

A Development of Logistics Management Models for the Space Transportation System

M. J. Carrillo, S. E. Jacobsen,
J. B. Abell, T. F. Lippiatt

35th
Year

Rand

R-3083-NASA

A Development of Logistics Management Models for the Space Transportation System

M. J. Carrillo, S. E. Jacobsen,
J. B. Abell, T. F. Lippiatt

September 1983

Prepared for the
National Aeronautics and
Space Administration

35th
Year



The research described in this report was sponsored by the National Aeronautics and Space Administration under Contract NAS10-10438.

ISBN 0-8330-0529-4

The Rand Publication Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

Published by The Rand Corporation

PREFACE

This report identifies the characteristics of logistics system capability assessment and stockage optimization methods that reflect the unique nature of the National Aeronautics and Space Administration's Space Transportation System's (STS) launch and recovery cycle. It presents the mathematical foundations of approaches to such methods and demonstrates their feasibility in the context of NASA's and the U.S. Air Force's need to develop a sound, well-formulated logistics support strategy for the STS program.

This is a final report on a research project, sponsored by NASA, the main purpose of which is to develop initial logistics methodologies relevant to NASA's STS. The report contains an extensive, non-technical summary that should be of interest to NASA and Air Force personnel involved in logistics support of the STS program. It should also be of interest to those concerned with technical aspects of logistics support and with mathematical derivations of the recommended methods.

SUMMARY

NASA and the U.S. Air Force currently face policy decisions of fundamental importance to the formulation of a sound, coherent, logistics support strategy for the Space Transportation System (STS). The importance of these decisions is clear. They not only involve large amounts of money, but they also will shape the essential character, quality, and cost-effectiveness of STS logistics support.

The policy decisions to be made include determination of the maintenance concept for the STS, i.e., the location, depth, and scope of component repair, levels of maintenance, and repair responsibility; the modes of transportation to be used for retrograde and serviceable component shipments; the proportion of component repairs to be done at each level of maintenance; the amount of investment in tools and equipment and its allocation; the amount of investment in reparable and consumable spare parts; and the spares stock level, by location, of each of the system's components.

Clearly, these decisions are interrelated. The computation of a spares posture depends on component characteristics, such as repair times and transportation times, that are the products of other policy variables, such as maintenance concept and level-of-repair decisions. NASA needs to understand how alternative maintenance concepts, choices of repair locations, repair level decisions, and transportation modes affect, for example, spares investment requirements and launch capability as a function of those investments. The complexity and interdependencies of the decisions suggest the need for a logistics

system capability assessment methodology that would enable NASA to evaluate policy alternatives as they affect a direct and meaningful measure of system performance, such as expected launch delay, and to do so in full light of the costs of those alternatives. Implicit in such a capability assessment methodology is the need for explicit representation of the relationship between system performance and spares investment level.

To assist NASA in meeting these needs, this report identifies the characteristics of an analytical modeling capability that would relate logistics support decisions and resource requirements to the capability to meet STS launch schedules, taking into account the unique characteristics of the shuttle program with its small fleet size and tight recovery and launch schedule. The report also presents feasible analytic approaches to both the capability assessment and the spares optimization problems. Because such analytic capability is only as good as the input data, the report discusses the quality and availability of data within NASA.

THE UNIQUE DIMENSIONS OF THE STS PROBLEM

NASA's STS Program differs sharply from previous NASA programs, such as Apollo. STS is NASA's first program with a relatively high launch rate of reusable vehicles; as a result, NASA has been faced with developing logistics policies to support a program that represents a significant departure from those supported previously. The analytical methodologies and decision aids available to support logistics decisions were developed for sustained military operations. But the STS Program, unlike the military, has a very small fleet size and a tight, but well-defined, launch schedule, both of which impinge upon the development of sound logistics policies.

Most of the capability assessment and spares optimization methods that are readily available were developed for military aircraft, which are usually deployed in large numbers and generate a relatively large number of sorties. A typical performance measure used by these methods is the number of mission-ready or mission-capable aircraft. Because one cannot predict when these aircraft must be mission capable, "ready" is taken to mean "ready at a random point in time."

Shuttle operations, by contrast, are characterized by a small fleet size (four or five) and a relatively low sortie rate of about 20 per year during full-scale operation. Thus, measures relating to launch delay rather than to the number "ready" are probably more relevant to shuttle operations. The problems of determining stockage requirements and logistics system capability using a measure of effectiveness directly related to launch delay have three distinguishing features:

1. The shuttle vehicle is required to be ready not at all points in time, but within a given number of time units from the beginning of prelaunch shuttle recovery process.
2. The prelaunch operations plan specifies a project network of activities to be carried out. Given this plan, it should be possible to identify the points in the schedule where demands for a particular part might occur.
3. The effect of a part shortage on launch delay depends not only on how long the shortage exists but also on (1) when in the schedule the demand occurs, (2) when the demand must be filled, and (3) the repair time of the part, which is a function of the basic repair level decision for the part in question.

The following example is presented to clarify the issues discussed above. It should not, of course, be construed as representing a typical shuttle prelaunch operation schedule.

Figure S.1 is a project network whose start is at node 1 and whose end is at node 7. The nodes represent events in time and the arcs represent activities. The orientation of each arc is from the lower numbered node to the higher numbered node. The numbers on the arcs represent activity times. For instance, activity (5,6) requires 5 time units and activity (3,4) requires 13 time units. The project network also reflects mandatory precedence relations. For instance, activity (3,4) may not begin until activity (2,3) is completed, and activity (6,7) may not begin until both activities (4,6) and (5,6) are completed. Note that the earliest possible completion time of this project is given by the length of the longer of the two paths beginning at node 1 and ending at node 7. The longest path in a project network is called the *critical path* and its length is the project duration. In the figure, the critical path follows the upper path and its total length is 44.

