



PROJECT AIR FORCE

THE ARTS

CHILD POLICY

CIVIL JUSTICE

EDUCATION

ENERGY AND ENVIRONMENT

HEALTH AND HEALTH CARE

INTERNATIONAL AFFAIRS

NATIONAL SECURITY

POPULATION AND AGING

PUBLIC SAFETY

SCIENCE AND TECHNOLOGY

SUBSTANCE ABUSE

TERRORISM AND
HOMELAND SECURITY

TRANSPORTATION AND
INFRASTRUCTURE

WORKFORCE AND WORKPLACE

This PDF document was made available from www.rand.org as a public service of the RAND Corporation.

[Jump down to document ▼](#)

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world.

Support RAND

[Purchase this document](#)

[Browse Books & Publications](#)

[Make a charitable contribution](#)

For More Information

Visit RAND at www.rand.org

Explore [RAND Project AIR FORCE](#)

View [document details](#)

Limited Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law as indicated in a notice appearing later in this work. This electronic representation of RAND intellectual property is provided for non-commercial use only. Permission is required from RAND to reproduce, or reuse in another form, any of our research documents.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE The Closed-Loop Planning System for Weapon System Readiness				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) RAND Corporation,1776 Main Street,PO Box 2138,Santa Monica,CA,90407-2138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 124	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This product is part of the RAND Corporation monograph series. RAND monographs present major research findings that address the challenges facing the public and private sectors. All RAND monographs undergo rigorous peer review to ensure high standards for research quality and objectivity.

The Closed-Loop Planning System for Weapon System Readiness

Richard Hillestad, Robert Kerchner, Louis W. Miller,
Adam Resnick, Hyman L. Shulman

Prepared for the United States Air Force

Approved for public release; distribution unlimited



PROJECT AIR FORCE

The research described in this report was sponsored by the United States Air Force under Contract F49642-01-0003. Further information may be obtained from the Strategic Planning Division, Directorate of Plans, Hq USAF.

Library of Congress Cataloging-in-Publication Data

The Closed-Loop Planning System for weapon system readiness / Richard Hillestad
... [et al.].

p. cm.

"MG-434."

"This report documents several ways in which the Air Force can improve its management of logistics."

ISBN 0-8330-3864-8 (pbk. : alk. paper)

1. United States. Air Force—Operational readiness. 2. Airplanes, Military—United States—Maintenance and repair. 3. United States. Air Force—Supplies and stores. I. Hillestad, R. J. (Richard John), 1942—

UG1103.C56 2006

358.4'18—dc22

2005031298

The RAND Corporation is a nonprofit research organization providing objective analysis and effective solutions that address the challenges facing the public and private sectors around the world. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

RAND® is a registered trademark.

© Copyright 2006 RAND Corporation

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from RAND.

Published 2006 by the RAND Corporation

1776 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, VA 22202-5050

4570 Fifth Avenue, Suite 600, Pittsburgh, PA 15213

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact

Distribution Services: Telephone: (310) 451-7002;

Fax: (310) 451-6915; Email: order@rand.org

Preface

The U.S. Air Force Spares Campaign and other analyses identified the need for a new, consistent approach for managing spares and repair requirements. The objective is to provide resources to achieve the levels of fully mission capable sorties and weapon system availability required in the Air Expeditionary Force (AEF) environment for both regional contingencies and major theater wars. The new approach is needed because the Air Force does not have a procedure to allocate limited funding for depot repair across weapon systems in a way that links the trade-offs to the projected readiness of each weapon system. In fact, no current system computes the readiness related to various levels of depot repair funding.

The current planning system for depot-level repairs develops a requirement based on historical repair demands. The funding of the requirement depends on other priorities within the Air Force budget process, and frequently the requirement is not fully funded. With the current planning systems and process, the implications of this shortfall cannot be easily estimated in terms of reduced sortie capability or weapon availability. This is, in effect, an “open-loop” system that cannot easily show decisionmakers the consequences of their repair budgeting decisions. A “closed-loop” system would allow decisionmakers to choose budget levels and feed back the readiness implications. Alternatively, it would allow decisionmakers to iterate among levels of readiness and required budgets to see what the budget consequences of desired levels of readiness would be.

During fiscal years 2003 and 2004, RAND completed and demonstrated a prototype decision support tool for allocating repair funds across weapon systems to meet sortie and availability goals given operational plans, logistics support infrastructure, and funding constraints. The prototype also supports the need to compare execution to the plan by providing a repair replanning mode.

This monograph describes the concept and design of the prototype, which is called the Closed-Loop Planning System, for application to depot-level repair of repairable components. The planning takes into account budgets, shop capacities, and operational goals. It permits decisionmakers to examine trade-offs of funding across weapon systems. The monograph is primarily a description and demonstration of a new approach to repair planning. As such it will be of interest to those currently involved in executing and improving that process.

This monograph is related to a variety of issues that have been the subject of RAND research over a considerable span of time. These include the Multi-Echelon Technique for Recoverable Item Control (METRIC) stockage model in the 1960s, which is the basis of the Air Force's readiness based leveling system, the pioneering work in the 1980s of Gordon Crawford that demonstrated the high level of variability in demands for spare parts, and the development of the Distribution and Repair in Variable Environments (DRIVE) model that was the precursor of the Execution and Prioritization of the Repair Support System (EXPRESS) in the 1990s. The following are two recent RAND reports of related interest:

- *How Should the U.S. Air Force Depot Maintenance Activity Group Be Funded? Insights from Expenditure and Flying Hour Data*, by Edward Keating and Frank Camm (MR-1487-AF, 2002). This study investigated the linkage between operating command activities and actual Air Force Materiel Command costs.
- *Effective Treatment of Logistics Resource Issues in the Air Force Planning, Programming, and Budgeting System (PPBS) Process*, by Frank Camm and Leslie Lewis (MR-1611-AF, 2003). This re-

port documents several ways in which the Air Force can improve its management of logistics resource issues in its PPBS process.

The research reported here was done within the project titled “Implementation of a Closed Loop Budgeting and Planning Process for Depot-Level Repair” and was sponsored by AF/A4I and conducted in the Resource Management Program of Project AIR FORCE. The project is responsive to recommendations resulting from the U.S. Air Force Spares Campaign.

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of the RAND Corporation, is the U.S. Air Force’s federally funded research and development center for studies and analyses. PAF provides the Air Force with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future aerospace forces. Research is performed in four programs: Aerospace Force Development; Manpower, Personnel, and Training; Resource Management; and Strategy and Doctrine.

Additional information about PAF is available on our Web site at <http://www.rand.org/paf>.

Contents

Preface iii

Figures xi

Tables xiii

Summary..... xv

Acknowledgments xxi

Abbreviations xxiii

CHAPTER ONE

Introduction 1

Context 1

The Goals of a Planning System for Depot-Level Repair..... 4

The “Closed-Loop” Depot-Level Repair Planning System 5

Organization of the Report 6

CHAPTER TWO

Current Processes for Planning Depot Repairs and the Need for
the Closed-Loop Methodology 7

The Sources of Demand for Depot Repair of Repairable Components 7

Stages of Planning and Execution 8

MAJCOM and AFMC Long-Term Planning 9

 A Closer Look at SIRS 11

 SIRS Does Not Consider Capacity Constraints on
 Repair Resources 11

 SIRS Does Not Recognize the Variance of Repair Demands for
 NRTS Parts 12

SIRS Does Not Provide a Relationship Between Budget Allocations for Repair and Readiness	12
Allocation of Obligation Authority to Sources of Repair	12
ALC Planning	13
Daily Execution	14

CHAPTER THREE

Formulation of the Depot Repair Planning Problem	17
The Overall Depot Repair Planning Problem	17
Why Optimize?	17
The Objective—Aircraft Availability	18
Measures of Aircraft Availability	18
Setting the Confidence Level	24
Demands for Depot Repair	24
The Decision Variables—the Components to Repair	25
Planning Period and Point of Evaluation	27
What Needs to Be Repaired?	27
Constraints	29
Budget Constraints	30
Carcass Constraints	31
Repair Capacity Constraints	34

CHAPTER FOUR

Stepping Through the Solution to the Optimization Problem	37
--	-----------

CHAPTER FIVE

The Prototype Database and Model for Depot Repair Planning	47
Database for the Prototype	47
Squadrons and a Scenario	48
Avionics LRUs and Test Stands	50
Asset Data	52
The Prototype Planning Model	53
Comparisons with Current Planning Models	64
Capability Differences	64
NSN Versus Work Center View	66
Cost Comparison with Constant Performance	67

CHAPTER SIX

Extensions and Next Steps	69
Some Possible Extensions to the Planning Model and Its Use of	
Data to Increase Its General Applicability	69
Planning for Future Years and Planning During the Execution Year	70
Optimizing Buy and Repair Decisions Together	70
Planning Depot-Level Repairs to Support a CIRF	72
Use of the Closed-Loop Planning System to Evaluate Spending Plans....	72
Replanning with the Closed-Loop Planning System	74
Next Steps.....	75
 Appendix	
A. Mathematical Description of the Planning Model	77
B. The Prototype Database	89
 Bibliography	95

Figures

S.1. Minimum Budget Determined by the Prototype to Achieve the Stated NMCS Goals When There Are No Repair Constraints	xvii
S.2. Adjustment of NMCS Goal to Fit Shop Capacity Constraints	xviii
S.3. Trade-Offs Between Weapon System NMCS at Constant Budget Levels	xix
1.1. The U.S. Air Force Working Capital Fund.....	2
3.1. NMCS Goals and Alternative Aircraft Groupings.....	21
3.2. Target NMCS Rates Broken Down by MDS, Theater, and Deployment Stage in a Major Theater Conflict	22
3.3. The “Pareto” Curve Showing How a Few High-Demand, High-Cost Components Account for Most of the Repair Load	26
3.4. Sources of Repair Requirements	28
3.5. The Catch-Up Repairs Needed to Achieve a Given Readiness Level	28
3.6. Total Repair Demands, Including Uncertainty Elements	29
3.7. The Process of Adjusting NMCS Goals to Determine a Budget for Repair	30
3.8. Budget Versus NMCS Trade-Off Curve Derived from the Closed-Loop Planning Model and Prototype Database	31
3.9. The Interaction Between Carcass Constraints and Achievable NMCS Goals	33

- 3.10. The Effect of Repair Capacity Constraints on the Cost to Achieve NMCS Goals 34
 - 4.1. Steps in the Optimization Process 37
 - 4.2. Setting the NMCS Goals 38
 - 4.3. Checking the Carcass Constraint 39
 - 4.4. Computing the Expected NSN Repair Requirement..... 41
 - 4.5. Checking the Work Center Capacity Constraints 42
 - 4.6. Computing the Joint Probability Distribution of Work Center Costs and Hours 43
 - 4.7. Optimizing the Allocation of Repairs to Work Centers by Marginal Analysis 44
 - 4.8. Observing the Cost of Depot-Level Repairs 45
 - 4.9. Repair and Cost Details for Depot-Level Planning 46
- 5.1. The Aircraft Components Used in the Prototype Database and Their Test Stand Assignment 52
- 5.2. Minimum Budget Determined by the Prototype to Achieve the Stated NMCS Goal When There Are No Repair Constraints 54
- 5.3. Budget Required as a Function of NMCS Goals When Those Goals Are Varied and Repair Is Unconstrained 54
- 5.4. Identification of Shop Capacity Constraints..... 58
- 5.5. Adjustment of the NMCS Goal to Fit Shop Capacity Constraints 59
- 5.6. Examination of the Cost of Repair for Different NMCS Goals with Different Levels of Overtime 60
- 5.7. Identification of Carcasses Needed to Achieve NMCS Goals 61
- 5.8. Trade-offs Between Weapon System NMCS Goals at Constant Budget Levels 62
- 5.9. Effects of Increased Flying Hours on NMCS and Budget..... 63
- 5.10. Detailed Repair Requirements by Shop, Aircraft Type, and Component 64
- 5.11. Additional Detailed Information for ALC and Shop Planning... 65
- 5.12. Cost Comparison Between Current Planning System and Closed-Loop Optimized Planning for Some NMCS Performance..... 68

Tables

5.1. Squadrons and Deployment Scenario 49

5.2. Totals of Aircraft Deployment Assignments..... 50

5.3. The Aircraft Components Used in the Prototype Database
and Their Test Stand Assignment 51

5.4. Marginal Cost Per Percentage of NMCS 56

5.5. Marginal Cost Per NMCS Aircraft 57

Summary

The U.S. Air Force does not have an effective way of allocating limited funding for depot-level repair across weapon systems and calculating the readiness implications of such allocations (p. 3). The RAND project discussed in this report directly addressed this problem by developing a methodology that estimates the effect of depot repair funding allocations on aircraft availability. We have called this the “Closed-Loop” Planning System because it provides this type of feedback, as opposed to the open-loop nature of the current planning system, which does not (p. 5).

The report describes the shortcomings in the current system to rationalize the development of the closed-loop methodology (pp. 9–15). It also illustrates the application of a prototype of the new planning system using a subset of real data from Air Force depot-level repair (Chapter Five). It compares the cost of achieving the same level of readiness with the current Air Force approach and the closed-loop methodology (pp. 64–68). In addition, it describes extensions of the methodology that would be useful for both long-term and short-term planning (p. 69).

Air Force depot-level component repair includes repairs of components removed at bases worldwide during flying operations and that cannot be repaired at either the base level or intermediate level. It also includes repairs of components needed to support programmed depot maintenance (PDM) and repairs contracted for with foreign militaries. This component repair operation at the depot level absorbed about \$3.1 billion in fiscal year 2003 (p. 2).

A review of the planning and budgeting processes for the depot level during the U.S. Air Force Spares Campaign in 2001 identified important disconnects in those processes. In particular, it was found that decisionmaking about depot repair budgets was done in the absence of information about how those decisions affected operational readiness. It also found that there were important inconsistencies between organizations (e.g., Major Commands and the Air Force Materiel Command) in assumptions and in their planning processes. Repair capacity was not considered in either process. Additionally, the U.S. Air Force Spares Campaign found that planning and execution were disconnected—the repair scheduling did not have adequate information about the planning goals, and it was unclear how to track the correspondence between the planning goals and execution of repairs (pp. 3–4).


The goal of our research was to define and demonstrate a methodology that could overcome several of these problems. In particular it connects the budgetary planning with its impact on operational readiness in terms of the influence on aircraft availability for missions. The Closed-Loop Planning System was thus developed; the system also produces a plan that considers repair capacity constraints, carcass constraints (that is, having something to repair), and uncertainty in the demands for repair. As far as we know, no current planning system includes these features (pp. 4–6).

The Closed-Loop Planning System solves the depot-level repair planning problem by starting with a statement of readiness goals in terms of end-of-planning-period aircraft availability. The availability is defined as the fraction of aircraft that are mission capable. In the case of the planning system, the important rate is the converse of this rate: aircraft that are not mission capable, supply (NMCS)—not mission capable because of supply. Readiness goals can be set by unit, theater, aircraft type, and command. The methodology then optimizes the mix of repairs to be planned for each shop to ensure, with high confidence, that the readiness goals can be achieved. Mathematically this is the same as ensuring that, with high confidence, the supply system can provide parts to the units to meet their readiness goals and that the depot-level repair shops can each provide the re-

paired parts to the supply system. The methodology identifies and optimizes within constraints to meet the availability goals in terms of shop capacity limits or carcass limits. Additional capacity in the form of overtime can be included in the optimization if necessary. The primary output is the budget necessary to achieve the readiness goals. In case the budget is not satisfactory, the user of the methodology can iterate the process by selecting units or aircraft to adjust readiness goals and then view the budget implications. It is also possible to program priorities for adjusting readiness goals to achieve a given budget level. Figure S.1 shows the basic decision-level inputs and outputs of the methodology. Conflict 1, Conflict 2, and CONUS (continental United States) represent a scenario for prioritizing readiness goals. The aircraft units to be deployed first are given the tightest goal, the units to deploy next a reduced goal, and the CONUS aircraft the most relaxed goal (Chapter Three).


Figure S.2 shows additional, marginal cost information that is available in terms of the additional depot repair costs necessary to make one more aircraft available, by aircraft type. Additional displays, illustrated in Chapter Five, show how the model can point out where capacity is exceeded, overtime to exceed that capacity, and confidence

Figure S.1
Minimum Budget Determined by the Prototype to Achieve the Stated NMCS Goals When There Are No Repair Constraints



NOTES: NMCS goal = full cannibalization of parts at time of deployment and full readiness spares package, with 90 percent confidence.

Figure S.2
Adjustment of NMCS Goal to Fit Shop Capacity Constraints




estimates that a shop will meet its repair goals. The model can also support detailed Air Logistics Center, shop, and workstation planning by producing detailed repair estimates for those entities (Chapter Five).

