


AD-A240 159



**AN ANALYSIS OF ISSUES RELATED
TO IMPLEMENTING THE DRIVE MODEL**

Report AF002R1

April 1991



Salvatore J. Culosi

91-09288



Prepared pursuant to Department of Defense Contract MDA903-90-C-0006.
The views expressed here are those of the Logistics Management Institute at
the time of issue but not necessarily those of the Department of Defense.
Permission to quote or reproduce any part - except for Government
purposes - must be obtained from the Logistics Management Institute.

LOGISTICS MANAGEMENT INSTITUTE
6400 Goldsboro Road
Bethesda, Maryland 20817-5886

91 8 28 094

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

REPORT DOCUMENTATION PAGE

Form Approved
OPM No. 0704-0188

Public reporting burden for this 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 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, and to the Office of Information and Regulatory Affairs, Office of Management and Budget, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE Apr 91		3. REPORT TYPE AND DATES COVERED Final	
4. TITLE AND SUBTITLE An Analysis of Issues Related to Implementing the DRIVE Model				5. FUNDING NUMBERS C MDA903-90-C-0006 PE 0902198D	
6. AUTHOR(S) Salvatore J. Culosi					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Logistics Management Institute 6400 Goldsboro Road Bethesda, MD 20817-5886				8. PERFORMING ORGANIZATION REPORT NUMBER LMI-AF002R1	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) WSMIS Program Office HQ AFLC/LMSC/SMWW Wright-Patterson AFB, Ohio 45433				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT A: Approved for public release; distribution unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report addresses procedural, analytical, and technical issues related to implementing the Air Force program for Distribution and Repair in Variable Environments (DRIVE). The DRIVE program is a software system developed to set priorities for depot-level repair and distribution of spares in a way that better meets the peacetime and wartime needs of Air Force units worldwide. The purpose of the DRIVE model is to relate the repair and distribution schedules to specific readiness and sustainability objectives and at the same time keep track of changing demand and asset conditions in the field. By tracking those changes and establishing the capability to modify repair and distribution plans accordingly, DRIVE will make the depot and distribution systems more responsive and proactive than they have been in the past. This report addresses synchronization and consolidation of the Production/Execution (Biweekly) and Quarterly Forecasting versions of the DRIVE model; inclusion of multiple repair sources; proper consideration of peacetime and wartime objective functions; alternative objective functions that would improve overall aircraft availability; benefits of redistribution; improved methods for scheduling the repair of shop replaceable units that are common to more than one line replaceable unit; and benefits of separate repair and distribution models.					
14. SUBJECT TERMS DRIVE, depot repair, distribution, maintenance priorities, priorities for repair and distribution, redistribution for aircraft availability				15. NUMBER OF PAGES 106	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	18. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL		

ACKNOWLEDGMENTS

This work was performed at the request of the Air Force Weapons System Management Information System (WSMIS) Program Office. That office helped to arrange essential meetings and provided points of contact within Headquarters, Air Force Logistics Command (HQ AFLC). The author would especially like to thank Major Curtis Ledgerwood, Captain Chris Moreland, and Ms. Deborah Gaston of the WSMIS Program Office for their guidance and direction.

The author also thanks Mr. Victor Presutti and his staff in HQ AFLC/XPS. In particular, Mr. Richard Moore was especially helpful in describing the intricacies of the current Distribution and Repair in Variable Environments (DRIVE) model and in analyzing alternative methods for implementing DRIVE.

Finally, the author thanks Dr. Craig Sherbrooke of LMI for his keen insights into the relative importance of repair and distribution and for providing the kernel of a model that was expanded and used to perform much of the analyses in this report.



Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

PREFACE

This report addresses procedural, analytical, and technical issues related to implementing the Air Force program for Distribution and Repair in Variable Environments (DRIVE). The DRIVE program was originally intended to operate in two modes: a Production/Execution DRIVE (also known as the Biweekly DRIVE) that schedules repair and distribution over the near term (2 to 4 weeks) and a Quarterly Forecasting DRIVE used by maintenance for planning and ordering parts.

The analyses in the beginning chapters focus on implementation issues. They demonstrate and verify methods for streamlining the DRIVE processes by consolidating the Production and Quarterly DRIVE programs and thereby ensuring they are compatible with each other. Those chapters also address possible procedures for dealing with multiple sources of supply.

The later chapters address analytical and technical issues. These analyses focus on improvements in the mathematical algorithms and objective functions that could be used to maximize aircraft availability within constrained resources.

We have directed our analyses toward those in the logistics community who are familiar with the Air Force DRIVE program and who must make it work. In some instances, we examine complex implementation issues that would be of interest to only those with a detailed knowledge of the algorithms and processes used by DRIVE. To that extent it has a very narrow audience. On the other hand, the executive summary captures the essence of the findings in these areas and would be of interest to all who are working to implement DRIVE.

Executive Summary

AN ANALYSIS OF ISSUES RELATED TO IMPLEMENTING THE DRIVE MODEL

The Air Force Logistics Command is implementing DRIVE – Distribution and Repair in Variable Environments – a software system developed to set priorities for depot-level repair and distribution of spares in a way that better meets the peacetime and wartime needs of Air Force units worldwide. Fundamentally, the purpose of the DRIVE model is to relate repair and distribution schedules to specific readiness and sustainability objectives and at the same time keep track of changing demand and asset conditions in the field. By tracking those changes and establishing the capability to modify repair and distribution plans accordingly, DRIVE will make depot and distribution systems more responsive and proactive than they have been in the past.

In implementing DRIVE, the Air Force Logistics Command has learned, not surprisingly, that some practical and operational questions have to be resolved before the Air Force can incorporate DRIVE ideas into its logistics system. In this report, we analyze seven of the most significant questions and recommend solutions to them:

- *Can we synchronize and consolidate the repair schedules from the Production/Execution (Biweekly) version of DRIVE with those from the Quarterly Forecasting version of DRIVE?* The answer is yes, with a Logistics Management Institute (LMI) procedure that we have successfully programmed as part of the Quarterly DRIVE system.
- *How do we handle the fact that we have multiple sources for depot repair? (The original DRIVE assumed a single depot-level repair point.)* We describe two approaches: one allows repair facilities in Europe and the Pacific to support customers only in those theaters, and the other permits all depot-level facilities to support all customers worldwide.
- *How do we develop repair and distribution schedules that consider both peacetime readiness objectives and wartime sustainability goals?* We propose a methodology for taking weighted, linear combinations of the peacetime and wartime "objective functions" developed for DRIVE.

- *Should we consider using any allocation procedures and objective functions that differ from those we used in early DRIVE development?* We describe alternative allocation rules and objective functions and the advantages and disadvantages of adopting them.
- *To what extent should we plan on redistributing assets among bases to complement DRIVE algorithms?* We demonstrate how redistribution among bases can increase aircraft availability significantly.
- *Can we improve DRIVE's ability to set priorities for repairing shop replaceable units (SRUs) that are common to more than one weapon system and thus improve the utilization of SRU repair capacity?* Yes, and we have developed a method for doing so. However, the method we used separates the repair and distribution functions currently combined in DRIVE.
- *Aside from providing the capability to solve the common SRU problem, does developing separate repair and distribution models have any other advantages?* We think so and list some below. Each of the separate models would be
 - ▶ more focused on a narrower part of the total problem with the opportunity to add features that could not be addressed previously (i.e., the common SRU problem);
 - ▶ simpler and easier to maintain or modify;
 - ▶ more adaptable to the community that will use it; and
 - ▶ more efficient to operate.

These questions and other DRIVE implementation issues are a reminder that the Air Force has not yet settled the debate on how its total logistics system – supply, maintenance, and distribution – should operate to maximize operational readiness. As the Air Force continues to gain experience with a phased implementation of DRIVE, we suggest that work on a repair-only DRIVE model and development of separate distribution/redistribution models be pursued as a way to inform that debate.

CONTENTS

	<u>Page</u>
Acknowledgments	iii
Preface	v
Executive Summary	vii
List of Tables	xi
List of Figures	xiii
Chapter 1. An Overview of This Report	1-1
Synchronization and Consolidation – Chapter 2	1-2
Multiple Repair Sources – Chapter 3	1-3
Peacetime and Wartime Readiness – Chapter 4	1-3
Alternative Objective Functions and Benefits of Redistribution – Chapter 5	1-3
Treatment of Common SRUs – Chapter 6	1-4
Benefits from Separate Repair and Distribution Models	1-4
Summary	1-6
Chapter 2. Synchronization and Consolidation of Production and Quarterly DRIVE Models	2-1
Introduction	2-1
The Synchronization Solution	2-3
The Consolidation Solution	2-6
Chapter 3. Incorporating Multiple Sources of Repair	3-1
Introduction	3-1
Solution 1 – Single Source of Supply	3-4
Solution 2 – Multiple Sources of Supply	3-7
Data Needs	3-8
Unique Repair Shop Designations/Capacities	3-8
LRU/SRU Source of Supply	3-8
Base/Depot Linkages	3-8
Depot Serviceable Assets/Due-ins/Awaiting Parts	3-9
The Advantages and Disadvantages of Each Solution	3-9

CONTENTS (Continued)

	<u>Page</u>
Chapter 4. Optimal Allocation of Serviceable Assets to Bases to Satisfy Peacetime and Wartime Targets	4-1
The Problem	4-1
A Proposed Solution	4-6
Chapter 5. Alternative Objective Functions	5-1
Introduction	5-1
Problems with the Current Process	5-2
LMI DRIVE-Like Model Assumptions and Definitions	5-2
Results of Analysis	5-6
A Comparison of Methods	5-7
Methods 1 and 2	5-7
Methods 3 and 4	5-11
The Effect of Selecting Optimum CAT Values	5-14
The Effect of Redistribution	5-16
An Assessment of Backorders	5-22
Conclusions and Recommendations	5-24
Chapter 6. Common SRU Problem: One Reason for a Repair-Only Model	6-1
The Common SRU Problem	6-1
The Effect of the Common SRU Problem	6-2
The Common SRU Solution as a By-Product of a Repair-Only Model	6-5
Biweekly Comparisons	6-6
Quarterly Comparisons	6-8
General Observations	6-8
Appendix A. The Repair Commonality Index	A-1 – A-2
Appendix B. Overview of LMI Repair Model	B-1 – B-15
Appendix C. Glossary	C-1 – C-2

TABLES

	<u>Page</u>
2-1. Sample Repair List from DRIVE	2-8
2-2. Sample Repair List from DRIVE Sorted on Quarterly Value	2-9
2-3. Sample Repair Lists for Quarterly and Biweekly Drive Analysis	2-9
3-1. Exemplary Sort List	3-5
4-1. Probability Table for Mean Demands of 2 for a Peacetime Objective	4-3
4-2. Probability Table for Mean Demands of 10 for a Wartime Objective	4-4
4-3. Sort Values Based on Larger of Peacetime or Wartime	4-5
4-4. A Notional Example of Sort Values for the Allocation of the Next Spare	4-7
5-1. A Comparison of Alternative Methods for Prioritizing Repair and Distribution of 41 Avionics LRUs at 40 Air Force Bases (No Redistribution)	5-23
5-2. A Comparison of Alternative Methods for Prioritizing Repair and Distribution of 41 Avionics LRUs at 40 Air Force Bases (With Redistribution)	5-23
6-1. Comparison of SRU Requirements With and Without Proper Treatment of Commonality	6-4
6-2. Biweekly Summaries (For Base DI Case)	6-10
6-3. Quarterly Summaries (For Base DI Case)	6-11

FIGURES

	<u>Page</u>
2-1. The Quarterly DRIVE Process	2-4
2-2. Quarterly Versus Biweekly Repairs for the First 2 Weeks (HQ AFLC Model/No Redistribution)	2-6
3-1. Schematic for Sources of Supply	3-2
3-2. Relationship Between ALC and Overseas Repair Cycles	3-9
5-1. Ability of Method 1 (HQ CATs) to Meet ENMCS Goals with No Redistribution	5-7
5-2. Cumulative ENMCS for Method 1 Using ENMCS Goals as CATs with No Redistribution	5-9
5-3. Ability of Methods 1 and 2 (HQ CATs) to Meet ENMCS Goals with No Redistribution	5-9
5-4. Cumulative ENMCS for Methods 1 and 2 with ENMCS Goals as CATs with No Redistribution	5-10
5-5. Repair Quantity for Methods 1 and 2 (HQ CATs) with No Redistribution	5-10
5-6. Ability of Methods 3 and 4 to Meet ENMCS Goals with No Redistribution	5-12
5-7. Cumulative ENMCS as a Function of Method with No Redistribution	5-12
5-8. Repair Quantity for Methods 3 and 4 with No Redistribution	5-13
5-9. Repair Quantity for Methods 1, 3, and 4 with No Redistribution	5-13
5-10. Impact of OPT CATs on ENMCS Method 1 with No Redistribution	5-14

FIGURES (Continued)

	<u>Page</u>
5-11. Impact of OPT CATs on ENMCS Method 2 with No Redistribution	5-15
5-12. Cumulative ENMCS as a Function of Method with No Redistribution	5-15
5-13. Repair Quantity for Methods 1 and 2 (OPT CATs) with No Redistribution	5-17
5-14. Repair Quantity for Methods 1 and 2 (OPT CATs) and Method 3 with No Redistribution	5-17
5-15. Cumulative ENMCS as a Function of Method with Redistribution	5-19
5-16. Method 1 (HQ CATs) Versus Method 2 (OPT CATs) in Meeting ENMCS Goals with Redistribution	5-19
5-17. Method 3 Versus Method 4 in Meeting ENMCS Goals with Redistribution	5-21
5-18. Repair Quantity for Methods 3 and 4 with Redistribution	5-21
5-19. Repair Quantity for Method 1 (HQ CATs – No Redistribution) and Methods 3 and 4 (with Redistribution)	5-22

CHAPTER 1

AN OVERVIEW OF THIS REPORT

The Air Force Logistics Command (AFLC) is currently implementing a program to develop the DRIVE (Distribution and Repair in Variable Environments) system. That system will be used to set priorities for the repair and distribution of spare parts in peacetime and during war so that Air Force units can meet its aircraft availability requirements. The system has the potential to increase significantly the Air Force's ability to meet its operational readiness goals because of DRIVE's following unique features:

- DRIVE is a near-real-time allocation system that combines the most recent estimates of supply and demand to determine priorities for repair and distribution in a way that compensates for the large uncertainties embedded in the long-term demand projections of the Air Force's D041 (Recoverable Consumption Item Requirements System) process. This near-real-time capability is an essential component for managing today's resources for the following reasons:
 - ▶ The supply resources we have today are the results of decisions made at least 18 months ago when we were trying to estimate what was going to happen today; we now know what happened and must take advantage of that knowledge.
 - ▶ We have better estimates of demands and less uncertainty in the demand projections because we are only projecting several weeks into the future rather than several years.
- Unlike the D041 process and the systems that flow from it, DRIVE can allocate repair and distribution resources against wartime requirements as well as peacetime requirements.
- DRIVE allocations for repair and distribution are based on specific quantitative assessments of a unit's capabilities measured against a unit's requirements; those assessments are far more discriminating than the current allocation processes. This capability can virtually eliminate the widespread abuse in the current system in which everything is given a high priority, which means that nothing has a high priority.

This report addresses the following seven areas for improving the DRIVE system so that it can better meet the Air Force's operational priorities during the execution of a logistics support program that is always resource constrained:

- Synchronization and consolidation of Production/Execution and Quarterly Forecasting DRIVE models
- Incorporation of multiple repair sources
- The balance between peacetime and wartime readiness
- Benefits from using alternative allocation procedures and objective functions
- Benefits of redistribution
- Proper treatment of common shop replaceable units (SRUs)
- Benefits from separate repair and distribution models.

SYNCHRONIZATION AND CONSOLIDATION – CHAPTER 2

The Air Force has two DRIVE models: a Production/Execution DRIVE (also referred to as the Biweekly DRIVE) that schedules repair and distribution over the near term (2 to 4 weeks) and a Quarterly Forecasting DRIVE used by maintenance for planning and ordering parts. For more than a year, the Air Force has been trying to synchronize the Production and Quarterly DRIVE models so that they are compatible. By compatibility, the Air Force means that the repair list from the Quarterly model should be the same as the sum of repairs from successive applications of the Production model that cover the same period. The Logistics Management Institute (LMI) has solved this problem with excellent results. In one case, the differences between the Production and Quarterly DRIVE results, using the LMI approach, were smaller by a factor of 10 than those of the next best alternative. The LMI approach has already been implemented in the current version of the Quarterly DRIVE model.

An outgrowth of this synchronization methodology is a procedure for consolidating the Production and Quarterly DRIVE models. This procedure will simplify the data processing, reduce the throughput time, and increase the efficiency of the DRIVE system.

MULTIPLE REPAIR SOURCES - CHAPTER 3

The current DRIVE model does not include the depot repair capabilities of the two overseas depots -- a problem that must be resolved if the Air Force is to take maximum advantage of its total depot repair capability. We propose two solutions: one restricts the European and Pacific depots to filling only requisitions from bases in their respective theaters, and the second permits "cross-fertilization" among all depots, with the European depot capable of resupplying Pacific bases and vice versa. We recommend that the Air Force pursue the second alternative because it provides the greater flexibility and better utilizes all the Air Force repair and distribution capability, albeit with some opportunity for increased distribution costs.

PEACETIME AND WARTIME READINESS - CHAPTER 4

The current DRIVE model uses a single wartime objective function; consequently, it does not adequately address the balance between peacetime and wartime readiness. We propose that DRIVE use a weighted, linear combination of the peacetime and wartime objective functions. The weighting factors would reflect both the operational priorities of the commanders and the changing world situation. In a relatively benign environment, weighting factors could be set to give the highest priority to forward-deployed bases and distribute shortages evenly across the remainder of the force. In periods of heightened tension, such as that during *Operation Desert Shield*, the weighting factors could be adjusted to reflect the wartime priorities, thereby maximizing aircraft availability against the wartime objective function and effectively redistributing shortages to the least important bases.

ALTERNATIVE OBJECTIVE FUNCTIONS AND BENEFITS OF REDISTRIBUTION - CHAPTER 5

The current DRIVE objective function is biased toward smaller bases (those with less than 24 aircraft); it tries to increase the availability of their aircraft at the expense of larger bases. Our analyses demonstrate that alternative procedures and objective functions for DRIVE would yield significant improvements in total aircraft availability. However, those improvements would come at the expense of slight reductions in aircraft availability at the smaller bases. Our analyses suggest that if we are willing to accept marginal reductions in availability at the smaller bases, we could increase significantly total aircraft availability across the force. On the other hand, if we insist on meeting specific by-base availability goals, we do so at the cost of

increases in total not mission capable (NMC) aircraft. That cost is lower and total NMC aircraft are significantly decreased when we take advantage of the enormous leverage that comes from redistribution of maldistributed spares. In one case, redistribution of spares reduced the number of NMC aircraft by almost 50 percent.

TREATMENT OF COMMON SRUs - CHAPTER 6

The DRIVE model does a poor job of setting priorities for the repair of SRUs that are common to more than one line replaceable unit (LRU). It overestimates the repair and distribution priorities for common SRUs at the expense of fewer repairs for other high-priority SRUs - a problem that results in inefficient use of scarce repair resources and lower aircraft availability. Proper treatment of common SRUs requires a revision of the design of the DRIVE model.

We have developed a repair-only prototype model that serves two purposes: it solves the common SRU problem and, through an assumption of perfect redistribution, recognizes the significant improvements that result from even minimal redistribution. This repair-only model is described briefly in Appendix B; the Addendum to Appendix B presents a comparison of results from the repair-only model with those from the current AFLC DRIVE model.

BENEFITS FROM SEPARATE REPAIR AND DISTRIBUTION MODELS

The principal advantages of developing separate repair and distribution models for DRIVE are as follows:

- Each model would be more focused on a narrower part of the total problem with opportunities to incorporate features that could not be addressed previously (e.g., the common SRU problem).
- The models would be simpler and easier to maintain and modify.
- Each model would be more adaptable to the community that will use it.
- The models would be more efficient to operate.