Suppose that there are two line replaceable units, LRU_1 and LRU_2 , that can fail and possibly delay the project. Assume that LRU_1 can fail at node 2 and that, if it does fail, it must be replaced before activity (5,6) can begin. We say that LRU_1 has node 2 as its *demand node* and node 5 as its *fill node*. Suppose that LRU_2 also has node 2 as its demand node, and has node 3 as its fill node. Finally, suppose that the repair time for LRU_1 is 16 and the repair time for LRU_2 is 5. Now, if there is no spare for either LRU, a failure of LRU_1 will delay the start of activity (5,6) by 10 time units but will not delay the project. That




Fig. S.1--An illustrative network with two items subject to failure and associated resupply times

is, the project duration is still 44 time units. However, if LRU_2 fails, the project duration will be increased by 3 time units to a total of 47, thus causing a delay. Under a typical ready-rate optimization, via marginal analysis, it is clear that there are conditions where spares of LRU_1 will be stocked, but perhaps none of LRU_2 , even though LRU_2 is the part most likely to cause a delay. For instance, if the failure rate for LRU_1 were larger than that for LRU_2 , and the cost of LRU_1 less than the cost of LRU_2 , the stockage derived by marginal analysis would be greater for LRU_1 than for LRU_2 . Moreover, aside from condemnation spares, the stockage of both items would cost more than is needed to minimize delay.

Additionally, if one is interested in deciding which repair times should be shortened, ready-rate or probability-of-sufficiency (POS)

optimization methods, if used to determine the tradeoff between stockage costs and repair time reduction costs, will generally lead to incorrect results. For instance, under the assumptions above, such an analysis will conclude that there is greater payoff in reducing LRU_1 repair time than there is in reducing LRU_2 repair time. Of course, the project network approach demonstrates that the exact opposite is true.

It is clear then that models developed for military aircraft operations are probably not appropriate for the shuttle, and that logistics system capability assessment and stockage optimization models for STS operations should have two important characteristics: (1) measures of effectiveness should be related to launch delay, and (2) explicit considerations should be given to the prelaunch task network.

CAPABILITY ASSESSMENT

There are two general approaches to developing a logistics capability assessment model that best represents the STS environment: Monte Carlo simulation and analytic queueing modeling. Each approach has its advantages and disadvantages. Simulation is very little limited by the amounts of detail that can be incorporated, but it usually requires a large input data base and frequently becomes slow and costly to run. In addition, because of random variations in any particular run, many computer runs are required to obtain a valid mean or standard deviation of any output measure. Therefore, especially for low-failure-rate-items, simulation may be unsuitable for use as a spare requirements methodology or for making sensitivity analyses. Indeed, the potentially threatening effects of low-failure-rate items on launch delay are usually difficult to expose through standard Monte Carlo simulation experiments.

Analytic queueing models are more difficult to derive and compromises usually have to be made about the amount of detail they can incorporate. But they definitely have an advantage because of their ease of use. A single run can generate means and standard deviations, and dealing with many types of spare parts is not a problem. These types of models, then, are often more suitable for logistics system capability assessment, spares requirements computation, and sensitivity analyses.

This report describes a new analytic queueing approach that relates stockage levels, repair level decisions, and the project network schedule of prelaunch operations directly to the probability distribution of launch delay. As a result, this approach can produce several measures based upon stock levels, repair, and transportation performance. These include expected delay and the probability of launch delay and its variance. Given appropriate inputs, it will also yield expected delay costs, as well as a rank-ordered list of those components most schedule-threatening.

The approach was developed under a strict set of assumptions about the shape and character of the prelaunch task network. It is, however, appropriate for networks similar to the Air Force Test and Evaluation Center (AFTEC) network of shuttle recovery tasks. Although the theoretical feasibility of this approach has been demonstrated, the approach requires full evaluation with a detailed shuttle network that reflects the demand and fill nodes for each component.

SPARES STOCKAGE OPTIMIZATION

The report describes two approaches to spares optimization under a budget constraint that have objective functions related to launch delay and that consider the prelaunch task network. The first is directly related to the approach used for capability assessment and is similarly restricted to networks such as those used by AFTEC. This approach is also relatively complicated and may require considerable computer time. The second approach is generalizable to any network and is relatively easier to use, but it reflects less information about the probability distribution of delay.

Because of the importance of launch delay as an objective function and in view of the simple network example discussed earlier, either of these approaches will perform better than the more conventional ready-rate or POS models. How much better remains to be evaluated, however, when more complete network and component characteristics data can be made available.

DATA ISSUES

The value of *any* logistics planning model is largely dependent on the quality of data used to run it. The data presently available to support logistics decisions were based on a comparability study of heavy aircraft components. That study provided estimates of maintenance demand rates (MDRs), i.e., the rates of component removals that generate demand for spares. There are two fundamental problems with these MDRs. The first is that they implicitly assume knowledge of component operating time, and operating time is not now routinely recorded. The second problem is not unique to the shuttle vehicle provisioning

problem, or even to NASA operations in general: initial estimates of component characteristics are notoriously unreliable. Thus, the orbiter initial provisioning problem is complicated by a severe paucity of useful, reliable information.

Discussions with NASA and its contractors have revealed no plan to use component removal or failure data from the first six orbiter flights to revise the initial MDRs. Furthermore, no systematized collection of component failure or removal data or operating time has been implemented. We infer from this that NASA intends to make logistics decisions on the basis of the initial estimates alone. A retrospective study of the F-16 initial provisioning problems, which is summarized in this report, suggests that, if the initial estimates alone do constitute the basis for those decisions, the performance of the resulting stockage posture could probably be achieved at dramatically less cost (or, conversely, given a specified investment level, the performance of the stockage posture could be dramatically improved) if an alternative strategy were used that took advantage of the body of theory that has emerged from the initial provisioning scenario.