Perhaps the most interesting use of the Closed-Loop Planning System is when it is necessary to make trade-offs in readiness across aircraft types to achieve budgetary goals. Figure S.3 shows such a trade-off curve between the F-15 C/Ds and the F-15 Es when the overall budget is kept constant (p. 62).

In the text, we also show a comparison of an approximate representation of a plan from the current Air Force Materiel Command planning system and the Closed-Loop Planning System. When we

Figure S.3
Trade-Offs Between Weapon System NMCS at Constant Budget Levels



NOTES: NMCS percentages include budget trade-offs between the two Mission Design Series, F-15 C/Ds and F-15Es, in Conflict 2 and CONUS. The entire budget equals \$124.32 million.

hold readiness constant, the budget required for the closed-loop plan is about 6 percent less.

In the last chapter of the report, we suggest a number of extensions that could be made to the Closed-Loop Planning System (pp. 69–74). It could be used for execution planning, in which it is necessary to start the model with the current state of the supply and repair system. In a similar vein, it could be used to track performance and replan during the execution year by tying it to the execution process and data. It could also be used to plan how and when to overcome repair capacity limitations through the purchase of additional equipment, manpower, or overtime.

Most important, the Closed-Loop Planning System is meant to help Air Force planners make decisions about budgets for depot-level repair with a true understanding of the readiness consequences of those decisions (p. 75). At a minimum, it should be integrated into the Spares Requirements Review Board process to help resolve depot-level repair budgeting issues. A broader goal would be to embed the closed-loop methodology directly within the Air Force Materiel Command planning process.

Acknowledgments

Lieutenant General Michael Zettler set this project in motion by initiating the U.S. Air Force Spares Campaign, and Brigadier General Robert Mansfield managed that campaign, leading to the recommendation (one of several) to perform this research. We acknowledge Mr. Grover Dunn (AF/ILI) for his support and oversight during the final year of this research. We thank Curt Neuman, AFMC/XP/SAO, and his staff for providing data and insight into the current D200 process as well as for reviewing our work. We also thank Maurice Carter, OO-ALC/IT (retired), Ogden AFB, for his support and insight into depot-level repair planning. We thank our reviewers, Jim Masters and Edward Chan, for their thoughtful comments and Sharon Drummond for manuscript preparation. Our posthumous gratitude goes to our coauthor Hy Shulman on this, his final publication. Hy's vision of providing the Air Force with a better way to plan for weapon system support at the depot based on improved readiness was the inspiration for this project.

Abbreviations

AAM	Aircraft Availability Model
ACC	Air Combat Command
AEF	Air Expeditionary Force
AETC	Air Education and Training Command
AFCAIG	Air Force Cost Analysis Improvement Group
AFMC	Air Force Materiel Command
ALC	Air Logistics Center
ANG	Air National Guard
APS	advanced planning and scheduling
ASM	Aircraft Sustainability Model
CIRF	centralized intermediate repair facility
COTS	commercial-off-the-shelf
DLR	depot-level repairable part
DMAG	Depot Maintenance Activity Group
DoD	Department of Defense
DRC	Dynamics Research Corporation
DRIVE	Distribution and Repair in Variable Environments
ECSS	Expeditionary Combat Support System
eLog21	Expeditionary Logistics for the 21st Century
EPM	EXPRESS Planning Module
ERP	Enterprise Resource Planning

EXPRESS	Execution and Prioritization of Repair Support System
FMS	foreign military sales
FTRSQ	fighter squadron
FW	fighter wing
FY	fiscal year
LRUs	line replaceable units
MAJCOM	Major Command
MDS	Mission Design Series, a type of aircraft
METRIC	Multi-Echelon Technique for Recoverable Item Control
MISTR	Management of Items Subject to Repair
MTW	major theater war
NMCM	not mission capable, maintenance
NMCS	not mission capable, supply
NRTS	not repairable at this station
NSN	national stock number
OA	obligation authority
OIMDR	organizational and intermediate demand rate
OWRM	other war reserve material
PACAF	Pacific Air Force
PDM	programmed depot maintenance
PMAI	Primary Mission Aircraft Inventory
QPA	quantity per application
RBL	readiness based leveling
RCC	repair cost center
RMS	Requirements Management System
RSP	readiness spares package
SIRS	Secondary Item Requirements Computation System
SMAG	Supply Management Activity Group

SPAWS	Single Priority Across Weapon Systems
SOR	source of repair
SOS	source of supply
SRRB	Spares Requirements Review Board
SRU	shop replaceable unit
TRSQD	training squadron
USAFE	U.S. Air Force Europe
VSL	variable safety level
WG	wing
WSMIS	Weapon System Management Information System

Introduction


Context

This report is about a methodology for more effectively planning and budgeting depot-level repairs of weapon system components. For those readers not familiar with it, depot-level repair (as opposed to base-level repair and intermediate-level repair) is the process of repairing those weapon system components that are shipped from the base or intermediate level to the depot because they cannot be repaired at those other locations. They are commonly referred to as the components that are not repairable at this station (NRTS) from bases and other repair facilities. Some parts cannot be fixed and are “condemned,” usually along with placing an order to a contractor for a new component. Those that can be repaired are sent to depot shops specializing in classes of repair, such as repair of avionics components. A given shop at a given depot may repair all of the NRTS components of a given class for a particular aircraft type worldwide. When there are too many repairs demanded of a shop, that shop can increase its capacity through overtime, but at some point the repairs become capacity limited and a backlog will grow. The depot repair budget is ultimately allocated to shops to pay for the resources to make the repairs. Once a component is repaired, it is put into the supply system, and that system allocates the “serviceable” component back to a base. The “planning” stage defines and allocates budgets to depots and ultimately to shops based on expected repairs. The “execution” stage involves the decisions about which specific components to repair given an existing demand and capacity limits of shops. This

may involve some reallocation of funds because of unexpected repair volume. The execution year is the current year and the planning years are the future years.

U.S. Air Force depot-level repair constitutes a sizeable business. Figure 1.1 shows the fiscal year (FY) 2003 U.S. Air Force working capital fund budget and the portion of this budget allocated to repair activity. The Depot Maintenance Activity Group (DMAG) is responsible for accomplishing the repairs, including contracted work. The Supply Management Activity Group (SMAG) acts as the intermediary between customers and the DMAG. It supplies parts to both the DMAG and customers (e.g., Major Commands [MAJCOMs]) and purchases repairs from the DMAG. The total DMAG planned sales were \$6.5 billion. Of that, the sales to the SMAG were planned to be \$3.1 billion. By way of contrast, the planned SMAG purchases of parts were planned to amount to one-third of the amount for the repair activity. In addition to sales to the SMAG, a major portion of DMAG activity is programmed depot maintenance (PDM), meaning overhauls of major weapon systems. Other, lesser sales are to other

Figure 1.1
The U.S. Air Force Working Capital Fund



SOURCE: U.S. Air Force, 2003b.

RAND MG434-1.1

services, to other Department of Defense (DoD) activities, and to foreign military sales (FMS). This report is about the non-PDM portion of the DMAG, the depot repair of repairable components of aircraft, missiles, and engines.

In 2001, Brig Gen Robert Mansfield initiated and supervised the U.S. Air Force Spares Campaign in which he directed five teams to identify “disconnects” in their respective topic areas and recommend corresponding remedies. From the vast number of issues raised, three generalizations relating to repair planning for depot-level repairable parts (DLRs) are of particular interest:¹

1. Decisionmaking is dominated by financial concerns and is often focused on commodities rather than weapon systems. The allocation of repair resources is generally done without a clear understanding of the impacts on weapon systems.
2. Beginning with MAJCOM budgeting and on through Air Force Materiel Command (AFMC) and Air Logistics Center (ALC) planning to the execution of repairs, there are many points of inconsistency across the processes, including differing views of what is worth spending resources on, varying goals and objectives, and alternative approaches to pricing repairs. Repair capacity is not a factor in planning above the ALC level.
3. Planning and execution processes are disconnected, with at best only weak linkages among them. There is a general lack of monitoring and feedback and no attempts to understand how the decisions made at one level are consistent with those at other levels or to understand when processes are deviating significantly from the higher-level plans.

¹ Original statements by RAND of the need for “closed-loop” repair planning processes predated the Spares Campaign by at least four years. But documents from the Spares Campaign provide a useful compendium of problems spanning a wide range of Air Force logistics activities. Each of the five teams produced a draft working paper. Please see United States Air Force, 2001a, b, c, d, e.

Ideally, each agency with a stake in depot-level repairs should be acting in ways that are consistent with a single overall plan that considers operational goals and logistics constraints, and the plan should be developed with an understanding of the sensitivities of matching operational goals with allocation of funds.

The Goals of a Planning System for Depot-Level Repair

To correct the problems noted above, a depot repair planning system should:

1. Link repair plans to the operational capability they are expected to provide. Operational capability goals should include those for “boiling peace,”² major conflict contingencies, and peacetime training. Generally, the repairs required to support an ongoing contingency or known deployments are not included in the planning process.
2. Explicitly take into account logistics constraints such as shop capacity. Weapon systems need a *balanced* supply of spares for their components. If constraints are ignored, not only are funds allocated that cannot be spent, but funds spent to repair other system components may not improve the weapon system’s availability.
3. Plan for the uncertainties in demand for repairs as well as the expected values. Although the supply system plans for pipeline uncertainties, this planning does not provide “protection” for the uncertainty of demands on the depot repair shops.
4. Provide decisionmakers with information on the trade-offs between weapon system performance and the budget allocation to weapon systems by relating expenditures for repair to operational factors. This information provides a basis for assessing the operational consequences when cuts and other changes to the repair budget are considered.

² “Boiling peace” refers to the turbulent condition of deployment and redeployment of forces around the world as tensions have arisen since the end of the Cold War.

5. Where applicable, provide the capability to consider a holistic budget strategy that includes investments in future capacity, e.g., additional test equipment along with funding near-term repair budgets.
6. For the year of execution, determine the best allocation of funds to the sources of repair as well as estimates of the repairs they will be making.
7. Provide a capability to monitor performance of the repair system and the health of the weapon system as compared with expectations of the plan and thereby relate repair execution to planning.

The “Closed-Loop” Depot-Level Repair Planning System

The current planning system for depot-level repairs develops a requirement based on historical repair demands. The funding of the requirement depends on other priorities within the Air Force budget process, and frequently the requirement is not fully funded. As noted in the previous section, the current planning systems and process do not supply the implications of this shortfall in terms of reduced sortie capability or weapon availability. Thus, this “open-loop” system cannot easily show decisionmakers the consequences of their repair budgeting decisions. A “closed-loop” system would allow decisionmakers to choose budget levels and would provide feedback on readiness implications. Or, it would allow decisionmakers to iterate between levels of readiness and required budgets to see what the budget consequences of desired levels of readiness would be. It would also close the loop between execution and planning by allowing evaluation of the current state of execution against the objectives of the plan. This report describes and demonstrates an approach that we believe satisfies these requirements. We call this approach the *Closed-Loop Planning System* for depot-level repairs.

There are two requirements for DLRs. The first requirement is for parts that fail during service and must be replaced. The second is parts needed to achieve specified readiness positions (“holes” in weapon systems and establishing war reserve stocks) by a future time.

We call these requirements, respectively, “keep-up” and “catch-up.” In general, unless there is a surplus of stock, or it is desired to reduce operational aircraft availability, the keep-up requirements must be satisfied, and these demands are random variables with associated uncertainties. Catch-up requirements are policy-driven variables, not subject to randomness, because decisionmakers can choose goals for holes in aircraft and war reserve parts. These, along with the initial state of the system (also deterministic), then define the catch-up requirement. Because the number of failures in operational flying (the keep-up quantity) is random, there is a risk that catch-up goals will not be met. Subject to financial and capacity constraints, the methodology we describe below allocates repair resources in an attempt to meet the policy-driven catch-up goals with high confidence. Through a process of iteration between the user and the model, the weapon system performance implications of various budget levels or “cuts” can be determined.

Organization of the Report

Chapter Two presents a brief critique and description of current repair planning and execution processes and suggests how they might benefit from the closed-loop methodology. Chapter Three describes the formulation of the planning problem that is the basis of our methodology, and Chapter Four describes the solution to the optimization problem. Chapter Five provides illustrations of the prototype database and model for depot repair planning and makes some numerical comparisons with the current process. Chapter Six suggests directions for building on this research and implementing an internally consistent planning process. Appendix A provides the mathematical details of the methodology, and Appendix B describes details of the prototype database used for testing.

Current Processes for Planning Depot Repairs and the Need for the Closed-Loop Methodology

This chapter reviews aspects of the Air Force's current processes for planning and carrying out the depot-level repairs of repairable parts. In the context of describing current processes, we also lay a foundation for the Closed-Loop Planning System methodology that is described in the following chapters.

The Sources of Demand for Depot Repair of Repairable Components

Demands for parts arise for several reasons:

- Demands from bases are generated when parts are removed from weapon systems in the course of daily operations and sent to the depot for repair and replacement.
- Parts are in pipelines and not available for needed installation in weapon systems. Pipelines, by definition, include unserviceable parts in transit to a depot and serviceable parts on the way to bases, parts in repair or waiting to be repaired at bases, and parts at a depot or contractor that are undergoing repair or in unserviceable condition and available for induction to a repair shop. These pipelines absorb a significant number of the Air Force's inventory of parts.
- Parts are needed for safety stock to buffer bases from the uncertainties of repair demands.

- Component repair demands are generated from the overhauling of aircraft and engines at a depot.
- Parts damaged beyond repair are condemned, possibly establishing requirements to acquire new replacements.
- Parts are needed to stock war reserves, both readiness spares packages (RSPs) and war reserve items held centrally (other war reserve material, OWRM).
- FMS users and other services require repair support.

Offsetting these demands are serviceable parts owned by the Air Force or expected to be acquired during the planning period of concern, future purchases of additional parts for future planning cycles (that would likely have delivery dates several years out), and depot repairs. If sufficient total numbers of parts, serviceable and unserviceable, are in the Air Force inventory, it is generally less expensive and quicker to meet needs by repairing unserviceable parts than acquiring new ones. Hence, the above-listed demands for parts generally lead to repair demands.

Constraints on meeting the repair needs are the total number of repairable components in the inventory (carcasses), repair capacity, and funds to pay for the repairs.

Stages of Planning and Execution

With respect to depot-level repair, the Air Force performs the following five hierarchical stages of planning and execution, starting with the longest planning horizon (38 quarters) down to daily execution, the shortest:

1. AFMC estimates “requirements” for buying new parts and repairing repairable weapon system components.
2. MAJCOMs plan and program funding to purchase serviceable spares from AFMC.
3. AFMC allocates funds for supporting weapon systems to sources of supply (SOSs) that in turn allocate obligation authority (OA)

to sources of repair (SORs).¹ Obligation authority is the authority to expend resources to perform repairs prior to “selling” the repaired component back to the SMAG.

4. Based on anticipated quarterly or semiannual repair programs, the ALCs make plans to have available the necessary resources: carcasses, capacity, repair parts, and funds.
5. Rather than simply carry out a preplanned program, managing daily execution of repairs is a control process, which is the underlying purpose of the Execution and Prioritization of Repair Support System (EXPRESS).

The closed-loop methodology we describe shortly should provide useful inputs to stages 1 through 4. Ultimately, it could also support tracking and replanning during the daily execution of repairs stage.

MAJCOM and AFMC Long-Term Planning

One of the issues raised in the U.S. Air Force Spares Campaign is that the first two stages in the list above are disconnected, with wide differences between views on what portions of the total need for depot-level repairs are to be funded by each organization. Basically, the MAJCOMs and AFMC use totally different planning methodologies.² The MAJCOMs plan and program the purchase of repairs needed for replenishment as a result of operations, essentially what we

¹ SOSs perform the item management functions and SORs are the repair shops or contractors who do the work. The SOS and SOR responsible for any particular item are not necessarily located at the same Air Logistics Center.

² This was raised by the Spares Campaign Programming & Financial Management Team (P&FM) under the heading “Funding Non-Sales Based DLR Requirements.” The following is a direct quote from the team’s draft, *Programming & Financial Management (P&FM) Team Background Papers* (United States Air Force, June 2001b): “The Air Force implemented a ‘spare is a spare’ doctrine in 1992, eliminating separate provisioning distinctions and funding appropriations for readiness spares. However, stock funding of DLRs and other major policy changes implemented since then have perpetuated the distinctions and yielded different funding results for sales- and non-sales based spares.”

term the “keep-up” need. From the MAJCOM perspective, the extra repairs to fill pipelines and provide war reserve spares do not fall within its scope of fiscal responsibility. The MAJCOM methodology is called “the AFCAIG process,” where the acronym stands for Air Force Cost Analysis Improvement Group. In application it amounts to a dollar cost per flying hour to the estimated flying hours—essentially a regression based on previous demands for repair and flying hours.

