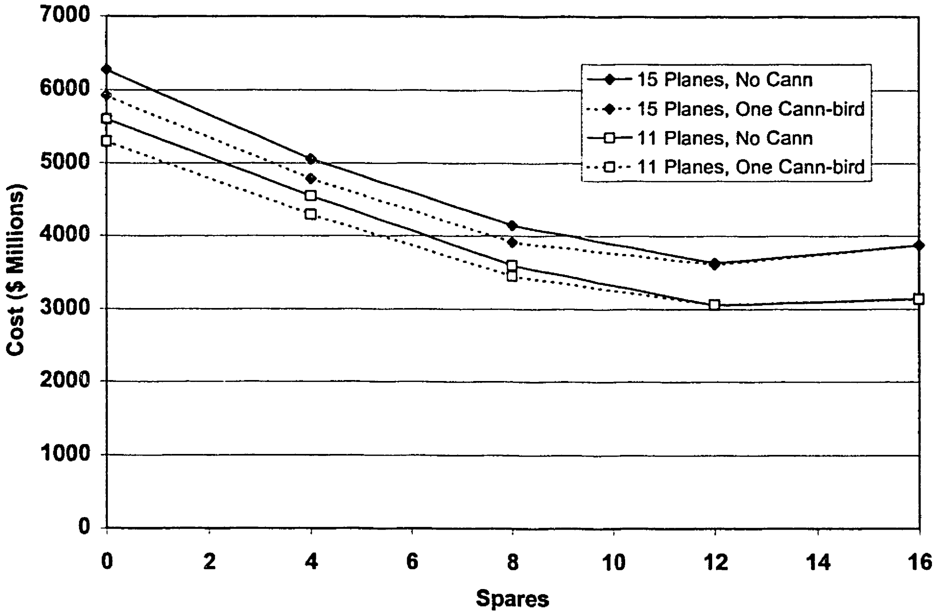

Figure 6-38 Cost Comparison: cost with no cannibalization vs. cost using one Cann-bird, for 15 aircraft and 11 aircraft.

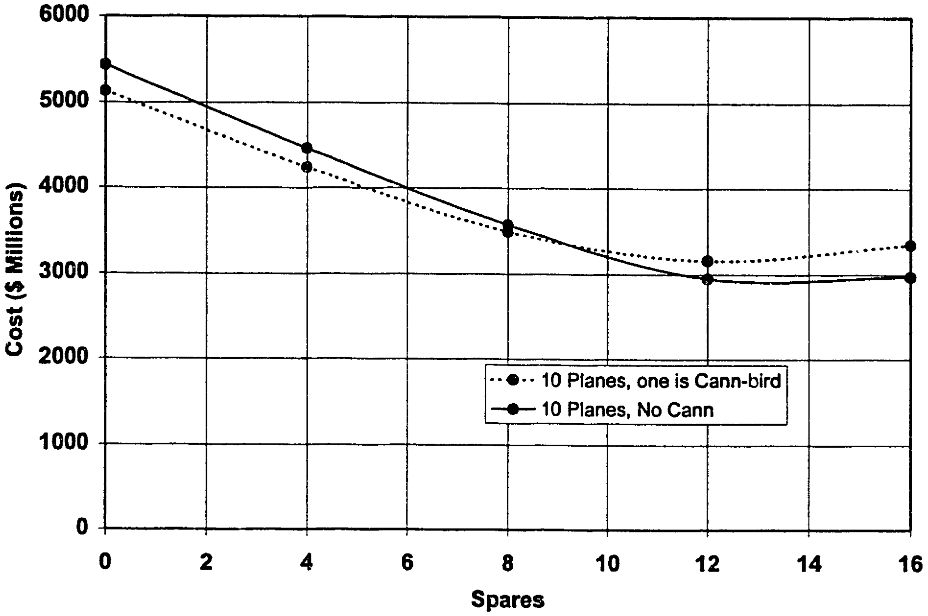

Figure 6-39 Cost Comparison: cost with no cannibalization vs. cost using one Cann-bird, for 10 aircraft.

6.3 SUMMARY

The model is able to illustrate tradeoffs between the number of aircraft, the number of spares, and the cannibalization policy. It also can be used to explore issues such as varying repair time and probability of failure. It can be used to show that within the given parameters of the model, cannibalization is generally a cost-effective policy.

Chapter 7 draws conclusions from these results and provides recommendations for further study.