The same F-16 study demonstrated that initial MDR estimates could be dramatically improved by revising them with sparse, initial operational data using Bayes-Lin techniques. On the basis of these results it suggested the NASA revise its current initial estimates using data from the early shuttle flights, and that it establish a program to continue the revision process as well as collecting or estimating operating time. It is also suggested that the uncertainty surrounding these MDR estimates be explicitly considered in the logistics planning process and models. The modeling approaches developed in this report accommodate this uncertainty.

MANAGEMENT ISSUES

NASA does not now have the data collection systems, the analytic modeling capability, or the management controls it needs for effective logistics management. NASA's current data systems, for example, do not allow systematic recording or estimation of subsystem or component operating times. Yet the only available estimates of component demand rates require at least estimated operating times to be useful. Furthermore, the data systems currently in use in the STS program are largely contractor-operated and lack interface. Spares requirements and logistics support policy recommendations are made by each major contractor independently, using models that may not be appropriate to the unique logistic support problem that the STS launch and recovery environment presents.

It seems to us important that NASA and the Air Force continue to develop, implement, and use the kinds of decision aids discussed in this report within their management framework. The following paragraphs offer some specific recommendations.

RECOMMENDATIONS

The models presented here demonstrate that it is feasible to develop improved analytical logistics modeling capabilities that will relate logistics support decisions and resource requirements to the capability to meet STS launch schedules. The evaluation and demonstration of the payoff to be gained from using these techniques remain to be done. It is recommended that these actions be accomplished in two phases because of the large amount of data that would be required for a full-scale evaluation. Recommendations concerning improved data

collection and parameter estimation are detailed below. These are important regardless of the logistics modeling methodologies that NASA may ultimately choose to use.

Improved Data Collection and Parameter Estimation

It is recommended that NASA:

1. Modify or design and implement an integrated data collection system that would routinely provide up-to-date component removal data, repair times, repair level distributions, retrograde shipment times, order-and-ship times, condemnation rates, procurement and repair costs, procurement lead times, and operating times or usage estimates.
2. Assemble whatever data are available (from either formal or informal systems) from previous STS flights and compare these data with those parameters currently being used for logistics planning and resource requirements computations. From this, judge whether revisions to initial estimates are required. If such is the case, as is likely, revise the initial estimates using the Bayes-Lin technique suggested in this report or a similar Bayesian technique. Continue this revision process as more flight experience is gained.
3. Estimate the uncertainty surrounding component removal rates and other logistics system performance parameters, and explicitly consider them in making logistics policy decisions (e.g., level of repair decisions) and in determining spares requirements.

Phase I: Initial Prototype Development and Evaluation

It is recommended that the evaluation and full-scale development and implementation of the logistics system capability assessment and spares optimization methodologies be carried out in two phases because of the difficulty in obtaining the necessary, detailed STS recovery task network and component data. The first phase would focus on the evaluation of these methodologies, using a limited set of representative components. If the outcome of the first phase is positive, the second phase would refine the techniques and implementation. For Phase I the following steps are recommended:

1. To the extent technically feasible, extend the methodologies presented in Secs. II and III to include non-Poisson processes with finite populations; multiple stockage points, including Vandenberg Air Force Base and other possible stockage sites; and variance-to-mean ratios other than unity.
2. Identify a subset of components that, to the extent possible, represent the population of all components from each of the projects (orbiters, boosters, tank and main engine). At the same time develop criteria to determine the range of components that should generally be considered in such models.
3. For that select subset develop the detailed network data corresponding to the prelaunch schedule of operations, including the demand and fill nodes for each component, and collect the most up-to-date component data.

4. Determine whether the network representation for these components is compatible with the assumptions inherent in these methodologies. If there are compatibility problems, develop and evaluate network editing techniques that could allow these methods to be used.
5. Evaluate and demonstrate the use of the capability assessment methodology using simulation for comparisons as appropriate.
6. Evaluate and compare the two stockage optimization techniques, in terms of launch delay or stockages costs, presented in this report (a) with each other and (b) with those techniques currently in use by NASA, using the capability assessment model or simulation as appropriate.

Phase II: Prototype Improvement and Implementation

If the Phase I evaluation results are positive, the following steps for Phase II are recommended:

1. Modify and improve the methodologies based on the Phase I results. In addition, improve them, to the extent feasible, so that they will be suitable for individual projects as well as for overall system assessment, will consider the availability of manufacturing assets, and will include indentured components (SRUs, etc.).
2. Define and assess their potential use for integrated logistics management of logistics operations, level of repair analyses, development of out-year requirements, and procurement and budget decisions.

3. Develop and implement a full-scale system.

The implementation of these several steps can be expected in the longer run to deliver significantly more cost-effective logistics support to the STS program than NASA's current plans.