By contrast, AFMC’s planning process encompasses each of the demands for parts enumerated in the first subsection. Computations supporting this planning are implemented in the D200A system, the Secondary Item Requirements Computation System (SIRS). SIRS estimates “buy requirements” and “repair requirements” by quarter, going out 38 quarters.³ It is run four times a year. In an attempt to reconcile the MAJCOM and AFMC views of spares requirements, the Air Force established the Spares Requirements Review Board (SRRB).⁴ It is at this point that the current process falls short in that it provides limited capability to understand the weapon system readiness implications of the negotiated solution.⁵ The methodology underlying the Closed-Loop Planning System model could offer a very useful tool for the SRRB process.

³ D200A, SIRS, is one part of D200, the Requirements Management System (RMS). The whole of RMS is documented in United States Air Force, 2003a (AFMCMAN 23-1), with a detailed step-by-step description of SIRS in Chapter 9.

⁴ We shall not comment on the success of SRRB. For one commentary, see Camm and Lewis, 2003, pp. 44–47, under the heading “The SRRB Cannot Address Most of These Issues.”

⁵ The AFCAIG process and AFMC process for arriving at repair requirements differ in fundamental ways. AFCAIG basically uses a regression of flying hours against repair costs to estimate repair costs in the future. It does not use item-level demand information. The AFMC process rolls up the requirements from historical item-level demand data. There is a continuing debate about the validity of the historical item-demand data for predicting future demands. The methodology we have proposed in this document also uses historical item-demand data (with uncertainty represented by the variance-to-mean ratio of 1) and is subject to the same questions about predictability of demands from historical data. Those questions remain to be answered in research beyond the scope of this project. We point out however, that both the AFCAIG process and AFMC process suffer from the key problem we note: They do not relate funding to readiness, while our methodology does.

A Closer Look at SIRS

We describe some aspects of the SIRS computation because it is the cornerstone of AFMC's planning and because we intend this discussion to set the stage for the need for our closed-loop modeling.

A detailed description of SIRS is beyond the scope of this report and can be found in other references (United States Air Force, 2003a). Basically SIRS determines the total demand for components by estimating the peacetime demand for repairs from historical removal rates and planned flying programs, determining the additional components necessary to fill pipelines and war reserve, and considering available serviceable assets and repairable components in the pipelines. Depot repair is then estimated from the expected NRTS fraction of the removals plus those repairable components in the pipeline to the depot, less expected "condemnations" (those components that cannot be replaced). There are basically three shortcomings with this process—it does not deal with constraints in repair, it does not consider uncertainty of demands for repair, and it does not provide the readiness impact of budget allocations that differ from those determined in the SIRS requirement.

SIRS Does Not Consider Capacity Constraints on Repair Resources

Depot-level repair is often constrained in various ways. Manpower and test equipment availability often limits repairs. This limitation can be overcome through the use of overtime, but that includes an additional cost burden. Another constraint is the availability of carcasses (repairable components) to repair. Because of condemnations, slow retrograde pipelines, or other shortages, there may not be enough components inducted to the repair shop to satisfy the immediate demands for repair. The SIRS computation does not recognize capacity constraints on repair resources. It assumes that any parts that arrive at the depot for repair and are not condemned can be repaired within a repair cycle time. This ignores the frequent realities of parts shortages and limited repair capacity. The closed-loop methodology includes repair capacity explicitly, which allows the additional management option of allocating budget to increase capacity in selected shops.

SIRS Does Not Recognize the Variance of Repair Demands for NRTS Parts

Randomness in demand on base asset positions requires safety stock to protect the bases from having aircraft that are missing parts because too many are temporarily in resupply. The randomness in demand has an additional effect. It induces a statistical variance in the number of parts that will need to be repaired over a given time interval, such as a quarter or year. This consequence of randomness is ignored in the SIRS computation even though the presence of this variance causes risk in aircraft availability. The Closed-Loop Planning System explicitly considers this variance in its repair budget estimates.

SIRS Does Not Provide a Relationship Between Budget Allocations for Repair and Readiness

The lack of linkage between repair budgets and readiness in the current process requires decisionmakers to make arbitrary budgeting decisions and hope for the best. This is the biggest drawback of SIRS. With the exception of the two items mentioned above, SIRS does a reasonable job of estimating the repair requirement based on historical information. However, when inevitable changes in funding occur, it is nearly impossible to determine what impact a partially unfunded requirement will have on readiness of units, aircraft types, or commands. The closed-loop methodology makes a linkage between repair budgets and readiness and permits an understanding of the readiness implications of such budgeting decisions.

Allocation of Obligation Authority to Sources of Repair

Within the execution year, the ALCs need to manage the application of budgeted OA to weapon systems. In the past, there have been occasions when an ALC had expended all of the year's OA long before the end of the fiscal year. In at least one instance where this occurred, a likely cause was the ALC's practice of repairing nearly every unserviceable part that came to it. To remedy the situation, a notion of "burn rate" was instituted whereby the amount of OA that can be

consumed over short intervals is limited. In this scheme not only is OA rationed by time, but also allocations are specified for weapon system and repair cost center (RCC) combinations.

Since the Closed-Loop Planning System deals directly in terms of funds for weapon systems and repair resources, it would be natural to develop the allocations as an extension to planning instead of continuing with the independent multistage process of allocating OA to ALCs, then to SOSs, and then to SORs.

ALC Planning

To be responsive to various customer needs, the depots need to take actions in advance of actual need to ensure that necessary resources will be there when required. In addition to funding, there is concern for capacity, for availability of parts, and that necessary contracts are in place. In the past, under the Management of Items Subject to Repair (MISTR) system, quarterly repair quantities were negotiated between item managers and shop chiefs. ALC planning was necessary then, but the need became more acute (and presumably, more difficult) with the introduction of EXPRESS. This is because EXPRESS makes repair highly responsive to current needs and, therefore, less predictable than a system that primarily repairs preplanned quantities.

As of this writing, there were two approaches to this kind of advance planning. The Ogden ALC has developed a system called the EXPRESS Planning Module (EPM) that is consistent with EXPRESS (Dynamics Research Corporation, 2000),⁶ which is a module of the Expeditionary Combat Support System (ECSS), a new Enterprise Resource Planning (ERP) system the Air Force is acquiring. The other approach, which is being tested at Oklahoma City ALC, is an

⁶ EPM incorporates a version of Distribution and Repair in Variable Environments (DRIVE) to produce a priority list that can then be truncated by the application of constraints with the Supportability Module, which is also a component of EXPRESS. EPM, like SIRS, does not take into account uncertainty in estimated keep-up quantities.

adaptation of a commercial advanced planning and scheduling (APS) system. It is not clear how APS, an item-level estimating process, might take into account the operational needs of aircraft availability. But in either case, the Closed-Loop Planning System methodology is capable of providing estimates of repair quantities to serve as the basis for ALCs' plans to have necessary resources in place.

Daily Execution

Daily decisions on induction of parts into repair and the allocation of serviceable components to customers are accomplished with the aid of EXPRESS. This report does not include specific recommendations with respect to EXPRESS, as that is ongoing research. However, we will indicate some topics under investigation.

Briefly, EXPRESS develops priority lists for induction and distribution based on very current worldwide asset positions. There are two priority-setting mechanisms. For parts whose demands are assumed to be related to flying hours, priorities are determined by the Distribution and Repair in Variable Environments (DRIVE) algorithm, which seeks to maximize the probability that all bases meet stated goals for supply that is not mission capable (NMCS) at the end of a close-in planning horizon (based on the nominal time it would take to get a part to a customer). For other parts, priorities are related to numbers of back orders.

DRIVE is a confidence optimization that assumes holes on aircraft are consolidated into the fewest number of airplanes at a base. This is consistent with the Closed-Loop Planning System methodology. EXPRESS and EPM also contain a clever device related to burn rates mentioned in the section on allocation of OA to SORs. The approach, called SPAWS for Single Priority Across Weapon Systems, addresses the fact that equal numerical values of DRIVE priorities, when derived for differing weapon systems, do not have equivalent implications for NMCS. Thus, when a shop repairs parts for more than one weapon system, a simple merging of the systems' priority lists would be inappropriate. SPAWS works by keeping the priority

lists separate and sequentially choosing parts from several lists in a way that maintains spending of OA in proportion to the planned allocations.

An issue that is not handled well by the current Air Force depot-level repair planning and control processes is that of replanning when there are shocks to the budget or dramatic changes in demand (from an unplanned warfighting contingency, for example). Budget shocks occur frequently, sometimes as a result of the need to shift funds to another activity and often as a result of new planning and programming. The Closed-Loop Planning System that we describe can serve to help with the replanning, showing the aircraft-availability consequences of the various ways that money for repair can be reallocated.

The next chapter describes the planning problem solved by the Closed-Loop Planning System.

Formulation of the Depot Repair Planning Problem

The Overall Depot Repair Planning Problem

The problem is to determine the best allocation of repair resources to achieve an aircraft availability objective by the end of a planning horizon. “Best” is defined as the lowest cost allocation of those repair resources to achieve the objective. This cost could be considered the required “budget” for achieving the desired availability goal. Alternatively, the problem could be defined as finding the best (highest) aircraft availability that can be achieved with a given level of budget. These two problems are interchangeable in our formulation. The methodology can be used either way—to maximize availability for a given budget/cost or to minimize the cost/budget to achieve target availability. In the description of the problem that follows, we first define the aircraft availability objective, then describe the driving variables (the demands for repair), discuss the decision variables (the repairs to be performed), and finally describe the constraints (budget, carcasses, parts, and repair capacity).

Why Optimize?

Given sufficient capacity and sufficient repair budgets, it would be possible to repair every component that failed as well as build up the safety stock and war reserve to meet all stated requirements. The Closed-Loop Planning System is built under the assumption that capacity and dollars are limited so that choices of what to repair must

be made in the planning process. Optimization is about making those choices rationally, with attention to operational performance objectives and the desires of decisionmakers. It can also draw attention to constraining factors. If decisionmakers can eliminate or mitigate critical constraints, they can generally develop repair plans that simultaneously provide better availability and budget reductions.

The Objective—Aircraft Availability

Measures of Aircraft Availability

Aircraft availability and its inverse, the rate that is not mission capable, have become commonly used measures of readiness for the Air Force. As a direct measure, availability is pretty simple—it just counts the number of aircraft that are available for operations. However, as a measure of supply or maintenance performance, or for planning resources or budgets, its meaning is not quite so simple. Even at the operational level, there are ambiguities. Whether or not an aircraft is considered available can depend on the current mission needs. For example, say that an electronic countermeasure system is not working, but it is not needed for a training mission. When some systems are out of order or do not function as required, but the aircraft can perform a subset of its intended missions, the aircraft is called *partially mission capable*.

In the planning process, when aircraft availability is used as a guide or predictor of how much is enough, we must use models and statistics to estimate the aircraft availability associated with each plan option. Ambiguities arise because various models and statistics estimate aircraft availability differently as a result of assumptions that are not generally agreed upon within the Air Force. An important assumption is whether or not parts are cannibalized across aircraft. With cannibalization, all shortages of parts are consolidated on the minimum number of airplanes. This means that aircraft availability will be dictated by the component part with the most shortages. If, on the other hand, there is an assumption of no cannibalization, then the shortages fall on the airplanes that experienced the problems. This

will generally result in more problem aircraft, compared with the cannibalization case. If we further assume that each aircraft is made nonoperational by one such problem part and that after it is “down” no other parts fail, then the number of aircraft that are not mission capable because of missing parts is equal to the total number of parts shortages counting every part.¹ This is a conservative assumption because the usual practice is to perform some cannibalization to consolidate shortages and to minimize the number of aircraft that are not mission capable, i.e., to maximize availability.

Currently, as described in Chapter Two, safety levels of stock for peacetime flying are computed under the assumption of no cannibalization. The argument for no cannibalization is that cannibalization takes time and resources and the Air Force does not want to build in a requirement to cannibalize aircraft. The counterargument is that computing parts required with an assumption of no cannibalization will lead to an estimate that provides the wrong mix of components when cannibalization is actually practiced, especially under budget limitations. The Air Force computes its RSP requirements for war-time deployments under an assumption of full cannibalization. It also manages day-to-day depot repair planning through the EXPRESS model, which sets priorities under a full cannibalization assumption.

Even if the cannibalization argument were resolved, various statistical measures of aircraft availability are used that differ in significant ways. Two such measures are the average or “expected” availability and the confidence (probability) that availability is at or above a given level of availability (at some specified time). The first of these is easier to compute when the model assumes no cannibalization. The second is more convenient under a full cannibalization assumption. We assume that full cannibalization is the better assumption, and our planning model sets a target aircraft availability or mission capable

¹ Components are usually defined as line replaceable units (LRUs) or shop replaceable units (SRUs). It is possible that an LRU is nonoperable because of one or more internal SRUs. It would be double counting to count a failed SRU and its parent LRU as contributing to more than one aircraft not being available. This illustrates that the NMCS calculations are a bit more complicated than stated above when considering such “indentured” parts. Also, it is possible that there is more than one failed part under a policy of no cannibalization.

rate and then attempts to allocate repair resources and dollars to achieve that rate with high confidence (90 percent). If budgets are high enough, we have found that it does not make much difference as to which assumption and model to use for planning resources. However, when budgets are constrained these measures and assumptions lead to quite different results.

Another issue with the availability measure is reflected in the question, whose availability? One can measure availability or mission capability rate for individual squadrons, for all aircraft at a location, for all aircraft in a theater, by Mission Design Series (MDS), or by wartime roles and deployment sequence of aircraft. Ensuring the availability of individual squadrons at a high level is more conservative and takes more resources than ensuring the average availability across a number of squadrons. These choices are not only budgetary, but involve operational considerations that cannot be incorporated into the planner, so we decided to leave it to the user to select the level of aggregation and differentiation for setting aircraft availability. The user can select target availability for each MDS or drop all the way down to goals for individual squadrons. It would be easy to provide the capability to intermix these options at the MDS or MAJCOM level.

Aircraft availability measures also distinguish the cause of nonavailability, maintenance, or supply. NMCS is the rate that is not mission capable because of supply and NMCM is the rate that is not mission capable because of maintenance at the basic level. Occasionally both factors operate simultaneously, and the classification becomes arbitrary. For depot repair planning, the NMCS is the measure of interest, because the depot is repairing components that are then distributed to the base supply system. When the depot does not deliver on time to fill a shortage and an aircraft is down for lack of that component, then the cause is the supply system.

Figure 3.1 illustrates NMCS goals set by MDS, by MDS and theater, and by MDS and deployment stage for demonstration of the database and scenarios used in our research prototype. Figure 3.2 is a more elaborate version closer to what might be present in an operational system.

Figure 3.1
NMCS Goals and Alternative Aircraft Groupings

MDS	NMCS	MDS	PACAF	USAFE	CONUS
F-16	10%	F-16	10%	10%	15%
F-15 C/D	10%	F-15 C/D	10%	15%	15%
F-15 E	15%	F-15 E	15%	15%	20%
Other	...	Other	...	...	...

MDS	Conflict 1	Conflict 2	Training
F-16 Block 40	0%	5%	15%
F-15 C/D	0%	5%	15%
F-15 E	0%	5%	15%

NOTE: PACAF = Pacific Air Force; USAFE = U. S. Air Force Europe.

RAND MG434-3.1

Figure 3.2 illustrates a less aggregate input of target availability.²

Note that there is an important reason why a single NMCS goal is not used for all weapon systems and theaters. As suggested earlier, the planning model was built under the assumption that one of the important tasks of planners is to make trade-offs between units, theaters, MDS, etc. when the budget for depot-level repair is not sufficient to perform the level of repairs to get all units to their target NMCS rates. As we shall show later, individualized NMCS goals provide one way that planners specify trade-off preferences. For instance, the first units likely to deploy in a conflict can be given lower NMCS goals than those units that would be left behind for training.

The meaning of the target NMCS rate requires a short explanation. If a total number of repairs are made, systemwide, how the repaired parts are distributed to each unit, and whether or not the parts go into aircraft or fill RSP, determines the NMCS value for each

² By this time, the reader has probably observed that we are using the not-mission-capable rate and aircraft availability interchangeably. Actually, the not-mission-capable rate is 100 percent minus availability.

Figure 3.2
Target NMCS Rates Broken Down by MDS, Theater, and Deployment Stage in a Major Theater Conflict

Aircraft and other end items	CONUS			PACAF			USAFE		
	Conflict 1	Conflict 2	Training	Conflict 1	Conflict 2	Training	Conflict 1	Conflict 2	Training
F-16	5	10	15	5	10	15	5	10	15
F-15	5	10	15	5	10	15	5	10	15
B-52	5	10	15	5	10	15	5	10	15
A-10	5	10	15	5	10	15	5	10	15
F-117	5	5	15	5	5	15	5	5	15
B-1B	5	10	15	5	10	15	5	10	15
B-2	10	10	10	10	10	10	10	10	10
C-130	15	15	15	15	15	15	15	15	15
...	...	...	...	...	...	...	...	...	...
F100 engine	10	10	10	10	10	10	10	10	10
F110 engine	15	15	15	15	15	15	15	15	15
...	...	...	...	...	...	...	...	...	...
Missiles	5	5	5	5	5	5	5	5	5
...	...	...	...	...	...	...	...	...	...
...	...	...	...	...	...	...	...	...	...