CHAPTER 7: CONCLUSION AND RECOMMENDATIONS FOR FURTHER STUDY

As eagerly as he constructed mathematical models, colleagues and former students said, Mr. Durand used the results as an aid to judgement rather than the final answer. 'Frequently this did not endear him to those enamored of a model' (said a former student)

Obituary for David Durand, MIT statistician,
in NY Times, March 1, 1996

7.1 CONCLUSION

This thesis used a high-level simulation model to illustrate characteristics and critical tradeoffs in the complex Air Force sustainment system. Specifically, this thesis modeled the sustainment system for a wing of C-5 aircraft and examined the high-level tradeoffs in readiness and cost between the number of aircraft, the inventory policy (that is, the number of spares), and the cannibalization policy. It also addressed tradeoffs due to the depot repair time and due to the probability of failure of components.

System models can be valuable tools for characterizing and understanding the Air Force sustainment system, can generate insight into system behavior, and can provide a valuable tool for learning. However, as with any modeling approach, there are tradeoffs in the way a system is represented. It is important to understand that there are limitations to the approach used in this thesis. The simple high-level simulation model used in this thesis makes abstractions, approximations, and assumptions to simplify the complex real world Air Force sustainment system. Many detailed issues in the real world system are not considered in the model. Thus, the results from this model must be used qualitatively and not quantitatively. The benefits to this approach are that the simple model can provide illustrations of general tradeoffs with a clarity that may be masked in a more complex model. The model is not data intensive and does not require the collection of large amounts of data from the Air Force. Rather, information from publicly available sources and reasonable approximation provides sufficient input. The simplicity of the model also allows a large number of tests to be completed quickly.

Perhaps the key point of this thesis is that it is important to consider the systemic effects of decisions and changes in the Air Force. Making decisions that seem to be beneficial on a local scale may in fact be detrimental to the overall system. Well meaning individuals or organizations may act in a counterproductive manner if they are not aware of their impact on the system. Furthermore, if performance metrics are not established with overall system performance in mind, individuals and organizations may actually be guided to act counterproductively. Understanding systemic behavior and the impact of decisions on the overall system is crucial.

Regarding the specific topics of this thesis, it is important to understand the behavior of the Air Force sustainment system and the tradeoffs between such factors as the number of aircraft in a wing, the number of spare parts, and cannibalization. While a small squadron of aircraft, a low spare parts stock level, or limited cannibalization may seem desirable from a cost viewpoint when considered individually, the combination of the three is generally not cost effective for the system as a whole if one considers the cost of missed missions due to the lack of MC aircraft. Tradeoffs must be made for the benefit of the system as a whole.

For example, due to the high cost of missing missions, adequate spares should be retained to ensure that a very high percentage of requested missions be flown. Although a high level of spares may not seem lean, if the costs of missing missions is high compared to the cost of spares, maintaining a high level of spares is cost effective. As shown in Figure 6-1 and Figure 6-4, in the case of no cannibalization or in the case where only one Cann-bird is used, the impact on the percentage of missions flown is greater from an additional set of spares than it is from an assembled aircraft. A large investment to increase the number of aircraft produces only a modest increase in the percentage of missions flown.

In order to maintain surge capacity, it is important to maintain a sufficient number of aircraft as well as adequate spares levels. As discussed in Chapter 6 and shown in Figure 6-30, although it may be most cost effective during peacetime to retain a number of aircraft just large enough to meet peacetime demands, in times of war the aircraft will not be able to meet the increased demands. For example, as shown in Figure 6-4, with adequate spares both a wing of 11 aircraft and a wing of 15 aircraft can meet the peacetime mission requirements of 10 flights per day. The wing of 11 requires a larger investment in spares than the wing of 15 aircraft. However, in the event of war, if the requirements jump to 13 flights per day, the smaller wing of 11 aircraft will clearly fail to meet the higher requirements. With 11 aircraft, a large investment was made in spares so that in peacetime the smaller wing of aircraft could meet demands. But in wartime the large number of spares does not amount to much when there are simply not enough aircraft. While a high level of capacity during peacetime may not seem lean, one must consider that the Air Force must also be prepared for wartime surge.

The Air Force must be efficient during peacetime yet prepared for the surge requirements of war, a challenge of conflicting expectations that requires many tradeoffs. Peacetime requirements and wartime requirements drive towards different behavior. The need to operate cost effectively during peacetime suggests fewer resources, while the need to be ready for war suggests more resources. It is apparent that metrics are needed that reflect this dichotomous situation to ensure that decision-makers are guided to act appropriately.

From an operations standpoint, cannibalization is highly inefficient due to the additional labor required. However, given reasonable cost assumptions, cannibalization was generally found to be a cost-effective means of increasing aircraft mission capability. As shown in Figure 6-12 through Figure 6-15, cannibalization increases the percentage of flights flown and the Mission Capable (MC) rate, especially when there are lower levels of spares. The drawbacks to cannibalization are that it requires more labor and thus has higher labor costs, and when the number of spares is high it can reduce the number of MC aircraft. The reduction in MC aircraft is due to the Cann-bird being grounded when it otherwise might be flying. As the cost of cannibalization was assumed to be low compared to the cost of a missed flight, it was shown (in Figure 6-38) that the use of a Cann-bird is a cost-effective policy. Indeed, for the given cost assumptions, maximum cannibalization was shown in Figure 6-34 to the most cost-effective policy. However, constraints on the amount of labor that can be provided preclude this from being a viable option in the real Air Force.