CONTENTS

PREFACE	iii
SUMMARY	v
FIGURES	xxi

Section

I. INTRODUCTION	1
The Unique Dimensions of the STS Program	2
An Illustrative Example	8
II. CAPABILITY ASSESSMENT	13
An Analytic Model for Capability Assessment	14
Development of the Model	16
Distribution of the Conditional Waiting Time	19
Distribution of the Unconditional Waiting Time	22
The Distribution of Delay	23
Performance Measures Derived from the Delay Random Variable	25
Adapting the Model To Deal with Data Variability	25
Adapting the Model To Deal with Finite Servers and a Finite Source Population	30
Summary	32
III. SPARES STOCKAGE OPTIMIZATION	35
Spares Optimization with the AFTEC LCOM Network	37
An Approach to More General Networks	47
Summary	54
IV. DATA ISSUES IN SPARES PROVISIONING	56
The F-16 Case	57
Summary and Recommendations	65

V. SUMMARY AND RECOMMENDATIONS	67
The Need for Capability Assessment	67
The Spares Stockage Optimization Problem	68
Data Issues	69
Management Issues	70
Recommendations	71
REFERENCES	75

FIGURES

S.1. An Illustrative Network with Two Items Subject to Failure and Associated Resupply Times	ix
1. An Illustrative Network	9
2. An Illustrative Network with Two Items Subject to Failure and Associated Resupply Times	10
3. An Abstraction of a Shuttle Project Network	18
4. A Network with Correlated Path Lengths	48
5. Typical Delay Distribution of an Item	50
6. Typical Path in a Project Network	52
7. Improvements in Availability Versus Cost Using the Bayes-Lin Technique	64

I. INTRODUCTION

The National Aeronautics and Space Administration (NASA) and the U.S. Air Force currently face policy decisions of fundamental importance to the formulation of a sound, coherent, logistics support strategy for the Space Transportation System (STS). The importance of these decisions is clear. Not only do they involve large amounts of money, but they will also shape the essential character, quality, and cost-effectiveness of STS logistics support.

The policy decisions to be made include the determination of the maintenance concept for the STS, i.e., the location, depth, and scope of component repair, levels of maintenance, and repair responsibility; the modes of transportation to be used for retrograde and serviceable component shipments; the proportion of component repairs to be done at each level of maintenance; the amount of investment in tools and equipment, and its allocation; the amount of investment in reparable and consumable spare parts; and the spares stock level, by location, of each of the system's components.

Clearly, these decisions are interrelated. The computation of a spares posture depends on component characteristics, such as repair times and transportation times, that are the products of other policy variables, such as maintenance concept and level-of-repair decisions. NASA needs to understand how alternative maintenance concepts, choices of repair locations, repair level decisions, and transportation modes affect, for example, spares investment requirements and launch capability as a function of those investments. The complexity and

interdependencies of the decisions suggest the need for a logistics system capability assessment methodology that would enable NASA to evaluate policy alternatives as they affect a direct and meaningful measure of system performance, such as expected launch delay, and to do so in full light of the costs of those alternatives. Implicit in such a capability assessment methodology is the need for explicit representation of the relationship between system performance and spares investment level.

To assist NASA in meeting these needs, this report identifies the characteristics of an analytical modeling capability that will relate logistics support decisions and resource requirements to the capability to meet STS launch schedules, taking into account the unique characteristics of the shuttle program with its small fleet size and tight recovery and launch schedule. The report also presents feasible analytic approaches to both the capability assessment and the spares optimization problems. Because such an analytic capability is only as good as the input data, a discussion of the quality and availability of data within NASA is also presented.

THE UNIQUE DIMENSIONS OF THE STS PROGRAM

NASA's STS Program differs sharply from previous NASA programs, such as Apollo. STS is NASA's first program with a relatively high launch rate of reusable vehicles; as a result, NASA has been faced with developing logistics policies to support a program that is significantly different from those it supported previously. The analytical methodologies and decision aids that are readily available to support logistics decisions have been developed for sustained military operations. The STS Program departs from such operations, however, for,

unlike the military, it has a very small fleet size and a tight, but well defined, launch schedule, both of which have an impact on the development of sound logistics policies.

Typically, military aircraft are deployed in fairly large numbers and generate relatively large numbers of sorties. For defense purposes, it is clearly desirable to hold "ready" as many of these aircraft as possible. Because one cannot predict when it may be necessary to have these aircraft mission capable, "ready" is taken to mean "ready at a random point in time."

In most spares requirements models, such as METRIC [1], VSL (Variable Safety Level, the Air Force's implementation of METRIC), MOD-METRIC [2], or Dyna-METRIC [3,4], supply performance is indirectly related to the requirement that aircraft be ready at a random point in time. Brooks, Gillen, and Lu [3] show that spares stockage results for models that emphasize inventory system performance (e.g., minimize expected backorders) or total system performance (e.g., minimize the expected number of grounded aircraft) yield essentially the same mixes of spares within a weapon system under the requirement of relatively high levels of readiness, at random points in time, or relatively high budget levels. In fact, the probability of experiencing no grounded aircraft is the same as the probability of no backorders; this is often called the "ready rate" of the system, or the operational rate. The system ready rate is the product of several probabilities, each of which is called the probability of sufficiency (POS) or ready rate of an item.

Such a measure of system performance may be quite appropriate in many situations but may not be appropriate to the STS. A given shuttle vehicle (orbiter, solid rocket booster and external tank), for instance,

will actually be on the ground a large percentage of the time; therefore, it is questionable whether STS stockage decisions should be based upon such measures of performance as system ready rate.

Determining spares requirements by stocking an item such that its POS strikes at least a prescribed level, or, in some sense, by maximizing the system ready rate (e.g., via marginal analysis [5,6]), is not an appropriate method for the STS.

A typical shuttle vehicle will be on the ground much of the time, and the times at which it will be required for launch will be determined by the launch schedule. In particular, it is not required that a shuttle vehicle be "ready" at a random point in time. We believe, therefore, that appropriate spares stockage postures and repair level decisions may be quite different from those currently emerging from the STS community.

The decisions that NASA must make are fundamentally important to the logistics support of the STS. Yet NASA is not now able to assess the effects of those policy decisions on launch capability or expected launch delays. Spares requirements and repair level decisions are meant to secure the launch schedule; therefore, such decisions should explicitly recognize this purpose. Current models are not appropriate for assessing such decisions because of their focus on steady state readiness rather than, say, expected launch delay.