RAND MG434-3.2

unit/organization/theater. The following paragraphs describe the distribution choices that are used in the closed-loop planner.

The first distribution choices are at the unit level. We earlier described the notions of “catch-up” and “keep-up.” Catch-up defines a deterministic quantity that has to do with meeting RSP goals, correcting shortfalls in stock, etc. Keep-up represents the repairs done to replenish failures while flying ongoing peacetime and contingency operations. The latter is a random quantity. Aircraft availability depends on how well these repair demands are met and whether or not serviceable components are removed from the RSP to temporarily cover parts shortages. Thus, in actuality there will be some parts missing from the RSP and some missing from the aircraft. For planning purposes, it is best to “protect” RSP by not considering the possibility of such temporary borrowing. Otherwise, the confidence of meeting NMCS goals would appear high, but the resulting budgets would require dangerous depletion of RSP stocks to meet those goals. To accomplish this “protection,” we virtually remove parts from aircraft to fill the RSP. The resulting number of holes is denoted as the

number of “effective” holes for the unit. Note that a specific NMCS goal implies a targeted number of “allowed” holes. We can then determine whether the NMCS goals can be met with a given level of repair (and repair dollars). The confidence of achieving the goals is the confidence of doing just enough repairs to both satisfy the random demands of keep-up and to make up any difference between the “effective” and “allowed” number of holes. This ensures that, in planning the level of repairs to achieve these rates, the repairs must also fill any shortfalls in the unit RSP kits.

There are also part distribution choices among units. We assume there is a priority to filling holes in aircraft that is dependent on the urgency of deployment. Holes are filled in units first to deploy, followed by holes in the follow-on units, and then holes in units not deploying. This priority can be changed in the model, but it is likely this would be the order of priority. This means however, that the main effects of shortages in repair will fall upon the nondeploying units. This should replicate the usual practice of “plusing-up” the deploying units at the expense of units left behind.

The target NMCS performance is evaluated at the time of deployment of the first units. This is assumed to be at exactly the end of the planning period, generally a one-year horizon. The model can accept other planning horizons, and for execution management this would in fact be required, but for budgeting we assumed that it is being used to define a budget for one year.

Ideally, it is preferable to use achievable sorties as the target objective. It is possible to devise ways to use a sortie objective in much the same way we used the availability objective, but those methods either require making very gross simplifying assumptions about sortie generation (x sorties per day per available aircraft, for example) or including operational factors (mission length, turnaround time, crew ratio, flying day length, etc.) that are very dependent on mission and scenario. These really have little bearing on depot repair performance and are likely to obscure the impact of good or bad planning for depot-level repair.

Another advantage of the aircraft availability measure is that other nonfighter aircraft end items repaired by the depot utilize similar measures—missiles, engines, large support aircraft, etc.

Setting the Confidence Level

We chose to set up the model to find the lowest cost solution to achieving the desired operational (NMCS) goal at high confidence. This confidence level is an input to the planning module, but we have normally run the model with 80 or 90 percent confidence. This means that the model must plan repairs to minimize the cost but must do enough repairs to achieve the desired probability (confidence) of not exceeding the target NMCS.

Demands for Depot Repair

As noted earlier, there are multiple sources of demand for depot-level repair of repairable components. These include repair demands from a base, from a centralized intermediate repair facility (CIRF), from parts removed from aircraft during depot periodic maintenance (PDM), and from aircraft of other services or foreign governments. As indicated in Figure 1.1, the monetary demands for repair are roughly 60 percent for aircraft parts, 25 percent for engine components, and 15 percent for “other” (missiles, etc.). The components repaired include LRUs and SRUs. Some of the depot repair dollars go to contractors who either provide a unique repair capability for some components or support the repairs done in the depot by adding capacity.

The estimated depot demand for repair is the sum of anticipated keep-up demand, replenishment of RSP, and repair demands from PDM aircraft, foreign governments, and other services. The PDM, foreign, and other service demands are estimated based on the previous year’s activity. The next paragraph discusses the estimation of the keep-up portion of demand.

In general, the demand for repair of aircraft components is related to flying activity and scales with flying hours (or more generally

with a linear combination of number of sorties and flying hours). Keep-up demand for engine components tends to be proportional to engine “cycles,” which depend not only on flying hours, but also on the type and intensity of sortie activity. The first step in the keep-up demand calculation is to calculate expected removals (of parts from aircraft). This is given as the product of the repair demand per flying hour and the flying hours planned by each organization.³ The removal estimate is multiplied by the fraction of the items that cannot be repaired at the base (the NRTS percentage) to yield an estimate of the average keep-up demand for each component at the depot.

Uncertainty about this estimate is due to several factors. First there is uncertainty in actual removal rates. A second uncertainty is due to deviations in the actual flying programs from those planned by the MAJCOMs. Recently this type of deviation has been driven by unforeseen contingencies and deployments. There is also uncertainty due to unknowns in the removal/break process. The removal process involves human decisionmaking and is sensitive to factors such as repair quality and defect rate in subcomponents. As a result, standard Poisson assumptions about failure processes do not always apply. Previous RAND work (Crawford, 1988) has shown that variance-to-mean ratios in these processes are often much higher than the unity relationship defined by Poisson failures. Although we shall assume a Poisson process in our description of the planning prototype, this is not a necessary requirement.

The Decision Variables—the Components to Repair


There are thousands of components to be repaired within the shops of the depot repair system. Fortunately it is not necessary to optimize the planning of repair of every single component type. Figure 3.3

³ This is highly simplified. The calculation ought to use a demand rate adjusted to each unit's mission profile, with the total demand being the sum of the products of each unit's demand rate and anticipated flying hours. In practice, the data to support such a detailed calculation are not available.

shows that the majority of repair dollars are associated with a relatively small number of components—the high-demand, high-cost components. This figure shows that 80 percent of the dollar volume of repair involves only about 20 percent of the components. Therefore, some repair capacity from low-cost, low-demand items can be used, based on their average removal rate, to optimize the repair of components that really drive costs and performance.

In the planning model, we select the important components to optimize from a Pareto analysis such as this and set aside the small amount of capacity and dollars for the remainder without including them in the optimization. The optimization determines how many of each type of the key components to repair in each shop to achieve the availability goals at high probability and low cost.

Figure 3.3
The “Pareto” Curve Showing How a Few High-Demand, High-Cost Components Account for Most of the Repair Load



Planning Period and Point of Evaluation

The planning module assumes that there is a planning period; this will generally be one year if the purpose is budgetary planning. For some other purposes, such as reallocation of budgets during a year, it may be desirable to plan for a shorter period. We assume that the point of evaluation of operational readiness is at the end of the planning period. It would be possible to plan to optimize for multiple points of evaluation, but these would then need to be “weighted” when trade-offs are to be made or they would have to be treated as constraints. The constraint approach is actually an illusion because the decisionmaker would be asked which constraint to relax if they could not all be met.


What Needs to Be Repaired?

The planning module assumes that repairs are done to perform keep-up with peacetime demands and to catch up on shortages for readiness, as illustrated in Figure 3.4.

It is not necessary to repair all components that fail if there are sufficient serviceable components available in the right location to achieve the desired level of mission capability. Basically, the repair plan is to repair all those components that fail plus those needed to replenish safety stock and RSP but not repair more than are necessary to achieve the desired operational performance. This is illustrated in Figure 3.5.

The desired final state is indicated by the left bar, which includes peacetime and wartime spares, serviceable assets, and no more than an allowed number of aircraft NMCS. The middle bar represents the desired final state in terms of serviceable spares and allowed aircraft holes. The initial state, at the beginning of the planning period, is indicated by the rightmost bar, showing initial total serviceable spares and aircraft holes. The catch-up repairs represent the


Figure 3.4
Sources of Repair Requirements



NOTE: MTW = major theater war.

RAND MG434-3.4

Figure 3.5
The Catch-Up Repairs Needed to Achieve a Given Readiness Level



RAND MG434-3.5


quantity needed to achieve the desired readiness. This is not the whole story, however, because repairs to keep up to peacetime flying must also be performed. This is shown in Figure 3.6.

In this figure, the rightmost bar includes repairs for the expected NRTS due to peacetime flying, the repairs necessary to catch up to pipeline and RSP requirements, and the repairs that must be planned to provide the desired confidence in achieving readiness goals given the uncertainty about the NRTS quantity. Note that this quantity is not accounted for in the current D200 planning system, which deals with stock pipeline uncertainties. Even if that stock is provided, the uncertainties in repair demand continue to exist and must be provided for in budgeting.

Constraints

A unique feature of the closed-loop repair-planning model is the ability to handle various constraints to depot-level repair. The constraints include those associated with budgets, carcass availability, and overtime. All of the constraints can be overridden if the user desires.

Figure 3.6
Total Repair Demands, Including Uncertainty Elements



Budget Constraints

In our implementation, the budget constraint is derived through iterative use of the model. Figure 3.7 illustrates the process.

In this diagram the user first sets the target NMCS, and the planning model then determines the cost of that target performance or identifies constraints to that performance. If other constraints such as the availability of carcasses, repair capacity, or overtime limit the performance, the user is asked to specify ways to override those limits—adding capacity, adding overtime, specifying purchase of spares—or adjusting the NMCS goals. Finally, if the cost of performing repairs exceeds the budget, then this also requires an adjustment of the NMCS goals. This iterative process, which necessarily involves the decisionmaker, gives a strong understanding of the relationship between the budgets required, the target or achievable NMCS goals, and various constraints to achieving those goals. Figure 3.8, derived with the closed-loop planning model, illustrates the repair cost associated with various levels of NMCS for the nondeployed

Figure 3.7
The Process of Adjusting NMCS Goals to Determine a Budget for Repair





Figure 3.8
Budget Versus NMCS Trade-Off Curve Derived from the Closed-Loop Planning Model and Prototype Database



NOTE: Data are for CONUS F-15 C/D NMCS aircraft.

RAND MG434-3.8

F-15 C/D units in our sample database and for repairs of the component parts in that database. For example, the model predicts that a budget constraint of \$120.7 million for the sample of parts would lead to a 15 percent NMCS rate for the nondeployed units in this dataset.

Carcass Constraints

When planning depot-level repairs, the repairs that can be achieved cannot exceed the number of carcasses that are available to repair. Several aspects of the supply and repair process limit the number of carcasses available. Repairable components may be in the retrograde pipeline from the units and not available at the depot. This includes those repairable carcasses that are held at the base and those in the transportation pipeline to the depot. At times, the repairable items are held at the base to cannibalize SRUs rather than sent immediately to

the depot.⁴ Some of the parts arriving for repair at the depot need to be condemned, removing a carcass from the system. The supply system attempts to anticipate this by planning for the purchase of components to keep up with the historical condemnation rate. However, if the historical rate does not predict future condemnations correctly, there may be a shortage of carcasses. Further sources of carcass shortages are shortages in the pipeline, safety stock, and RSP. These may be the result of insufficient prior-year budgets, imbalances across the Air Force in parts distribution, new requirements due to increased OPTEMPO (operational tempo), etc. Although it is generally desirable to provide serviceable parts by performing repairs, it will be necessary to purchase additional components when the carcasses to fill these holes do not exist.


In the planning model, we first make the assumption that there is about the same number of repairable components in the preplanning period pipeline (and that will arrive during the planning period) as near the end of the planning period. The latter would arrive at the depot after the end of the planning period. Thus, we assume the system is “in balance” with respect to pipelines and thus does not lead to a carcass shortage from this cause. If the planning module is used during the execution year, then it will be necessary to remove this approximation and deal with what is actually in the pipelines. For planning, which may be up to one to five years in the future, this approximation should be reasonable.

The pipeline shifting approximation just described allows us to assume that all component removals during the planning period that require depot-level repair provide carcasses for repair, with the exception of condemnations. Condemnations are handled by a model input rate that should be based on history. With these assumptions, we can deal deterministically with the carcass constraint for RSP and safety stock. The additional carcasses needed for these are simply the

⁴ In three-level repair, in which both the base and the depot repair LRUs, the average depot repair of an LRU can be quite a bit more extensive than that at the base because the LRU the depot gets can be filled with failed SRUs or holes through the process of SRU cannibalization.

current RSP requirement, plus the safety stock requirement, less the on-hand spares, and less the NMCS goal. The last term must be multiplied by the number of the individual part on each aircraft under an assumption of full cannibalization. Appendix A provides the actual equation used. It should be noted that this is a deterministic quantity: Statistical fluctuations in removal rates are exactly matched by equal fluctuations in the number of carcasses. If the calculation indicates a carcass shortage, it can be remedied either by planning to buy new parts (this is okay if the planning is one or more years in the future) or reducing the NMCS goal. Basically, the planning model asks to have this carcass constraint resolved in order to achieve targeted NMCS. An option is to run the planning model and ignore this constraint. Figure 3.9 illustrates an unachievable set of goals due to car-

Figure 3.9
The Interaction Between Carcass Constraints and Achievable NMCS Goals




ness constraints. It also shows the required carcasses and the NMCS goals achievable if the carcass constraints are not resolved.

Repair Capacity Constraints

Capacity to repair can constrain the output of certain depot shops and ultimately affect aircraft availability. If one shop is seriously constrained while others are not, then, without someone taking steps to improve the constrained shop, the other shops may be spending repair dollars without actually helping to achieve aircraft availability goals. Figure 3.10 illustrates this situation. The constrained repair curve shows the fall off in NMCS goals achievable when the repair constraint is encountered. Clearly the additional budget can be spent, but it does not buy additional aircraft availability. An option is provided to overrun a capacity constraint by allowing the user to allocate overtime hours to the shop. This adds to the cost of repair but may lead to a feasible solution. Figure 3.10 show the effect of adding overtime at two different cost levels.

Figure 3.10
The Effect of Repair Capacity Constraints on the Cost to Achieve NMCS Goals




NOTE: Data are for CONUS F-15 C/D NMCS aircraft.
RAND MG434-3.10

Identifying the capacity constraints in shops can be complex. Constraints may be due to lack of specialized personnel or test/repair equipment. Parts may need various stages or stations for repair, and the capacity constraint may be at only one of these. Parts of different kinds going through a single shop may flow through different types of stations and require different skills. In our modeling of repair capacity, we have assumed that it is possible to define entities within a shop associated with groups of components for which the capacity can be determined. In avionics shops, these are typically the test stands and the personnel who run those stands. In the hydraulics shops, it is also the combination of specialized groups of personnel and some test equipment that defines the capacity for groups of components. The capacity is also defined by the number of hours of operation during regular operations as well as in overtime situations.

This completes the description of the problem in terms of the objectives, the decision variables, and the constraints. The next chapter describes the approach to the optimization problem.

Stepping Through the Solution to the Optimization Problem

Figure 4.1
Steps in the Optimization Process




NOTE: NSN = national stock number.

RAND MG434-4.1

Figure 4.1 shows the steps involved in the planning-problem solution. The shaded steps are those that must be performed by the

user of the planning system. The software algorithm performs all others. We describe each of these steps below.

Figure 4.2
Setting the NMCS Goals




RAND MG434-4.2

The NMCS goals should be provided to the planning algorithm user in the aircraft groupings desired by the user. Figure 4.2 illustrates a simple matrix of goals set up in the planning module prototype. In reality this might include many more types of aircraft and end items as well as different categories such as MAJCOM, theater, etc., as illustrated earlier in Figure 3.2.

Note that there are several blocks for user input to adjust goals in the process. In this case, the goals have been set to give priority (lowest NMCS) to aircraft designated for conflict 1, moderate priority for conflict 2, and lowest priority to aircraft remaining in CONUS. These adjustments are required when the goals cannot be

met because the cost is too high, because there are insufficient carcasses, or because there is insufficient repair capacity in one or more shops.