The LMI repair-only prototype model provides all of the above advantages as well as retaining those of the current DRIVE model. It makes efficient use of scarce repair capability from three perspectives:

- It better allocates SRU repair capacity because it calculates properly the support requirements of SRUs that are common to more than one LRU.

- It better utilizes repair capability for LRUs and SRUs by building a list of spares that are in short supply *worldwide* rather than those that are in short supply at a *specific location*. That capability can be provided with the current DRIVE model by "turning on the switch" that redistributes serviceable assets among bases. This explicit redistribution is not necessary in the LMI repair-only prototype, thereby making it simpler and more efficient.
- As a collateral benefit, the LMI repair-only prototype tends to provide a more stable repair list for a maintenance system that thrives on low volatility in repair scheduling. A repair-only model that not only considers the needs of the force but also puts stability into the repair scheduling would seem to strike the right balance between the supply community that "wants what it wants when it wants it" and a maintenance community that must consider the real problems of production efficiency.

The LMI repair-only model intentionally begs the question of how to distribute spares because we believe distribution decisions are more appropriately addressed

- later in time when we know what has actually been repaired and not just what has been assigned priorities for repair;
- at a level at which we have better visibility of the day-to-day needs of the force; and
- with knowledge of how redistribution can enhance the combat capability of the force.

With regard to the first point, we can use the current DRIVE model. After repairs are made, we can rerun the current DRIVE model using available serviceable assets and zero repair capacity. In that mode, the current DRIVE model will not merely distribute serviceables but would also perform calculations needed to determine a repair list we neither want nor need. This makes the current DRIVE model less efficient as a distribution model.

On the second point – the need to make distribution decisions with better visibility of force requirements – DRIVE needs to interface with other systems at lower echelons such as TRADES (Theater/Region Allocation/Distribution Execution System) and MASS [MICAP (mission capability) Asset Sourcing System] that distribute spares to the force. A separate distribution will make that interface simpler.

Further, we need a more capable distribution model than the current DRIVE model, one that will be better able to complement a repair-only model that assumes perfect redistribution. Because the LMI repair-only model assumes perfect redis-

tribution, its use could lead to lower peacetime availability to the extent that assets are really maldistributed and *not* routinely redistributed in peacetime. That situation would be ameliorated through the wise use of lower echelon systems such as MASS and TRADES that will have the capability to redistribute assets. If maldistribution persists, however, we could redistribute these assets during periods of warning before a conflict and thus achieve greater combat capability than we would have if we repaired items only because they were maldistributed. Thus, a critical component of a new distribution model would be the development of redistribution "packages" that *could* be executed on short notice *as necessary*.

SUMMARY

This report documents the solutions to some implementation issues for DRIVE and suggests approaches to others. Much, however, remains to be done if DRIVE is to be effective. DRIVE is a big step forward in Air Force logistics because of its near-real-time planning horizon and its wartime sustainability orientation. These strengths, however, often place DRIVE in conflict with existing systems – systems that are not going to disappear over night. Development of the repair-only prototype should continue and the development of a separate distribution model should begin. These initiatives can address some of the remaining technical issues and can serve to address questions of policy on how the total logistics system should operate to maximize the operational readiness of the forces.

CHAPTER 2

SYNCHRONIZATION AND CONSOLIDATION OF PRODUCTION AND QUARTERLY DRIVE MODELS

INTRODUCTION

The fundamental purpose of the DRIVE system is to develop for the maintenance depots a priority list of spares that should be repaired and distributed to the units to meet more efficiently and effectively the operational needs of the Air Force within total budget constraints. Because of the wide variability in demands for items repaired in a particular maintenance shop, the Air Force realized that weapon system availability could be improved significantly if the maintenance shops could be made more responsive to the day-to-day needs of the force. Thus, DRIVE was intended to operate with a time horizon of about 2 weeks.

At the same time, the current maintenance planning process requires an estimate of repairs for a longer time horizon in order to ensure the availability of the consumable items needed to repair assets projected for induction into maintenance depots. That longer time horizon has traditionally been 90 days. Against this backdrop, the logistics community began to talk about two DRIVE models:

- A Production DRIVE model that has often been referred to as the 2-week or Biweekly DRIVE.
- A Quarterly DRIVE model that sets priorities for repair and distribution for 90 days.

The process of implementing the DRIVE model was to begin with the development of the Quarterly DRIVE and then follow up with the Production DRIVE model. Early in that process, the question of compatibility arose; that is, would results from successive runs of the Production model, when added, yield the same answers as the Quarterly model for the same repair period?

Compatibility demands that, under a specified distribution of demands across all bases for 90 days, the sum of repairs from six consecutive Biweekly (15-day) runs

of the Production model¹ should be the same as the repairs from one run of the Quarterly (90-day) model. If the totals were significantly different, we would know that an inherent bias exists in the way these two models operate – a factor that would essentially introduce even more variability into the system than that generated by the uncertain demand patterns.

A caution before we move on: The above definition does not imply that six Biweekly DRIVE runs ought to give the same answers as the Quarterly run *in the execution of the program*. It merely claims that *if the demands over time occur as predicted*, then the answers from each model should be the same. The answers from the Quarterly model should not – and will not – be the same as the sum of six Biweekly analyses if the demands predicted in each model do not in fact occur. In fact, the inspiration for developing DRIVE was the realization that demand forecasting is not, and never will be, accurate “enough” and that demands will never occur as we predict. Thus, we need flexible and responsive maintenance, frequently reviewing actual demands and aircraft status, to react to varying and unpredictable demands.

This issue is also referred to as synchronization of Quarterly and Production DRIVE. We will not belabor the earlier attempts to synchronize these two models; Headquarters (HQ) AFLC and the RAND Corporation are familiar with the details. We will instead present the approach that has proven to be the most successful in obtaining compatibility and synchronization.

Another compatibility issue concerns consolidation. As HQ AFLC began to implement the Quarterly DRIVE model and was about to begin implementation of the Production DRIVE model, it naturally desired to get the Production and Quarterly results from a single run of the model. If it could do so, the Production run would be a special output from the Quarterly run and the entire process would be accelerated.

The solutions for both these problems are presented in the following sections.

¹At the end of each Biweekly run, the asset posture for each item at each base would be incremented by the repairs made for that base and decremented by the expected failures at that base. The model would then be run against that new asset posture. At the end of six runs, the repairs would be totaled and compared with a Quarterly run that used that same total repair capacity as the sum of the six Biweekly repair capacities.

THE SYNCHRONIZATION SOLUTION


Synchronizing the Quarterly and Production DRIVE models requires that we use a common objective function in both models and adjust the input data for the Quarterly DRIVE run.

We begin by summarizing a selective list of steps that must be followed in running the Production model. These steps are the areas that will be modified in the Quarterly run to make it compatible with the Production run:

- Specify the repair period covered by the Production run period. Let that period be noted as time $T1$; in Figure 2-1, $T1$ is equal to 15 days. If the order-and-shipping time (OST) is $T2$ and the war period is $T3$, then the total time horizon (TH) for this Production run is $T1 + T2 + T3$. The total failed assets sent to the depot (DI) are calculated using this time horizon.
- Calculate the total assets of each item, $S1$, that each base will have at the end of the time horizon if no failures occur from this point forward and no additional serviceable assets are sent to this base.
- Calculate marginal benefits ("sort values"), as currently done in the DRIVE model, using DI as the mean number of "due-ins" and the starting supply level of $S1$. The sort value calculated in the DRIVE model is the expected "benefit" of repairing and distributing the next spare (LRU or SRU) to the indicated base divided by the "repair cost" of providing that spare. The benefit is the relative increase in the probability² of achieving a specified number of available aircraft and the repair cost is the number of labor hours needed to repair the item. For serviceable items, the repair cost is set to 1.0.
- Carcasses are constrained to those that are currently at the depot and available for repair; let that number be noted as $C1$.
- Total man-hours of repair capability for the Production run is defined as MHP (man-hours for Production).

Making the Quarterly run compatible with multiple applications of the steps in the Production process requires that we essentially move the Production run "downstream" $90 - T1$ days. In Figure 2-1, since $T1$ is 15 days, the process is moved

²For technical reasons, the benefit in the sort value is actually the logarithm of the relative increase in this probability.



Note: NRTS = not reparable this station.

FIG. 2-1. THE QUARTERLY DRIVE PROCESS

downstream 75 days. The following steps for the Quarterly process parallel those for the Production process:

- The time horizon for the objective function remains the same (*TH*) with the total number of failed assets sent to the depot equal to *DI* as noted previously.
- The starting asset posture for the Quarterly run (*S2*) is *S1* minus all the expected failures classified NRTS for the period 90 – *T1*. Thus, in the example shown in Figure 2-1, we would subtract 75 days of failures from *S1* to get *S2*, the new starting assets posture for the Quarterly DRIVE run.
- The sort values are then calculated as currently done in the DRIVE model, using *DI* as the mean number of due-ins and *S2* as the starting supply level.
- Carcass availability is increased by the failures that occurred during the 90 – *T1* days. This is the same number that was subtracted from the asset posture in the second step above. Thus, the total carcasses for an item in the Quarterly run (*C2*) is equal to *C1* + *S1* – *S2*.
- Total man-hours of repair capability available for the Quarterly run is defined as *MHQ* (man-hours for Quarterly) and is the total repair hours available for 90 days. It is given by the following formula:

$$MHQ = \frac{MHP \times 90}{T1}$$

This solution emerged from an analysis using a DRIVE-like model that LMI developed. When six Biweekly runs were made and compared with the results of a Quarterly run using the above approach, we obtained identical repair lists. In fact, we obtained identical results (perfect compatibility) in 11 of 12 excursion analyses. One case was "off" by about 3 percent. [By "off" we mean that the sums of the absolute differences in repair quantity divided by the total repair quantity across all items was a little over 0.03.]

The LMI approach is now being used by HQ AFLC. Before it was implemented, it was tested by HQ AFLC/XPS and while the approach did not give identical results, it was found to be the most compatible one by a wide margin. The Production and Quarterly lists were not identical in the HQ AFLC test because the DRIVE model at HQ AFLC has several constraints that are not present in the LMI DRIVE-like model. Nonetheless, the HQ AFLC results were excellent. In one case, the sum of the squared deviations of repairs across 250 items for the LMI method was only 1/20th that of the closest competitor.³

On the other hand, our analysis shows that the *sequence* of repairs from the six Biweekly runs differs significantly from that of the Quarterly runs (see Figure 2-2). In Figure 2-2, the items are sorted; items requiring the largest number of repairs over the 90-day period are on the left and those requiring the lowest number of repairs on the right. Each point on the curve is the *difference* in number repaired for the following two cases:

- The number repaired in the first 2 weeks from the first Biweekly run
- The first 2 weeks worth of repairs from the Quarterly run.

We found that the Quarterly run gives highest repair priority to those items that have the largest daily demands and the lowest supply posture. This fact is particularly relevant when we considered options for consolidating the Quarterly and Production runs discussed in the next section (i.e., we could not take the first 2 weeks of repair from the Quarterly analysis and call it the Production run).

³The details of these comparisons can be obtained from Mr. Richard Moore of HQ AFLC/XPS.




FIG. 2-2. QUARTERLY VERSUS BIWEEKLY REPAIRS FOR THE FIRST 2 WEEKS
(HQ AFLC model/no redistribution)

THE CONSOLIDATION SOLUTION

The process for developing a consolidated Quarterly and Production DRIVE analysis arises from the solution we just presented for synchronizing these processes. A prerequisite for understanding the proposed solution for "consolidation" is an understanding of the current process that is used to develop a separate Quarterly and Production analysis. We outline that process briefly below:

- The current DRIVE model first calculates a sort value for each possible allocation of an item in an LRU family on the basis of the appropriate objective function and asset posture for that run. Since this calculation is based on an installation's asset positions and carcass availability, we would get different sort values and a different priority of repair and distribution depending on whether we are making a Quarterly or Production run. The analysis presented in the previous section on synchronization shows that while the total repairs are the same, the order is very different.
- The current DRIVE model then merges the allocations for all LRUs and SRUs on the basis of the sort value for that run.
- The analyst then uses a postprocessor program called "LINE DRAW" to assign workloads to the individual workstations at the depot; the program

essentially moves down the list of priorities and stops when the total capacity for a shop is exceeded.

The solution proposed for consolidation would be to calculate two sort values in the main program. They would be used in the postprocessors to develop a Quarterly and Production result. This process is outlined below.

We begin with the Quarterly version of the DRIVE model that operates with the adjustments proposed above in the synchronization process. This model would be modified to calculate a second sort value — one that will be used in the postprocessor to develop a Production list of priorities. The first sort value would be a Quarterly one; the second would be a Production sort value. The sort value for Production would use the same objective function but with a supply level at each base that is not reduced by the base NRTS for 90 — T1 days. Also, available asset carcasses would be reduced by these NRTS demands.

The following notional example shows how the process should work. Table 2-1 is a sample repair list that would be generated by the DRIVE model as it processed one LRU family at a time. Note the two sort values for each base record: a Quarterly value and a Biweekly value (Production period). Table 2-2 is the result of a descending sorting of this file using the Quarterly sort value.

In our notional example, we assume further that the Biweekly repair capacity is 2 labor units and that each repair costs 1 unit. Thus, over a 90-day period, 12 labor units are available to repair 12 items. The dashed line in Table 2-2 is the cutoff line for the 12 repairs for this LRU shop.

The two repairs that will be done in the first Biweekly period are determined by sorting the base records in Table 2-2 on the basis of the Biweekly sort values. We have not done that type of sorting but, rather, in Table 2-2, we have "boxed" the two items with the greatest values. The total repairs by LRU type for the Quarterly and Biweekly (Production) analysis are summarized in Table 2-3.

At a briefing, HQ AFLC generally agreed that this approach offers the best promise for consolidating the two DRIVE models. The approach is scheduled for testing in the Production model, and if it is validated, it will be implemented.

TABLE 2-1
SAMPLE REPAIR LIST FROM DRIVE

LRU	Base number	Sort value	
		#1 - Quarterly value	#2 - Biweekly value
A	1	10.0	5.0
	3	9.5	6.0
	2	9.0	5.1
	3	8.4	5.8
	4	6.5	5.2
	2	6.0	4.5
	1	5.8	4.6
	⋮	⋮	⋮
B	1	9.4	7.0
	2	9.2	6.5
	4	7.2	5.0
	⋮	⋮	⋮
C	1	8.7	6.8
	2	7.0	6.1
	⋮	⋮	⋮
D	4	8.5	4.3
	1	5.6	4.0

TABLE 2-2

SAMPLE REPAIR LIST FROM DRIVE SORTED ON QUARTERLY VALUE

LRU	Base number	Sort value	
		#1 - Quarterly value	#2 - Biweekly value
A	1	10.0	5.0
A	3	9.5	6.0
B	1	9.4	7.0
B	2	9.2	6.5
A	2	9.0	5.1
C	1	8.7	6.8
D	4	8.5	4.3
A	3	8.4	5.8
B	4	7.2	5.0
C	2	7.0	6.1
A	4	6.5	5.2
A	2	6.0	4.5
A	1	5.8	4.6
D	1	5.6	4.0

TABLE 2-3

SAMPLE REPAIR LISTS FOR QUARTERLY
AND BIWEEKLY DRIVE ANALYSIS

LRU	Quarterly repairs	Biweekly repairs
A	6	0
B	3	1
C	2	1
D	1	0

CHAPTER 3

INCORPORATING MULTIPLE SOURCES OF REPAIR

INTRODUCTION

The Air Force repairs depot-level reparable at each of the five Air Logistics Centers (ALCs) and at two overseas depots [one in the European theater at Royal Air Field (RAF) Kimball and one in the Pacific theater at Kadena Air Base, Japan]. The current DRIVE model considers only the five ALCs in the United States and does not yet have the capability to set priorities for the workload at the repair shops for the overseas depots. In fact, this problem is really much broader in that the current DRIVE model considers only organic repair capability and does not integrate contracted depot repair. This deficiency in the current DRIVE program is referred to as the "multiple-sources-of-repair" problem.

Before proposing solutions for this problem, we present some of the "physics" of how these multiple sources of repair interface with the units. At each Air Force base, each reparable spare (LRU and SRU) is assigned a main source of supply. When an LRU or SRU fails at a base and cannot be repaired locally at base or intermediate repair, it is sent to one of the depots for repair. (Such a failure is referred to as not reparable this station or NRTS.) In general, RAF Kimball is the source of supply for the European bases *for the subset of reposables that RAF Kimball can repair*. That does *not* mean that *all* spares denoted as NRTS from European bases have RAF Kimball as their main source of supply; each spare has its own designation for source of supply at each base. Figure 3-1 is a schematic example for the interfaces that could exist among bases and their sources of supply for an LRU family that consists of LRU A and SRU A1.

In Figure 3-1, the solid and dashed lines represent the main source of supply (MSOS) for each item. The solid lines identify the main sources of supply for LRU A; the dashed lines show the MSOS for SRU A1. Depot number 1 (D1) is the overseas depot with name OS1; D2 is the designation for an ALC depot that can also repair these items; and D3 is the other overseas depot (OS2). Note that Depot D3 cannot




FIG. 3-1. SCHEMATIC FOR SOURCES OF SUPPLY

repair SRUs and must get them from the ALC. That arrangement adds a further wrinkle: *depots could have other depots as a source of supply.*

Under normal procedures, requisitions for a specific spare are filled by the main source of supply. Thus, for example, for LRU A and SRU A1, the following conditions would apply:

- Bases B1 and B2 would normally get resupplied by Depot D1 (we emphasize again that for other LRUs and SRUs, the linkages could be different for these bases).
- Bases B3, B4, and B5 get their requisitions filled by Depot D2.
- Bases B6 and B7 get their LRU A requisitions filled by Depot D3 and their SRU A1 requisitions filled by Depot D2.
- Depot D3 gets its SRU A1 requisitions filled by Depot D2.

In the current DRIVE model, these detailed linkages are missing. All bases are assumed to be resupplied by a single source at the ALC. The effect is that the current DRIVE repair list for the ALCs in the United States is suboptimized to the extent that the overseas depots are scheduling the same items for repair during the same repair cycle. This problem must be addressed before the DRIVE system can become fully operational.

The solution to this problem depends on a critical assumption about how the Air Force wants to operate its logistics support structure in the future. That is, what constraints ought to be imposed on resupply? Who may fill a requisition and how

much flexibility is needed or desirable in the logistics system? Do we want to permit RAF Kimball to resupply *routinely and as necessary* Pacific bases that have higher priority for spare parts than the European bases and vice versa? Alternatively, do we want to prevent such "cross-fertilization" and constrain the system to a "single source of supply"?

Since the Air Force may well decide on some form of hybrid solution that combines the best features of each approach, we propose two solutions in this chapter: one is based on a single source of supply, and the other permits multiple resupply sources as necessary. They are outlined briefly below.

- *Single Source of Supply.* In this approach, we assume that all requisitions for a specific item (LRU or SRU) are made to a single source of supply *and resupplied from that source*. No cross-repair or redistribution occurs among depots.
- *Multiple Sources of Supply.* In this approach, all repair capability is assumed to be aggregated in one location. An item is repaired and distributed and then parceled out to the separate locations based on individual capacities. This decomposition of repair and distribution would be conducted in a way that minimizes the amount of "cross-fertilization." In this approach, every depot can resupply every other base or depot regardless of the flow of requisitions.

The subsequent sections of this chapter provide the following:

- *Solution 1 – Single Source of Supply.* A more detailed, macro solution based on the assumption of a single source of supply
- *Solution 2 – Multiple Sources of Supply.* Similar information for a macro solution based on the assumption of a multiple source of repair
- *Data Needs.* A definition of the additional data needs to enable the DRIVE program to solve the problem
- *The Advantages and Disadvantages of Each Solution.* An exposition of the advantages and disadvantages of each approach.

The first three of these subsections have been prepared for those who are thoroughly familiar with the current processes and design of the DRIVE model. They describe specific ways for organizing and processing data to achieve the desired effect. Those who are less familiar with these processes are provided with an assessment of the advantages and disadvantages of each approach in the final subsection.