This report demonstrates that the full development of an assessment methodology that corrects shortcomings of current models is not only feasible but of great importance to NASA in understanding the effects on launch delay of the various logistics policy decisions currently being made. We do not necessarily recommend that the methodology be adopted

by the various NASA contractors, but rather that NASA itself undertake full development and use of this methodology to better understand the effects of spares requirements and repair level decisions upon launch availability and the required costs to support a specified level of launch availability.

Repair level analyses usually attempt to formulate a minimum-cost repair level decision for a given item or group of items. Such analyses are usually not connected to measures of system performance, and are usually only loosely connected to spares stockage methods. In the STS environment, as discussed above, the spares stockage methods currently in use are not connected to measures of system performance. In this report, we develop spares requirements methods that are directly related to system performance, an obvious improvement over methods currently in use. Moreover, with our methods, the effects of repair level decisions on system performance can be evaluated directly, as can their effects upon spares costs, given a specified level of system performance. Current methods cannot do this.

Since the various shuttle vehicles are not required to be ready at a random point in time, current methods of estimating the effects of spares requirements and repair level decisions upon system performance, or the costs required to achieve a given level of system performance, are not applicable. To illustrate the point somewhat simply, consider a system composed of only 100 components and assume that the item POS is the same for each item. If each item POS is 0.95, then the system ready rate is $(0.95)^{100} = 0.0059$. On the other hand, if the decisionmaker wishes the system ready rate to be at least 0.95, then, if all the item POSs are equal, each must be at least $(0.95)^{1/100} = 0.99949$. However,

such a POS may be expensive to attain; moreover, it may be unnecessary. For instance, a failed item may be repaired before the launch schedule is threatened. The example points out again that it is inappropriate to try to determine stock levels of individual items, or appropriate spares investment levels, without methods that explicitly relate such levels to system performance.

Shuttle operations, during the operational phase, will consist of four or five vehicles making approximately 20 flights per year. Thus, STS operation is distinctly different from that of typical military aircraft. As a result, logistics decision tools need to differ somewhat from those developed by and for the Air Force.

Typical prelaunch operations of a given shuttle vehicle include inspect, transfer, assemble, test, fuel, and checkout. These activities take place at several locations and frequently result in the discovery of malfunctions, which may in turn result in demands for replacement parts for the orbiter, boosters, or external tank. Often such parts are expensive, and, because repair facilities are often remote from the location of demand, relatively long resupply times may occur. Generally, replacement parts are available from stock on hand, local repair, cannibalization of other vehicles, remote repair by the manufacturer, or depot-level repair.

During prelaunch operations, total demands for parts for a given shuttle vehicle are likely to decrease as the prelaunch schedule approaches the launch date, since a large number of the malfunctioning items will have been discovered earlier in the prelaunch preparation process. On the other hand, the potential penalty for not having a spare of a demanded item or for not being able to rapidly repair tends

to increase as the launch date approaches. However, it is not correct to concentrate only on the countdown phase just before launch, or on any other particular phase of the prelaunch schedule, to make appropriate stockage and repair level decisions. For instance, the later phases of prelaunch operations provide some information for stockage decisions but are relatively useless for appropriate remote repair level decisions, because time has become increasingly critical and most repairs undertaken at that point, to fill a part demand, are more likely to delay the shuttle launch. Similarly, if repair times are in fact quite short, the early phases of the prelaunch schedule provide little information for stockage decisions. Clearly, the entire prelaunch schedule of operations must be considered carefully to make useful stockage and repair level decisions.

The problems of selecting stockage and repair level policies and evaluating them, in terms of a measure of effectiveness directly related to launch delay, have three distinguishing features.

1. The vehicle in question is required to be ready not at a random point in time, but merely within a given number of time units from the beginning of prelaunch operations.
2. The prelaunch operations plan specifies the project network of activities to be carried out. Therefore, given the prelaunch operations plan, it is possible to identify the points in the schedule where demands for a particular part might occur. For instance, a demand for a certain type of valve may occur during checkout of the propulsion system but not during checkout of the guidance system.

3. The effect of a part shortage on delay depends not only on the length of time the shortage exists but also on (1) where in the prelaunch schedule of operations the demand occurs, (2) where in the schedule the demand must be filled, and (3) the repair time of the part, which, in turn, is the result of the basic repair level decision for the part in question.

AN ILLUSTRATIVE EXAMPLE

The following simple example is presented to clarify the issues discussed above. It should not, of course, be construed as representing a typical shuttle prelaunch operation schedule.

Figure 1 is a project network, the start of which is at node 1 and the end of which is at node 7. As usual, the nodes represent events in time and the arcs represent activities. The orientation of each arc is from the lower numbered node to the higher numbered node. The numbers on the arcs represent activity times. For instance, activity (5,6) requires 5 time units and activity (3,4) requires 13 time units. The project network also reflects mandatory precedence relations. For instance, activity (3,4) may not begin until activity (2,3) is completed, and activity (6,7) may not begin until both activities (4,6) and (5,6) are completed. Note that the earliest possible completion time of this project is given by the length of the longer of the two paths beginning at node 1 and ending at node 7. The longest path in a project network is called the *critical path* and its length is the project duration. In Fig. 1, the critical path follows the upper path and its total length is 44.