Figure 4.3
Checking the Carcass Constraint



RAND MG434-4.3


As discussed in the last chapter, the repair shop can repair only those broken components (carcasses) that already reside in the shop at the beginning of the planning period. We make the assumption that all keep-up components, which failed as a result of ongoing operations during the planning period, will reach the shop. The remaining number of carcasses needed is the catch-up repair requirement, which is the difference between the RSP requirement plus initial holes in aircraft, less the serviceable stock on hand, less the allowed holes in aircraft based on the target NMCS. If, for example, the serviceable stock on hand initially matches the RSP requirement and the initial

holes in aircraft for the item equals the allowed holes based on the NMCS criterion, then there is no carcass constraint. On the other hand, in the same situation, if the RSP is short by any number of this item, then there is a carcass shortage equal to this number. Similarly, if stock initially on hand matches the RSP, but the aircraft have more than the allowed holes for this item, then there is a carcass shortage. In these cases, additional carcasses must be obtained to achieve a feasible solution. In equation form, the number of carcasses needed, for each component i , is

$$\begin{aligned} \text{Carcasses}(i) \text{ needed} = & \text{Max}(0, \text{RSP}(i)) + \text{Holes}(i) - \\ & \text{Spares}(i) - Q_i \times \text{NMCS}, \end{aligned}$$

where Q_i is the quantity per aircraft. Because this calculation is deterministic, it can be computed prior to running the optimization problem. If the system is short carcasses, it cannot achieve its NMCS goals without first addressing the carcass problem. The user solves this by either planning to procure more stock for the item or adjusting the NMCS goal until there is no carcass shortage. Figure 4.3 illustrates this step in the algorithm along with the possible user adjustments to spares or goals.


Figure 4.4
Computing the Expected NSN Repair Requirement



RAND MG434-4.4

The expected number of repairs of each item to achieve the NMCS goal is the net of the expected removals plus the catch-up quantity (RSP requirement less the spares on hand, plus the holes existing, less the NMCS allowed holes for the part). The latter is the number of allowed NMCS aircraft multiplied by the number of the part on each aircraft. The actual repair requirement will be distributed around this expected repair requirement, as illustrated in Figure 4.4.


Figure 4.5
Checking the Work Center Capacity Constraints



RAND MG434-4.5

Given the expected number of repairs, the algorithm makes a preliminary check on the shop capacity constraints. This is an approximate check based on a multiple of the repair requirement standard deviation that is generally fairly close in predicting whether or not a shop's capacity will be exceeded. This check allows the user to choose to add capacity through overtime prior to seeing the optimization results. The user can also choose at this point to adjust the NMCS goal or to ignore the shop capacity constraint. Figure 4.5 illustrates this step.


Figure 4.6
Computing the Joint Probability Distribution of Work Center Costs and Hours



RAND MG434-4.6

The next step translates the required repairs in shops to the hours and costs to achieve those repairs. Because the required repairs have a random component, this creates a joint probability distribution of hours and costs for each work center. The optimization problem is to determine where to set the budget for each shop to achieve enough repairs to provide the desired high confidence of satisfying the NMCS goal for each aircraft type. Figure 4.6 illustrates the work center problem.


Figure 4.7
Optimizing the Allocation of Repairs to Work Centers by Marginal Analysis



RAND MG434-4.7

Marginal analysis solves the optimization problem for multiple shops. This process allocates a small amount of budget to each shop, one at a time; computes the improvement in NMCS confidence achieved; and, after completing this for each shop, determines which will provide the most NMCS confidence improvement per dollar expended. This shop then gets the increment, and the process repeats until the desired confidence is achieved or a capacity constraint leads to almost no improvement per dollar expended. Figure 4.7 delineates this process. Appendix A describes the optimization process in more detail.


Figure 4.8
Observing the Cost of Depot-Level Repairs



RAND MG434-4.8

The result of the optimization leads to a result (or set of results if the NMCS goals are varied). The user observes these results and decides whether or not to iterate the process further by adjusting goals, etc. If the cost/budget is satisfactory and the desired goals are achieved, the iterative process can be terminated. Figure 4.8 illustrates a set of results obtained by varying the NMCS goal.

Figure 4.9
Repair and Cost Details for Depot-Level Planning



RAND MG434-4.9

When the cost and NMCS performance goal are satisfactory, the planning algorithm can provide the depot and shops detailed information about the repairs to be done by NSN, shop, and aircraft type. Some of these outputs are illustrated in Figure 4.9.

The next chapter illustrates the use of the prototype-planning module on a prototype database.

The Prototype Database and Model for Depot Repair Planning

The ideas of the previous chapter are embodied in a closed-loop prototype that has two purposes. It provides a necessary test platform to support research leading to a suitable optimization algorithm. Second, it satisfies the need to have a demonstration vehicle so that people can understand the kinds of questions that could be answered and trade-offs that could be made. For both purposes, it is desirable to have a parsimonious scenario so that runs can be made quickly and so that easily comprehended computer displays can be created. On the other hand, the scenario has to be sufficiently rich that there is something significant to test and display. We first discuss the data used to exercise the prototype and then illustrate its execution by showing inputs and outputs. The third part of the chapter compares results from the prototype with a representation of the current planning process with the same data.

Database for the Prototype

To satisfy the two purposes mentioned above, we chose to work with avionics LRUs on three fighter aircraft. The aircraft are F-15 C/D, F-15 E, and F-16 C, Block 40 (B40). Fighter aircraft provide a good starting point for research because their populations are large and their operations are easier to characterize than those of large aircraft. Since a main purpose of the Closed-Loop Planning System is to help the Air Force make informed trade-offs, including both F-15 C/Ds

and F-15 Es in the prototype is useful because they share many LRUs and compete for repair capacity. The F-16s were included because their military purpose is similar to that of the F-15 Es, and decision-makers might be interested in trading off funding and military capability across the two weapon systems.

Avionics LRUs provide an interesting set of components for prototyping. They are expensive, critical to the operation of the aircraft, and constitute a significant proportion of RSP kits. Across the three MDS, there are a reasonably large number of LRUs (70 subgroup masters are in the prototype database) involving nine types of test stands. Each kind of test stand can repair several different types of LRUs. Our “shop-based model,” as opposed to NSN-based algorithms, recognizes the advantage of the flexibility inherent in repair resources that can be used for several different kinds of repair.

Having made these choices, we obtained data from a variety of Air Force sources including D041,¹ EXPRESS,² MERLIN, and a previous RAND report on supporting F-15 avionics (Peltz et al., 2000).

Squadrons and a Scenario

We created a planning scenario that postulated a year of peacetime operations followed by two major conflicts (major theater wars—MTWs), one involving U.S. Air Force Europe (USAFE) and the other, Pacific Air Force (PACAF). The squadrons that fly each of three MDS were identified and assigned to the first or second conflict or were designated not deploying.³ The squadrons, MDS, number of aircraft, and deployment assignments are shown in Table 5.1, and

¹ Although D200 would have been more current, we had easier access to recent D041 data, which was used mainly to identify NSNs and corresponding work unit codes. Work unit codes are necessary for identifying parts, but the fields in the EXPRESS database are not populated.

² Appendix B describes a schema for a relational database, similar to the EXPRESS database, for use with the prototype.

³ No classified sources were consulted; assignments were made simply on the basis of what seemed logical based on peacetime locations.

Table 5.2 summarizes the total number of aircraft assigned to each deployment category.

The Dynamics Research Corporation (DRC) provided 14 months' worth of flying-hour and sortie history for each of the MDS by wing. These data were used as guides for developing planned peacetime flying hours to support estimates of demand for the model prototype. In a real application, such estimates would be for the future and would be provided through a classified system such as the Scenario Subsystem in EXPRESS.

Table 5.1
Squadrons and Deployment Scenario

Squadron	Wing	Location	MDS	PMAI	Command	Deploy
58FTRSQ	33FW	Eglin	F-15 C/D	24	ACC	MTW 1
60FTRSQ	33FW	Eglin	F-15 C/D	24	ACC	MTW 1
27FTRSQ	1FW	Langley	F-15 C/D	24	ACC	MTW 1
71FTRSQ	1FW	Langley	F-15 C/D	24	ACC	MTW 1
94FTRSQ	1FW	Langley	F-15 C/D	18	ACC	MTW 2
390FTRSQ	366WG	Mt. Home	F-15 C/D	18	ACC	MTW 2
15CTNG	57WG	Nellis	F-15 C/D	10	ACC	Not Deploying
391FTRSQ	366WG	Mt. Home	F-15 E	18	ACC	MTW 1
335FTRSQ	4FW	Seymour	F-15 E	24	ACC	MTW 1
336FTRSQ	4FW	Seymour	F-15 E	24	ACC	MTW 2
15ETNG	57WG	Nellis	F-15 E	5	ACC	Not Deploying
333FTRSQ	4FW	Seymour	F-15 E	36	ACC	Not Deploying
522FTRSQ	27FW	Cannon	F-16 B40	18	ACC	MTW 1
412FTRSQ	388FW	Hill	F-16 B40	24	ACC	MTW 1
34FTRSQ	388FW	Hill	F-16 B40	18	ACC	MTW 1
69FTRSQ	347FW	Moody	F-16 B40	18	ACC	MTW 1
68FTRSQ	347FW	Moody	F-16 B40	18	ACC	MTW 1
523FTRSQ	27FW	Cannon	F-16 B40	18	ACC	MTW 2
466FTRSQ	388FW	Hill	F-16 B40	24	ACC	MTW 2
493TFSQD	48FW	Lakenheath	F-15 C/D	24	USAFE	MTW 1
492TFSQD	48FW	Lakenheath	F-15 E	24	USAFE	MTW 1
494TFSQD	48FW	Lakenheath	F-15 E	24	USAFE	MTW 1
510FTRSQ	31FW	Aviano	F-16 B40	18	USAFE	MTW 1
555FTRSQ	31FW	Aviano	F-16 B40	18	USAFE	MTW 1
12FTRSQ	3WG	Elmendorf	F-15 C/D	18	PACAF	MTW 2
19FTRSQ	3WG	Elmendorf	F-15 C/D	24	PACAF	MTW 2

Table 5.1—Continued

Squadron	Wing	Location	MDS	PMAI	Command	Deploy
44FTRSQ	18WG	Kadena	F-15 C/D	24	PACAF	MTW 2
67FTRSQ	18WG	Kadena	F-15 C/D	24	PACAF	MTW 2
90FTRSQ	3WG	Elmendorf	F-15 E	18	PACAF	MTW 2
18FTRSQ	354FW	Eielson	F-16 B40	18	PACAF	MTW 2
36FTRSQ	51FW	Osan	F-16 B40	24	PACAF	MTW 2
114FTRSQ	173FW	Kingsley	F-15 C/D	15	ANG	MTW 1
112FTRSQ	180FW	Toledo	F-16 B40	15	ANG	MTW 1
188FTRSQ	150FW	Kirtland	F-16 B40	15	ANG	MTW 2
124FTRSQ	132FW	Des Moines	F-16 B40	15	ANG	MTW 2
125FTRSQ	138FW	Tulsa	F-16 B40	15	ANG	MTW 2
1FTRSQ	325WG	Tyndall	F-15 C/D	24	AETC	Not Deploying
2FTRSQ	325WG	Tyndall	F-15 C/D	24	AETC	Not Deploying
95FTRSQ	325WG	Tyndall	F-15 C/D	19	AETC	Not Deploying
61FTRSQ	48FW	Luke	F-16 B40	76	AETC	Not Deploying

NOTES: PMAI = Primary Mission Aircraft Inventory; FTRSQ = fighter squadron; TRSQD = training squadron; FW = fighter wing; WG = wing; ACC = Air Combat Command; ANG = Air National Guard; AETC = Air Education and Training Command.

Table 5.2
Totals of Aircraft Deployment Assignments

MDS	Conflict 1 (USAFE)	Conflict 2 (PACAF)	No Deployment (Training)	Total Aircraft
F-16 B40	111	129	76	316
F-15 C/D	135	126	77	338
F-15 E	90	78	5	173

Avionics LRUs and Test Stands

The prototype database includes 70 distinct LRUs. As summarized in Table 5.3, 30 are F-16 parts, each of which is repaired on one of four test stands at Ogden. The remaining 40 are F-15 parts, repaired on five test stands at Warner Robbins. Twenty-three of the 40 are common to both the C/D and E series of F-15. Dealing with common items tends to be a problem with any algorithm that purports to optimize, which is one reason why both F-15 series are included in the

Table 5.3
The Aircraft Components Used in the Prototype Database and Their Test Stand Assignment

Test Stand	F-15 C/D Unique	F-15 E Unique	F-15 Common	F-16 B40	Total NSNs on Stand
OO Computer				12	12
OO RF				4	4
OO Display				7	7
OO Pneumatic				7	7
WR TISS	6		5		11
WR Display		2	7		9
WR METS		6	9		15
WR EARTS	2		1		3
WR Microwave	1		1		2
MDS total	9	8	23	30	70


NOTE: Only LRUs with a significant dollar volume (demand rate $\times$ cost) are represented.

design. The Ogden ALC provided a list of the F-16 B40 LRUs and their test stands. The included F-15 LRUs involved the following analysis: An initial list was developed by filtering a D041 database for the two MDS using work unit codes and program select codes to identify the flying-hour-driven avionics LRUs. This list was checked against Table B.1 in Appendix B of the RAND F-15 avionics study (Peltz et al., 2000).

Table 5.3 shows the number of LRUs cross-matched to the various F-15/F-16 avionics test stands. From data provided by AFMC/XP and DRC, the total annual repair requirement represented by the LRUs in the prototype database represents about one-third of the total funded F-15 repair requirement (\$110 million out of \$315 million), as shown in Figure 5.1.

The Closed-Loop Planning System requires capacity of repair resources in terms of hours of availability over the planning period. Correspondingly, hours of capacity consumed by the repair of an LRU are a required data element. In the case of avionics LRUs, capacity consumed by a repair would be time on a test stand. For the

Figure 5.1
The Aircraft Components Used in the Prototype Database and Their Test Stand Assignment



NOTE: Fiscal year 2003 data provided by AFMC/XP and DRC.

RAND MG434-5.1

prototype data, we used the `repair_hours` field in the `dbo_part_depot` table in the EXPRESS database.⁴

In a real-world implementation of closed-loop planning, other estimates of capacity may be required. To exercise the prototype, however, we assigned capacities to the various test stands based roughly on assumed numbers of test stands of each type multiplied by assumed working hours in the planning period. Our purpose was to match demand and capacity in a way that produced “interesting” cases for testing the process.

Asset Data

Additional data on assets relating to both the depot and bases are available in the EXPRESS database. Depot-related LRU-specific in-

⁴ These data may be total labor hours rather than test stand hours, which are not the same since generally two people work on a test stand. Ultimately, if the DRILLS system is implemented, there will eventually be an empirically observed source of these data.

formation includes serviceable stock on hand, unserviceable LRUs awaiting induction, and LRUs in process. Data on LRUs associated with bases include stock levels, RSP authorizations, assets on hand, and LRUs missing from aircraft. We did not count assets in retrograde and forward pipelines because the fields for that information were not populated in the version of the EXPRESS database we had available.

The Prototype Planning Model


This part of the chapter will be used to illustrate the application of the prototype Closed-Loop Planning System for various problems, utilizing the previously described prototype database.⁵

The easiest application of the system is to determine the depot repair cost to support the NMCS goal. We have arbitrarily set F-16 B40, F-15 C/D, and F-15 E goals for aircraft depending on their war-time assignment in the prototype scenario. As can be seen in Figure 5.2, the goal is the most stringent, 0 percent NMCS, for the first deployment, assumed to be Conflict 1. The goals are relaxed to 5 percent NMCS for Conflict 2 and 15–20 percent for the nondeploying-CONUS units. These goals and a 90 percent confidence requirement for meeting the goals lead to a total budget requirement of \$120.70 million when the Closed-Loop Planning System determines the optimal set of repairs.

Figure 5.3 illustrates the effect on the model-predicted budget of varying the F-15 C/D NMCS requirement for the CONUS (nondeploying) forces. Driving the NMCS goal for these aircraft to 0 percent increases the budget requirement to \$122.6 million.

⁵ We have developed two versions of the prototype. One is implemented in an Excel workbook with input data and some calculations saved in worksheets. The other prototype is a windows application (programmed in Visual Basic) that runs off an Access database. The purpose of the Visual Basic version was to create an implementation that could be readily transferred to the Air Force. The database is a plausible abstract of the EXPRESS database. The figures in this illustration came from exercising the Excel version.


Figure 5.2
Minimum Budget Determined by the Prototype to Achieve the Stated NMCS Goal When There Are No Repair Constraints



NOTES: NMCS goal = full cannibalization of parts at time of deployment and full readiness spares package, with 90 percent confidence.

RAND MG434-5.2

Figure 5.3
Budget Required as a Function of NMCS Goals When Those Goals Are Varied and Repair Is Unconstrained



RAND MG434-5.3

Outputs of the closed-loop model make it possible to determine the cost to remove each additional NMCS aircraft (that is, take an aircraft out of the NMCS category) or to change the NMCS by 10 percent. This is illustrated in Tables 5.4 and 5.5. Table 5.4 provides the marginal cost in terms of percentage of NMCS aircraft, whereas this number is put in the context of number of aircraft in Table 5.5.