SOLUTION 1 – SINGLE SOURCE OF SUPPLY

The current DRIVE model calculates a sort value for each spare in an LRU/SRU family. The sort value is the relative increase in the probability of meeting the target aircraft availability for this LRU/SRU at a particular base divided by the cost of providing this capability (measured in repair hours). This sort value considers the asset posture of both the base and the repair facility (depot). Since we have more than one repair depot and since one repair depot can be a source of supply for another repair depot, three changes are required to the current solution:

- An LRU or SRU sort value for a base must consider the supply posture of its main source of supply.¹
- If a base's MSOS for an LRU is a depot (say, Depot 1) that cannot repair SRUs but must get them from another depot (say, Depot 2), we need to know which depot is the MSOS for the SRU that is needed by Depot 1.
- The sort list must be able to make allocations to other depots as well as to the bases. The following expands on this new requirement.

The sort list is that group of records produced by the main DRIVE program that calculates the sort values for each LRU family. For the purposes of this discussion, we define some of the elements for each record of the sort list as follows:

- *ALLOC* – *allocation index*. This is the number of the allocation and is only used here to group the actions necessary to implement a given allocation (records with the same allocation number will also have the same sort value).
- *NSN* – *national stock number*.
- *IND* – *indenture level*. 1 for LRU; 2 for SRU.
- *DEST* – *destination of the item in that record*. This is the code for a base or a depot.

¹The current treatment of common SRUs requires that the assets at each source of supply be "allocated" to each parent LRU that uses the SRU. That allocation is based on relative demands for the SRU generated by the parent LRUs. If SRU A1 is common to two LRUs (LRU A and LRU B), each with the same demands and quantity per application (QPA) for SRU A1, then any serviceable assets of SRU A1 will be divided 50/50 between LRUs A and B. This division of serviceable assets would be poor if LRU A is in good supply and LRU B needs all the serviceable assets of the SRU. Since the current model works one LRU family at a time, we cannot know a priori the relative importance of these LRUs. The result is that we may well be misallocating SRU repair capability – a deficiency in the current model that is eliminated with the LMI repair-only model (see Chapter 6).

- **MSOS** – *main source of supply code for the item*. This is normally a depot location. (If an ALC is the source of supply, it can be viewed as the central point for organic and contractor repair. Total repairs for this source could then be distributed between organic and contractor repair.)
- **SHOP CODE** – *shop code*. This entry, in conjunction with the MSOS, will uniquely define the source of repair. For example, a particular type of repair shop (say, SH1) may exist at two depots with different capacities. However, Shop Type SH1 at Depot D1 is a distinctly different shop than Shop SH1 at Depot D2. This shop code, in conjunction with the MSOS code, will constitute unique shop identifiers. In Table 3-1, "Svc" in the shop code column means that there is a serviceable item available for distribution to the indicated place.
- **SORT VALUE** – *sort value*. This is the relative increase in the probability of achieving the aircraft availability objective for a base divided by the repair cost to provide this capability.

TABLE 3-1
EXEMPLARY SORT LIST

ALLOC	NSN	IND	DEST	MSOS	SHOP CODE	SORT VALUE
1	SRU A1	2	B6	D2	Svc	60.0
2	LRU A	1	B6	D3	SH1	59.0
2	SRU A1	2	D3	D2	SH2	59.0
3	LRU A	1	B1	D1	SH1	58.0
3	SRU A1	2	D1	D1	Svc	58.0
4	SRU A1	2	B2	D1	Svc	57.0
5	LRU A	1	B3	D2	Svc	56.0
6	LRU A	1	B4	D2	SH1	55.0
6	SRU A1	2	D3	D2	SH2	54.0
7	SRU A1	2	B2	D1	Svc	53.0
8	SRU A1	2	B5	D2	DH2	52.0

Table 3-1 is an example of a sort list for an LRU family that consists of LRU A and SRU A1. In addition to using an identical allocation number to identify actions associated with the same allocation, we use a blank line to separate these entries for

clarity. The base/source-of-supply linkages for this item are those defined in Figure 3-1.

Referring to Figure 3-1 and Table 3-1, we see that the first allocation sends a serviceable SRU from Depot D2 to Base B6. It either supplies an SRU for an LRU that is awaiting parts (AWP) or reduces projected backorders of SRUs, which in turn affects LRU availability. In this allocation, the sort value did not use any repair cost since a serviceable item was available at D2.

In Allocation 2, we repair LRU A at Depot D3 in Shop SH1 and send it to Base B6. However, Depot D3 needs an SRU from Depot D2 (the third entry in Table 3-1). Since Depot D2 no longer has serviceable items (it sent its last one to Base B6 in Allocation 1),² an SRU is inducted for repair at Depot D2 in Shop SH2 *but it must be sent to Depot D3*.

- Note that Depot D1 has some serviceable SRU A1 items (see Allocations 3, 4, and 7 in Table 3-1), but they cannot be used since we are in the single source of repair mode.
- Similarly, Depot D2 has a serviceable LRU A item (see Allocation 5 where the serviceable item goes to Base B3) but cannot send it to Base B6 for the same reason.

The main difference between the above approach and that of the current DRIVE program is that in the former we needed additional information to know who is responsible for filling SRU needs for LRU A. In the current program, only one source of supply would be available for SRU A1. In Allocation 2 (Table 3-1), we had to know that LRUs fixed at Depot D3 get their SRUs from Depot D2. Thus, when repairing LRUs at the depot, we can no longer use a single depot record for the LRU indenture structure. Each of the base records for the LRU family must identify the main source of supply for LRUs and SRUs. Only in this way can we know what to put in the *DEST*, *MSOS*, and *SHOP CODE* columns for the SRU A1 needs for Allocation 2 (third entry in Table 3-1).

The postprocessor will then prepare a repair and distribution list for each depot on the basis of its repair shop capability and available carcasses. This requires a unique shop repair code for each depot that can repair the same item. For example, in

²Other SRU A1 items may be available, but they may be allocated to LRU B (see Footnote 1).

the above case we would workload the following unique shops: D1-SH1, D1-SH2, D2-SH1, D2-SH2, and D3-SH1.

SOLUTION 2 – MULTIPLE SOURCES OF SUPPLY

The second solution permits resupply from any source. Thus, we begin by consolidating all assets from the separate depots into one location. Pseudo-repair shops are defined as the sum of individual repair shops at all depots that can repair the item. Sort values are then calculated as done today for each LRU/SRU family; in this case, however, an MSOS and a pseudo-source of supply for the LRU/SRU are indicated. The workload is assigned in the postprocessor described below using three passes through the merged repair and distribution list as follows:

- *Pass 1.* In the first pass through the merged "repair and distribution" output file, assign the workload for all the pseudo-repair shops as though all repair was collocated. Serviceable assets going to a base are directed from the MSOS until its assets are exhausted; after that, serviceable assets are redistributed (done in the third pass).
- *Pass 2.* In the second pass, serviceable assets used to repair LRUs at the depots are treated at the same time the workload is directed to the shops as follows:
 - ▶ Decompose the workloads for the pseudo-shops into workloads for specific shops based on the MSOS for each base/item combination; if a repair is required for a base (B1, for example) and repair capacity at its MSOS (Depot D1) is used up, we cross-repair between depots. *Write this record to a new output file, which will be processed in a third pass.*
 - ▶ In the above process, when we schedule a repair of an LRU at a specific depot, we need to ensure that sufficient SRUs are available for that repair.
 - If a serviceable asset was available when the sort value was calculated, we would first check the assets at the MSOS. If none remain, we need to provide the asset from another depot. *In this case, we write the information to another file.* In the third pass, we will see which depot has the serviceable assets and direct them to this depot.
 - If a serviceable SRU asset was not available and one needed to be repaired, we would first try to assign the workload to the repair shop for the main source of repair. If we have exceeded capacity for the main source of repair, we need to repair an SRU at one depot and send it to another depot. *Write this information to another file and process it in the third pass.*

- **Pass 3.** In the third pass, we take two actions:
 - ▶ Redistribute the remaining serviceable assets that were not yet distributed above
 - ▶ Finish assigning the workload to those shops to which no workload has been assigned and redistribute these repaired assets to the appropriate depot or base.

DATA NEEDS

The two solutions – those involving a single source of supply and multiple sources of supply – require the following additional information on the supplies and maintenance capacity at each of the overseas depots.

Unique Repair Shop Designations/Capacities

For each depot, we need to develop unique repair shop codes. For the multiple-sources-of-supply solution, we need to develop pseudo-repair-shop codes that link all similar depot repair shops. For example, assume

- Shop D1-SH1 at Depot D1 can repair LRUs A and B from Bases B1 and B2;
- Shop D2-SH1 at Depot D2 can repair LRUs A, B, and C from Bases B3, B4, and B5 and LRU C from Bases B1 and B2; and
- Shop D3-SH1 at Depot D3 can repair LRUs A and C for Bases B6 and B7, and LRU B from Bases B6 and B7 cannot be repaired at Depot D3 and must go to Depot D2.

In that case, these three repair shops must be grouped into a single pseudo-repair shop so that we can consider redistribution of repair among depots for Solution 2.

LRU/SRU Source of Supply

For each base, we must define the MSOS for each LRU and SRU. We also need to specify which depot provides SRU support to the LRU repair shops.

Base/Depot Linkages

The average order-and-shipping time (OST) from each base to each depot is required to determine the appropriate time horizon and repair cycle period for the depots. For example, if the ALC OST is 23 days to a particular base but the OST from an overseas depot to that base is 5 days, then the repair cycle period for the overseas

depot must be 18 days longer than the repair cycle time used at the depot. This adjustment is required because we are trying to determine the optimum asset posture for each base at a specific supply target point in time (see Figure 3-2), which is assumed to be the beginning of the war period.




FIG. 3-2. RELATIONSHIP BETWEEN ALC AND OVERSEAS REPAIR CYCLES

Since this calculation of repair cycle time can be done for each base, the repair cycle time used in the model would be the largest of these values across all bases linked to the source of supply.

Depot Serviceable Assets/Due-ins/Awaiting Parts

We need serviceable asset data at all the depots. Those data include information on serviceable SRUs that may be in the maintenance inventory center (MIC). We also need information on due-ins at the overseas depots (reparables that have been inducted into the maintenance depot but have not yet been repaired) and the number of these due-ins that are AWP at the depots at the same level of detail as provided for the ALC depots.

THE ADVANTAGES AND DISADVANTAGES OF EACH SOLUTION

The solutions proposed here intentionally address opposite ends of a spectrum: no cross-fertilization among depots (Solution 1) and unrestricted cross-fertilization as necessary (Solution 2). As noted earlier, some hybrid solution is surely to evolve. The discussion presented in this section is intended to help HQ AFLC decide how far it wants to go in either direction.

For one thing, the multiple-sources-of-supply approach requires that the objective functions used in the DRIVE model capture adequately the relative priorities

among theaters because it may well move assets from one theater to another based on what the model sees as the priorities for spare parts.

The single-source-of-supply approach (Solution 1) would maintain the integrity of repair and supply at each depot with no redistribution. It is more decentralized than Solution 2.

On the other hand, if we have a good representation of theater priorities in DRIVE (and we believe that is possible), Solution 2, the centralized approach, would more effectively meet the Air Force needs since it takes the view that all the depots support the whole Air Force. It takes maximum advantage of all supply and repair capability through limited redistribution as necessary. It is more complex in terms of postprocessing but would optimize the utilization of all repair capability and better meet availability objectives across the Air Force. Admittedly, such a capability would come at the cost of some redistribution. However, since we already have a capability to resupply bases by air, we ought to use it to its greatest advantage.

In both solutions, the workload would be assigned at the headquarters and passed down to the individual depots. That procedure could present a problem if the depots are required to meet these goals without question. Policies and procedures would be required to permit some decentralized control over repairs at the overseas depots. These procedures would need to be crafted carefully so that they give flexibility to the overseas depots on the one hand and have enough constraints so that the theaters cannot thwart "essential" redistribution on the other.

CHAPTER 4

OPTIMAL ALLOCATION OF SERVICEABLE ASSETS TO BASES TO SATISFY PEACETIME AND WARTIME TARGETS

THE PROBLEM

The current DRIVE model calculates a sort value based on a single objective function that combines the peacetime and wartime demands for bases with a wartime mission and includes only the peacetime demands for bases without a wartime mission. The use of a single objective function for peacetime and wartime raises some questions about what "availability targets" to use in the current DRIVE algorithm to maintain the "proper" balance between peacetime and wartime readiness.

The HQ AFLC/XPS staff described to us the current procedures used in the DRIVE model to set base availability targets to balance peacetime and wartime objectives. DRIVE attempts to maximize the following:

- The probability of satisfying a peacetime availability goal of 100 percent at the end of the peacetime planning horizon
- The probability of satisfying a wartime availability goal of 85 percent at the end of the wartime planning horizon.

Typically, the wartime goal is more difficult to meet (it has larger demand rates and a longer time period); thus, DRIVE was originally configured to use the wartime objective as its sole goal for bases with both peacetime and wartime missions.

RAND discovered that when an item has a large quantity per application (QPA),¹ the DRIVE model does not order enough of the items to be repaired to do a good job in peacetime since the number of allowable holes in wartime is large (allowable wartime holes = $0.15 \times \text{fleet size} \times \text{QPA}$). Consequently, RAND suggested that before running DRIVE, the model compute the following:

- The probability of meeting the peacetime mission [$P(p)$] with assets that would be available at the end of the peacetime time horizon if no other assets are sent

¹The QPA is the number of LRUs on an aircraft or SRUs on an LRU.

- The probability of meeting the wartime mission $[P(w)]$ with expected assets at the end of the wartime period.

If $P(w) > P(p)$, then the allowable holes for wartime are reduced until $P(w) \leq P(p)$.

In effect, this procedure amounts to raising the wartime target from 85 percent to something larger on a particular item at a particular base. It ensures that the peacetime mission will not be ignored, but as more stock is added, the allocation becomes inappropriate. To understand why, consider the following example.

Suppose that the peacetime pipeline is 2 and the wartime pipeline is 10 (the wartime pipeline typically is larger). Let the wartime allowance for holes on this item be 9, and assume that $P(w) > P(p)$. Further assume that if the allowable wartime holes are reduced to 6, then $P(w) \leq P(p)$. Now, as we add stock to this base, its probability of meeting the peacetime mission is increasing much more rapidly than the probability of meeting its wartime mission (because the peacetime pipeline is smaller).

Thus, this procedure puts too much emphasis on the peacetime mission at first (because it has reduced the allowable wartime holes from 9 to 6) and then puts too much emphasis on the wartime mission as we add stock. It is much like a hunter shooting ahead of ducks half of the time and then behind ducks the other half in hopes that on average he will hit the ducks.

We need a procedure for setting the target for a base with no wartime mission. Consider a numerical example consistent with the data assumed above for one item:

- The peacetime mean over the planning horizon is 2 and the wartime mean is 10.
- Demand probabilities are Poisson.
- We have no stock.
- The peacetime allowance for holes is 0.
- The wartime allowance for holes is 9.

Table 4-1 shows relevant data for mean demands of 2 and will be used in discussions about the peacetime demands and objective function. Table 4-2 is the counterpart of Table 4-1 but for mean demands of 10 to be used in addressing the wartime objective function. The columns of each table are described next.

In the tables, s is the number of demands over the period and is shown in Column 1. Column 2 [labeled $Pr(s)$] is the probability of having s or fewer demands over the period. It can also be interpreted as the probability of j or fewer backorders with a supply of $s - j$ spares. For example, $Pr(6)$ can represent any of the following probabilities:

- Having 0 or fewer backorders with a stock level of 6
- Having 2 or fewer backorders with a stock of 4
- Having 6 or fewer backorders with a stock of 0.

Column 3 is the change in $Pr(s)$ as the demands increase by 1. Column 4 is the logarithm of $P(s + 1)/P(s)$. It shows the sort value used in the DRIVE allocation where the cost is 1. The sort value in Column 4 is the benefit achieved when allocating the $s + 1$ spare.

From Table 4-1 we see that the probability of meeting the peacetime objective (no holes) with 0 stock is 0.1353 (first row with $s = 0$). On the other hand, Table 4-2 shows that the probability of meeting the wartime mission (9 holes or less) with 0 stock is 0.4579.

TABLE 4-1

PROBABILITY TABLE FOR MEAN DEMANDS OF 2 FOR A PEACETIME OBJECTIVE

s	$Pr(s)$	$\Delta Pr(s)$	$\text{Log}[Pr(s + 1)/Pr(s)]$
0	0.1353	0.1353	0.4771
1	0.4060	0.2707	0.2218
2	0.6767	0.2707	0.1027
3	0.8571	0.1804	0.0435
4	0.9473	0.0902	0.0162
5	0.9834	0.0361	0.0053
6	0.9955	0.0120	0.0015
7	0.9989	0.0034	0.0004
8	0.9998	0.0009	0.0001
9	1.0000	0.0002	0.0000

TABLE 4-2

PROBABILITY TABLE FOR MEAN DEMANDS OF 10 FOR A WARTIME OBJECTIVE

s	Pr(s)	$\Delta \text{Pr}(s)$	$\text{Log}[\text{Pr}(s + 1)/\text{Pr}(s)]$
0	0.0000 ^a	0.0000 ^a	1.0414
1	0.0005	0.0005	0.7439
2	0.0028	0.0023	0.5720
3	0.0103	0.0076	0.4518
4	0.0293	0.0189	0.3605
5	0.0671	0.0378	0.2878
6	0.1301	0.0631	0.2284
7	0.2202	0.0901	0.1794
8	0.3328	0.1126	0.1386
9	0.4579	0.1251	0.1049
10	0.5830	0.1251	0.0774
11	0.6968	0.1137	0.0554
12	0.7916	0.0948	0.0383
13	0.8645	0.0729	0.0254
14	0.9165	0.0521	0.0161
15	0.9513	0.0347	0.0098
16	0.9730	0.0217	0.0057
17	0.9857	0.0128	0.0031

^a This number is 0 to this level of significance. Its real value is $e^{-10} = 0.0000454$.

The current DRIVE procedure tells us to reduce the wartime allowance for holes so that the probability of meeting the wartime objective [$P(w)$] with the starting asset position (0 in this case) is at or below the probability of meeting the peacetime objective (0.1351). Since the second column of Table 4-2 can be interpreted as the probability of having s or fewer backorders (holes) with a stock of 0, we see that $Pr(s)$ is less than or equal to 0.1351 when $s = 6$ [$P(w) = Pr(6) = 0.1301$].

The current DRIVE focuses on this wartime probability and the relative change in this wartime probability as we repair more of this item. With an adjusted wartime objective that would require 6 or fewer backorders, the DRIVE model will use the fourth column of Table 4-2 as a sort value starting with Row 7 where $s = 6$. It would use a sort value of 0.2284 for the first spare, 0.1794 for the second, 0.1386 for the

third, etc. This sequence of sort values is listed in Column 4 of Table 4-3 with the heading "adjusted wartime."

TABLE 4-3
SORT VALUES BASED ON LARGER OF PEACETIME OR WARTIME

Spare number	Actual peacetime	Actual wartime	Adjusted wartime	Larger of actual peacetime or actual wartime
1	0.4771	0.1049	0.2284	0.4771 P
2	0.2218	0.0774	0.1794	0.2218 P
3	0.1027	0.0554	0.1386	0.1027 P
4	0.0435	0.0383	0.1049	0.0435 P
5	0.0162	0.0254	0.0774	0.0254 W
6	0.0053	0.0161	0.0554	0.0161 W
7	0.0015	0.0098	0.0383	0.0098 W
8	0.0004	0.0057	0.0254	0.0057 W
9	0.0001	0.0031	0.0161	0.0031 W
10	0.0000	0.0016	0.0098	0.0016 W

If one wanted to allocate based on the objective function that gave the greater sort value for peacetime or wartime, then the sort values would be those in the fifth column of Table 4-3. A "P" or "W" is appended to the value in the last column to indicate the objective function that gives the largest sort value. The second and third columns of Table 4-3 are the actual sort values for each additional spare against the actual peacetime allowance of 0 (Column 2) and a wartime allowance of 9 (Column 3).