Suppose that there are two line replaceable units, LRU_1 and LRU_2 , which can fail and possibly delay the project. In particular, assume that LRU_1 can fail at node 2 and, if it does fail, that it must be replaced before activity (5,6) can begin. We say that LRU_1 has node 2 as its *demand node* and node 5 as its *fill node*. Suppose that LRU_2 also has node 2 as its demand node, and that it has node 3 as its fill node. Finally, suppose that the repair time for LRU_1 is 16 and the repair time for LRU_2 is 5. Figure 2 represents the situation. Note that if there is no spare of either LRU, then a failure of LRU_1 will delay the start of activity (5,6) by 10 time units but will not delay the project. That is, the project duration is still 44 time units. However, if LRU_2 fails, the project duration will be increased by 3 time units to a total




Fig. 1 — An illustrative network




Fig. 2 — An illustrative network with two items subject to failure and associated resupply times

of 47. Under ready-rate optimization, it is clear that there are conditions where spares of LRU_1 will be stocked, but perhaps none of LRU_2 . For instance, if the failure rate for LRU_1 were larger than that for LRU_2 , and the cost of LRU_1 less than the cost of LRU_2 , then the stockage would be greater for LRU_1 than for LRU_2 . Moreover, aside from condemnation spares, the stockage of both items would spend more than is needed to minimize delay (i.e., additional project duration time).

The following will illustrate the point. Let the failure rates and unit costs of LRU_1 and LRU_2 be, respectively, $\lambda_1 = 0.25$, $C_1 = 10$ and $\lambda_2 = 0.05$, $C_2 = 20$. Suppose that the resupply times are $t_1 = 16$ and $t_2 = 5$, and assume a budget of 40. A simple marginal analysis [5] shows that the amounts of each item to stock so as to maximize the ready rate, subject to the budget constraint, are $S_1 = 4$ and $S_2 = 0$. Thus, the

ready rate optimization stocks four units of the item whose failure does not, and none of the item whose failure does, threaten to lengthen the project duration.

Additionally, if one is interested in deciding which repair times should be shortened, then ready-rate or POS optimization methods, if used to determine the tradeoff between stockage costs and repair time reduction costs, will generally lead to incorrect results. For instance, under the assumptions above, such an analysis will conclude that there is greater payoff in reducing LRU_1 repair time than there is in reducing LRU_2 repair time. Of course, the project network approach demonstrates that the exact opposite is true.

Generally, an LRU type may have more than one demand node. For instance, assume that there are two units of LRU_2 . One of these units has node 2 as its demand node and node 3 as its fill node, and the other unit has node 6 as its demand node and node 7 as its fill node. If there is one spare unit of LRU_2 in stock, then, neglecting remove and replace times, one or two failures of LRU_2 cannot delay the project. To see this, note that, if the LRU_2 at node 2 does not fail, then the spare of LRU_2 can be used to cover the possible failure of the other LRU_2 at node 6. On the other hand, if the LRU_2 at node 2 does fail, it will be replaced by the spare, and the failed unit will be repaired in 5 time units. If there is one spare unit of LRU_2 , the failed LRU_2 will be repaired in time to become an available spare to cover the possible failure of the other LRU_2 component at node 6. Such is not the case if the repair time of LRU_2 is 25 time units, for, if there is one spare LRU_2 , the project will be delayed by four time units owing to failures of LRU_2 at nodes 2 and 6.

The above simple example captures the flavor of the complicated relationships among the prelaunch schedule of operations, repair times, and stock levels, and their effects on project delay (i.e., launch delay). Moreover, since repair level analyses affect the component repair times of failed components, repair level decisions should fully recognize these complicated interactions. Herein lies some of the power of capability assessment.

In the remainder of this report, we discuss the development of models that focus on launch delay in contrast to optimization models that limit themselves to ready rate. Specifically, in Sec. II we discuss capability assessment applications, and in Sec. III the spares optimization problem. Section IV presents a discussion of data issues in initial provisioning. A summary and recommendations are provided in Sec. V.

II. CAPABILITY ASSESSMENT

Two general approaches to developing a logistics capability assessment model to represent the STS environment are Monte Carlo simulation [7,8] and analytic queueing modeling. Each has its advantages and disadvantages. Simulation is, to an extent, not limited by the amounts of detail that can be incorporated; but it usually requires a large input data base and frequently becomes slow and costly to run. In addition, because of the random variation in any particular run, many computer runs are required before a valid mean or standard deviation of any output measure can be obtained. Therefore, especially for low-failure-rate items, this approach may become unsuitable for use as a spare requirements methodology or for doing sensitivity analyses because of the many runs required for a detailed examination of just one of the many parts in a logistics system. Indeed, the potentially threatening effects upon launch delay of low-failure-rate items are usually difficult to expose through standard Monte Carlo simulation experiments.

Analytic queueing models are more difficult to derive, and compromises usually have to be made about the amount of detail they can incorporate. When trying to derive an analytic model, one is forced to seek an abstraction of the relationships among the various system components. This is a useful exercise on its own and is easy to overlook in simulation. If available, analytic models definitely have an advantage because of their ease of use. One single run can obtain means and standard deviations, and dealing with many parts is not a

problem. These models, then, are more suitable for capability assessment, spares requirements computation, and sensitivity analyses.

There are some simulation models in use by either NASA or the Air Force for limited purposes (e.g., JFK's ARTEMIS, MFSC's BOSIM and extensions, AFTEC's LCOM shuttle simulation); but, for the reasons stated above, we believe that a general analytic capability assessment model is needed which represents the unique characteristics of the STS environment.

AN ANALYTIC MODEL FOR CAPABILITY ASSESSMENT

The following discussion describes in non-mathematical language the characteristics, capabilities, and limitations of the model developed in subsequent pages. The model is intended to serve NASA's need for capability assessment. Although it is specific to a particular network structure, i.e., that of the AFTEC LCOM network [12,13], it can be extended to other, more complex networks. It is not, however, a general network model. Its implementation would require validation as well as additional data describing certain characteristics of the NASA project network. It is, nevertheless, a promising approach that we believe deserves additional research because of its potential for quantifying the effects of alternative policy decisions and clarifying the interrelationships among them in terms of both cost and launch capability.