The model was used to examine the impact of changing the depot repair time. As shown in Figure 6-33, reducing the depot repair time would decrease the optimum number of spares and could thus reduce the total cost of spares. This in part forms the basis for such programs as the Air Force's former Lean Logistics program and the current

Agile Logistics program. The cost-effectiveness of reducing depot repair time depends on how the cost of accomplishing the depot repair time reduction compares to the savings in spares.

The model also briefly explored the issues of aging and the changing probability of component failure. It was shown in Figure 6-29 that for a fixed number of aircraft and spares, the percentage of flights made and the MC rate did not decrease linearly with an increasing probability of component failure. The percentage of flights made and MC rate were shown to remain high until a threshold value for failure probability was reached. Once this threshold was passed, percentage of flights and MC rate decreased rapidly before leveling off at a low value. As probability of failure can be roughly correlated to aging, this suggests that decreases in performance could appear rapidly as parts reach a threshold age.

The issue of relative costs is important to consider when evaluating tradeoffs in the system. The results from the model are often highly dependent on the relative impact of the various costs in the system, including the cost of aircraft, the cost of spares, the cost of missed missions, and the cost of cannibalization. For example, if the cost of cannibalization became high enough relative to the cost of missed missions, then cannibalization would no longer be a cost effective option. If the ratio of the cost of spares to the cost of aircraft changes, the cost-optimal number of spares changes. The costs, especially the cost of missed missions, are difficult to estimate. The cost of a missed mission is difficult to estimate because it is dependent on the value of the mission, which may not have an assigned value or may have different values from different perspectives. For example, if a mission to provide medical supplies and ammunition to combat troops was missed and as a result casualties were higher, the cost would be difficult to determine. Thus it is important to understand the assumptions for cost and the sensitivity of the results when evaluating the results from the model.

It is also important to realize that the results in this thesis are dependent on many other assumptions and approximations. The results can not be taken quantitatively, but provide qualitative insights. When comparing real world systems to this model it is important to understand and check the applicability of the assumptions.

7.2 RECOMMENDATIONS FOR FURTHER STUDY

There are many opportunities for further study, both by using this model and by exploring the issues of the Air Force sustainment system in another manner. Many of the simplifying assumptions made in this model could be addressed. There are many valid directions for further study that could be quite useful. A few that the author considers interesting will be discussed here.

It may in some cases be useful to expand the current model to incorporate new ideas, but perhaps adding complexity would diminish the clarity of the model. For the sake of retaining simplicity it may be better to create separate models to address certain issues. Perhaps these separate models could be collected into a suite of models and be integrated with a post-processor.

One interesting option for further study would be to look more closely at the dichotomous expectations of peacetime efficiency and wartime readiness placed on the Air Force. In this model, the daily requested number of flights was assumed to be constant. The model could be enhanced to allow a probabilistic number of requested flights to simulate varying demand for missions. For example, perhaps 95% of the time the number of flights is 10, while 5% of the time the number of requested flights surges to 13. The costs of missing a flight could be varied as well, with perhaps the flights during surge carrying a much higher missed mission cost.

Another interesting possibility for further exploration is the consideration of a constrained labor force. In this model, it was assumed that the maintenance personnel would have sufficient capacity to finish all of the required repairs in time for the next day's missions. In reality this is not necessarily the case. A new model with a different approach may be required in order to address this issue adequately. The model would probably need to consider readiness on a more continuous time scale, and take into account NMCM (Not Mission Capable due to Maintenance) time as well as NMCS (Not Mission Capable due to Supply) time. A model that could take into account labor capacity and the effects of labor workload on fatigue and subsequently on performance could be rather interesting. It might also be useful to include such issues as the effects of labor workload on the retention rate of personnel.

This model assumed that all of the modeled parts could be cannibalized and that cannibalizations were always successful. In reality a fraction of cannibalization result in breaking the part that is being cannibalized. This increases the cost of cannibalization beyond just the labor costs, as repair costs as well as missed mission costs are potentially increased. A model that could examine this issue would be important support for this thesis. The attractiveness of cannibalization as shown in the model results of this thesis may be diminished slightly.

The model used in this thesis assumed that there was no attrition of aircraft or condemnation of spare parts. In doing so, this model eliminated consideration of the need to occasionally order new aircraft and parts. A model that could address the ordering process could be interesting, as the ordering process would change the calculation of total costs and could introduce problems such as unpredictable or long order lead times.