So far we have not imposed shop capacity constraints. The closed-loop prototype has been designed to estimate, before performing an optimization, whether or not any test stand or shop will exceed its capacity when it repairs a number of components equal to the expected repair demand and three standard deviations about this mean. If this estimated volume exceeds the repair capacity, then the model worksheet shows negative numbers in the row labeled excess capacity to indicate the approximate percentage of capacity shortfall. This situation is illustrated in Figure 5.4, in which the F-16 shops can provide only a 72 percent confidence that the goals could be met.

At this point the goal could be adjusted until the excess capacity row shows all zeros or positive numbers. This was done and illustrated in Figure 5.5. By adjusting the F-16 goal for Conflict 2 to 5 percent (from 0 percent) and to 15 percent (from 5 percent) for the CONUS units, the shop capacities are not exceeded, and the F-16 confidence increases to the desired 90 percent for meeting these new goals.

It is possible to fine-tune this adjustment by first estimating the goals to make the shops not exceed their capacity, then run the optimization. If capacities are not exceeded, then the goals can be adjusted to be slightly tighter and the optimization rerun. This can continue until some shop or test stand exceeds its capacity. It is also possible to allow the shops to invoke overtime to avoid the capacity limits. The model will then attempt to minimize the cost of achieving the target goals at high confidence, but it will use overtime when capacities are exceeded and include the cost of this overtime in the cost of repair. Figure 5.6, shown earlier as Figure 3.10, is a graph created by running the Closed-Loop Planning System with various goals and allowing overtime. The constrained repair curve does not allow over-

Table 5.4
Marginal Cost Per Percentage of NMCS

Aircraft	Budget Allocated (in millions of dollars)	Percentage of NMCS Aircraft			Millions of Dollars Per 10 Percent of NMCS		
		Conflict 1	Conflict 2	CONUS	Conflict 1	Conflict 2	CONUS
F-16 B40	\$18.5	0%	5%	20%	\$0.92	\$1.10	\$0.63
F-15 C/D	\$49.5	0%	5%	15%	\$2.09	\$1.97	\$1.20
F-15 E	\$52.6	0%	5%	15%	\$2.47	\$1.10	\$1.11
Total (rounded)	\$120.7						

Table 5.5
Marginal Cost Per NMCS Aircraft

Aircraft	Dollars Per Additional Aircraft	Number of NMCS Aircraft			Total Number of Aircraft			Total Number of Aircraft
		Conflict 1	Conflict 2	CONUS	Conflict 1	Conflict 2	CONUS	
F-16 B40	\$106,950	0	6	15	111	129	76	316
F-15 C/D	\$155,328	0	6	12	135	126	77	338
F-15 E	\$273,733	0	2	6	90	42	41	173

Figure 5.4
Identification of Shop Capacity Constraints

	Conflict 1	Conflict 2	CONUS
F-16 B40	0%	5%	15%
F-15 C/D	0%	5%	15%
F-15 E	0%	5%	15%

Results	Work center										
	OO COMP	OO RF	OO DISPLAY	OO PNEUM	WR TISS	WR DISP	WR METS	WR EARTS	WR MICRO	HYDRL 1	HYDRL 2
Estimated capacity (estimated percentage) (regular hours)	-2	-5	-3	-1	2	0	0	0	0	3	-6
Regular hours available	40,518	19,258	20,571	15,163	32,878	1,852	11,802	23,300	7,361	27,030	19,450
Overtime hours used	0	0	0	0	0	0	0	0	0	0	0
Overtime (% of regular hours)	0	0	0	0	0	0	0	0	0	0	0
Overtime cost (\$ millions)	—	—	—	—	—	—	—	—	—	—	—
Budget (\$ millions)	5.44	3.82	4.90	2.27	42.10	0.51	21.96	21.29	10.34	5.80	3.75
Work center confidence	0.922	0.920	0.917	0.985	0.978	0.994	0.978	0.976	0.988	0.988	0.937

Mission Design Series (MDS) confidence	F-15	0.906
	F-16	0.718
MDS repair cost (\$ millions)	F-15	102.00
	F-16	20.18
Total repair cost (\$ millions)	122.18	

Carcasses needed to buy		Number of carcasses	Number of components	Total cost (\$ millions)
	F-16 B40	0	0	—
	F-15 C/D	0	0	—
	F-15 E	0	0	—
Dollars per additional mission capable aircraft	F-16 B40	106,950.47		
	F-15 C/D	155,328.00		
	F-15 E	273,733.52		

Figure 5.5
Adjustment of the NMCS Goal to Fit Shop Capacity Constraints





Figure 5.6
Examination of the Cost of Repair for Different NMCS Goals with Different Levels of Overtime



NOTE: Data are for CONUS F-15 C/D NMCS aircraft.

RAND MG434-5.6

time. Note that because of capacity constraints, beyond a certain point larger budgets do not buy improvement in NMCS. The lowermost curve shows what the budget/NMCS trade-off would be if there were no capacity constraints. The two curves that lie in the middle permit capacity to be purchased at fixed dollar amounts per hour.

Earlier we described another constraint to repair, namely the availability of carcasses to repair. For various reasons the ALCs may not have sufficient carcasses on which to perform repair. The closed-loop prototype can determine this number based on repairable assets in the system and the RSP and NMCS goals. Figure 5.7 illustrates the carcass constraint output of the system. As indicated in the figure, the goals for NMCS must be adjusted upward to 5 percent in Conflict 2

Figure 5.7
Identification of Carcasses Needed to Achieve NMCS Goals

Target NMCS goals			
	Conflict 1	Conflict 2	CONUS
F-16 B40	0%	5%	15%
F-15 C/D	0%	0%	15%
F-15 E	0%	0%	15%

Carcasses needed to buy		Number of carcasses	Number of components	Total cost (\$ millions)
	F-16 B40	0	0	—
	F-15 C/D	10	3	2.91
	F-15 E	7	4	2.46

Achievable NMCS goals			
	Conflict 1	Conflict 2	CONUS
F-16 B40	0%	5%	15%
F-15 C/D	0%	5%	15%
F-15 E	0%	5%	15%

NN	Shortage	Total cost (dollars)
014467839-2	0	—
013619240-3	0	—
012247827-3	0	—
012308578-3	0	—
013841108-3	1	249,416.29
012368438-3	0	—
012905835-3	0	—
012876182-3	0	—
011799699-3	0	—
012112335-3	0	—
014467839-3	2	1,650,811.75
012363657-2	0	—
012238179-2	0	—
012731990-2	0	16,028.68
012320680-3	0	—
013429774-3	0	—
014449008-3	0	—
013062073-3	0	—
012731990-3	0	—
012400136-3	2	384,138.68
011234126-2	0	—
012778913-2	3	1,226,811.48
011007363-2	6	1,664,649.69
012778913-3	0	—
010512886-2	0	—
013732801-2	0	—
013732801-3	1	174,845.02

RAND MG434-5.7


for the F-15s to avoid demanding repairs that will be carcass constrained. A carcass constraint in the year of execution is a strong constraint in that it is not possible to resolve this constraint other than by relaxing the NMCS goal, unless new serviceable parts can be ordered and delivered during the current planning period. In future planning years, it should be possible to purchase the additional components needed. Note that each part purchased avoids only one repair. Figure 5.7 shows the breakdown of the parts (and their costs) that would need to be ordered to achieve feasibility with respect to the carcass constraint.

It is possible to use the Closed-Loop Planning System to examine trade-offs in availability related to trade-offs in depot-level repair budgets allocated to particular aircraft types. As previously mentioned, providing decisionmakers with the capability to perform such trade-offs under tight budget situations was one of the most important motivations and challenges for the development of this proto-

type. Figure 5.8 shows the effect of keeping the entire budget constant for the F-15s while adjusting the NMCS goals for the F-15 C/Ds and the F-15 Es in our scenario.


The Closed-Loop Planning System can also be used to examine other potential cost and performance impacts of depot-level repair. Included in Figure 5.9 is a curve resulting from a 5 percent increase in flying hours for the CONUS F-15 C/Ds. Such increases in flying hours could be the result of additional training demands or an unanticipated contingency operation. Neither of these is uncommon. The important thing is that the implications of such an increase in demand for depot-level repairs, aircraft availability, and budgets can be estimated quickly using the model.

Figure 5.8
Trade-offs Between Weapon System NMCS Goals at Constant Budget Levels



NOTES: NMCS percentages include budget trade-offs between the two Mission Design Series, F-15 C/Ds and F-15Es, in Conflict 2 and CONUS. The entire budget equals \$124.32 million.

Figure 5.9
Effects of Increased Flying Hours on NMCS and Budget



NOTE: Data are for CONUS F-15 C/D NMCS aircraft.


RAND MG434-5.9

Once the aircraft availability and budget allocations have been decided upon, the work breakdown in terms of ALC and shop repair requirements and dollars allocated can be obtained from the closed-loop prototype. Figure 5.10 shows the breakdown two ways, in terms of allocation of a shop budget to aircraft types and in terms of the percentage of budget (and alternatively, work hours) by component expected in the shops for repair during the planning period. This breakdown is based on an expected value division of the repair dollars for a shop proportional to the component repair demands.


Finally, additional information for shop and ALC planning is available from the model's front page in terms of capacity available, overtime planned, budget expected, allocation to aircraft type, expected performance, etc. This is illustrated in Figure 5.11.

Figure 5.10
Detailed Repair Requirements by Shop, Aircraft Type, and Component

**Test center budget breakdown
by MDS, Shop 5
(millions of dollars)**



**Test center budget breakdown
by component,
Shop 5**



RAND MG434-S.10

Comparisons with Current Planning Models

In Chapter Two we discussed some of the models and systems used in planning depot-level repair. We are now in a position to be more explicit about the differences between them and the closed-loop planning methodology.

Capability Differences

The D200 system provides the primary AFMC estimates of repairs on a quarterly basis, projected out for 36 months. Included in the D200 calculation are the RSP requirements process and the Aircraft Availability Model to compute peacetime safety stock levels. The EPM provides a prioritized list of repairs for 90 days from which repair parts can be estimated. This list is constrained by estimated throughput of the various shops and budget allocations. The APS system, a commercial-off-the-shelf (COTS) application, is expected to perform the functions of the EPM and is undergoing a test phase within certain ALCs and for some applications.

Figure 5.11
Additional Detailed Information for ALC and Shop Planning

	Conflict 1	Conflict 2	CONUS
F-16 B40	0%	5%	20%
F-15 C/D	0%	5%	15%
F-15 E	0%	5%	15%

Results	Work center										
	OO COMP	OO RF	OO DISPLAY	OO PNEUM	WR TISS	WR DISP	WR METS	WR EARTS	WR MICRO	HYDRL 1	HYDRL 2
Estimated capacity (estimated percentage) (regular hours)	6	2	5	16	2	0	0	0	0	3	1
Regular hours available	40,518	19,258	20,571	15,163	32,878	1,852	11,802	23,300	7,361	27,030	19,450
Overtime hours used	0	0	0	0	0	0	0	0	0	0	0
Overtime (% of regular hours)	0	0	0	0	0	0	0	0	0	0	0
Overtime cost (\$ millions)	—	—	—	—	—	—	—	—	—	—	—
Budget (\$ millions)	5.05	3.63	4.57	1.91	42.10	0.51	21.96	21.29	10.34	5.80	3.54
Work center confidence	0.979	0.981	0.979	0.995	0.978	0.994	0.978	0.976	0.988	0.988	0.964

<i>Mission Design Series (MDS) confidence</i>	F-15	0.906
	F-16	0.902
<i>MDS repair cost (\$ millions)</i>	F-15	102.00
	F-16	18.71
<i>Total repair cost (\$ millions)</i>	120.70	

<i>Carcasses needed to buy</i>		Number of carcasses	Number of components	Total cost (\$ millions)
	F-16 B40	0	0	—
	F-15 C/D	0	0	—
	F-15 E	0	0	—
<i>Dollars per additional mission capable aircraft</i>	F-16 B40	106,950.47		
	F-15 C/D	155,328.00		
	F-15 E	273,733.52		

There are a number of differences between the Closed-Loop Planning System and the current planning process. First, based on the objective of the research, the Closed-Loop Planning System has been designed to show the effects of budget trade-offs and priorities in depot-level repair on aircraft availability and across MDS. The current systems do not provide aircraft availability output associated with depot repair planning and, consequently, do not show the operational consequences of changes in budget allocations for depot-level repair.

The second major difference is the attention to shop capacity constraints allowed in the Closed-Loop Planning System. We have shown that ignoring such constraints can lead to planning of repairs that make no improvement in aircraft availability because some critical parts cannot be repaired in constrained shops. Of course, the capacity constraint can be ignored or overridden with overtime allocations. The importance of considering the constraint is that it allows managers to focus on capacity problems before they restrict the availability of aircraft, and this can be done in the planning stage. This is an important function of the recently inaugurated Spares Requirements Review Board. The Closed-Loop Planning System provides a capability needed by that board.

There is an inconsistency between some of the elements of the current process. The RSP model, EXPRESS, and EPM all assume that cannibalization of components takes place to consolidate “holes” in aircraft. The Aircraft Availability Model assumes that there is no cannibalization. The Closed-Loop Planning System is consistent with the RSP and EXPRESS models in that it assumes that cannibalization will take place. We have noted earlier in this report that when the normal practice is to cannibalize components, planning under the assumption that there is no cannibalization provides an incorrect and suboptimum mix of components.

NSN Versus Work Center View

We noted in Chapter Two that the current planning processes do not account for the uncertainty in repair demands when computing the repair requirement—only the mean value estimates are used. Ignoring this uncertainty can leave a shop underbudgeted when the uncertain

demands are factored in. The Closed-Loop Planning System includes the repair uncertainty in its budget and repair estimates. The Closed-Loop Planning System consolidates unit repair demands to the shop level to allow “pooling” of the uncertainty. It promises, in effect, to provide parts to units at high confidence when demands occur, but it does not promise to provide additional “safety repair stock” for each unit. We think this perspective is the appropriate way to handle the uncertainty and project performance associated with depot-level repairs. The main idea is that we are aggregating “requirements” over units (flying units, that is) and over NSNs in a shop and relating all that to weapon system availability. Current processes do not do that. Aggregation is good because the pooled demand provides a smaller relative uncertainty (expect some units to have larger-than-expected and some smaller-than-expected DLR demands). The cost of DLR for stated performance at a unit is likely to be overstated when there is a safety level at each unit (which in turn means a high confidence of every NSN meeting its goals—one could call this approach the NSN perspective). The costs were 30 percent higher in one example tested. The other advantage of the work center perspective relative to this NSN perspective is that it provides direct outputs for planning at the work center and shop levels and provides flexibility in execution.

Cost Comparison with Constant Performance


Figure 5.12 compares costs for providing the same NMCS performance with calculations representative of current planning methods and the closed-loop method. The same NMCS performance is achieved with 6 percent fewer budget dollars using the Closed-Loop Planning System. The figure shows the mix of planned repairs derived under each system. These estimates were done for an unconstrained depot repair system. In the case of repair constrained by work center capacity or carcasses, we expect that the planning difference between the current system and the Closed-Loop Planning System would be considerably larger.

The analysis illustrates the benefit of pooling risk across items repaired at the same shop. We chose a set of weapon system NMCS goals and calculated the budgets necessary to meet the goals at 90

percent confidence using the Closed-Loop Planning System. Alternatively, we calculated budgets with a heuristic that is representative of current planning methods. With existing planning systems, shop repair budgets are crafted as the sum of individual item budgets. Repair for each item is budgeted to meet all expected demands, plus a small amount for protection against uncertainty in demands. In this exercise, we budgeted each item so that the joint likelihood of being able to repair all items needed to meet the NMCS goals was 90 percent.

The Closed-Loop Planning System takes a weapon system approach to budgeting for uncertain demands. The main savings of the Closed-Loop Planning System in this example appears when significant risk pooling can be leveraged within shops repairing multiple items. In Figure 5.12, the largest budgets and largest difference between the two plans is for the TISS shop at Warner Robins because it has the largest workload.

Figure 5.12
Cost Comparison Between Current Planning System and Closed-Loop Optimized Planning for Some NMCS Performance



Extensions and Next Steps

Some Possible Extensions to the Planning Model and Its Use of Data to Increase Its General Applicability

During this project we designed and prototyped the Closed-Loop Planning System that computes the depot-level repairs that will support both a peacetime flying requirement and catch-up to additional requirements driven by desired safety stock, war reserve stock, and operational aircraft availability goals expressed by planners. However, there are really several different types of planning, planning horizons, and repair organizational structures. For example, *execution planning* refers to the near-term, day-to-day decisions about what to repair now and to whom to send the repair part. The need to replan during the execution year can occur when there are reallocations of budget dollars or additional dollars for repair are made available. *Replanning* is also necessary when new contingency operations arise and when unforeseen repair demands or constraints occur. Generally, repair is less expensive than the purchase of new LRUs. However, there can be circumstances such as capacity limitations in some repair shops or shortages of carcasses in which it is necessary to plan to buy LRUs. In such cases the optimal solution to the *buy-and-repair planning problem* is different from one that would be determined by optimizing buys and repairs separately. It is possible to use (or start to use) the Closed-Loop Planning System as merely an *operational evaluator of other repair plans* rather than using it to define the plan. In this way, it might be used as an evaluator of plans created by an APS system or

the D200. We use this chapter to address each one of these extensions/uses of the Closed-Loop Planning System.