By comparing the last two columns of Table 4-3, we see that the DRIVE model, which uses the adjusted wartime values, understates the return for Spares 1 and 2 relative to the peacetime objective and overstates the value of each spare thereafter. Even though the RAND "fix" makes repair of this item more attractive for the first two spares using sort values of 0.2284 and 0.1794 (adjusted wartime), respectively, instead of 0.1049 and 0.0774 (actual wartime), it still understates the real returns for the peacetime objective functions, which are 0.4771 and 0.2218. On the other hand, the reverse happens with Spares 3 and 4 with a bias in favor of the peacetime objective: the adjusted wartime values are 0.1386 and 0.1049, which are greater than

both the actual peacetime and the actual wartime. In fact, this bias continues to the point that when Spare 7 is allocated, the model thinks the probability of wartime mission accomplishment is 0.8645 (line for $s = 13$ in Table 4-2) but it really is 0.9730 (line for $s = 16$ in Table 4-2).

This problem becomes even more important under changes now being considered in War Readiness Spares Kit (WRSK) computations. More than one wartime requirement may have to be considered; e.g., a Day 7 requirement and a Day 30 requirement. Thus, DRIVE may need to look at three different objective functions simultaneously.

A PROPOSED SOLUTION

In the preceding discussion, we seemed to suggest that the solution to the peacetime and wartime balance problem lies in calculating multiple objective functions and then selecting the largest of the group. Such an approach is certainly better than the current solution for one item as we demonstrated in the previous section. However, when we introduce more than one item, selection of the larger of two objective functions could ignore an item that is second best when looking across items. The following notional example demonstrates this point.

Consider the data in Table 4-4. There we have three items — A, B, and C — that already have a certain number of spares allocated. Item A has N spares, Item B has M spares, and Item C has I spares. The table identifies the notional returns for peacetime, wartime, and the sum of peacetime and wartime when allocating one additional spare for each item. For example:

- The $(N + 1)$ spare of Item A will yield 5 units for wartime, 20 units for peacetime, and 25 units for peacetime and wartime (the sum of peacetime and wartime assumes equal weighting for peacetime and wartime, which could be varied depending on the operational priorities at the time).
- After the $(N + 1)$ spare is allocated, the sort values for the $(N + 2)$ spare of Item A is 4 for wartime, 18 for peacetime, and 22 for peacetime and wartime.

We now assume that we can only repair two items in the next DRIVE period. Which two should we repair?

If we base our allocation on the largest of the returns for peacetime or wartime, we would first repair one of Item A (based on a return of 20 for peacetime). After that

TABLE 4-4

A NOTIONAL EXAMPLE OF SORT VALUES FOR THE ALLOCATION OF THE NEXT SPARE

Item	Numbered spare	Notional return		
		Wartime	Peacetime	Wartime and peacetime
A	$N + 1$	5	20	25
	$N + 2$	4	18	22
B	$M + 1$	19	5	24
	$M + 2$	18	4	22
C	$I + 1$	15	15	30
	$I + 2$	14	14	28

one is allocated, the next largest peacetime or wartime allocation comes from repairing one of Item B with a return of 19 for wartime. The sum of the returns for these two allocations would be a total of 49 units: 25 for ($N + 1$) spare of Item A (5 for peacetime and 20 for wartime) and 24 for the ($M + 1$) spare of Item B.

If, on the other hand, we set priorities for repair on the basis of the sum of the two objective functions (the last column), the first repair would be for the ($I + 1$) spare for Item C with a combined return of 30 units. After that allocation, the second repair would be for the ($I + 2$) spare for Item C with a return of 28. The combined return in that case would be 58 units as compared with the 49 units obtained by selecting the larger of the two objective functions.

Therefore, we propose that the DRIVE program use a weighted sort value that reflects the relative priorities for peacetime and wartime. In that case, we could assign weighting factors for peacetime and wartime either globally or by base.

- In a relatively benign environment, we may want to distribute "unavailability" across the force uniformly except at selected forward-deployed bases.
- During a crisis, we can change the relative weights for peacetime and wartime by base, giving higher relative weight for wartime to the war bases and higher peacetime weight to the CONUS training bases.

CHAPTER 5

ALTERNATIVE OBJECTIVE FUNCTIONS

INTRODUCTION

The search for alternative objective functions for DRIVE is principally motivated by the need to repair and distribute spare parts in a way that recognizes the operational priorities for available aircraft across the entire Air Force. This need requires that we provide available aircraft to units in a priority sequence that reflects the operational needs of the force. This need implicitly means that one may want to get one more additional aircraft at Unit x instead of at Unit y even though it may cost less to get one more available aircraft at Unit y .

This kind of a priority is difficult to model in the traditional ways. Furthermore, the analytic community has not yet developed a method for translating the commander's operational priorities into the appropriate availability targets for the current DRIVE model. In fact, the analyses presented in this chapter raise serious concerns about the utility of the current DRIVE objective function to meet tomorrow's operational priorities.

In the first place, the current operations priorities are no longer defined in terms of a probability of having a minimum number of aircraft down — the current DRIVE objective function. Air Force WRSK calculations are now based on expected aircraft availability; thus, operations priorities should be defined in terms of expected availability, and to do so, we need to move closer to maximizing aircraft availability — a criterion that is not now embedded in the current DRIVE model.

Clearly, we ought not abandon the current objective function unless we have a suitable replacement or at least one that is more robust. (A more robust objective function and allocation process is one that gives better results more often than not.)

This chapter explores the utility of the current objective function and process in meeting aircraft availability goals and introduces alternative objectives and processes that offer greater potential for meeting operational priorities.

PROBLEMS WITH THE CURRENT PROCESS

The current DRIVE objective function does not optimize aircraft availability; nobody ever claimed it did. It maximizes the probability of having no more than a specified number of aircraft down. We also find that the current objective function is biased in favor of small air bases. That bias leads to lower overall aircraft availability across the force than would be the case if the bias were removed. Why, then, should we use it?

The simple answer is that it is a holdover from the way WRSK calculations were done in the past and, quite frankly, it is a vast improvement over earlier methods used to set priorities for repair. It is also a convenient objective function when full cannibalization is permitted. In the full-cannibalization case, the current objective function has a mathematical form that makes it separable in the sense that one could determine the optimum sort value for one LRU without knowing the supply availability of other LRUs.¹

The current DRIVE model is designed around having a separable objective function. It processes one LRU family at a time, developing a sort value for each possible allocation. The LRU and SRU lists are sorted in a postprocessor, and those items that give the biggest "bang per buck" are repaired and distributed to the units.

The aircraft availability objective function for full cannibalization is not separable. Consequently, it cannot be accommodated in the current DRIVE program without a significant redesign of the entire model. However, the following analyses suggest the existence of more robust alternative objective functions and processes that provide better results more often than not.

LMI DRIVE-LIKE MODEL ASSUMPTIONS AND DEFINITIONS

LMI developed a DRIVE-like model to perform the analyses presented in the next section. This section describes this model and its assumptions. It also introduces the definitions of alternative allocation methods addressed in our analyses.

¹The aircraft availability objective function used in the Aircraft Availability Model is also separable in the case in which no cannibalization is permitted. That case is not considered here nor in the DRIVE model because in wartime, cannibalization is assumed to be a fact of life. A no-cannibalization logic is judged to be unrealistic for most LRUs in the WRSK.

With the exceptions noted below, the LMI model is essentially the same as the most recent version of the Ogden prototype DRIVE model. It uses the same input data file (PPOUT.DAT) and sets priorities for the repair and distribution of spares at each base. The key differences are:

- We do not consider any SRUs in this analysis; thus, the analysis ignores the LRUs that are AWP at the bases.²
- No LRU carcass constraints are imposed.
- All LRUs are repaired in a single shop with a repair capacity equal to the "keep-up" requirement, i.e., the capacity needed to repair everything that breaks during a period.
- The model is designed to run for six Biweekly periods. At the end of each period, base asset postures are adjusted by
 - ▶ Adding those repairs allocated to that base in the previous period
 - ▶ Subtracting expected demands for the previous period.
- The model calculates the number of expected not mission capable-supply (ENMCS) aircraft at each base that results from the repair and distributions made by the model.
- The ENMCS calculations for the expected number of aircraft "down" assumes that all 41 LRUs are on all the aircraft at each base. We do not distribute the expected backorders (EBOs) for each LRU to each aircraft in proportion to the usage for that aircraft.³

In addition to the above, the LMI model has the capability to examine the improvements in aircraft availability that result from:

- The use of alternative objective functions
- Iterative selection of base-specific readiness target values required as input to the models.

The current DRIVE model and the LMI DRIVE-*like* model require that the user enter as an input for each base the number of aircraft that are allowed to be not

²We do not see this as a serious constraint for an analysis that is attempting to understand the physics of the allocation problem. We hypothesize that the treatment of SRUs would have essentially made more LRUs available and improved aircraft availability across the board in roughly the same proportion as presented in the analysis.

³This assumption has the effect of understating slightly the total number of down aircraft resulting from this set of LRUs and could therefore lead to an underestimate of the benefits for those alternatives that reduce total ENMCS.

mission capable-supply (NMCS). The number of allowable NMCS aircraft is referred to as the model's CAT value (cannibalization threshold). It has been traditionally set equal to the difference between the primary aircraft authorized (PAA) and the direct support objective (DSO) for a unit. If a base has more than one unit, the input CAT value to the HQ AFLC DRIVE model for that base is the sum of differences between the PAA and DSO for each unit at the base.

With this definition of terms, it is important not to confuse a CAT value with an implicit DSO. If, for example, a unit requires a DSO of 23 out of 24 aircraft on Day 7 of the war, then one may wish to use a CAT value of 1 as the target value input to the DRIVE model for that base (this base has only one unit assigned to it). On the other hand, one may wish to raise or lower the input CAT values to the DRIVE model to give one unit (base) with a DSO higher priority than another unit (base) with the same DSO. If, in this situation, we get the desired effect by giving one unit (base) a target input CAT value of 4 and the other unit a CAT of 0, *we have NOT changed the DSO for those units*. The DSOs remains unchanged; they are set by operational requirements. We have merely changed an input value to a model that is trying to achieve some specific objective.

Against this backdrop, we now proceed to discuss the alternative objective functions used in this report. The analyses presented consider the current method (Method 1) and three alternative methods for calculating a sort value. They are defined as follows:

- *Method 1.* The sort value (SV) for Method 1 is based on maximizing the probability of having only CAT aircraft NMCS. The sort value for Method 1 is the change in the logarithm of P divided by the "cost to repair" the LRU, where P is the probability of having no more than CAT aircraft down because of this LRU.

$$SV 1 = \frac{\Delta \log(P)}{\text{repair cost}} \quad [\text{Eq. 5-1}]$$

- *Method 2.* The sort value for Method 2 is based on trying to equalize the probability (P) across all bases; therefore, it *tends* to allocate to the base with the lowest value for P . It was motivated by the realization that the Method 1 sort value is biased toward small bases. When a zero CAT value is used in Equation 5-2, it becomes the reduction in EBOs per repair cost.

$$SV 2 = \frac{(1-P)}{\text{repair cost}} \quad [\text{Eq. 5-2}]$$

- *Method 3.* This sort value is approximately equal to the expected number of aircraft down because of this LRU and is given by the following equation:

$$SV 3 = \max\left(\frac{EBO}{QPA}\right) \quad [\text{Eq. 5-3}]$$

where QPA is the number of LRUs of this type per aircraft. This method reduces the backorders of the LRU that is expected to contribute most to the unavailability of the aircraft regardless of the repair cost to achieve this objective.

- *Method 4:* This method is the same as Method 3 except the sort value is divided by the number of aircraft (NAC) at the base and serves as a proxy for allocating to the base with the lowest fraction of available aircraft.

$$SV 4 = \max\left(\frac{EBO}{QPA \times NAC}\right) \quad [\text{Eq. 5-4}]$$

In addition to the above methods for setting sort values, our analysis showed that it is also possible to get improved aircraft availability across all bases if the CAT values for each base are chosen appropriately. We refer to this set of CATs as "OPT CATs" or optimum CATs. The OPT CATs were found by running the model iteratively:

- At the end of each run, we calculated the EBOs for each item at each base.
- The CAT value for a specific base was then reset to the largest of all the LRU EBOs at that base and the model was rerun with the new CAT values.
- The process ended when the CAT values stopped changing.

Note that the selection of OPT CATs applies only to Methods 1 and 2 since Methods 3 and 4 are independent of CAT and repair cost.

The rationale for this procedure is to avoid repairing items that do not improve overall aircraft availability. Thus, after one run of the model, we look at each base and determine which LRU is contributing most to unavailability. That LRU is the one with the largest EBOs. Thus, we know the total number of down aircraft at this base will be no smaller than the number of backorders. By setting the allowable

number of down aircraft equal to this minimum number, we essentially adjust the allocation process by optimizing in the neighborhood of what is possible with the given number of spares in the system and not necessarily where we would like to be.

For example, the original input target CAT values are normally based on war-time requirements for available aircraft that are, in turn, used to size the WRSK. If the Air Force has not fully funded all its WRSK requirements, which is normally the case, an unconstrained input target CAT value consistent with the WRSK requirement will automatically force us to optimize on the wrong part of the curve. Put another way, we need to know, a priori, how much availability is possible in the system given our prior investment decisions and then we must develop a scheme for distributing that availability around the force. The process outlined above is an attempt to address the question by using the models iteratively.

RESULTS OF ANALYSIS

We compared alternative methods for repairing and distributing spares in three dimensions:

- We first examined changes in aircraft availability at each base and in the aggregate across all bases.
- We then compared the repair lists that resulted from the allocations to determine the extent to which the ENMCS changes were mostly a distribution problem. The extent to which they are such a problem would be manifested by large differences in total ENMCS and only small changes in the by-item repair quantity.
- Finally, we compared the total backorders for each alternative — a measure often used to assess the relative extent of cannibalization.

We present the results of our analysis in the following subsections:

- *A Comparison of Methods.* This subsection compares the by-base and total ENMCS and repair lists for Methods 1 through 4 using the D + 30 ENMCS goals⁴ for inputs in Methods 1 and 2.
- *The Effect of Selecting Optimum CAT Values.* This subsection compares ENMCS and repair lists when optimum input target values are used for Methods 1 and 2.

⁴In this analysis, we used the current ENMCS goals of 15 percent of PAA.

- *The Effect of Redistribution.* This subsection demonstrates the effect that redistribution has on the ability of the alternative allocation processes to minimize total ENMCS and meet individual base ENMCS goals.
- *An Assessment of Backorders.* This subsection gives an overview of how total backorders vary as a function of method with and without redistribution.

A Comparison of Methods

Methods 1 and 2

Figure 5-1 shows the ENMCS by base that results when allocations for repair and distribution are based on Method 1 (current DRIVE algorithm) using the same CAT values currently used in the HQ AFLC DRIVE model. These values equate to 85 percent of PAA for bases with a wartime mission (wartime bases) and 0 percent for those with no wartime mission.




FIG. 5-1. ABILITY OF METHOD 1 (HQ CATs) TO MEET ENMCS GOALS WITH NO REDISTRIBUTION

These input CAT values are referenced in two ways and are shown by the monotonically decreasing solid line in Figure 5-1:

- As the HQ CAT values when we talk about inputs for use with Methods 1 and 2
- As individual base ENMCS goals when we compare the actual ENMCS achieved by a specific allocation method.

In Figure 5-1, the bases are sorted on the ENMCS goals for each base essentially sorting the bases by the largest base with a wartime mission. The large wartime bases have the lower base number and a high ENMCS goal (Base 1 has an ENMCS goal of 20). Bases numbered 15 through 32 are the smaller wartime bases; bases 33 through 40 are the peacetime bases with an ENMCS goal of 0.

Figure 5-1 shows the ENMCS that result at each base when the current DRIVE logic and input data are used to set priorities for the repair and distribution of 41 LRUs to 40 Air Force bases. It shows that all of the larger bases, except Base 6, have fewer down aircraft than would be permitted by the base ENMCS goals. On the other hand, all the smaller bases have more down aircraft than specified by the ENMCS goals. Thus, with the exception of Base 6, the larger bases meet their goals and the smaller bases do not.

Figure 5-2 is a plot of the cumulative ENMCS across all these bases; the total ENMCS is 171.7 aircraft.

Figures 5-3 and 5-4 are the same as Figures 5-1 and 5-2 except we have added curves to represent the allocations resulting from using Method 2 with the HQ CAT values as target ENMCS. Figure 5-4 shows that the cumulative ENMCS for Method 2 (165.2) is smaller than that for Method 1 (171.7); however, the new objective function has a bias in favor of larger bases, which is shown in Figure 5-3.

Figure 5-3 shows the curve for Method 2 at or below the curve for Method 1 for the first 21 bases and only slightly above the Method 1 curve for the remaining bases. That is, Method 2 improved availability at the large bases but did not significantly detract from the smaller bases.

Figure 5-5 compares the repair lists for Methods 1 and 2. In that figure and all the figures that show repair lists, the LRUs were given numbers based on a descending sort on the numbers of LRUs repaired over a 90-day period for one of the




FIG. 5-2. CUMULATIVE ENMCS FOR METHOD 1 USING ENMCS GOALS AS CATs WITH NO REDISTRIBUTION




FIG. 5-3. ABILITY OF METHODS 1 AND 2 (HQ CATs) TO MEET ENMCS GOALS WITH NO REDISTRIBUTION




FIG. 5-4. CUMULATIVE ENMCS FOR METHODS 1 AND 2 WITH ENMCS GOALS AS CATs WITH NO REDISTRIBUTION




FIG. 5-5. REPAIR QUANTITY FOR METHODS 1 AND 2 (HQ CATs) WITH NO REDISTRIBUTION

cases. The only purpose of this kind of sort is to make the graphs easier to read with the larger numbers on the left and smaller numbers generally on the right.

As a further aid in assessing the extent of commonality between two different repair lists, we have developed a repair commonality index (RCI). The RCI is a percentage between 0 and 100 with 0 percent meaning that the two lists have no repairs in common and 100 percent meaning that both repair lists are identical. Appendix A provides the mathematical definition for RCI and presents Table A-1, which summarizes the RCI values for a number of combinations of repair lists generated by our analyses. Many of the numbers in Table A-1 are referenced in the body of this report.

The RCI for the two repair lists plotted in Figure 5-5 is 90.2 percent, which indicates that these two lists have about 90 percent of their repairs in common. Given the small changes in the ENMCS for these two cases, the similarity of these repair lists is no surprise.

Methods 3 and 4

Figure 5-6 is the analog of Figure 5-3. It compares the by-base ENMCS goals with those resulting from allocations using Methods 3 and 4. Method 3 allocates items to the LRU/base combination with the largest number of aircraft down because of this LRU, and Method 4 allocates items to the LRU/base combination to equalize the percentage of available aircraft at each base. Figure 5-7 adds the cumulative ENMCS information for these methods to those already plotted in Figure 5-4.

This experiment showed that Methods 3 and 4 reduce total ENMCS by between 6 percent and 12 percent (see Figure 5-7) and have a definite bias in favor of large bases and against small bases, with Method 3 showing a greater tendency in that direction (see Figure 5-6). The differences in repairs are shown in Figure 5-8. There, we see larger differences in repairs (an RCI of 81 percent) showing that some of the improvement is coming from what we repair as well as where it is distributed.

Figure 5-9 adds the repair list for Method 1 (HQ CATs) to Figure 5-8. It shows that a significant difference in repairs is generated by these processes. The RCI for Method 1 (HQ CATs) versus that for Method 3 and Method 4 is 71.2 percent and 71.0 percent, respectively.