Depot repair times in the model used in this thesis were assumed to be deterministic and constant within the model. In reality there are large differences in repair time between parts of different type, and great variation in repair time for parts of the same type. A model that could include these differences and variability might have some benefit.

This thesis assumed that the modeled parts all had the same probability of failure, the same depot repair time, and the same cost. A model that could consider individual parts that had different properties would open up many possibilities for exploration. One issue that could then be considered is the cost effectiveness of stocking low demand spares versus that of relying on a Cann-bird for parts that fail infrequently.

As mentioned before there are many areas for further development of the model used in this thesis and the topics addressed in this thesis. It is hoped that this thesis will provide a useful stepping stone for further study.

BIBLIOGRAPHY

*Logistics Executive Summary.*436th AW/512th AW, Dover AFB. June 1998.

USAF Fact Sheet: C-5 Galaxy. http://www.af.mil/news/factsheets/C_5_Galaxy.html. .

"World Military Aircraft Inventory." *Aviation Week & Space Technology*, January 11, 1999 1999, 261.

Abell, John B. , and H. L. Shulman. *Evaluations of Alternative Maintenance Structures.* R-4205-AF. Santa Monica, CA: RAND. 1992.

Adams, John L., John B. Abell, and Karen E. Isaacson. *Modeling and Forecasting the Demand for Aircraft Recoverable Spare Parts.* R-4211-AF/OSD. Santa Monica, CA: RAND. 1993.

Air Combat Command, USAF. *Air Combat Command Director of Logistics Ten Year Lookback Standards and Performance FY88-FY97*USAF HQ ACC/LGPA. November 4, 1997.

Air Force Instruction (AFI) 65-503. *US Air Force Cost and Planning Factors, Table A19-1.* April 1, 1999.

Air Mobility Command, USAF. *1998 Air Mobility Master Plan*HQAMC.

Air Mobility Command, USAF. *HQ AMC Orientation.* Scott AFB, IL: HQ AMC/LG. August 12, 1998.

Air Mobility Command, USAF. *HQ AMC Orientation Briefing*AMC/LG, HQ. August 12, 1998.

Boulding, Kenneth E. *Economics as a Science.* New York: McGraw-Hill, 1970.

Brown, B. B. *Characteristics of Demand for Aircraft Spare Parts.* R-292. Santa Monica, CA: RAND. July 1956.

Crawford, Gordon B. *Variability in the Demands for Aircraft Spare Parts: Its Magnitude and Implications.* R-3318-AF. Santa Monica, CA: RAND. 1993.

Department of the Air Force, U.S. *U.S. Air Force White Paper on Long Range Bombers.* March 1, 1999.

General Accounting Office, United States. *Air Force Supply: Management Actions Create Spare Parts Shortages and Operational Problems.* GAO/NSIAD/AIMD-99-77GAO. April 1999.

Hillestad, R.J. *Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control.* R-2785-AF. Santa Monica, CA: RAND. March 1982.

Isaacson, Karen E. , and Patricia M. Boren. *Dyna-METRIC Version 5: A Capability Assessment Model Including Constrained Repair and Management Adaptations.* R-3612-AF. Santa Monica: RAND. 1988.

Isaacson, Karen E. , and Patricia M. Boren. *Dyna-METRIC Version 6: An Advanced Capability Assessment Model.* R-4214-AF. Santa Monica, CA: RAND. 1993.

Office of the Under Secretary of Defense. *National Defense Budget Estimates for FY 1999*U.S. Department of Defense. March 1998.

Pyles, Raymond A. , and Hyman L. Shulman. *United States Air Force Fighter Support in Operation Desert Storm.* MR-468-AF. Santa Monica, CA: RAND. 1995.

Raymond, Neil V. *Lean Sustainment Initiative Baseline Seminar Logistics Overview*GRC International, Inc. May 1997.

Romero, Maj. Craig. *Quarterly Agile Logistics Pipeline Analysis Update.* http://38.178.7.130/alconference/briefings.htm. August 1998.

Senge, Peter M. *The Fifth Discipline: The Art & Practice of The Learning Organization.* 1st ed. New York: Doubleday, 1990.

Slay, F. Michael, Tovey C. Bachman, Robert C. Kline, T.J. O'Malley, Frank L. Eichorn, and Randall M. King. *Optimizing Spares Support: The Aircraft Sustainability Model.* AF501MR1. McLean, VA: Logistics Management Institute. October 1996.

United States Air Force. *Posture Statement 1999.* http://www.af.mil/lib/afissues/1999/posture/index.html. 1999.