Planning for Future Years and Planning During the Execution Year

In defense planning, budgeting plans are made for the next fiscal year, called the execution year, and for four or more years following that one. The primary difference between planning for these different periods is that the time horizon for the execution year can be too short to plan some long lead time activities such as the purchase of additional shop capacity (people or equipment) and the purchase of LRUs. Another way the execution year constrains the solution is that it is closer to the current state of the world and may need to consider ongoing or known contingencies or deployments as well as the current state of pipelines and repair facilities. As one moves into the future, it is possible to assume more flexibility in correcting current imbalances as a result of prior management actions and budgets.

Optimizing Buy and Repair Decisions Together

Usually it is optimal to repair an LRU rather than buy one because of the difference in cost. However, the depot shop can repair only the broken LRUs that arrive at the depot. When LRUs are condemned as “unfixable” or sit in retrograde pipelines (for example, when a base does not quickly return its broken LRUs to the depot after drawing a serviceable LRU from base stock), then there can be a shortage of assets to repair even though there is a known demand for the repair. Such shortages of LRU carcasses can lead to a need to purchase serviceable LRUs. The closed-loop model permits the calculation of carcasses needed to achieve the catch-up requirement and meet the aircraft availability goal. If these carcasses are not returned for repair or new parts are not purchased, then there is a limit on meeting the catch-up goal and requisite aircraft availability. Repairs cannot substi-

tute because there are not enough LRUs to repair. This then, should be the first stage of the combined buy and repair process, and it will lead to either a purchase of LRUs or a compromise in the availability goal. There may also be a shortage of LRUs to repair because of condemnations during the planning year. The expected number of condemnations should be added to the aforementioned carcass shortage. Although the condemnations are a stochastic quantity, it is probably reasonable to buy to the expected mean value as is done currently. Thus the buy quantity is the number of carcasses needed for catch-up plus the mean condemnations from all sources.

There is one situation in which it is possible that a joint optimization of repair and buy decisions would be necessary. When there are sufficient carcasses and some shop has a capacity constraint, limiting the planning of repairs to achieve an availability goal, then it is possible to substitute for the constrained repairs with the addition of serviceable LRUs in a buy program. However, this is a very expensive way to overcome the repair capacity constraint. Each purchased LRU substitutes at great cost for just one repair. In general, it will be better to plan to expand the shop capacity by the use of overtime (the needed overtime is calculated by the closed-loop model) or the purchase of additional manpower or equipment within the shop. If serviceable LRUs were to be purchased, then it would be possible, through iterations of the closed-loop model and a marginal analysis algorithm to determine the best number of LRUs to purchase to minimize the total buy and repair cost. For example, by iteratively changing the number of each serviceable LRU and running the closed-loop model for minimizing repair cost, then the best number of LRUs to purchase when combined with the repair cost could be determined. After that item is added, the process would repeat until either a budget constraint was achieved or aircraft availability performance was judged satisfactory.

Planning Depot-Level Repairs to Support a CIRF

A centralized intermediate repair facility is generally an in-theater location that performs intermediate-level repairs. The in-theater location allows it to be potentially more responsive to the particular needs of the theater locations supported as well as to provide shorter pipelines while producing economies of scale that cannot be achieved at the operating locations. Broken LRUs are removed from aircraft and shipped to the CIRF to be repaired. At the CIRF, some LRUs are repaired and some are shipped on to the depot to be repaired. In this latter case, the demands for repair at the depot are the NRTS demands from the CIRF. The closed-loop model would use this NRTS demand in determining the optimal set of repairs. The problem is that the effect on aircraft availability depends not only on the repairs at the depot but on the repair priorities and responsiveness of the CIRF. However, this is not different from our assumptions made about base-level repair. In the absence of a CIRF, the depot sees only the NRTS parts from bases and has to make the assumption that base-level repair is responsively repairing or servicing the aircraft. We make the same assumption about the CIRF, that parts that are not NRTS are repaired in a priority scheme and responsive manner that is consistent with the availability goals used for depot-level repair planning.

Use of the Closed-Loop Planning System to Evaluate Spending Plans

The Closed-Loop Planning System model could be used to merely evaluate spending plans created by another planning process regarding the effects on aircraft availability. The current process, D200, does not provide a measure of aircraft availability except that related to spares purchases in the Aircraft Availability Model (AAM) and in the execution process, EXPRESS. The COTS system, APS, under test by the Air Force also does not include an aircraft availability measure related to repair planning. Using the plans as inputs and providing

the resulting availability as an output, the Closed-Loop Planning System model could provide this function for either system. The number of repairs of each LRU would be input to the closed-loop model, and it would output the aircraft availability that could be achieved at high (90 percent) confidence for this repair plan. This process would require either a manual adjustment of goals input to the closed-loop model until the high confidence is achieved or an algorithmic variation of the goals until the confidence is achieved. The reason for these adjustments is that there are trade-offs across aircraft and across aircraft units that can be made with respect to availability. Some units could have goals of high availability and some lower goals in order to meet the high confidence level. If these choices are to be completely ad hoc and made by senior decisionmakers, then the closed-loop model could be used to examine choices. If a set of rules can be established with priorities for compromises and trade-offs, then the process can be done more algorithmically. We have tested various rules such as this with the prototype and have verified that the closed-loop model would function very well as a plan evaluator via either approach.

Such a plan evaluation could also examine the implications of shop capacity and carcass constraints. The closed-loop model will automatically estimate whether the planned repairs call for increased capacity, overtime, or leaving unused capacity at each shop entity. It will also display the carcasses needed to achieve the goals of the plan. These can be checked against the planned purchases of serviceable LRUs.

Finally, a natural evolution from the plan evaluator is to use the closed-loop model to examine excursions from the plan in terms of impact on aircraft availability and capacity. This could be done by manually iterating the plan or by asking the Closed-Loop Planning System to adjust some of the repairs, i.e., allowing some of the optimization process to suggest alternatives.

Replanning with the Closed-Loop Planning System

There are often reasons to replan the allocation of depot-level repairs. New contingencies happen, unexpected constraints or demands arise, budget reallocations occur, etc. The Closed-Loop Planning System can be used to support such replanning. Several elements distinguish replanning from the exercise of building the initial plan. First, if the replanning is for the execution year during the execution year, then the time horizon must be changed to the time remaining in the execution year. This is a variable in the planning module; so this is easy to do. Of course, if the time horizon is very short there will not be much flexibility to improve or change availability rates through depot repair.

The second element to distinguish the replanning activity during the execution year from initial planning is that the actual pipeline quantities should be used in determining the stock levels and number of carcasses available, which means that data should be drawn from the execution control dataset used by EXPRESS. These data can be used to estimate the parts in the pipeline that will arrive during the planning period and the on-hand LRUs to be repaired.

The third aspect of replanning is the possible need to constrain the new solution to *not* deviate too far from the original plan in various ways. Such constraints can apply to the level of activity in repair shops, LRUs that must be repaired according to plan, and allocations of repair to aircraft types. Each of these constraints can be accommodated in the model by inputting the part of the solution that must remain fixed and allowing the Closed-Loop Planning System to optimize (build to an availability level) the remaining repairs. The primary issue is an implementation one. Interfaces must be set up to allow the choice of constraints and to direct the acquisition of the appropriate real-time and previous plan data. We have not examined this interface issue in the development of the prototype other than to verify that the model could be run in a more constrained mode as described here.

Next Steps

The Air Force does not have an effective way to allocate limited funding for depot repair across weapon systems or to link repair resources to projected readiness. Nor is there a system that computes the readiness related to various levels of depot repair funding. This research project was undertaken to determine how to make the links between repair funding and readiness as measured by aircraft availability. We have created a successful design for a depot-level repair planning model that uses a readiness objective as the goal and that allows the user to examine trade-offs across weapon systems. That is, a decisionmaker may favor the readiness of one weapon system over another or favor some set of operational units over another in the planning process, which usually involves the allocation of limited repair dollars. The model uses these priority inputs to achieve the lowest-cost plan to meet the goals or the best plan that can be generated given a budget constraint.

After developing the concept, the next part of the research involved testing it with a prototype that used real repair and demand data. This report describes the model and the prototype with examples that use a sample of real data. Based on our tests, we are convinced that the approach is valid and that value can be achieved by moving toward the implementation of the Closed-Loop Planning System within the Air Force. The first step in that implementation is to perform some “operational” testing. That is, put the model up on an Air Force platform, tie it to Air Force data, and then run the system in parallel to the processes it will supplant or augment. After this testing is completed successfully, it will be necessary to develop the detailed operational concept that defines responsibilities and system interfaces. It should be recognized that embedding this type of system within an organization such as the Air Force takes time and energy. The Air Force EXPRESS model took many years from concept to implementation to become the feedback control system for the Air Force depot. Time is required to set up and run the tests, to digest the data from the tests, to adjust the model based on testing results, and to retest after the changes are made. During the implementation

period, there needs to be continuing support from Air Force leadership (across the frequent leadership changes) and an acceptance of the concept by the myriad organizations that make up the planning and budgeting process, especially during periods when running parallel systems increases the workload of those organizations. Expectations regarding time to implementation must be set realistically and with careful thought given to the degree of testing and integration required. Successfully navigating the testing and implementation process should lead to the ability of the Air Force, for the first time, to relate the detailed planning of depot-level repair to aircraft readiness and to make readiness-based trade-offs when managing limited repair budgets.

Mathematical Description of the Planning Model

This appendix offers a description of the algorithm employed by the prototype planning model. The next paragraph sets the stage by introducing some nomenclature and general comments about the planning problem to be solved by the algorithm.

We are concerned with planning the repairs that depot-level shops will accomplish over a planning period spanning an interval $[0, T]$. For our purpose, a *shop* is defined to be a collection of resources, such as a group of identical avionics test stands, devoted to repairing a set of parts. For each depot repairable part, the plan takes into account several kinds of requirements, which for modeling purposes are either random or deterministic. The random component, called “keep-up,” is the replacement of parts that fail during the planning period. The deterministic element, called “catch-up” is the number of repairs needed to arrive at a target number of serviceable parts by time T starting from the asset position at time 0, without counting keep-up demands. The target asset position is calculated to satisfy needs for war reserve, safety stocks, and filling pipelines and embodies an allowance for parts missing from weapon systems consistent with specified numbers of NMCS aircraft. When executed, a plan can result in failure to achieve the target asset position at the end of the planning period if the keep-up demands over the planning period are large. Were that to occur, what to do about it in terms of holes in weapon systems or parts missing from war reserve stocks is an operational choice that the planning system does not resolve; it is concerned only with whether or not the goal is met.

We wish to consider shop capacities (when adequate data are available) and financial budgets by weapon system. The fundamental choice in formulating a plan is the allocation of dollars to support weapon systems in shops. Because the keep-up demand is random, the choices are reflected in probabilities that catch-up quantities will be achieved, implying that the resulting number of NMCS systems will not exceed their specified NMCS limits.

A variety of optimization problems can be formulated within this framework. Our choice is to minimize costs subject to constraints that specify probabilities of meeting weapon system NMCS goals, where the probabilities depend on dollar allocations and shop capacities. An alternative might be maximizing probabilities subject to constraints on funding and capacities. But this is less desirable because it requires combining all weapon systems' probabilities into a single objective function. We have illustrated in the main portion of this document how a decision support system based on the optimization procedure allows decisionmakers to iterate performance to achieve a budgetary goal.

Although data on asset positions originate at the unit (squadron or wing) level, these data are aggregated so that keep-up demands, asset positions, and NMCS allowances for each kind of part are aggregated up to total worldwide values. This reflects an assumption that the execution system (e.g., EXPRESS) will do a good job in allocating parts. Also, this assumes that parts will be redistributed to deploying units as needed to fill holes in their weapon systems and RSP kits. The planning process accepts requirements for war reserve (RSP) as a component of the catch-up quantities. The objective of the planning system is to get ready for deployments by time T .

The Catch-Up Requirement as a Decision Variable

For now we assume that the DLRs are aircraft components whose demands over the planning period arise from base-level operations. A *unit* is a collection of aircraft of the same type (MDS) operating from the same base. Let j be an index of units and i refer to a specific kind

of part. For each unit that has aircraft using type i parts, the Air Force has established (using processes outside of the closed-loop planning process) a stock level, $L_{i,j}$, that includes a safety level and expected size of pipelines, and a number of components, $R_{i,j}$, to be in the unit's RSP. If at time 0 there are $M_{i,j}$ type i parts missing from aircraft at unit j , filling the holes and meeting the stock level and RSP requirement at unit j would require $L_{i,j} + R_{i,j} + M_{i,j}$ parts. On the other hand, there may not be enough parts to meet the total requirement. Let N_j be the number of aircraft-allowed NMCS at unit j and q_i be the number of copies of the part on the aircraft (quantity per application—QPA). Assuming that parts of type i are cannibalized when necessary, the requirement is reduced by an amount $q_i N_j$. The total end of planning period requirement for type i parts is

$$\sum_j (L_{i,j} + R_{i,j} + M_{i,j} - q_i N_j), \quad (1)$$

where the summation is over units that have aircraft using type i parts.

Offsetting this number of parts are the serviceable parts in the inventory at time 0, including $S_{i,j}$ serviceable parts at or in transit to base j and serviceable parts D_i at the depot. The total catch-up quantity C_i is

$$C_i = \sum_j (L_{i,j} + R_{i,j} + M_{i,j} - q_i N_j - S_{i,j}) - D_i. \quad (2)$$

In Equation (2), all the quantities except the N_j are either observed status at time 0 or are determined by existing Air Force processes. The N_j are regarded by the planning system as decision variables that can be manipulated by decisionmakers in performing trade-offs. From this point of view C_i is a function of a vector of N_j values. A more complete notation would be $C_i(\mathbf{N}_i)$, where $\mathbf{N}_i$ is the set of allowable not-mission-capable aircraft at bases that use part i . (For the prototype, we grouped units by MDS and one of three deployment

situations, which is only an illustration of how one might want to think about scenarios and is not essential to the overall idea.)

The Random Keep-Up Requirement

Denote by X_i all the repairable parts (excluding condemnations) removed from systems and sent to the depot during the planning period $(0, T)$. The planning algorithm operates with the probability distributions of the X_i . The Air Force has been dealing with this estimation problem for a long time. For parts whose removal processes are judged to be linked to flying hours, D200 maintains demand rates, which are assumed to hold worldwide. The prototype follows the same procedure. Let FH_j be the anticipated flying hours that unit j will perform during the planning period. From D200 we obtain the removal rate per 100 flying hours (*OIMDR*) and the fraction of removals that go to the depot (*NRTS* rate). The expected number of type i parts to be repaired at the depot is

$$\overline{X}_i = \left(NRTS_i \times q_i \times OIMDR_i \times \sum_j FH_j \right) / 100. \quad (3)$$

Were the distribution functions of removals Poisson, the variances would be equal to the means. In the prototype, we have made that assumption and used normal distributions, which is justified through the central limit theorem.¹ Thus, we assume that the keep-up demands for depot repairs are normal distributions with means and variances given by Equation (3).

¹ A variety of data-collection efforts over the years has shown that sample variances of demands over time periods are larger than sample means, which would not be the case if demand processes were stationary Poisson. Some Air Force systems operate with a “variance-to-mean” ratio calculated from exponential functions. See Crawford, 1988.

The Role of Available Repairable Carcasses

A useful way to view the distinction between the catch-up and keep-up components of the overall demand for repairs is that keeping up—i.e., a repair will be done for every part that fails in service—is obligatory unless the asset is in long supply. Then, any failure of the plan, say, because the keep-up demand that eventuates is very high, will appear as a failure to meet the catch-up goal. This would result in either more NMCS weapon systems than the stated allowances, or parts missing from war reserves, or both. Whatever the keep-up demand turns out to be, there will always be a matching number of repairable carcasses generated. (This may not be strictly true because parts removed and sent to the depot may not arrive in time to be repaired during the planning period or there may be some condemnations.) Nevertheless, a necessary condition for a plan to be feasible is that the number of repairable carcasses at the depot or in retrograde at the beginning of the planning period must be at least equal to the catch-up quantity $C_i(N_i)$ because the catch-up is the number of *additional* serviceable parts needed, independent of the keep-up.

Given a census of unserviceable parts, the planning system can report violations of this “carcass constraint.” How to proceed with violations is up to users of the planning system. One might relax NMCS goals, or simply ignore the violation.