FIG. 5-6. ABILITY OF METHODS 3 AND 4 TO MEET ENMCS GOALS WITH NO REDISTRIBUTION






FIG. 5-7. CUMULATIVE ENMCS AS A FUNCTION OF METHOD WITH NO REDISTRIBUTION



**FIG. 5-8. REPAIR QUANTITY FOR METHODS 3 AND 4
WITH NO REDISTRIBUTION**



**FIG. 5-9. REPAIR QUANTITY FOR METHODS 1, 3, AND 4
WITH NO REDISTRIBUTION**

We can compare all four methods by referring to Figures 5-3, 5-6, and 5-7. There, we can see how Methods 3 and 4 reduce total ENMCS (relative to Methods 1 and 2) by increasing slightly the ENMCS at smaller bases while reducing significantly the ENMCS at the larger bases.

The Effect of Selecting Optimum CAT Values

In this subsection, we show how the optimum (OPT) selection of input CAT values (described earlier) further improves aircraft availability across the force. Figures 5-10 and 5-11 show how the selection of OPT CATs affects ENMCS goals under Method 1 (Figure 5-10) and under Method 2 (Figure 5-11). The cumulative ENMCS at the bases is shown in Figure 5-12 where we have added these last two cases to Figure 5-7. Figure 5-10 shows that optimum selection of CAT values reduces total ENMCS by about 20 percent relative to the baseline cases that use HQ CATs (see Figure 5-4).




FIG. 5-10. IMPACT OF OPT CATs ON ENMCS METHOD 1 WITH NO REDISTRIBUTION




FIG. 5-11. IMPACT OF OPT CATs ON ENMCS METHOD 2 WITH NO REDISTRIBUTION




FIG. 5-12. CUMULATIVE ENMCS AS A FUNCTION OF METHOD WITH NO REDISTRIBUTION

Figure 5-10 shows what is now coming to be a familiar shift in favor of large bases (except for Base 6) with small increases in ENMCS at smaller bases. In Figure 5-10, Method 1 gives poor results for Base 6. Apparently, a number of LRUs at Base 6 are in short supply but do not "win" when Method 1 is used. That phenomenon results in larger backorders and a larger CAT value through the iterative process used in this analysis. Figure 5-11 shows that Method 2 (OPT CATs) is more robust in this regard. It "sees" the "bad actors" at Base 6 and allocates to it. And, yes, Method 2 (OPT CATs) further degrades the small bases and improves the availability of larger bases *with still a further reduction in total ENMCS* as is shown in Figure 5-12.

From an ENMCS point of view, it is clear that we ought to move away from the current process (Method 1 HQ CATs) since all the other methods give lower overall ENMCS values. However, from a by-base perspective, the current method comes closer to the targets for the smaller bases. We must decide whether we are willing to do a little worse at some bases that are not meeting their targets in order to realize a significant increase in total aircraft "up," albeit at bases that "don't need them."

Figure 5-13 shows small differences in the repair lists for Methods 1 and 2 using OPT CATs (an RCI of 85.5 percent), thereby suggesting that the difference in ENMCS between these two cases is the result of who gets the spare (distribution) as well as what gets repaired. On the other hand, since these repairs are significantly different from those resulting from Method 3 (see Figure 5-14), we can see that repair plays a significant role in reducing ENMCS. The RCI for Method 3 versus Method 1 (OPT CATs) is 58.1 percent. The RCI for Method 3 versus Method 2 (OPT CATs) is 65.1 percent.

The Effect of Redistribution

Any allocation process that removes constraints must by definition do no worse than the previous allocation. Thus, we would expect to do better if we do not constrain each base to its starting asset posture but redistribute all these assets to the bases that need them most. Two questions must be answered by this subsection:

- How much better do we do? The improvement that can be achieved is really a function of the amount of maldistribution in the system.
- When assets are not maldistributed, is any method more robust than the others or does any process meet overall aircraft availability and by-base ENMCS goals significantly better than the others?




FIG. 5-13. REPAIR QUANTITY FOR METHODS 1 AND 2 (OPT CATs) WITH NO REDISTRIBUTION




FIG. 5-14. REPAIR QUANTITY FOR METHODS 1 AND 2 (OPT CATs) AND METHOD 3 WITH NO REDISTRIBUTION

The redistribution process was conducted as follows:

- Before allocating any items to repair and distribution, we collected all base-serviceable items and placed them at the depot.
- We then allocated these serviceable items to the bases using the specific method we were considering.
- After that, we began the allocations using available repair capacity.

Obviously, many combinations of cases can be examined in the manner shown in Figures 5-1 through 5-14. Ideally, we should be able to discern differences from a figure in which all cases were displayed. However, we clearly have too many cases to display all of them. Thus, we begin by comparing the cumulative ENMCS curves for redistribution in Figure 5-15 with those for no redistribution from Figure 5-12. We note the following:

- The cumulative ENMCS curves for the redistribution cases in Figure 5-15 are all lower than their counterparts in Figure 5-12 (not a surprising result).
- The use of the current ENMCS goals as input targets for Method 1 (the current DRIVE algorithm) and Method 2 gives significantly *lower* overall availabilities than those resulting from the other four methods.
- Over all bases, Method 3 (allocating items to the base with the largest number of aircraft down because of the item) gives the lowest ENMCS. Method 4 (allocating items to the base with the largest fraction of total aircraft down because of the item) is second best.
- The above observations suggest the following: If assets are distributed correctly in the real world, then we may do substantially better (from an ENMCS point of view) if we use Method 3 or 4.

Figure 5-15 tells only half the story; we must now look at the by-base results to see how the methods, with redistribution, meet the individual base ENMCS goals. Figure 5-16 compares by-base ENMCS for the following two cases:

- Method 1 (HQ CATs) with the ENMCS goals used as input CAT values and a total ENMCS of 140.3 aircraft
- Method 2 (OPT CATs) using optimum CAT values and yielding a total ENMCS across all bases of 103.8.

Method 1 (HQ CATs), the current method, does a much better job of meeting the ENMCS goals of the smaller base and the bases with no wartime mission than does Method 2 (OPT CATs). While the ENMCS at the larger bases is higher using the




FIG. 5-15. CUMULATIVE ENMCS AS A FUNCTION OF METHOD WITH REDISTRIBUTION




FIG. 5-16. METHOD 1 (HQ CATs) VERSUS METHOD 2 (OPT CATs) IN MEETING ENMCS GOALS WITH REDISTRIBUTION

current DRIVE method, it is still below the input ENMCS goals for those bases (except for Base 6). Conversely, Method 2 (OPT CATs) favors large bases, sacrificing availability at the smaller bases and those with no wartime mission to get a 26 percent reduction in down aircraft.

In Figure 5-17, we examine the by-base ENMCS resulting from Methods 3 and 4 (the two methods that Figure 5-15 shows are the best from an overall ENMCS point of view). The significant points from Figure 5-17 are as follows:

- Method 3 comes closer to meeting the ENMCS goals for the smaller bases with a wartime mission (Bases 14 through 32) than does Method 2 (OPT CATs) (see Figure 5-16). It is also not much worse than the current DRIVE method (see Figure 5-16) in meeting the ENMCS goals of these bases. However, Method 3 still penalizes the bases with no wartime mission. On the up side, as shown in Figure 5-15, Method 3 imposes that penalty while reducing total ENMCS by almost 38 percent!
- Method 3 has the effect of making the number of down aircraft the same across all bases (see Figure 5-17). That is ameliorated some by Method 4, which provides generally larger ENMCS at the larger bases and lower ENMCS at the smaller bases.
- Compared with Method 1 (HQ CATs), Method 4 reduces total ENMCS by a little over 32 percent (as opposed to a 38 percent reduction for Method 3). However, Method 4 seems to strike the proper balance in meeting wartime goals. In fact, it does a marginally better job than the current method for the small wartime bases but does significantly better for the larger bases (compare Figures 5-16 and 5-17). On the other hand, Method 4 still does not meet the ENMCS goal of 0 for the peacetime bases.

A comparison of repairs is shown in Figure 5-18, which is the with-redistribution analog of Figure 5-13. There, we again see small differences in the repair list generated by Methods 3 and 4 when both methods use the same redistribution assumption. In this with-redistribution case, the RCI for Methods 3 and 4 is 86.7 percent; the RCI for these two methods in the no-redistribution case was 81.0. On the other hand, the RCI for Method 3 with and without redistribution is 57.1 percent and that for Method 4 is 54.2 percent (see tabulation in the addendum to Appendix B). This suggests that redistribution, when items are maldistributed, has more of an effect on what to repair than does the method used to allocate priorities for repair.




FIG. 5-17. METHOD 3 VERSUS METHOD 4 IN MEETING ENMCS GOALS WITH REDISTRIBUTION




FIG. 5-18. REPAIR QUANTITY FOR METHODS 3 AND 4 WITH REDISTRIBUTION

When we vary method and assumptions about redistribution, we find, as expected, still larger differences in number of items repaired. An attempt to portray these differences is shown in Figure 5-19 where we have added to Figure 5-18 the no-redistribution repair list for the Method 1 (HQ CATs) (heavy solid line). The RCIs for the no-redistribution Method 1 (HQ CATs) case versus Methods 3 and 4 are 47.5 percent and 48.3 percent, respectively.




FIG. 5-19. REPAIR QUANTITY FOR METHOD 1 (HQ CATs - NO REDISTRIBUTION) AND METHODS 3 AND 4 (WITH REDISTRIBUTION)

An Assessment of Backorders

Until now, we have ignored the question of total backorders. However, total backorders is often used as a measure of the relative number of cannibalizations actions that will occur in each case. Tables 5-1 and 5-2 present a summary of total ENMCS, total EBOs across all items at all bases, and average aircraft availability across the force for each of the cases discussed in the previous subsection. Table 5-1 presents summaries for the no-redistribution cases and Table 5-2 for the redistribution cases.

Table 5-1 shows that total EBOs goes down when we search for OPT CAT values for Methods 1 and 2. In these no-redistribution cases, we not only reduce ENMCS but

TABLE 5-1

**A COMPARISON OF ALTERNATIVE METHODS FOR PRIORITIZING REPAIR AND DISTRIBUTION
OF 41 AVIONICS LRUs AT 40 AIR FORCE BASES**

(No redistribution)

Parameters	Method					
	HQ CATs		3	4	OPT CATs	
	1	2			1	2
ENMCS	171.7	165.2	150.2	161.4	136.4	131.3
Total EBOs	626.4	601.2	515.4	555.0	503.7	463.3
Availability (%)	89.0	89.4	90.3	89.6	91.2	91.6

Note: Total aircraft = 1,555.

TABLE 5-2

**A COMPARISON OF ALTERNATIVE METHODS FOR PRIORITIZING REPAIR AND DISTRIBUTION
OF 41 AVIONICS LRUs AT 40 AIR FORCE BASES**

(With redistribution)

Parameters	Method					
	HQ CATs		3	4	OPT CATs	
	1	2			1	2
ENMCS	140.3	138.2	88.7	94.9	106.7	103.8
Total EBOs	617.0	605.8	201.2	228.1	388.3	281.0
Availability (%)	91.0	91.1	94.3	93.9	93.1	93.3

Note: Total aircraft = 1,555.

we also reduce the number of cannibalization actions that are likely to occur. That is also true for Methods 3 and 4.

In the redistribution cases (Table 5-2), the total number of EBOs, relative to the no-redistribution cases in Table 5-1, is reduced across the board as expected. What was not expected was the very small reduction in total EBOs between these two tables for Methods 1 and 2 using the HQ CATs. An explanation for this may be that

the total Air Force assets are sufficient to provide better availability than required by the ENMCS goals. This explanation is obviously supported by the lower ENMCS goals that we get from the other cases. Thus, the larger CAT values used in these cases essentially permit backorders at each base that are not permitted when we use ENMCS goals that are more closely tied to capability.

Another important point is that redistribution sharply reduces total ENMCS and EBOs in Methods 3 and 4.

CONCLUSIONS AND RECOMMENDATIONS

Our analyses offer many ways for improving significantly the aircraft availability across the total Air Force. We must now address the issue of selecting the method to recommend for DRIVE.

When LRUs are maldistributed, the procedures that search for OPT CAT values for each base will give the highest total aircraft availability regardless of whether we use the current DRIVE method (Method 1) or the method that tries to allocate items to the base with the lowest probability of achieving its objective (Method 2). That kind of iterative, searching process, however, is neither practical nor feasible for a production process. HQ AFLC cannot have a production process that requires iterative runs of a model. Admittedly, the iterations can be made invisible by performing them internally within the model and stopping at convergence, but they would not only increase the complexity of the model's design but would also increase run time by at least a factor of 10.

On the other hand, the other two methods addressed in this study (Methods 3 and 4) are independent of CAT value input. They do not need to know where we are in terms of ENMCS capability and they do a better job than the current method, which uses the HQ CAT values. And with redistribution, Methods 3 and 4 perform the best, minimizing ENMCS and EBOs and coming closest to meeting the by-base ENMCS goals.

In light of the above, it appears that HQ AFLC should move toward an allocation method that uses LRU backorders in a way that minimizes the number of aircraft (Method 3) or the fraction of aircraft (Method 4) at a base that are expected to be NMCS because of an LRU. We have intentionally omitted any discussion of how to

distribute, which is a more complicated matter that was touched upon in the introduction of this chapter and has not yet been resolved.

Also, our analyses show the enormous leverage of redistribution when assets are maldistributed. Consequently, the Air Force should seriously consider a repair allocation process that assumes perfect redistribution of serviceable assets. We hold this view for the following reasons:

- If assets are significantly maldistributed, they really ought to be redistributed. We expand on this point below.
- If assets are not significantly maldistributed, the assumption of perfect distribution is not an important one. In this instance, we propose that other systems operating closer to the unit (for example, TRADES⁵ or MASS⁶) should have the capability to make any smaller redistributions at a time and a place that are more appropriate.

We now turn our attention to the first point made above. If, in fact, seriously maldistributed items never get redistributed during peacetime, then an assumption of perfect distribution would be false. Such an assumption, however, may not be the disaster one might expect at first glance for the following reasons:

- The failure to repair an item for Base A because Base B has an excess of that item will result in other systems such as MASS finding those assets and transferring them, when they are needed to make an aircraft mission capable.
- During periods of increased tension, the Air Force can take those redistribution actions that are required to maximize the combat capability of the total force.

Under such a situation, we could accept marginally lower peacetime readiness; however, analyses show that the needed redistributions prior to combat would increase significantly the number of aircraft that would otherwise be available. These redistribution actions could be calculated every 30 days and become part of our wartime readiness plans or used selectively to make certain high-leverage redistributions in peacetime.

Taking the above approach leaves the depot free to do what it does best — help solve the availability problem at the macro level by repairing those spares that are

⁵TRADES = Theater/Region Allocation/Distribution Execution System.

⁶MASS = MICAP (mission capability) Asset Sourcing System.

most critical to the total Air Force from a worldwide perspective rather than from a by-base perspective. We believe that the Air Force repair capability should not be used to solve a maldistribution problem. Redistribution is much faster and effective than *repair and distribution*. The details of how to redistribute ought to be determined at the lower levels in a controlled, decentralized manner — a process that should be the focus of future analysis efforts.

CHAPTER 6

COMMON SRU PROBLEM: ONE REASON FOR A REPAIR-ONLY MODEL

THE COMMON SRU PROBLEM

A number of SRUs are common to many different LRUs, and the proper treatment of those common SRUs in inventory models throughout the Air Force has been a continuing problem. The inventory theory is well known and documented.¹ The Air Force Aircraft Availability Model is probably the model that treats this problem the best; it recognizes the indenture structure of the common SRUs, collects its demands as a single SRU, then distributes the shortages of the common SRU among the parent LRUs in proportion to the demands of the parent LRUs.

The current DRIVE model essentially ignores the common SRU problem. In the DRIVE model, SRUs that are common to more than one LRU are treated as separate and distinct SRUs with separate demands and supply. The demands for each of these "artificial"² items are those generated by the LRU parent. The supplies at the base and depot for one of these artificial SRU items are equal to the prorated share of the total supplies at these locations based on the distribution of demands across all the parent LRUs.

For example, if three different LRUs at a base generate demands for the same SRU and these individual demands are 1, 2, and 3 over a specific time horizon, then the DRIVE model will operate under the assumption that there are three separate SRUs (all with the same stock number) with demands of 1, 2, and 3, respectively. If the total base assets for the item are 12, then these stocks will be allocated to each of the artificial SRUs in a 1:2:3 proportion (2 for the first, 4 for the second, and 6 for the third). This same process will be continued at the depot with SRU demands from all

¹One source for a mathematical solution to the multi-echelon, multi-indenture inventory problem is *Vari-METRIC: A Multi-Echelon, Multi-Indenture Model with EOQ Items*, Craig C. Sherbrooke, Final Report, for HQ USAF/XPS under Contract F33600-86-M6240, January 1987.

²Since the DRIVE model creates additional, unique SRU items that really do not exist in practice, we refer to these SRU creations as artificial SRUs.

bases from its LRUs parents summed and the total depot assets of this SRU distributed proportionally among the artificial SRUs.

THE EFFECT OF THE COMMON SRU PROBLEM

The principal effect of the common SRU problem is that the current DRIVE model wants to repair more of the common SRUs at the expense of SRUs that are unique to a single LRU. The model will increase the availability of common SRUs while it detracts from what would otherwise be greater availability for the other SRUs. As a result, the current method either

- reduces aircraft availability across the force for a given repair capacity or
- requires additional repair capacity to achieve a desired aircraft availability.

To quantify this problem, we examined the current Ogden Air Logistics Center (ALC) data base of avionics components for the F-16 aircraft. Among the components were 41 LRUs and 209 SRUs. The 209 SRUs had the following commonality characteristics:

- 126 were unique in that they had only one parent LRU.
- 60 were common to two LRUs.
- 10 were common to three LRUs.
- 13 were common to four LRUs.

Thus, 40 percent of the SRUs are common to two or more LRUs. Although this percentage is probably higher than that for nonavionics items, it does suggest that commonality is too important to be ignored.

Our quantification of the effect of DRIVE's poor treatment of common SRUs proceeded as follows:

- We built a small data base with the same characteristics as those of the Ogden ALC data base for the F-16 avionics components:
 - ▶ Four LRUs and 20 SRUs, which constitute roughly one-tenth the LRUs and SRUs in the Ogden data base. The 20 SRUs had the following characteristics: 12 were unique, 6 were common to two LRUs, 1 was common to three LRUs, and 1 was common to four LRUs.
 - ▶ Forty identical bases, each with a 15-day order-and-shipping time (OST).

- ▶ All LRUs reparable at the base, with a base repair time of 4 days.
- ▶ No SRUs repaired at the base; all SRUs repaired at the depot, with a 30-day depot repair pipeline.
- In one case, we calculated a cost-versus-availability curve for cases in which each common SRU was treated as a separate SRU.
- In a second case, we calculated the cost-versus-availability curve for cases in which common SRUs were properly indentured to each of their parent LRUs using the Vari-METRIC analysis contained in the Aircraft Availability Model.³

We found that at typical high availability rates, proper treatment of commonality reduces LRU and SRU requirements by 12 percent each. Admittedly, an analysis of how much to buy (procurement cost versus availability) differs somewhat from setting priorities for repairing items in an environment that is constrained by a specific supply of spares. However, the analysis is not so much different that we can ignore the significant reductions in requirements.

In the first place, in its procurement of peacetime spares, the Air Force currently uses methodologies that consider SRU commonality in some level of detail. When we ignore these commonality considerations in repair during execution, we will not be achieving the required availability for the unique SRU items. The lower inventory of common SRU stocks that results from using our current procurement models will tell the DRIVE model to repair more of the common items since they are "seen" as separate items that (seem to) need higher availability (than is *actually* needed). A simple example will clarify this.

Assume the following:

- All demands are Poisson.
- A specific SRU is common to four LRUs.
- The expected SRU pipeline generated by each of the four LRU parents is one, yielding a total pipeline of four for this common SRU.
- A detailed analysis of requirements that properly treated commonality requires a stockage level of nine (against an aggregate expected pipeline of four) to attain some specified aircraft availability.