Womack, James P. , Daniel T. Jones, and Daniel Roos. *The Machine That Changed the World: The Story of Lean Production.* 1st Harper Perennial ed. New York: Harper Collins, Original Published by Rawson Associates, 1990.

APPENDIX A

USAF Fact Sheet[60]
C-5 Galaxy

MISSION

With its tremendous payload capability, the gigantic C-5 Galaxy, an outsized-cargo transport, provides the Air Mobility Command intertheater airlift in support of United States national defense.

FEATURES

The C-5 is one of the largest aircraft in the world. It can carry outsized cargo intercontinental ranges and can take off or land in relatively short distances. Ground crews can load and off load the C-5 simultaneously at the front and rear cargo openings since the nose and aft doors open the full width and height of the cargo compartment. It can also "kneel down" to facilitate loading directly from truck bed levels. Other features of the C-5 are:

- High flotation landing gear with 28 wheels sharing the weight.
- Full width drive-on ramps at each end for loading double rows of vehicles.
- An automatic trouble-shooting system that records and analyzes information and detects malfunctions in more than 800 test points.
- Able to take off fully loaded within 8,300 feet (2,530 meters) and land within 4,900 feet (1,493 meters).

The C-5 is similar in appearance to the smaller transport aircraft, the C-141 Starlifter, although the C-5 is much larger. Both aircraft have the distinctive high T-tail, 25-degree wing sweep and four turbofan engines mounted on pylons beneath the wings.

The Galaxy carries nearly all of the Army's combat equipment, including such heavy oversized items such as its 74-ton mobile scissors bridge, from the United States to any theater of combat on the globe.

Four TF-39 turbofan engines power the big C-5. They are pylon-mounted and rated at 41,000 pounds thrust each. They weigh 7,900 pounds (3,555 kilograms) each and have an air intake diameter of more than 8.5 feet (2.6 meters). Each engine pod is nearly 27 feet long (8.2 meters).

The Galaxy has 12 internal wing tanks with a total capacity of 51,150 gallons (194,370 liters) of fuel -- enough to fill more than six regular size railroad tank cars. The fuel load

---

[60] Taken from: *USAF Fact Sheet: C-5 Galaxy*, http://www.af.mil/news/factsheets/C_5_Galaxy.html

weighs 332,500 pounds (150,820 kilograms) allowing the C-5 with a load of 204,904 pounds (92,207 kilograms) to fly 2,150 nautical miles, offload, and fly to a second base 500 nautical miles away from the original destination -- all without aerial refueling. With aerial refueling, crew endurance is the only limit to the aircraft's range.

BACKGROUND

Lockheed-Georgia Co. delivered the first operational Galaxy to the 437th Airlift Wing, Charleston Air Force Base, SC, in June 1970. C-5s are stationed now at Altus AFB, OK; Dover AFB, DE; Travis AFB, CA, and Kelly AFB, TX. AMC transferred some C-5s to the Air Reserve starting with Kelly AFB, in 1985; followed by Stewart Air National Guard Base, NY, and Westover Air Reserve Base, MA.

In March 1989, the last of 50 C-5B aircraft were added to the 76 C-5A's in the Air Force's airlift force structure. The C-5B includes all C-5A improvements as well as more than 100 additional system modifications to improve reliability and maintainability. All 50 C-5B's are scheduled to remain in the active-duty force, shared by comparably sized and collocated Air Force Reserve associate units.

The C-5, C-17 Globemaster III and C-141 Starlifter are partners of AMC's strategic airlift concept. The aircraft carry fully equipped, combat-ready military units to any point in the world on short notice then provide field support required to help sustain the fighting force.

GENERAL CHARACTERISTICS

Primary function: Outsized cargo transport
Contractor: Lockheed Georgia Co.
Power Plant: Four General Electric TF-39 engines
Thrust: 41,000 pounds, each engine
Wingspan: 222.9 feet (67.93 meters)
Length: 247.1 feet (75.3 meters)
Height: 65.1 feet (19.84 meters) (at tail)
Cargo Compartment: height, 13.5 feet (4.11 meters); width, 19 feet (5.79 meters); length, 143 feet, 9 inches (43.8 meters)
Takeoff/Landing Distances: 8,300 feet (2,530 meters) takeoff fully loaded; and 4,900 feet (1,493 meters) land fully loaded
Pallet Positions: 36
Speed: 518 miles per hour (.68 Mach)
Range: 6,320 nautical miles (empty)
Crew: 7 (pilot, co-pilot, two flight engineers and three loadmasters)
Unit Cost: C-5B, $184.2 million (FY96 constant dollars)
Date Deployed: June 1970 (operational)
Inventory[61]**:** Active force, 81; ANG, 13; Reserve, 32.