The Joint Distribution of Cost and Capacity Within a Shop

The index k will be used to indicate a shop (as defined above) and $I(k)$ the set of parts that are repaired by shop k . The planning algorithm operates with joint probability distributions of dollars and hours needed to be expended within shops in order to meet the catch-up and keep-up quantities of the parts repaired in the respective shops. These random variables are denoted by B_k and H_k . Let the cost of repairing a type i part be π_i and the capacity consumed in hours η_i . The total cost of accomplishing the required repairs in shop k is

$$B_k = \sum_{i \in I(k)} \pi_i (C_i + X_i), \quad (4)$$

and the capacity needed in shop k is

$$H_k = \sum_{i \in I(k)} \eta_i (C_i + X_i). \quad (5)$$

The corresponding expected values, $\overline{B_k}$ and $\overline{H_k}$, are obtained by the same formulas with the X_i replaced with their expected values, $\overline{X_i}$. Since we are assuming that the variances of the X_i are also equal to the means,²

$$\begin{aligned} Var(B_k) &= \sum_{i \in I(k)} \pi_i^2 \overline{X_i} \\ Var(H_k) &= \sum_{i \in I(k)} \eta_i^2 \overline{X_i}, \\ Cov(B_k, H_k) &= \sum_{i \in I(k)} \pi_i \eta_i \overline{X_i}. \end{aligned} \quad (6)$$

Denote the corresponding standard deviations and correlation by the superscripted symbols σ_k^b , σ_k^h , and ρ_k^{bh} . With these means and covariance matrixes, we can define bivariate normal probability distributions:³

$$F_k(b_k, h_k) = prob\{B_k \leq b_k, H_k \leq h_k\}. \quad (7)$$

² In addition to assuming that the variances of demands over the planning period are equal to the means, these estimates of variance are also low because the repair times and costs are being treated as constants rather than random variables.

³ To evaluate bivariate normal distribution functions, we have followed the method described by Tong (1990).

The Optimization Problem

Given a set of weapon system NMCS goals—which determine the C_i through Equation (2) and then the probability functions through Equations (4) through (6)—and limits on shop capacities, there are trade-offs between dollar costs and probabilities of meeting the goals. One could seek to maximize the probability of meeting all goals as a function of dollar budget levels.⁴ Alternatively, one could minimize the cost of achieving specified probabilities of meeting goals. We favor the latter in the belief that people have difficulty in relating to choices of probability values and that it would be better to declare probability goals and try to achieve them as cheaply as possible. The simplest version of this formulation would be to declare an overall probability P as the goal and assume that the shop capacities are fixed. Referring to the distribution functions defined by Equation (7), the shop capacities, b_k , are fixed, and the shop budgets, b_k , are decision variables. Formally, the problem is to determine the set of b_k as

$$\min \sum_k b_k$$

$$\text{subject to:} \quad \prod_k F_k(b_k, b_k) \geq P.$$

This formulation, however, is not totally satisfactory because the single constraint includes all the various weapon system types in an analysis. It is preferable to have a probability constraint for each weapon system, which would facilitate making cost/performance trade-offs for individual systems. To accomplish this, the probability that a weapon system's availability goal is met is taken to be the product of the probabilities associated with the all the shops that support the weapon system. With this approach, the probability function

⁴ This kind of probability maximization is done by the DRIVE algorithm in EXPRESS.

for a shop that supports multiple systems will appear in more than one constraint.

We have defined $I(k)$ as the set of parts repaired in shop k . In addition, we define $I_w(m)$ as the set of parts on weapon system m and $K(m)$ as the set of shops that support weapon system m . $K(m)$ would then be

$$K(m) = \{k : I(k) \cap I_w(m) \neq \emptyset\},$$

that is, the set of shops such that the intersection of the set of parts repaired by the shop and the set of parts on weapon system m is not empty. The expanded formulation of the optimization problem is

$$\min \sum_{k \in K(m)} b_k$$

$$\text{subject to: } \prod_{k \in K(m)} F_k(b_k, h_k) \geq P_m \forall m.$$

Although the constraints permit varying the probabilities across weapon systems as indicated by the subscript on the right-hand sides, it is likely that a constant value across systems would be used. Alternative formulations of the planning problem are of interest, but we will discuss approaches to solving this programming next.

Solving the Programming Problem

As is the case for many of the allocation processes employed in Air Force logistics, some form of marginal analysis can be used.⁵ The constraints are converted to sums by taking logarithms. In this case, we would assume that initially the b_k (and also the h_k) are set to values above—1.5 standard deviations in Equation (8)—their expected val-

⁵ The earliest careful proof known to the authors is Fox (1966).

ues. We consider incrementing the b_k by increments α , say 0.05, times their standard deviations. So-called “sort values” are formed by

$$SV_{kt} = \frac{[\ln \Phi_k(1.5 + \alpha t) - \ln \Phi_k(1.5 + \alpha(t-1))]}{\alpha \sigma_k^b}. \quad (8)$$

In this equation, the Φ_k are standardized versions of the distribution functions F_k with the second arguments suppressed (because they are constant), and t is a count of budget increments of $\alpha \sigma_k^b$ dollars. Because we are taking logarithms of distribution functions with arguments greater than their expected values, the $\ln \Phi_k$ are concave functions. If we allocate dollars sequentially to the various shops, taking at each step the shop with the largest sort value, the budget will be spent in the most efficient manner in terms of increasing the probabilities.

Although we have written logic to solve the problem with marginal analysis, we have found using linear programming to be more expedient and flexible. A linear program can be set up in the following way:

Assume we adopt 1.5 standard deviations and 0.05 as the standard deviation step size in Equation (8). We select a number of steps indicated by τ . We let the number of shops be κ and compute the κ by $\tau + 1$ matrix $W_k(t)$ as

$$\begin{aligned} W_k(0) &= \ln \Phi_k(1.5) \\ W_k(t) &= \ln \Phi_k(1.5 + 0.05t) - W_k(t-1). \end{aligned}$$

The variables of the linear program are a corresponding matrix $Y_k(t)$ for $t = 1, \dots, \tau$ of nonnegative variables with upper bounds of 1.0. The nonzero values of these variables indicate how many budget increments of size $\alpha \sigma_k^b$ are to be included in the plan for shop k . The linear program can be written as

$$\begin{aligned}
& \min \sum_k \alpha \sigma_k^b \sum_{t=1}^{\tau} Y_k(t) \\
& \text{subject to: } \sum_{k \in K(m)} \sum_{t=1}^{\tau} Y_k(t) W_k(t) \geq \ln(P_m) - \sum_{k \in K(m)} W_k(0) \forall m \\
& Y_k(t) \leq 1 \quad \forall k, (t = 1, \dots, \tau).
\end{aligned}$$

A problem with this formulation arises from the fact that dollars are allocated to shops, whereas one of our primary motivations is to provide the ability to associate support budgets to NMCS goals for individual weapon systems. If individual shops were dedicated to a single weapon system, this would not be an issue. To associate budgets with weapon systems, a possibility is a compromise that calls for artificially decomposing shared shops into several shops, each dedicated to a single weapon system. Of course, this would dilute the benefit of being able to observe capacity constraints and would not recognize the flexibility that a shop has to utilize its capacity where most needed. On the other hand, the current execution system does operate with weapon-system-within-shop budget allocations and uses the allocations in the SPAWS algorithm in EXPRESS.

Buying Capacity

A variation of the above is to include the possibility of buying additional capacity, which might be interpreted as planning for overtime. Given the cost per unit of added capacity in a shop, one can associate a dollar cost to a small increment of capacity in a shop analogous to the way in which increments of budget in a shop are handled in the earlier formulation. Developing the W matrix is more complicated, and a companion matrix, $U_k(t)$, of cost increments is also generated. For a shop, k , at each step through the t increments, three possibilities are considered: Increase the budget in the shop, increase the capacity, or increase both. For each, divide the increment in the logarithm of

probabilities over the previous step by the cost incurred. For that step, $W_k(t)$ is the logarithm of the increased probability, and $U_k(t)$ is the associated cost. The objective function to be minimized then becomes

$$\sum_k \sum_{t=1}^{\tau} U_k(t) Y_k(t).$$

The Prototype Database

Description

The prototype planning system obtains information from a Microsoft Access database containing the tables described below. Almost all of the information is in the EXPRESS database, although combining several fields in EXPRESS tables creates some fields here. The two tables tblGoals and tblMDS exist for the convenience of users making multiple model runs. The table giving squadron-level operational plans, tblProgram, is a stand-in for the output of the EXPRESS scenario subsystem. In some cases, examples of entries or comments follow field names. Fields named MDSd200 are used to identify aircraft types. NSN and NIIN are used interchangeably to refer to a specific kind of part.

tblBaseMDS

There is one record for every SRAN/MDS combination. Find MDS at a SRAN or all the SRANs that have a given MDS.

Fields:

SRAN_MDS (Key)	FB2027_F016C
SRAN	FB2027
MDSd200	F016C
Name	Hill

tblBaseNSN

There is one record for every combination of SRAN, NSN, and MDS. This is the source of current data on assets specific to the base

and aircraft. This is an abstraction of the EXPRESS table `dbo_part_base` and is very large.

Fields:

SRAN_NIIN_MDS (Key)	FB2027_010397817__F016C
SRAN	
NIIN	
MDSd200	
BaseName	
RetrogradePipeDt	(Dt means “delay time”)
ServiceablePipeDt	
Level	(Stock level or requisition objective)
RSP	(Authorized, not on hand)
Spares	(Serviceable POS and RSP on hand)
Holes	(Missing from aircraft at the base)
NumServiceablesInTransit	

tblDepotAssets

There is one record for every NSN. This is the source of current data on depot assets. The three fields are intended to account for all parts at depots regardless of condition codes or other fine distinctions.

Fields:

NIIN
Depot_Stock_OH
Awaiting_Induction
Parts_in_Work

tblGoals

There is one record for each MDS. This keeps user-supplied aircraft availability goals as fractions for three “contingencies.” The entries can be changed through the user’s interface.

Fields:

MDSd200 (Key)	
NMCSfrac_C0	(first contingency)

NMCSfrac_C1	(second contingency)
NMCSfrac_C2	(nondeploying units)

tblMDS

There is one record for each MDS. A Yes/No indicator can be used to mark MDS to include in runs to save the user from always having to specify them.

Fields:

Name	F-16 B40
MDSd200 (Key)	F016C
Include	Yes/No

tblNSN

There is one record per NSN. This is the source of static information.

Fields:

NIIN	
QPA	(Quantity per application)
RepairResCode	YDCMET
DollarsPerRepair	
HoursPerRepair	
Failures_Per_100_FH	
Description	
WUC	

tblProgram

There is one record per squadron. It contains data on flying activity for developing estimates of demands for parts. This version is unclassified.

Fields:

Squadron	522FTRSQ
Wing	27FW
Location	Cannon
SRAN	FB4855
AnnualFH	
AnnualSorties	
MDSd200	

tblUnit

There is one record per squadron. The MTW_Assignment field associates the squadron in question with one of the three contingencies whose goals are in tblGoals.

Fields:

Name	522FTRSQ
Location	Cannon
PAI	
MTW_Assignment	(0, 1, 2)
SRAN	
MDSd200	

tblWorkCenter

There is one record per “shop.” This is the source of capacity information and relates codes with common designations. The Capacity field is the primary data item. It may be changed through the user’s interface.

Fields:

Name	Computer
RepairResCode	MBRK9G
SupportLocationName	Ogden
Capacity	
NumTestStands	
ShiftHours	

Calculating Catch-Up, Keep-Up, and Carcasses

In the prototype, the catch-up quantity for a part is calculated as LRU_{s_Needed} less Total_Serviceables and Allowed_Holes. In terms of the prototype database these are

- $LRUs_Needed \text{ for an NSN} = \text{sum over SRANs and MDS of the totals of the Level, RSP, and Holes fields from tblBaseNSN.}$

- Total_Serviceables for an NSN = sum over SRANs and MDS of Spares plus NumServiceablesInTransit from tblBaseNSN plus Depot_Stock_OH from tblDepotAssets.
- Allowed_Holes is developed by first combining information from tblUnit and tblGoals to compute how many aircraft are allowed NMCS at each base (SRAN). Then records from tblBaseNSN are processed to accumulate the allowed missing for each part using the QPA information from tblNSN.

The expected keep-up demand for each part is calculated from information in tblProgram, tblBaseNSN, and tblNSN. The first table provides the projected annual flying program hours for the various base/NSN combinations. Running through tblBaseNSN permits accumulating the projected flying hours for each part, with the QPA information supplied from tblNSN. Once the projected flying hours for the parts to be included in the plan have been developed, the mean and variance of the capacities and dollars required are accumulated for each shop, along with the covariances, using information from tblNSN.

A way to view the carcass constraint is to suppose that it would be possible to repair as many parts as one would like overnight with no capacity or cost constraints. Then, could the catch-up be met by tomorrow? The only reason for it not being met would be that there are not enough broken parts available to repair. To check on the feasibility of achieving the catch-up goal for a kind of part is to add up the unserviceable parts. These would be the Awaiting_Induction and Parts_in_Work fields in tblDepotAssets. Parts in retrograde pipelines should also be included, but at the time our database was developed the corresponding field in the EXPRESS database was not populated.

Bibliography

- Abell, John B., Louis W. Miller, Curtis E. Neumann, and Judith E. Payne, *DRIVE (Distribution and Repair in Variable Environments): Enhancing the Responsiveness of Depot Repair*, Santa Monica, Calif.: RAND Corporation, R-3888-AF, 1992.
- Camm, Frank, and Leslie Lewis, *Effective Treatment of Logistics Resource Issues in the Air Force: Planning, Programming, and Budgeting System (PPBS) Process*, Santa Monica, Calif.: RAND Corporation, MR-1611-AF, 2003.
- Crawford, Gordon B., *Variability in the Demand for Aircraft Spare Parts: Its Magnitude and Implications*, Santa Monica, Calif.: RAND Corporation, R-3318-AF, 1988.
- Dynamics Research Corporation, *EXPRESS Planning Module (EPM), EPM Technical Description, Requirements Description*, Andover, Mass.: Dynamics Research Corporation, E-8799U, May 2000.
- Fox, Bennett, "Discrete Optimization via Marginal Analyses," *Management Science*, Vol. 13, No. 3, November 1966.
- Headquarters, Air Force Materiel Command, Directorate of Item Management, *Readiness Based Leveling (RBL) Production System Process Description*, Ohio: Wright-Patterson Air Force Base, April 2003.
- Keating, Edward G., and Frank Camm, *How Should the U.S. Air Force Depot Maintenance Activity Group Be Funded? Insights from Expenditure and Flying Hour Data*, Santa Monica, Calif.: RAND Corporation, MR-1487-AF, 2002.

- Miller, Louis W., and John B. Abell, *DRIVE (Distribution and Repair in Variable Environments): Design and Operation of the Ogden Prototype*, Santa Monica, Calif.: RAND Corporation, R-4158-AF, 1992.
- O'Malley, T. J., *The Aircraft Availability Model: Conceptual Framework and Mathematics*, Washington, D.C.: Logistics Management Institute, 1983.
- Peltz, Eric, H. L. Shulman, Robert S. Tripp, Timothy L. Ramey, Randy King, and John G. Drew, *Supporting Expeditionary Aerospace Forces: An Analysis of F-15 Avionics Options*, Santa Monica, Calif.: RAND Corporation, MR-1174-AF, 2000.
- Sherbrooke, Craig C., *METRIC: A Multi-Echelon Technique for Recoverable Item Control*, Santa Monica, Calif.: RAND Corporation, RM-5078-PR, 1966.
- , *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*, New York, N.Y.: John Wiley & Sons, Inc., 1992.
- Slay, F. Michael, et al., *Optimizing Spares Support: The Aircraft Sustainability Model*, McLean, Va.: Logistics Management Institute, AF501MR1, 1996.
- Tong, Y. L., *The Multivariate Normal Distribution*, New York: Springer-Verlag, 1990, pp. 193–194.
- United States Air Force, *Requirements for Secondary Items, the Users Manual for D200A (SIRS) & D200N (CSIS)*, Ohio: Wright-Patterson Air Force Base, AFMC Manual 23-1, 2003a.
- , *Working Capital Fund FY 2003 Budget*, 2003b.
- , *Spares Campaign Plan, Potential Changes to Air Force Spares Command & Control Processes*, Draft Working Paper, June 2001a.
- , *Spares Campaign Plan, Programming & Financial Management (P&FM) Team Background Papers*, Draft Working Paper, June 2001b.
- , *Spares Campaign Plan, Recommended Changes to Air Force Requirements Allocation, Execution, and Distribution Processes*, Draft Working Paper, June 2001c.
- , *Spares Campaign Plan, Recommended Changes to Air Force Spares Requirements Determination Process*, Draft Working Paper, June 2001d.
- , *Spares Campaign Plan, Recommended Changes to Air Force Supplier Relationships*, Draft Working Paper, June 2001e.