³See Craig C. Sherbrooke, *Vari-METRIC: A Multi-Echelon, Multi-Indenture Model with EOQ Items*, Final Report, for HQ USAF/XPS under Contract F33600-86-M6240, January 1987.

Under these assumptions, the optimum solution yields a supply availability of 0.9918 and EBOs of 0.01226 for the SRU. When this problem is decomposed into four SRUs each with an expected pipeline of one, a supply level of 16 (four for each of the decomposed SRUs) would yield an individual supply availability of 0.9963 and EBOs of 0.00435. The probability of no stock out occurring for any of these four individual SRUs is 0.9852 (0.9963^4) and the EBOs is 0.01740 (4×0.00435). These statistics are summarized in Table 6-1.

TABLE 6-1
COMPARISON OF SRU REQUIREMENTS
WITH AND WITHOUT PROPER TREATMENT OF COMMONALITY

Parameters	Proper commonality	Improper treatment of commonality	
		Decomposed SRU	Sum of individual SRUs
Number of items	1	1	4
Mean pipeline level	4	1	4
Stockage level	9	4	16
Supply availability	0.9918	0.9963	0.9852
EBOs	0.01226	0.00435	0.01740

Thus, Table 6-1 shows that even with 16 spares, the improper treatment of commonality (last column in Table 6-1) leads the DRIVE model to the belief that it is achieving lower supply availability and greater backorders than are actually required with a sparing level of nine (the proper commonality column in Table 6-1). That improper treatment leads to a 77 percent increase in requirements *and the model then believes we are not doing as well*. In fact, the real availability with 16 spares in the system under commonality is greater than 0.99999 and the EBO is smaller than 0.000002.

The current DRIVE model will try to repair more of the common items, thinking it has low supply availability when in fact it has higher availability. This repair is done at the expense of the unique items.

Another way to see how this improper treatment manifests itself in the current model is to consider the filling of requisitions for SRUs to "free up" LRUs that are

AWP. Since serviceable depot SRUs are distributed to artificial SRUs at the depot, it is possible that an SRU will be repaired at the depot when that depot has a serviceable one on the shelf.

Assume the following:

- SRU 1 is common to two LRU units (A and B).
- Demands for LRU A are two times those for LRU B, and the requirements for SRU 1 from each failure of the LRU units are the same.
- The depot has nine serviceable SRU 1 units.
- Nine LRU A units are AWP at bases because they need SRU 1 units and no LRU B units are AWP for lack of an SRU 1 unit.

Under this situation, six of the nine serviceable SRU 1 units at the depot would be allocated to artificial SRU 1 (LRU A) and three to artificial SRU 1 (LRU B). The DRIVE allocation process would see the six serviceable SRU 1 units (LRU A) for six of the nine LRU A units that are AWP for SRU 1 but would not see the other three SRU 1 units that are reserved for LRU B. It would conclude that it needs to repair some SRU 1 units and may allocate repair resources to do so when, in fact, adequate SRU 1 units are available (the three SRU 1 units reserved for LRU B).

THE COMMON SRU SOLUTION AS A BY-PRODUCT OF A REPAIR-ONLY MODEL

Proper treatment of common SRUs in DRIVE requires a redesign of the current model. The new, revised model would no longer process one LRU at a time but would process a family of LRUs and SRUs defined as that set of items that includes all LRUs and their SRUs that have any SRU in common.

We have developed the logic and computer programs for the proper treatment of common SRUs. However, our prototype is a "single-base" model — one that could look at a single Air Force base or could treat the entire Air Force as one big base. It was built on the idea that some redistribution was essential within the Air Force if we wanted to take maximum advantage of scarce repair and distribution resources.⁴

This single-base model does not directly address the question of distribution. Base information is aggregated into a single base in a way that compensates for the

⁴The case for some redistribution is strongly suggested by the findings in Chapter 5.

deficiencies that are inherent in assuming that all aircraft and assets are at a single base with full cannibalization permitted. This model is currently referred to as the LMI repair-only prototype model. It is described briefly in Appendix B.

The LMI repair-only prototype model is designed to run on the same PPOUT.DAT data file used in the current DRIVE program. We have made two runs of this model for both a biweekly and quarterly analysis. The difference in each of these two runs turns on what serviceable assets are assumed to be available at the depot:

- The "Base DI" run (see the addendum to Appendix B) is intended to match the current assumptions in the HQ DRIVE model. This run assumes that all base due-ins of LRUs that are not AWP are available and that none of the depot due-ins (LRUs and SRUs) are available.
- The "All DI" run assumes that all base and depot LRUs that are not AWP and all depot due-ins of SRUs are available (we had no way of knowing how many depot due-ins for SRUs were AWP).

The raw data comparisons are also shown in the addendum to Appendix B. Time and resources did not permit a detailed and extensive analysis of these raw data. A brief comparison of the biweekly and quarterly runs and some general observations follow.

Biweekly Comparisons

The assumption about available assets (the Base DI case versus the All DI case) has a significant impact on the repair lists for both LRUs and SRUs in the biweekly run. However, this problem is not a major one; we merely need to have accurate information on what is going to come out of repair so that the priorities for the next period are more reflective of the real situation. To assume that none of the due-ins comes out of repair is as erroneous as the assumption that all come out.

In comparing the Base DI run with the HQ AFLC results, we find significant agreement on those items that neither model required to be repaired. Of the 250 national stock numbers (NSNs) considered (41 LRUs and 209 SRUs), both models excluded 137 items from the repair lists (12 were LRUs and 125 were SRUs).

After the exclusions, 29 LRUs remained; of those, the LMI model selected 11 for repair and the current DRIVE model selected 28. The LRUs selected for repair by the LMI repair-only prototype model include fewer stock numbered items (NSNs) but

more repairs of the NSN items that were selected. That approach reflects the fact that the repair-only prototype model assumes perfect redistribution; the LMI repair-only model is repairing items that are needed most in the Air Force and not necessarily those that are needed at a particular base.⁵

- The 11 LRUs that the LMI repair-only model selected to be repaired had an average of 10 repairs per NSN item.
- The 28 LRUs the current DRIVE model selected to be repaired had an average of 4 repairs per NSN item.

After exclusion of 125 SRUs from the original 209, 51 SRUs were selected for repair by the LMI model and 58 by the current DRIVE model. Here, we see greater differences between the number of cases in which some items are selected for repair by one model and not by the other.

The above comparisons show that in all but one case, if an LRU was selected for repair by the LMI model, it was also selected for repair by the current DRIVE model. That was not true with SRUs where only about 50 percent of the items selected for repair by the LMI model were also selected by the current HQ AFLC model.

We did not have enough time and information to identify which reason (perfect redistribution or common SRU treatment) accounted for which differences. We would expect both assumptions to play a significant part but we do not know which contributed more to the differences. However, we were able to determine that the LMI repair-only model is biased toward eliminating AWP time for LRUs by sending SRUs to the bases. This bias comes from the redistribution assumption embedded in the LMI model and is discussed further below.

When the LMI model determines that the "system" needs another LRU, it first checks to see the system has an LRU that is AWP. With a perfect redistribution assumption, the model could be selecting an SRU to be repaired and sent to Base A to release an LRU that is really needed by Base B. On the other hand, such a situation has several positive aspects:

- We are able to get another needed LRU into the total system in the cheapest way possible although it needs to be sent from Base A to Base B.

⁵The current DRIVE model is able to redistribute assets among bases, but that ability was not exercised because we had neither sufficient time nor sufficient resources.

- We use LRU shop capacity to repair other critical LRUs and not those that can be repaired at the bases.
- With regard to setting priorities for repairs of SRUs, those that free up LRUs that are AWP will be used immediately when they get to the base rather than sit on the shelf at another base in anticipation of a need. This condition has several positive effects associated with the financial aspects of the stock fund that is now managing reparable:
 - ▶ The cash position of the stock fund increases because of the immediate sale of the SRU at Base A.
 - ▶ The receiving base is happy to receive the SRU so that it can repair the LRU and return it to supply where it receives a credit for the repair.⁶

Quarterly Comparisons

Quarterly comparisons showed considerably more agreement on the items that ought to be repaired but this is not surprising for two reasons:

- Any maldistributions tend to take care of themselves over time making the LMI assumption of perfect distribution less critical.
- Items with large daily demands must be continually repaired and appear on the repair list of both models.

On the other hand, very large differences still must be explained. The addendum to Appendix B shows that the LMI model generally selects for repair fewer of those SRUs that are common to more than one LRU. We must review the comparisons further to determine the reasons for other differences, and in that review, we may well uncover deficiencies in the LMI analysis/approach and may also identify additional shortcomings in the current DRIVE model.

General Observations

The LMI model produces a convenient "Constraint Matrix" summary that permits one to determine which repair shop is the binding constraint in the repair of high-priority items. The constraint matrices for the Base DI case for the LMI

⁶When an LRU is AWP for an SRU, the Air Force stock fund credits the customer's account with the net price of the LRU that is equal to the standard price less the LRU depot repair cost. When the SRU is provided to the customer and the LRU is no longer AWP for this SRU, the customer account is charged the net price of the LRU (the LRU carcass is "sold" back to the customer). After the item is repaired in base maintenance, it is sold to base supply for the standard price. The net effect of the above transactions is a net gain of the LRU repair cost in the customer's account.

biweekly and quarterly runs are shown in Tables 6-2 and 6-3. Entries in each horizontal row indicate unused capacity (top part of Table 6-2) or fraction of unused capacity (lower part of Table 6-2) for the shop in that column *when the shop shown in the first column in that row exceeded its capacity*. Negative numbers mean capacity was exceeded by that amount. For example, Table 6-2 shows the following:

- When the capacity of Shop RF was exceeded by 2 hours, Shop CI had 81 hours of unused capacity and Shop D had 1,088 hours of unused capacity.
- When the total capacity of Shop CI was exceeded by 0.8 percent, Shop PP used only 21.7 percent of its total capacity.

These tables show that the single most important shop is Shop RF, which repairs LRUs. The most critical SRU shop is Shop A. The least critical shops are the DI and M shops. From the above, it appears that we could probably close these shops for about a year and not feel any impact, which brings us to a related point.

The tables show that in some cases, more than 100 percent of capacity is used in the A and D shops that repair SRUs. How could this happen if the model was run in a capacity-constrained mode? It occurs because the model is trying to schedule the workload at LRU Shop DI. To do so, it must schedule the repairs of SRUs in Shops A and D, when those shops have already been scheduled to repair higher priority items. Thus, we are confronted with a choice:

- Do we maintain the model priorities for repairing SRUs and not repair SRUs needed by maintenance people in Shop DI? If we do this, we are not utilizing these DI maintainers very efficiently; all they do is generate AWP.
- Do we give the SRU repairs needed for LRU repairs in Shop DI higher priority by dropping lowest priority items off the bottom of the A and D list, recognizing that these SRU repairs really have higher readiness impact on the force?

In the LMI analysis reflected in the repair summaries, we have taken the latter view. We presumed that if the people are there to work, they ought to be doing useful work. If they are not needed, they should be cross-trained into another shop and the input capacities data changed to reflect reorganization of the depot.

TABLE 6-2
BIWEEKLY SUMMARIES
(For Base DI case)

<i>Unused Capacity (hours) at Shop</i>							
Shop	RF	CI	PP	A	D	DI	M
RF	- 2	81	360	592	1,088	395	1,684
CI	- 2	- 4	282	459	1,031	395	1,684
PP	- 2	- 4	- 1	20	850	395	1,684
A	- 2	- 4	- 1	- 7	850	395	1,684
D	- 2	- 4	- 1	- 14	- 4	356	1,145
DI	- 2	- 4	- 1	- 318	- 86	- 2	673
M	- 2	- 4	- 1	- 336	- 86	- 2	- 12

<i>Fraction of Total Capacity Used at Shop</i>							
Shop	RF	CI	PP	A	D	DI	M
RF	1.004	0.843	0.000	0.487	0.188	0.000	0.038
CI	1.004	1.008	0.217	0.603	0.231	0.000	0.038
PP	1.004	1.008	1.003	0.983	0.366	0.000	0.038
A	1.004	1.008	1.003	1.006	0.366	0.000	0.038
D	1.004	1.008	1.003	1.012	1.003	0.099	0.346
DI	1.004	1.008	1.003	1.275	1.064	1.005	0.615
M	1.004	1.008	1.003	1.291	1.064	1.005	1.007

Notes: LRU shops are RF, CI, PP and DI; SRU shops are A, D, and M.

TABLE 6-3
QUARTERLY SUMMARIES
(For Base DI case)

<i>Unused Capacity (hours) at Shop</i>							
Shop	RF	PP	A	D	CI	DI	M
RF	- 14	1,790	3,480	6,564	1,660	2,370	10,368
PP	- 14	- 7	1,826	4,957	1,112	2,370	8,798
A	- 14	- 7	- 9	3,994	391	2,370	7,649
D	- 14	- 7	- 9	- 5	13	2,178	4,646
CI	- 14	- 7	- 9	- 5	0	2,171	4,600
DI	- 14	- 7	- 1,987	- 465	0	- 5	1,174
M	- 14	- 7	- 1,987	- 471	0	- 5	- 4

<i>Fraction of Total Capacity Used at Shop</i>							
Shop	RF	PP	A	D	CI	DI	M
RF	1.005	0.171	0.498	0.184	0.463	0.000	0.013
PP	1.005	1.003	0.737	0.383	0.640	0.000	0.162
A	1.005	1.003	1.001	0.503	0.873	0.000	0.272
D	1.005	1.003	1.001	1.001	0.996	0.081	0.558
CI	1.005	1.003	1.001	1.001	1.000	0.084	0.562
DI	1.005	1.003	1.287	1.058	1.000	1.002	0.888
M	1.005	1.003	1.287	1.059	1.000	1.002	1.000

Notes: LRU shops are RF, CI, PP and DI; SRU shops are A, D, and M.

APPENDIX A

THE REPAIR COMMONALITY INDEX

In Chapter 5 of the main text of this report, we address alternative methods for establishing priorities for the repair and distribution of spare parts and how the allocations from those alternative methods affect aircraft availability. In analyzing differences among the methods, it is useful to know whether the differences in aircraft availability from two alternative methods or processes should be attributed to differences in repair or to differences in distribution. To this end, we have developed the following repair commonality index (RCI) that best defines the extent to which the same repairs are recommended by two lists:

$$RCI = \frac{\sum_{i=1}^{N_{items}} \min[R(1,i), R(2,i)]}{\sum_{i=1}^{N_{items}} \max[R(1,i), R(2,i)]} \times 100$$

where

- $R(1,i)$ = number of Item i repaired in one alternative,
- $R(2,i)$ = number of Item i repaired in the second alternative, and
- N_{items} = the total number of line replaceable unit (LRU) items.

Table A-1 summarizes the RCIs for a number of the alternatives presented in Chapter 5. There, we also define the various methods (Method 1 through Method 4) and the terms "HQ CATs" (headquarters cannibalization threshold) and "OPT CATs" (optimum CATs). In Table A-1, the numbers circled along the diagonal are the RCIs for the two repair lists generated by the *same method* but *different redistribution* assumptions. Thus, for example, the 62.0 in the upper-left corner of the table is the RCI for the following two cases:

- A repair list generated when using Method 1 (HQ CATs) and *no redistribution*
- A repair list generated when using Method 1 (HQ CATs), *with redistribution*.

TABLE A-1
SUMMARY OF SELECTED REPAIR COMMONALITY INDICES
(Percent)

Method	HQ CATs				OPT CATs	
	1	2	3	4	1	2
1 (HQ CATs)	62.0	90.2	71.2	71.0	70.7	70.6
2 (HQ CATs)	86.4	60.6	73.7	73.0	70.9	72.6
3	47.5	43.2	57.1	81.0	58.1	65.1
4	48.3	43.9	87.6	54.2	54.0	58.0
1 (OPT CATs)	67.2	62.3	48.1	49.2	71.6	85.5
2 (OPT CATs)	70.2	66.9	48.4	50.8	87.1	70.4

} *No redistribution*

} *With redistribution*

The indices above and to the right of the diagonal in Table A-1 are the RCIs for the two repair lists generated by *different methods* but with the same *no-redistribution* assumption. Thus, the 70.6 in the top row, last column is the RCI for the following two cases:

- A repair list generated when using Method 1 (HQ CATs) and *no redistribution*
- A repair list generated when using Method 2 (OPT CATs) and *no redistribution*.

Conversely, the indices below and to the left of the diagonal in Table A-1 are the RCIs for the two repair lists generated by *different methods* but with the same *with-redistribution* assumption. Thus, the 70.2 in the bottom row, left-most number column is the RCI for the following two cases:

- A repair list generated when using Method 1 (HQ CATs), *with redistribution*
- A repair list generated when using Method 2 (OPT CATs), *with redistribution*.

APPENDIX B

OVERVIEW OF LMI REPAIR MODEL

This appendix offers a broad overview of the process used and an appreciation for the advantages and disadvantages of the LMI repair-only model for the Distribution and Repair in Variable Environments (DRIVE) model. It presumes that the reader is generally familiar with the current DRIVE model and the data it uses as input. It is not intended to be a full and comprehensive documentation of the model.

The LMI repair-only model uses the same input file – PPOUT.DAT – as that used by the Headquarters, Air Force Logistics Command (HQ AFLC) DRIVE model. This appendix is presented in five parts:

- A discussion of how demands are calculated and aggregated
- The allocation logic
- Model options
- The model's objective function
- Model options used in the LMI analyses.

The addendum to this appendix compares the results from the LMI model with results from the HQ AFLC DRIVE model.

DEMAND INFORMATION

For each base, the model calculates expected number (u) and variance (v) of due-ins for each item over the time horizon:

- Time horizon = repair cycle time + order-and-shipping time (OST) + war period
- Variance is determined by the following formula:

$$v = VMR \times u$$

where

$$\begin{aligned} \text{VMR (variance-to-mean ratio)} &= 1 + 0.14 \times \sqrt{AD}, \text{ and} \\ AD &= \text{annual demands for the item.} \end{aligned}$$

- Expected due-ins also include those items in the base repair pipeline.

The model then aggregates the expected due-ins and their variance of these due-ins by summing these parameters over all bases to yield for the force a single expected number of due-ins (u) and the variance of those due-ins (v).

Supply data for all bases, including intransit assets, are aggregated. The model has options for including base and depot assets that are already inducted into maintenance and not awaiting parts (AWP).

The model uses the following data to estimate the line replaceable units (LRUs) that are AWP at the bases:

- The number of shop replaceable unit (SRU) holes at each base
- The quantity per application (QPA) for the SRU on each of its parent LRUs (commonality)
- A distribution of holes to LRUs based on the fraction of the demands for the SRU that are generated by the LRU parent (commonality)
- An assumption that all SRU holes are consolidated into the smallest number of LRUs at each base.

In the single-base aggregation, the total number of LRUs that are AWP at the bases are summed with the SRU holes in each LRU identified. In the single-base aggregation, the LMI model does not consolidate all SRU holes into the smallest number of LRUs. That procedure is consistent with the way the depot treats SRU holes: each LRU has a specific set of SRU holes that must be filled in order to repair it.

On the other hand, when the model is allocating SRUs to the bases to reduce base SRU backorders, which in turn affects LRU availability, it does assume that all SRU backorders are consolidated into the smallest number of LRUs. That assumption clearly understates impact of SRU backorders on LRU availability across the force. Consequently, when LRUs and SRUs compete (as is the case when we are making procurement decisions), the LMI model will tend to select fewer SRUs for

repair relative to the number of LRUs it selects for repair. However, since LRUs and SRUs do not compete for resources in DRIVE (they are generally repaired in different shops), the allocation problem boils down to finding those SRUs that are in the worst shape relative to each other. Thus, our assumption that all SRU backorders are consolidated into the smallest number of LRUs does not have a serious effect.