---

[61] The numbers from the outdated Factsheet were updated with information from "World Military Aircraft Inventory," *Aviation Week & Space Technology*, January 11, 1999 1999, 261.

APPENDIX B

The purpose of this appendix is to explain how to run the model and to present the model code. A description of the model is given in Chapter 5, and a description of model results is given in chapter 6.

The model is coded with Matlab, which is needed in order to run the model. The main challenge to running the model will be to get it saved into an appropriate location so that Matlab can run it. The instructions for this will for the most part be left to the user's specific help manual, if such help is needed. In general, the file for this model, called SustainmentModel, must be placed in the Matlab directory or bin file. Once the model is saved in the appropriate location, the model is executed simply by typing the filename SustainmentModel into the Matlab window.

When the model begins to run, it will request inputs. For the most part, the requests are self explanatory, but we will step through them for the sake of clarity.

The first request is for the number of aircraft. Simply enter the number, for example 15, and hit the enter key.

The second request is for the number of parts that compose an aircraft. This is not the number of spares, but the number of parts that make up a "complete" aircraft in the model. In this thesis, 25 parts were used.

The third request is for the number of spares. This is the number of initial spares for each of the 25 parts. Only one number can be entered. In the thesis the number was between 0 and 16.

The fourth request is for the number of missions requested per day. Simply enter the number, such as 10.

The fifth request is for the probability of failure. Each part has this probability of failure on a given flight. The number must be entered in decimal form, such as 0.05 for 5 percent.

The sixth request is for the depot repair time, which is simply the number of days.

The seventh request is for the cannibalization policy. For no cannibalization enter 0. For one Cann-bird, enter 1. For unrestricted cannibalization, enter 2.

The eighth request is for the number of trials. This is the number of times you want the simulation to be run. Simply enter the number. The results from the multiple trials are averaged together. For this thesis, 10 trials were used for each run.

The ninth and final request for input is for the duration of each simulation run. This is the number of days in the model that you want to consider. In this thesis, a value of 360 was chosen to approximate one year. Thus if there were 10 trials, there would be 10 trials of 360 days each.

Once the inputs are entered, the model will begin running. It will output the number of the trial that it is running, so that the user can be aware of its progress. When the model is finished running, it will output the results to the screen. The user can then copy them in whatever manner he or she chooses. In running the model for this thesis, it was helpful to enter the results into a spreadsheet for post-processing. The model will output the following:

- The percentage of flights made and its standard deviation over the number of trials.
- The Mission Capable rate and its standard deviation.
- The average number of failures over the duration of simulation and its standard deviation.
- The average number of cannibalizations (not including the cannibalizations due to the Cann-recovery process), and its standard deviation over the number of trials.

In the case of unrestricted cannibalization the model will output:

- The average maximum cannibalization level. The maximum number of aircraft during each simulation that must be gone through to find an available part to cannibalize is averaged.
- The maximum cannibalization level. The maximum number of aircraft, over all the trials, that must be gone through to find an available part to cannibalize.

In the case of one Cann-bird the model will output:

- The average number of cannibalizations required for Cann-recovery, and the standard deviation over the number of trials.

The model is simple to execute, and is simple to edit as well. The file can be edited with almost any text processor. Users who would like to modify the code will only need to learn how to use Matlab.