Given a network structure similar to that of the AFTEC LCOM network, which can be represented as two major sets of tasks essentially in parallel followed by a third, the model will compute the probability of launch delay, the probability that the delay will be less than any specified value, and the expected length of delay as functions of spares

stockage levels, repair times, and transportation times. The computations are network-specific. They take explicit account of task times and slack times in the network. Thus, the computations enable the model to estimate the effects of changes to the network as well as policy alternatives.

In addition to requiring descriptions of tasks in the network and their expected durations and interdependencies, the model requires data that describe the points in the network where parts of each type can be demanded, and the points in the network at or before which parts of each type are required to preclude subsequent task delays. Data describing parts characteristics (demand rates, costs, repair times, etc.) are also required.

The model uses a steady-state approximation to launch delay by focusing on the delay of a typical shuttle vehicle that has just entered the prelaunch schedule of operations. Of course, as stated above, the probability distribution of this delay depends critically upon important parameters of the logistics system such as demand rates, spares stockage levels, and repair times.

Additionally, the model derived in this report is a non-cannibalization one. We chose to develop such an initial model for two main reasons: (1) there seems to be a great deal of uncertainty with respect to those items that are interchangeable from one shuttle vehicle to another, and (2) a computational model that takes explicit account of cannibalization will, with respect to relatively low-failure-rate items, recommend that cannibalization be routinely considered and utilized. We feel, therefore, that although cannibalization is a management strategy available to NASA, its incorporation into a computational model is inappropriate at this time.

The discussion that follows describes the development of the model in necessarily mathematical terms. The general reader is referred to the end of this section for a summary discussion of the analytical approach developed thus far, and for recommendations for further development and evaluation.

DEVELOPMENT OF THE MODEL

Consider a queueing system in which part failures cause arrivals of customers and completions of part repairs cause departures. Let the state of the system, $Z_i(t)$, be the number of parts of type i in the system at time t , e.g., waiting for repair or in repair. $\{Z_i(t), t \geq 0\}$ is a stochastic process with state space $\{0, 1, 2, \dots\}$.

Let the process have Poisson input with identically distributed exponential interarrival times with mean $1/\lambda'_i$. The repair time random variable, Y_i , has a general distribution with finite mean $1/\mu_i$. Also assume that the arrival and repair processes are independent.

We will assume for now an infinite number of servers and an infinite calling population, though results can be obtained similarly for other cases. Let S_i be the stock level for part type i . If X_i denotes the state $Z_i(t)$ of the system at a random point in time, then by Palm's theorem [10,11], in the steady state, X_i has a Poisson distribution with mean λ'_i/μ_i .

We proceed by finding the distribution of waiting time (until a part is available) W_i given that a failure of part i was just detected. We then obtain the distribution of U_i , the unconditional waiting time. Groups M_j , $j = 1, 2, \dots$, of part types are selected so that their part failures (if any) are discovered at the same time and required later to

be operational at the same time. We then find the distribution of U^j , the aggregate waiting time for the parts in the group M_j ; that is,

$$U^j = \max\{U_k, k \in M_j\}.$$

Each group of parts, M_j , has a fixed delay, K_j , in addition to the random waiting time, U^j . In addition, we define K_0 to be a constant representing certain characteristics of a NASA shuttle network, as discussed below.

Some of the groups of random arcs are arranged in parallel with respect to other groups in the NASA project network, and some in series. After consideration of the AFTEC network for the LCOM model [12,13] and the STAR reports [14], we are optimistic that the actual duration of the NASA shuttle network can be accurately represented by the delays in any number of the groups of arcs in parallel, plus the delay in one group in series. (Refer to Fig. 3.)

From the AFTEC and STAR nets, the $\{M_j\}$, and hence the corresponding $\{U^j\}$, correspond, in increasing order, to orbiter unscheduled maintenance before the orbiter integration test; to SRB unscheduled maintenance before and during SRB and ET stacking; and to unscheduled maintenance before and during the SRB/orbiter mating and shuttle integration test. It is clear that there is also a group of deterministic arcs, corresponding to scheduled maintenance tasks that could be represented by constants, $\{K_j\}$. K_0 corresponds to payload processing; K_1 , K_2 , and K_3 correspond to scheduled maintenance tasks.

With groups M_1 and M_2 in parallel and M_3 in series, the delay random variable, D , is given by

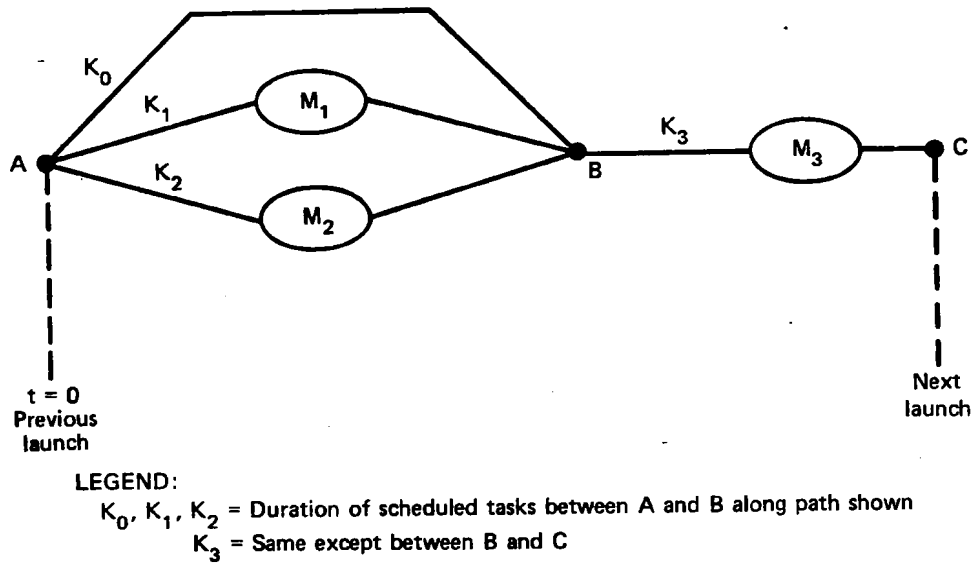


Fig. 3 — An abstraction of a shuttle project network

$$D = \max\{K_0, K_1 + U^1, K_2 + U^2\} - K^* + U^3,$$

where the constant $K^* = \max\{K_0, K_1, K_2\}$. Note that $K^* + K_3$ is the duration of the project network if no failures are discovered.