THE ALLOCATION LOGIC

In the LMI model all allocations for all items are performed in a single program; there is no separate postprocessor program. The allocation logic procedure is outlined below:

- Select the item that has the best sort value regardless of whether it is an LRU or SRU. More details on this sort value are presented subsequently in this appendix in the section entitled "Objective Function."
- After the workload at the LRU shop is scheduled, do not consider any LRUs that are scheduled for repair at that LRU shop *unless those LRUs are AWP at the base and can be repaired by sending SRUs* either from depot stock or by repairing an SRU (provided there is capacity remaining in the SRU shop).
- When this allocation process is completed, ensure that all LRUs selected for repair have sufficient SRUs. The way the model now operates, an LRU could be inducted for repair even though no repair capacity remains for an SRU that the inducted LRU needs. In that case, the model makes room in the SRU repair shop to repair the needed SRUs and drops those SRUs with the lowest priority from the repair list. This is done with two passes through the sorted repair list:
 - ▶ The first pass collects LRU repairs and those SRU repairs needed to complete the LRU repairs.
 - ▶ In the second pass, SRUs are scheduled for induction in priority sequence until the workloads in the SRU shops are completely scheduled.

MODEL RUN OPTIONS

Options for running the model are listed below:

- *Biweekly or quarterly option.*
- *MICAP (mission capability) option:*
 - ▶ Ignore all MICAP input data, assuming there are no MICAPs.

- ▶ Process all existing MICAPs as highest priority at the beginning of the process.
- ▶ Let existing MICAPs compete with all backorders and have the optimization determine which MICAPs get eliminated.
- *Capacity constraint on repair shops (yes or no).* Even if no capacity constraints exist, shop capacities must be input. In the unconstrained mode, the model stops when the workload for the last shop has been scheduled while allowing all the other shops to repair items in excess of the input values.
- *Carcass constraints (yes or no).* In all cases, SRU carcasses that are accumulated at the depot are added to the available carcass pool.
- *Job-route option.* If this option is used, the model does not require the availability of SRU carcasses to schedule repair of an inducted LRU. Otherwise, an SRU carcass must be available for repair.
- *SRU constraints on LRU repair.* When an SRU must be scheduled for repair before an LRU can be scheduled for repair because the depot has no serviceable SRUs, the user has the following options:
 - ▶ Use a sort value that assumes the SRU needed to repair the LRU can be repaired; or
 - ▶ Eliminate this LRU from consideration if the model is running in the capacity-constrained mode and SRU repair capacity has been exceeded or if the model is running in the carcass-constrained mode and no carcasses are available and the job-route option is off.
- *LRU workload priority.* This option should not normally be used. When it is on, the user has the option of scheduling the workloads at LRU shops as highest priority before allowing the optimization to select an SRU to reduce SRU backorders at the bases. When this option is off (the preferable setting), the model selects the item with the best sort value regardless of indenture. The only benefit of this option is that it cuts the run time by more than half (from 3.0 minutes to 1.5 minutes). Preliminary runs show that the on option does not have a significant effect on the LRU repairs at the depot; however, it does have an effect on the depot SRU repairs and, in general, results in higher LRU backorders at the bases.
- *MIC (maintenance inventory center) asset utilization option.* When this option is used, SRU assets in the maintenance shops are made available to the depot to clear base AWP or to reduce base SRU backorders.
- *Repair cost option.* By turning this option on, the user can ignore the repair cost in the calculation of the sort value and allocation process. The sort value then becomes the change in the objective function.

- *Base LRU due-in option.* If the option is on, the base LRU assets are augmented by those base due-ins that are not AWP.
- *Depot LRU due-in option.* If this option is on, the depot LRU assets are augmented by those depot due-ins that are not AWP.
- *Depot SRU due-in option.* If this option is on, the depot SRU assets are augmented by SRU depot due-ins.

OBJECTIVE FUNCTION

The objective function for the sort value is the weighted expected backorders (EBOs). The weighting method and the rationale for it are described below:

- What we want to be able to do is maximize aircraft availability subject to the following:
 - ▶ Cannibalization.
 - ▶ Aircraft priorities (F-15 aircraft may be of more relative importance than C-130s).
 - ▶ The number of aircraft allowed to be down in peacetime and wartime in the requirement process.
 - ▶ The availability of spare parts and repair capacity. (We may want to have an aircraft availability of 90 percent but we do not have enough assets — spares and repair capability — to achieve that value.)
- At the same time, we want to have a model that takes a reasonable amount of time to run. For example, we discarded the idea of calculating the marginal increase in aircraft availability associated with inducting one of each item in a cannibalization environment because it would lead to unacceptably long running times.
- We have considered using a "no-cann" logic that makes the problem "separable" and simpler, but we discarded that idea because of the large body of opinion that says that the no-cann logic is inappropriate especially when we buy War Readiness Spares Kits (WRSK) on the cann logic.
- The current method of using a direct support objective (DSO) target is convenient but could lead to misallocations in a constrained environment. In Chapter 4 of the main text, we show the optimum cannibalization thresholds (OPT CATs) that increase the aircraft availability across the board.
- We have done other analyses (not addressed in this report) in which the allocation for a single base using weighted backorders gives much the same results as a detailed "greedy algorithm" that minimizes the number of expected not mission capable-supply (ENMCS). The sort value for the

weighted EBO objective function that was used in that analysis is given below:

$$\frac{(\Delta EBOs) \times EBOs}{(QPA \times RepCost)}$$

where

- EBO = the expected backorders,
- QPA = the quantity per application, and
- $RepCost$ = repair cost of item.

This equation gives a good approximation for optimizing aircraft availability because it tries to improve the item that is keeping the aircraft down. If we have larger backorders on Item A relative to Item B but Item B is very cheap to repair, the current HQ AFLC DRIVE process will repair Item B because we get the largest increase in the probability of getting x or fewer backorders per repair cost. Item A, however, may really be keeping the aircraft down and should be repaired at any cost.

- Based on these earlier analyses, the current LMI model uses the following sort value for LRUs:

$$\sum_i \frac{SortV_i}{RepCost}$$

where

- i = the index on aircraft,
- $SortV_i$ (the sort value) = $\Delta EBO_i \times (EBO_i / QPA_i) / CAT_i$,
- EBO_i = expected backorders for aircraft i ,
- QPA_i = LRU quantity per application on aircraft i , and
- CAT_i = number of allowable aircraft down.

- The sort value for the SRUs is consistent with the above equation and is captured through its impact on the parent LRUs through the cannibalization logic for SRUs.

MODEL OPTIONS USED AND ASSUMPTIONS MADE IN THE LMI REPAIR-ONLY ANALYSES

Addendum 1 to this appendix is a comparison of the biweekly and quarterly repair lists for the seven avionics shops at Ogden Air Logistics Center (ALC), Utah. The case information used in these runs follows:

- The number of hours available for each shop are

CI	D	A	RF	M	DI	PP
515	1,340	1,155	480	1,750	395	360

- LRUs are given low priority if SRU capacity is exceeded or SRU carcasses are unavailable.
- MICAP on switch = 0: do not process MICAPs.
- ILRU on switch = 0: do not schedule LRU shops first.
- Carcasses are constrained.
- Shop capacity is constrained.
- Two standard deviations for safety level for SRUs in MIC are used.
- Biweekly DRIVE is run.
- Aircraft availability targets are

F-16A	F-16B	F-16C	F-16D
0.950	0.950	0.950	0.950

- Repair costs are used.
- MIC assets to depot supply are not used.
- SRUs for LRUs are repaired as necessary.
- The addendum has two cases presented for considering due-in (DI) assets at the depot:
 - ▶ *Base DI.* This case does not include any depot due-ins; it only considers base due-ins that are not AWP.
 - ▶ *All DI.* This case includes all due-ins that are not AWP at the base and depot and all due-in SRUs.

ADDENDUM

A COMPARISON OF RESULTS FROM THE LMI REPAIR MODEL AND HQ AFLC REPAIR AND DISTRIBUTION MODEL FOR BIWEEKLY AND QUARTERLY RUNS

Summary for Shop CI									
Index	Master NSN	Description /IM ES WUC	Biweekly			Quarterly			
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
1	1270010453976WF	FIRE COMP HW VH74CA0	0	0	1	0	0	1	
107	1270012223829WF	XFCC HW VH74CA0	0	0	0	2	2	4	
178	6605010463533WF	FC NAV PAN HW VX74DD0	50	14	11	132	81	87	
186	6605010876645WF	INU 74DA0 HW VC74DA0	0	3	9	19	45	32	
196	6610010397817WF	ACCELER AS HW VE14AF0	0	0	1	0	0	1	
210	6610011230046WF	ECA 14FB0 HW BA14FB0	0	0	1	1	0	5	
224	6610011480712WF	ECA C/D HW BA	0	0	0	0	0	0	
226	6615010427834WF	GYRO HW VE14AG0	0	0	1	3	0	6	
229	6615011273160WF	PANEL HW BA14AD0	0	20	1	29	48	24	
235	6615011297445WF	PANEL TRIM HW BA14AE0	11	0	18	128	42	125	
236	6615011611592WF	FLT CTL CO HW BA14AA0	0	0	0	0	0	0	
247	6615011720136WF	FLCC HW BA14AA0	0	0	0	1	0	1	
249	6615012203851WF	FL CTL CTR HA BA14AA0	0	7	0	4	22	8	
250	6625011146771WF	RECORD ASY HW VZ14AL0	0	7	0	23	15	15	

Summary for Shop RF									
Index	Master NSN	Description /IM ES WUC	Biweekly			Quarterly			
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
11	1270010932174WF	ANTENNA RA HW VN74AA0	0	0	0	0	0	0	
17	1270010932256WF	RADAR XMTR HA VP74AC0	0	0	0	0	0	0	
32	1270011022962WF	LOW PWR RF HA V574AB0	0	0	1	0	0	1	
41	1270011022963WF	LOW PWR RF HA V574AB0	0	0	1	0	0	3	
43	1270011022965WF	LOW PWR RF HA V574AB0	0	0	3	0	4	12	
45	1270011022966WF	LOW PWR RF HA V574AB0	0	0	7	0	0	17	
90	1270011464630WF	ANTENNA HW VN74AA0	11	11	4	71	69	50	

Summary for Shop DI									
Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
23	1270010946872WF	RCP 74AH0 HA VV74AH0	1	1	1	9	0	37	
26	1270010948505WF	HUD PDU HA VF74BA0	0	0	1	0	0	16	
47	1270011229955WF	HUD ELECT HA VF74BC0	6	9	6	28	17	14	
124	1270012740543WF	HUD EU ADF HA VF74BC0	1	0	2	0	0	8	
157	5841010963945WF	DISP 74EA0 HW VG74EA0	18	6	11	50	38	47	
164	5841010964833WF	RDR E74EB0 HW VG74EB0	5	5	3	44	60	30	

Notes: NSN = national stock number, IM = item manager, ES = equipment specialist, WUC = work unit code, and alloc = allocation

Summary for Shop PP									
Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
62	1270011336494WF	DIG SIG PR HA VA74AD0	0	0	0	0	0	0	
93	1270012099982WF	RADC HA VA74AF0	8	8	2	37	37	24	
105	1270012122990WF	DSP ADF HA VA74AD0	0	0	1	0	0	1	
116	1270012733858WF	OCU RADC HA VT	0	0	0	0	0	0	
121	1270012733859WF	S2 RADC HA VT	0	0	1	0	0	3	
126	1270012827914WF	ADF RADC HA VT	0	0	0	0	0	0	
127	1280011091499WF	MRIU 75DB HA WA74DB0	0	0	2	0	0	17	
130	1280011216879WF	SCP 75DA0 HA WA75DA0	5	0	5	27	9	32	
133	1280012248924WF	XCIU HA WA	0	0	0	0	0	0	
146	1280012804855WF	CIU-S2 HA WA	0	9	3	17	33	16	
154	1290010800203WF	CRIU 75DE0 HA WA75DE0	0	0	1	0	3	2	
176	5999010803978WF	JRIU 75DD0 HA WA75DD0	0	0	1	19	25	1	
197	6610010891018WF	COMPUTR CA HW NA51FA0	0	0	7	0	0	28	
227	6615010427835WF	PNE SENSOR HW BA14FC0	2	0	2	9	9	7	

Summary for Shop D										
Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator	
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.		
			Base DI	All DI		Base DI	All DI			
2	1270010653427WF	PROCESSOR1 HW VH74CAA	0	1	0	0	0	0		
4	5998010657867WF	CCA MULTIP HW VH75CAH	0	0	0	0	0	0		
5	5998010672075WF	CCA DIGITA HW VH74CAG	0	0	0	0	1	0		
6	5998010672076WF	CCA MEMORY HW VH74CAF	0	0	0	0	0	0		
7	5998010687880WF	CCA-MIA HW VH74CAJ	0	0	0	0	0	0		
8	5998010694390WF	CONVERTER HW VH74CAC	0	3	1	0	1	1		
9	5998010702774WF	CCA COMMONHW VH74CAD	0	0	0	0	0	0		
10	5998010794168WF	CCA PROC HW VH74CAB	0	0	0	0	0	0		
21	1285010847356WF	DIGIBUS BD HA VP74ACD	0	0	0	0	0	0		
24	5998010726306WF	BOARD ASSY HA VA74AHA	0	3	0	6	4	8		
25	5999010696483WF	BOARD ASSY HA VA74AHD	1	2	0	5	3	7		
39	5998010993253WF	CONTR BD HA VA74ABA	17	0	13	36	15	63	c4	
49	1270010848495WF	OCCCLUSION HA VF74BCA	0	0	0	0	0	0	c	
50	1270010972254WF	INTRFC CTL HA VF74BCJ	0	0	0	0	0	0	c	
52	1270011007336WF	SYM GEN B HA VF74BCM	0	3	0	3	8	0		
53	1270011230069WF	EPROM HA VF74BCN	0	0	0	10	2	1	c	
54	5998010795412WF	ARITHM PCB HA VF74BCB	0	5	0	7	9	0		
55	5998010795413WF	DIGTL IN A HA VF74BCC	1	2	1	8	6	3	c	
56	5998010795414WF	DIGTL IN B HA VF74BCD	1	5	0	6	9	3	c	
57	5998010799816WF	PROCESSOR HA VF74BCK	0	4	0	10	8	0		
58	5998010799817WF	SYM GEN A HA VF74BCL	0	4	0	2	5	0	c	
59	5998011006496WF	CLOCK ASSY HA VF74BCE	1	0	0	5	2	1	c	
63	1270010795395WF	BOARD ASSY HA VA74ADH	0	0	0	0	0	1	c	
64	1270011657143WF	BOARD ASSY HA VA74ADS	0	1	0	0	4	1	c	
65	5998010777476WF	BOARD ASSY HA VA74ADB	0	0	1	0	0	3	c	
66	5998010777477WF	BOARD ASSY HA VA74AEB	0	0	0	0	0	3	c	
67	5998010777478WF	BOARD ASSY HA VA74AEC	0	0	0	0	0	1	c	
68	5998010777480WF	CFAR NO 3 HA VA74ADL	0	0	0	0	0	1	c	
69	5998010777481WF	CFAR NO 2 HA VA74ADJ	0	0	0	0	0	1	c	
70	5998010777482WF	CFAR NO 1 HA VA74ADT	0	0	0	0	0	1	c	
71	5998010777483WF	CFAR NO 0 HA VA74AEE	0	3	1	0	4	2	c	

Notes: NSN = national stock number; IM = item manager; ES = equipment specialist; WUC = work unit code; alloc = allocation; c = common to two LRUs; c3 = common to three LRUs; and c4 = common to four LRUs.

Summary for Shop D (continued)

Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
72	5998010777484WF	BOARD ASSY HA VA74ADR	0	0	0	0	2	2	c
73	5998010777485WF	AU SCRATCH HA VA74AEA	0	0	1	0	0	8	c
74	5998010777487WF	AU OUTSIDE HA VA74ADW	0	0	1	0	0	3	c
75	5998010777488WF	AU INSIDE HA VA74ADX	0	0	0	0	0	1	c
76	5998010777489WF	AU INSIDE HA VA74ADY	0	0	0	0	0	1	c
77	5998010777492WF	BOARD ASSY HA VA	0	0	0	0	0	1	c
78	5998010777494WF	BOARD ASSY HA VA74ADJ	0	0	0	0	0	1	c
79	5998010777495WF	XY BD ASSY HA VA74ADM	0	0	1	0	0	4	c
80	5998010777496WF	BOARD ASSY HA VA74ADF	0	0	1	0	0	5	c
81	5998010777497WF	BOARD ASSY HA VA74ADD	0	0	1	1	1	4	c
82	5998010778165WF	CT I-O HA VA74ADE	0	0	0	0	0	1	c
83	5998010783390WF	BOARD ASSY HA VA74ADQ	0	3	2	0	5	4	c
84	5998010795393WF	BOARD ASSY HA VA74ADG	0	8	1	4	10	2	c
85	5998010795396WF	BOARD ASSY HA VA74ADF	0	0	0	0	0	1	c
86	5998010950251WF	CONV CTRL HA VA74ADL	0	2	1	0	3	2	c
87	5998010950252WF	A-A TAR ME HA VA74ADP	0	4	0	0	4	1	c
88	5998010950253WF	A-A TAR GE HA VA74ADN	0	8	1	5	11	2	c
89	6130011336495WF	BOARD ASSY HA VA74ADA	0	0	0	0	3	0	c
94	1270010983541WF	BOARD ASSY HA VA74AFD	0	0	0	5	4	0	c4
95	1270010995200WF	ARITH BD HA VA74AFA	0	0	0	6	0	10	c4
96	1270011013774WF	BOARD ASSY HA VT74AFC	0	0	0	0	0	0	c3
97	1270011030374WF	BOARD ASSY HA VT74AFB	0	0	1	0	0	4	c
98	1270012100977WF	EPROM BD 1 HA VA74AFL	0	0	3	15	5	9	
99	1270012100978WF	EPROM BD 2 HA VA74AFM	0	0	1	4	3	3	
100	5998011013776WF	BOARD ASSY HA VT74AFE	9	0	2	21	11	11	c3
102	5998011035943WF	16K RAM BD HA VT74AFF	12	14	8	27	28	6	c4
106	5998012129127WF		0	0	0	0	0	0	
108	1270012221936WF	CCA NLD HW VH	0	2	0	2	2	0	
110	1270012223849WF	CCA PIF HW VH	0	0	1	2	0	6	
111	1270012229491WF	INTR AV HW VH	0	0	0	1	1	1	
112	5998012221935WF	CCA-BCI HW VH	3	2	0	4	3	1	
113	5999012221932WF	CCA ADC HW VH	0	0	0	0	0	0	
114	5999012222933WF	XFCC PART HW VH	0	0	0	3	3	2	
115	5999012223850WF	CCA-P2F HW VH	0	0	0	0	0	2	
117	1270011012774WF		0		0	0	0	1	
118	1270011012776WF		0	0	0	0	0	1	
119	1270012098985WF	EPROM BD 1 HA VT	0	0	0	1	0	1	
120	1270012098986WF	EPROM BD 2 HA VT	0	0	0	0	0	1	
122	5999012739728WF	16K EP BD HA VT	0	0	0	0	0	1	
123	5999012739729WF	32K EP BD HA VT	0	0	0	1	5	2	
125	5999012746290WF	EPROM BR D HA VF	1	1	1	0	4	2	
128	1280011083415WF	CIRCUIT CD HA WA75DBC	0	0	0	64	158	155	
129	5999010722440WF	CKT CARD HA WA75DOB	0	4	30	275	300	249	c3
131	1280010722439WF	CKT CD ASY HA WA75DA0	0	0	5	0	0	8	
134	1280010726305WF	CKT CD ASY HA WA75DCB	0	8	0	18	20	6	c
135	1280010785539WF	CKT CARD A HA WA75DCH	0	4	1	9	15	14	c
136	1280010970602WF	CKT CD ASY HA WA75DCG	0	0	0	4	0	2	c
137	1280011013020WF	CKT CD ASY HA WA75DCS	1	7	1	15	19	8	c
138	1280011119800WF	CIRCUIT CD HA WA75DCF	23	0	4	55	20	27	c
139	1280012223136WF	CIR CARD HA WA	0	4	0	0	8	0	
140	1280012223862WF	CIR CARD HA WA	0	1	0	0	4	0	
141	5999011404494WF	CIRCUIT CD HA WA75DCC	3	0	3	25	12	14	c
142	5999012633336WF	HA WA75DC	0	0	0	0	0	0	
143	5999012662329WF	HA WA75DC	0	0	0	0	0	0	