The model code is included below. The % symbol denotes a comment line.

```
%SustainmentModel.m

%This model accompanies the 1999 MIT SM TPP thesis of Luis %Tsuji
%copyright August 6, 1999
%by Luis Tsuji

clear all;
%We begin by collecting all of the inputs.  The inputs of course can be
%hard wired into the code for ease of use. For example, the first few
%inputs could be coded as:
%NumPlane=15;
%NumParts=25;
%inputspares=12;
%It depends on the users preference whether to ask for input or to
%hard wire it.
NumPlane=input('Enter the number of aircraft:');
NumParts=input('Enter the number of parts that compose an aircraft');
inputspares=input('Enter the number of spares for each part');
MissionReq=input('Enter the number of missions requested per day');
probfail=input('Enter the probability of failure');
inputrepair=input('Enter the depot repair time');
CannPolicy=input('Enter 0 for No Cann, 1 for One Cann-bird, and 2 for
Unrestricted Cann');
numtrials=input('Enter the number of trials you wish to run for each
simulation');
timelimit=input('Enter the simulation duration in days');


%Next we begin initializing all our variables...
if CannPolicy==1         %if we use one cann-bird
   cannbird=NumPlane;    %we set the last aircraft to be the cannbird
end
Pfail=zeros(1,NumParts);
for part=1:NumParts
   Pfail(part)=probfail;
   TotalSpares(part)=inputspares;
   RepairTime(part)=inputrepair;
end
flightsmade=zeros(numtrials,1);
totalflightcount=0;
MaxCannLevel=zeros(numtrials,1);
CannSwitch=zeros(numtrials,1);
SumReadySwitch=zeros(numtrials,1);
CannActions=zeros(numtrials,1);
```

```
Failures=zeros(numtrials,1);

%Run a number of trials and average results
for trial=1:numtrials
   %more initialization
   trialnumber=trial    %this outputs the current trial
                        % number while the model is running.
   ready=ones(NumPlane,NumParts);
   PlaneReady=ones(NumPlane,1);
   totalflightcount=0;
   stock=TotalSpares;
   MissionCapable=zeros(1,timelimit);
   dailyflightcount=zeros(1,timelimit);
   OrderArrival=zeros(NumParts,timelimit+max(RepairTime));
   OrderAmt=zeros(NumParts,timelimit+max(RepairTime));
   withdrawal=zeros(NumParts,timelimit+max(RepairTime));
   time=0;
   cannactions=0;
   TroublePart=0;

   %for each day
   for time=0:timelimit-1
      time=time+1;
      %check for part shipment arrivals, clear orders
      for part=1:NumParts
         if OrderArrival(part,time)==1
            stock(part)=stock(part)+OrderAmt(part,time);
         end
      end
      planesdown=0;
      flying=zeros(NumPlane,1);

      %The CANN Stuff!
      if CannPolicy==1 %for one Cann-bird
         if rem(time,45)==0 %if it is time to switch cann-
                            %birds then count the switch-canns
            CannSwitch(trial)=CannSwitch(trial)+NumParts-
                              sum(ready(cannbird,1:NumParts));

            SumReadySwitch(trial)=SumReadySwitch(trial)+sum(ready(cannb
                                  ird,1:NumParts));
         end

         if sum(PlaneReady)<MissionReq %if the total planes
                                       %are less than needed for day.
            %Want to determine most damaged plane...
            for plane=1:NumPlane %for each plane
               NumGoodParts(plane)=sum(ready(plane,1:NumParts));%count
                                                            %good parts
               NumHoles(plane)=NumParts-NumGoodParts(plane);
                                                       %count bad %parts
            end
            [Y,FixOrder]=sort(NumHoles);     %order the aircraft by
                                             %broken parts
            [y,CannOrder]=sort(NumGoodParts); %order by good parts
```

```
      for plane=1:NumPlane  %for all of the planes
         if sum(PlaneReady)<MissionReq  %if don't have enough
                                       %ready planes for missions.
            %need to determine if plane can be fixed...
            %that way we don't do a lot of unnecessary cann.
            if PlaneReady(FixOrder(plane))==0  %if the plane is
                                               %broken somewhere
               readyTEMP=ready;                %we use dummy
                                               %variables to check
               PlaneReadyTEMP=PlaneReady;      % and see if the
                                              %plane can be fixed.
               cannactionsTEMP=cannactions;
               for part=1:NumParts
                  if readyTEMP(FixOrder(plane),part)==0
                     if readyTEMP(cannbird,part)==1
                        readyTEMP(FixOrder(plane),part)=1;
                        readyTEMP(cannbird,part)=0;
                        cannactionsTEMP=cannactionsTEMP+1;
                     end
                  end
               end                             %if we can fix it
                                                %then...do it.
               If sum(readyTEMP(FixOrder(plane),1:NumParts))=
                                                       =NumParts
                  PlaneReady(FixOrder(plane))=1;
                  ready=readyTEMP;
                  cannactions=cannactionsTEMP;
               end
            end
         end
      end
   end
end

if CannPolicy==2         %if we allow maximum cannibalization...
   if sum(PlaneReady)<MissionReq %if the total of ready planes is
                                 %less than needed.
      %want to determine most damaged plane...
      for plane=1:NumPlane
         NumGoodParts(plane)=sum(ready(plane,1:NumParts)); %count
                                             %good parts on plane.