Notes: NSN = national stock number, IM = item manager, ES = equipment specialist, WUC = work unit code, alloc = allocation, c = common to two LRUs, c3 = common to three LRUs, and c4 = common to four LRUs

Summary for Shop D (continued)

Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
144	5999012672425WF	CIR CARD HA WA	0	0	0	0	0	0	
147	1280010944080WF	CKT CD ASY HA WA75DCE	0	0	0	6	5	13	
148	1280011004375WF	CT CD ASY HA WA75DCA	1	0	0	13	5	8	
149	1280011043066WF	CKT CARD HA WA75DCJ	0	0	0	5	9	10	
150	1280011156216WF	WCKT CD AY HA A75DCJ	7	8	4	18	18	12	
151	5999012801754WF	CIR CARD HA WA75DCU	0	2	0	4	5	3	
152	5999012801755WF	CIR CARD HA WA	0	0	0	3	3	0	
153	5999012802869WF	CIR CARD HA WA75DCT	0	3	0	12	14	13	
155	1280010714919WF	STATUS LOG HA WA75DEC	0	0	3	0	0	4	
156	1280010838515WF	CK CD ASSY HA WA75DED	0	0	2	0	2	4	
165	5841010954942WF	CIRCUIT AY HW VG74EBL	0	0	0	0	0	0	
166	5998010656683WF	CCA MUX ME HW VG74FBE	0	0	0	0	0	0	
167	5998010657865WF	MUX INTERF HW VG74EBB	0	0	0	0	0	0	
168	5998010677665WF	MOV SYM GE HW VG74EBJ	0	0	0	0	0	0	
169	5998010830486WF	CCA CPU PR HW VG74EBF	0	0	0	0	0	0	
170	5998010861366WF	CCA PROC C HW VG74EBG	1	0	0	7	5	1	
171	5998010861367WF	CCA SELF T HW VG74EBH	4	6	0	14	14	0	
172	5998010861368WF	CCA FIXED HW VG74EBK	0	0	0	0	0	0	
173	5998010954941WF	DIGITAL MU HW VG74EBD	0	0	0	6	4	0	
177	1280011127124WF	CIRCUIT CD HA WA75DDC	0	0	0	0	46	13	
179	5998010831364WF	LOGIC NO 2 HW VX74DDB	2	0	2	9	0	10	
180	5998010831365WF	LOGIC NO 1 HW VX74DDA	17	2	7	31	18	24	
181	5998010844614WF	MUX CD #1 HW VX74DDD	18	0	10	35	3	23	
182	5999010843485WF	MUX CD #2 HW VX74DDE	16	7	12	34	23	24	
183	5999010845103WF	LOGIC #3 HW VX74DDC	0	0	2	11	0	13	
187	5999010754766WF	MUX 1 HW VC74DAL	0	0	1	9	17	5	
191	5999011835277WF	CPJ HW VC74DAK	26	11	14	83	61	73	
192	5999011835278WF	INPT OPT 2 HW VC74DAH	13	7	7	30	28	31	
195	6605010764730WF	IN OUT 3 HW VC74DAJ	0	0	2	9	0	10	
198	5999010744138WF	ARITH ASSY HW NAS1FAC	0	0	0	0	0	5	
199	5999010796295WF	MUX BIAS HW NAS1FAF	0	0	2	1	0	9	
200	5999010846163WF	CIRCUIT CD HW NAS1FAB	0	0	1	0	0	6	
203	6610010617716WF	SCR OUTPT HW NAS1FAJ	0	0	3	3	0	6	
204	6610010617717WF	PARA CNVTR HW NAS1FAK	2	0	12	33	10	31	
205	6610010617718WF	DIG CONVTR HW NAS1FAL	0	0	0	0	0	0	
206	6610010617719WF	DIG CONVTR HW NAS1FAM	0	0	6	26	10	23	
208	6610010744139WF	PROG CONTL HW NAS1FAD	0	0	0	0	0	0	
209	6610010744140WF	TIME CONTL HW NAS1FAE	0	0	0	0	0	2	
220	5999011153930WF	CK CD ASSY HW BA14FBS	0	0	0	0	0	2	c
222	5999011253912WF	CK CD ASSY HW BA14FBR	0	0	0	0	0	0	c
223	6610011990760WF		0	1	1	0	1	1	c
231	5998010720096WF	SEQNR HW BA11ADD	5	10	6	18	19	11	
233	5998011311324WF	TORQUE HW BA14ADE	0	0	0	1	2	0	
234	5999010709400WF	CRT CARD HW BA14ADC	0	6	4	12	17	15	
244	5999011329473WF	CCA LOGIC HW BA14AAG	0	11	5	17	21	12	c3

Notes: NSN = national stock number; IM = item manager; ES = equipment specialist; WUC = work unit code; alloc = allocation; c = common to two LRUs; c3 = common to three LRUs; and c4 = common to four LRUs

Summary for Shop A

Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
3	1270010665979WF	POWER CONV HW VH74CAK	0	0	0	0	0	0	
12	1270010600697WF	PHASE SHIF HW VN74AAB	2	0	0	6	0	4	c
13	1270010615082WF	BUFFER,RES HW VN74AAC	7	0	0	21	3	17	c
14	1270010778077WF	AZ EL COMP HW VN74AAD	7	0	2	21	34	67	c
16	6130010778197WF	ANT PWR SU HW VA74AAE	0	0	0	0	0	0	
19	1270010830473WF	POWER SPLY HA VP74ACG	0	0	0	0	0	0	
22	5998011153249WF	BOARD ASSY HA VP74ACB	0	0	0	0	0	1	
27	1270010754768WF	PANEL CONT HA VF74BAC	0	0	2	0	0	16	
28	1270010754864WF	PDU HVPS HA VF74BAE	0	0	0	0	0	2	
29	5998011169554WF	VIDEO HA VF74BAF	0	0	0	0	0	0	
30	5998011178597WF	VIDEO DRV HA VF74BAJ	0	0	0	0	0	0	
31	6130010781818WF	PWR SUPPLY HA VF74BAG	0	0	0	0	0	0	
34	1270010770836WF	PWR SUPPLY HA VS74ABN	0	4	6	2	23	13	c4
38	1270011070140WF	SAMPLE ASS HA VS74ABC	0	5	6	0	44	32	c4
40	5998011015290WF	CONTROL BC HA VS74ABB	0	4	5	0	37	20	c4
48	1270010795459WF	EU LVPS HA VF74BCZ	1	0	0	8	4	1	c
51	1270011006495WF	ANALOG INP HA VF74BCG	0	0	0	5	0	1	c
60	5998011007335WF	ANLOG OUTP HA VF74BCH	2	0	0	10	5	0	c
61	5998011069835WF	LOSS CARD HA VF74BCU	0	0	0	13	0	0	
92	6130011464687WF	ANT PWR SU HW VN74AAE	0	0	0	0	0	0	
101	5998011034653WF	BOARD ASSY HA VT74AFJ	2	0	0	7	8	4	c4
103	5998011045867WF	CIRCUIT CD HA VT74AFH	5	0	6	13	0	14	c4
104	6130011034551WF	PWR SUPPLY HA VT74AFK	2	0	0	20	4	14	c4
109	1270012223783WF	PWR CONV HW VH	0	0	0	0	0	3	
132	1280011202059WF	CARD AY HA WA75DAE	0	0	2	6	0	27	
145	6130011001875WF	PWR SUPPLY HA WA75DCM	5	2	12	93	118	133	c
158	5841010756723WF	POW SUPPLY HW VG74EAF	5	0	3	44	16	7	
159	5841010659504WF	VERT DEFL HW VG74EAB	3	0	0	12	7	1	
160	5841011493220WF	HVPS FOCUS HW VG74EAG	0	0	0	2	6	0	
161	5960010844987WF	CRT ASSY HW VG74EAE	19	39	5	86	110	58	
162	5998010866714WF	CIRCUIT CD HW VG74EAA	4	0	0	13	7	2	
163	5998010999472WF	CKTCDASY HW VG74EAC	8	8	0	35	37	0	
174	5998010999473WF	CKT CD AY HW VG74EBA	7	25	3	51	69	16	
175	6130011146824WF	PWRSUPAY HW VG74EBM	25	26	16	78	90	47	
184	6130010785507WF	PWR SUPPLY HW VX74DDG	9	0	3	27	9	22	
185	6130010845069WF	PWR SPY #1 HW VX74DDF	3	0	1	13	5	13	
188	5999010754865WF	AC DC PWR HW VC74DAP	0	0	2	11	15	11	
189	5999010770725WF	PLAT ELEC HW VC74DAF	1	6	13	73	85	78	
190	5999010951078WF	HSI CCA HW VC74DAR	17	0	21	62	0	62	
193	6130010942037WF	DC/DC CARD HW VC74DAQ	0	0	1	4	5	9	
194	6605010754770WF	SERVO MOD HW VC74DAN	0	0	0	4	19	18	
201	6130010599017WF	POWER SUPP HW NA51FAQ	0	0	1	0	0	9	
202	6610010599121WF	ANLG OUTPT HW NA51FAH	0	0	0	1	0	7	
207	6610010690109WF	CKT ASSY HW NA51FAA	0	0	0	0	0	8	
211	5999010720084WF	A OF A MON HW VE14FBC	0	0	0	0	0	1	c
212	5999010720085WF	CCA LEF C/L HW VE14FBD	0	0	0	0	0	9	
213	5999010720086WF	CK CD ASSY HW BA14FBE	0	0	2	0	0	6	c
214	5999010720087WF	AUTOPLT #1 HW BA14FBF	0	0	0	0	0	0	
215	5999010720088WF	GUN COMP HW BA14FBG	0	0	0	0	0	1	c
216	5999010720089WF	PWR SUPPLY HW BA14FBJ	0	0	1	0	0	6	c
217	5999010720090WF	SLE TS REL HW BA14FBK	0	0	0	0	0	0	c
218	5999010720092WF	A OF A CMP HW BA14FBQ	0	0	0	0	0	0	c
219	5999010902639WF	CIRCUIT BD HW BA14F9A	0	0	0	0	0	2	
221	5999011169555WF	CK CD ASSY HW BA14FBM	0	0	3	0	12	5	c

Notes: NSN = national stock number, IM = item manager, ES = equipment specialist, WUC = work unit code, alloc = allocation, c = common to two LRUs, c3 = common to three LRUs, and c4 = common to four LRUs

Summary for Shop A (continued)

Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
225	5999011673719WF	CIRCUIT CD HW BA	0	0	0	0	0	0	
228	6615010714139WF	PRNTD CRCT HW VE14FCA	0	0	0	0	0	0	
230	5998010720093WF	DISC HW BA14ADA	0	0	3	9	1	6	
232	5998010720098WF	CKT CD HW BA14ADF	0	0	2	8	5	4	
237	5999011347206WF	CCA PITCH2 HW BA14AAD	0	1	1	0	9	7	c
238	5999011610638WF	CCA YAW HW BA14AAA	0	0	6	0	13	25	c3
239	5999011610641WF	CCA AUX HW BA14AAK	0	0	0	0	0	0	
240	5999011611146WF	CCA FLAP HW BA14AAF	0	0	0	0	0	1	c3
241	5999011611147WF	CCA HORIZ HW BA14AAE	0	1	0	1	9	3	c3
242	5999011611678WF	POWER SUP HW BA14AAH	0	0	4	2	1	16	c3
243	5999011626671WF	CCA ROLL HW BA14AAB	0	8	6	1	22	30	c3
245	6615010996864WF	CKT CD ASY HW BA14AAL	0	0	0	0	0	0	c3
246	6615011611679WF	CCA PITCHI HW BA14AAC	0	0	0	0	0	0	c
248	5999010996870WF	CKT CD ASY HW BA14AAD	0	0	0	0	0	0	

Summary for Shop M

Index	Master NSN	Description / IM ES WUC	Biweekly			Quarterly			SRU commonality indicator
			LMI model		HQ AFLC alloc.	LMI model		HQ AFLC alloc.	
			Base DI	All DI		Base DI	All DI		
15	1270010956768WF	DIGIBUS AS HW VA71AAA	0	0	0	0	0	3	c4
18	1270010830398WF	PRESS VSL HA VP74ACH	5	0	0	98	111	45	
20	1270010976096WF	DETECT ASY HA VP74ACC	11	8	3	25	22	27	
33	1270010770710WF	REC ASSY HA VF74ABD	26	0	26	129	71	213	
35	1270010854653WF	REF SOURCE HA VS74ABG	2	16	5	12	0	13	c4
36	1270010993205WF	PHASE LOCK HA VS74ABJ	0	2	13	3	0	39	
37	1270010993206WF	LOW NOISE HA VS74ABE	27	6	25	69	48	95	
42	1270010770724WF	REF SOURCE HA VS74ABG	11	41	5	59	54	18	
44	1270010860716WF	REF SOURCE HA VS74ABG	0	0	5	8	37	27	c4
46	1270010759674WF	REF SOURCE HA VS74ABG	0	0	4	4	35	17	
91	1270011465363WF	DIGIBUS AS HW VN74AAA	11	9	4	16	8	5	

Notes: NSN = national stock number; IM = item manager; ES = equipment specialist; WUC = work unit code; alloc = allocation; c = common to two LRUs, c3 = common to three LRUs, and c4 = common to four LRUs

CAPACITY CONSTRAINTS MATRICES

BIWEEKLY SUMMARIES								QUARTERLY SUMMARIES							
For Base DI Case								For Base DI Case							
Unused Capacity (hours) at Shop								Unused Capacity (hours) at Shop							
Shop	RF	CI	PP	A	D	DI	M	Shop	RF	PP	A	D	CI	DI	M
RF	-2	81	360	592	1,088	395	1,684	RF	-14	1,790	3,480	6,564	1,660	2,370	10,368
CI	-2	-4	282	459	1,031	395	1,684	PP	-14	-7	1,826	4,957	1,112	2,370	8,798
PP	-2	-4	-1	20	850	395	1,684	A	-14	-7	-9	3,994	391	2,370	7,649
A	-2	-4	-1	-7	850	395	1,684	D	-14	-7	-9	-5	13	2,178	4,646
D	-2	-4	-1	-14	-4	356	1,145	CI	-14	-7	-9	-5	0	2,171	4,600
DI	-2	-4	-1	-318	-86	-2	673	DI	-14	-7	-1,987	-465	0	-5	1,174
M	-2	-4	-1	-336	-86	-2	-12	M	-14	-7	-1,987	-471	0	-5	-4
Fraction of Total Capacity Used at Shop								Fraction of Total Capacity Used at Shop							
Shop	RF	CI	PP	A	D	DI	M	Shop	RF	PP	A	D	CI	DI	M
RF	1.004	0.843	0.000	0.487	0.188	0.000	0.038	RF	1.005	0.171	0.498	0.184	0.463	0.000	0.013
CI	1.004	1.008	0.217	0.603	0.231	0.000	0.038	PP	1.005	1.003	0.737	0.383	0.640	0.000	0.162
PP	1.004	1.008	1.003	0.983	0.366	0.000	0.038	A	1.005	1.003	1.001	0.503	0.873	0.000	0.272
A	1.004	1.008	1.003	1.006	0.366	0.000	0.038	D	1.005	1.003	1.001	1.001	0.996	0.081	0.558
D	1.004	1.008	1.003	1.012	1.003	0.099	0.346	CI	1.005	1.003	1.001	1.001	1.000	0.084	0.562
DI	1.004	1.008	1.003	1.275	1.064	1.005	0.615	DI	1.005	1.003	1.287	1.059	1.000	1.002	0.888
M	1.004	1.008	1.003	1.291	1.064	1.005	1.007	M	1.005	1.003	1.287	1.059	1.000	1.002	1.000
For All DI Case								For All DI Case							
Unused Capacity (hours) at Shop								Unused Capacity (hours) at Shop							
Shop	RF	PP	CI	A	DI	D	M	Shop	RF	PP	CI	A	D	DI	M
RF	-2	152	383	945	395	1,221	1,750	RF	0	1,044	2,299	5,455	7,399	2,370	10,500
PP	-2	-5	244	638	395	963	1,750	PP	0	-9	846	3,394	6,287	2,370	8,813
CI	-2	-5	-2	391	395	834	1,750	CI	0	-9	-1	2,502	4,994	2,161	8,261
A	-2	-5	-2	-9	395	767	1,733	A	0	-9	-1	-4	3,139	1,574	7,955
DI	-2	-5	-2	-154	-3	73	1,644	D	0	-9	-1	-35	0	1,437	7,843
D	-2	-5	-2	-170	-3	-4	1,627	DI	0	-9	-1	-1,319	-239	-17	5,083
M	-2	-5	-2	-170	-3	-4	-10	M	0	-9	-1	-1,342	-247	-17	0
Fraction of Total Capacity Used at Shop								Fraction of Total Capacity Used at Shop							
Shop	RF	PP	CI	A	DI	D	M	Shop	RF	PP	CI	A	D	DI	M
RF	1.004	0.578	0.256	0.182	0.000	0.089	0.000	RF	1.000	0.517	0.256	0.213	0.080	0.000	0.000
PP	1.004	1.014	0.526	0.448	0.000	0.281	0.000	PP	1.000	1.004	0.726	0.510	0.218	0.000	0.161
CI	1.004	1.014	1.004	0.661	0.000	0.378	0.000	CI	1.000	1.004	1.000	0.639	0.379	0.088	0.213
A	1.004	1.014	1.004	1.008	0.000	0.428	0.010	A	1.000	1.004	1.000	1.001	0.603	0.336	0.242
DI	1.004	1.014	1.004	1.133	1.008	0.946	0.061	D	1.000	1.004	1.000	1.005	1.000	0.394	0.253
D	1.004	1.014	1.004	1.147	1.008	1.003	0.070	DI	1.000	1.004	1.000	1.190	1.030	1.007	0.516
M	1.004	1.014	1.004	1.147	1.008	1.003	1.006	M	1.000	1.004	1.000	1.194	1.031	1.007	1.000

Notes:

- LRU shops are RF, CI, PP, and DI. SRU shops are A, D, and M.
- Entries in each row indicate unused capacity (first table) or fraction of unused capacity (second table) for the shop in that column when the shop shown in the first column in that row exceeded its capacity. (Negative numbers mean capacity was exceeded by that amount.)

APPENDIX C

GLOSSARY

AFLC	=	Air Force Logistics Command
ALC	=	Air Logistics Center
AWP	=	awaiting parts
CAT	=	cannibalization threshold
CONUS	=	continental United States
D041	=	Recoverable Consumption Item Requirements System
DRIVE	=	Distribution and Repair in Variable Environments
DSO	=	direct support objective
EBO	=	expected backorder
ENMCS	=	<i>expected not mission capable-supply</i>
EOQ	=	economic order quantity
HQ	=	Headquarters
HQ CAT	=	Headquarters cannibalization threshold
LMI	=	Logistics Management Institute
LRU	=	line replaceable unit
MASS	=	MICAP Asset Sourcing System
MHP	=	man-hours for Production
MHQ	=	man-hours for Quarterly
MIC	=	maintenance inventory center
MICAP	=	mission capability
MSOS	=	main source of supply
NAC	=	number of aircraft

NMC	=	not mission capable
NMCS	=	not mission capable-supply
NRTS	=	not reparable this station
NSN	=	national stock number
OPT	=	optimum
OPT CAT	=	optimum cannibalization threshold
OST	=	order-and-shipping time
PAA	=	primary aircraft authorized
QPA	=	quantity per application
RAF	=	Royal Air Field
RCI	=	repair commonality index
SRU	=	shop replaceable unit
SV	=	sort value
TH	=	time horizon
TRADES	=	Theater/Region Allocation/Distribution Execution System
USAF	=	U.S. Air Force
VMR	=	variance-to-mean ratio
WRSK	=	War Readiness Spares Kit
WSMIS	=	Weapon System Management Information System