         NumHoles(plane)=NumParts-NumGoodParts(plane); %count
                                                    %broken parts.
      end
      [Y,FixOrder]=sort(NumHoles);            %determine fix order.
      [y,CannOrder]=sort(NumGoodParts);       %determine cann
                                              %order.

      for plane=1:NumPlane                    %for all planes
         if sum(PlaneReady)<MissionReq
            %need to determine if plane can be fixed...
            %that way we don't do a lot of unnecessary cann.
            if PlaneReady(FixOrder(plane))==0
               readyTEMP=ready;
```

```
               PlaneReadyTEMP=PlaneReady;
               cannactionsTEMP=cannactions;
               MaxCannLevelTEMP=MaxCannLevel;
               for part=1:NumParts
                  if readyTEMP(FixOrder(plane),part)==0
                     for cannbird=1:NumPlane
                        %check to see if cann can fix plane...
                       if PlaneReadyTEMP(CannOrder(cannbird))==0
                        if readyTEMP(CannOrder(cannbird),part)==1

                           readyTEMP(FixOrder(plane),part)=1;
                           readyTEMP(CannOrder(cannbird),part)=0;
                          cannactionsTEMP=cannactionsTEMP+1;

                          if cannbird>MaxCannLevelTEMP(trial)
                             MaxCannLevelTEMP(trial)=cannbird;
                             TroublePartTEMP(trial)=part;
                          end
                          break %if make cann, stop looking for a
                                   %cannbird
                         end
                       end
                     end
                  end
              end                      %if cann works, do it...
          if sum(readyTEMP(FixOrder(plane),1:NumParts))==NumParts
                  PlaneReady(FixOrder(plane))=1;
                  ready=readyTEMP;
                  cannactions=cannactionsTEMP;
                  MaxCannLevel(trial)=MaxCannLevelTEMP(trial);
                  TroublePart(trial)=TroublePartTEMP(trial);
          end
         end
        end
      end
   end
end

if CannPolicy==1
   PlaneReady(cannbird)=0;   %it is assumed,in the case of one
                             %Cann-bird,
end                          %the Cann-bird does not fly

for plane=1:NumPlane
   if PlaneReady(plane)==1 %if plane is ready to fly
      MissionCapable(time)=MissionCapable(time)+1; %count MC
                                                      %time.
      if dailyflightcount(time)<MissionReq    %if plane is needed
                                      %to meet flight requirement.
         flying(plane)=1; %if plane ready and needed, then it
                            %flies.
         dailyflightcount(time)=dailyflightcount(time)+1; %count
                                            %the number of flights
         for part=1:NumParts %each part
            if rand<Pfail(part)       %see if the part fails...
```

```
                      ready(plane,part)=0; %part broken, plane not ready
                     OrderArrival(part,time+RepairTime(part))=1;%submit
                                 %part for repair,determine return time
                     OrderAmt(part,time+RepairTime(part))=
                                 OrderAmt(part,time+RepairTime(part))+1;
                                 %count number of parts
                    if stock(part)>0  %if part is in stock
                        stock(part)=stock(part)-1;%take part from stock
                        ready(plane,part)=1; %replace part.
                     else
                        PlaneReady(plane)=0; %plane is not ready
                     end
                  end; %if rand
               end %for part
            end; %if plane needed
         else %if plane not ready, it doesn't fly, but gets repaired
               when stock arrives.
            PlaneReady(plane)=1; %by default, assume plane can be fixed
            for part=1:NumParts %check all parts
               if ready(plane,part)==0 %if part broken
                  if stock(part)>0  %if part is in stock
                     stock(part)=stock(part)-1; %take part from stock
                     ready(plane,part)=1; %place part on plane
                  else
                     PlaneReady(plane)=0; %plane is not ready
                  end %if stock
               end %if ready=0
            end %for part
         end %if PlaneReady=0
      end %for plane
   end %time
   flightsmade(trial)=sum(dailyflightcount); %compute the number of
                                              %flights made
   MCrate(trial)=mean(MissionCapable)/NumPlane;
   CannActions(trial)=cannactions;
   Failures(trial)=sum(sum(OrderAmt));
end %trial
avgflightsmade=mean(flightsmade);
PercentFlightsMade=avgflightsmade/(MissionReq*timelimit)*100
StandardDeviationFlights=std(flightsmade)/(MissionReq*timelimit)*100
MissionCapableRate=mean(MCrate)*100
StandardDeviationMC=std(MCrate)*100
AvgFailures=mean(Failures)
stdFailures=std(Failures)
AvgCanns=mean(CannActions)
stdCanns=std(CannActions)
if CannPolicy==2
   AvgMaxCannLevel=mean(MaxCannLevel)
   MaxMaxCannLevel=max(MaxCannLevel)
end
if CannPolicy==1
   AvgCannSwitch=mean(CannSwitch)
   stdCannSwitch=std(CannSwitch)
end
```

THESIS PROCESSING SLIP

FIXED FIELD: ill. ______ name ______

index ______ biblio ______

COPIES: [✓ Archives] Aero [✓ Dewey] Eng Hum

Lindgren Music Rotch Science

TITLE VARIES: ☐ ______

NAME VARIES: ☑ Christophe

IMPRINT: (COPYRIGHT) ______

COLLATION: 125 P

ADD: DEGREE: ______ DEPT.: ______

SUPERVISORS: ______

NOTES:

cat'r: date:

page: 36

DEPT: T&P

YEAR: 1997 DEGREE: S.M.M.O.T.

NAME: TSUJI, Luc C.