The payoff of this approach is derived from the presence of the $\{K_j\}$, which generate slack times for some of the $\{M_j\}$. These slack periods allow time to get parts from repair/resupply, not just from stock on hand, without delaying the project network. This also emphasizes the potential payoff of expedited repair and transportation for selected parts.

The elements of $\{K_j\}$ are obtained from the scheduled tasks in the AFTEC LCOM network (slightly modified) by finding the longest path between two nodes in a directed network, which is the basis of the critical path method of project networks (see Ch. 5 of [15], for example), and which enabled us to derive the graph of Fig. 3.

DISTRIBUTION OF THE CONDITIONAL WAITING TIME

Consider the above stochastic process in steady state. This is equivalent to assuming that a group of shuttle vehicles has been flying sorties for some period of time at the frequency expected in the operational period, currently estimated to be 20 or more flights per year.

We want to find the cumulative distribution of the waiting time, W , until a part is available to satisfy a demand, given that its failure was just discovered. To simplify our notation, we drop the subscript i corresponding to part type i in this subsection. Thus we wish to find $\text{Prob}[W \leq w]$. The following priority rule will be assumed for the current model.

Priority Rule. Parts entering the repair process lose their vehicle identity, and vehicles take advantage of the spares stock protection on a first-come, first-served basis.

Thus, since Y denotes the repair time and S the stock level,

$$\text{Prob}[W \leq w] = \text{Prob}[Y \leq w] \quad \text{if } S = 0.$$

That is, when there is no stock protection, waiting time is equal to repair time; otherwise we condition on the actual number in the pipeline, X , before the arrival. If there is some stock on hand, there is no wait for parts; that is,

$$\text{Prob}[W \leq w \mid X = x] = 1, \quad \text{if } 0 \leq x < S;$$

otherwise,

$$\text{Prob}[W \leq w \mid X = x] = G_{jk}(w)$$

where $j = x + 1$, $k = S$, and $x \geq S > 0$, and G_{jk} is defined as follows:

$G_{jk}(w)$ = Probability of at least one transition into state k
(i.e., k in repair/resupply) within w time units,
given that state j was just entered due to an
arrival and the system is in steady state.

Note that we have used the priority rule stated above, in that a vehicle will not have its demands satisfied unless all demands ahead of it are satisfied. It follows that

$$\text{Prob}[W \leq w] = \text{Prob}[X < S] + \sum_{x=S}^{\infty} G_{x+1,S}(w) \text{Prob}[X = x]$$

Now, the probability mass function of X is Poisson with mean λ'_i/μ_i , since we are looking at the system at a random point in time. On the other hand, the distribution G_{jk} is difficult to obtain, because it is affected by subsequent incoming failures. For the problem at

hand, however, many failure rates are rather small. So we obtain an approximation, H_{jk} , $k < j$, by assuming no further arrivals; thus, $H_{jk}(w) \leq G_{jk}(w)$; i.e., H_{jk} is conservative. In using this approximation with no more arrivals, the repair rate becomes proportional to the number currently in repair, i.e., μ ; if we were to allow additional arrivals, the repair rate would be greater. With this approximation, the failure that just arrived can be repaired according to the distribution of Y ; however, items already undergoing repair have a remaining repair time given by the equilibrium distribution for the distribution function of Y (see Takacs [11], p. 161):

$$A(w) = \mu \int_0^w \text{Prob}[Y > y] dy$$

We now derive an explicit expression for H_{jk} . First, let R_{nj} denote the distribution function of the j th order statistic, $j = 1, 2, \dots, n$, of the equilibrium distribution, A ; n is the sample size. Then,

$$R_{nj}(w) = \sum_{k=j}^n \binom{n}{k} A^k(w) [1 - A(w)]^{n-k} \quad \text{for } A(w) < 1,$$

where $\binom{n}{k}$ is the binomial coefficient. Note that

$$R_{nj}(w) = 1 \quad \text{for } A(w) = 1.$$

Furthermore, let $R_{nj} = 1$ for $j = 0$. It follows that

$$H_{jk}(w) = \text{Prob}[Y \leq w] R_{j-1, j-1-k}(w) + \text{Prob}[Y > w] R_{j-1, j-k}(w),$$
$$k = 0, 1, \dots, j-1.$$

Note that, for Y exponential, the equilibrium distribution, A , is equal to the distribution function of Y . In this case, $H_{jk} = R_{j, j-k}$. In the case of exponential repair, $\{Z(t), t \geq 0\}$ becomes a special class of continuous-time Markov chains called the birth and death process, for which a substantial amount of literature exists. For Y deterministic, $A(w)$ is uniform, reaching unity when w reaches $1/\mu$.

DISTRIBUTION OF THE UNCONDITIONAL WAITING TIME

Let $\{N_i(t), t \geq 0\}$ be the counting process of cumulative discovered failures by time t of parts of type i for the vehicle in question, and let $t = 0$ at the time of the last launch. Assume that the present vehicle had all of its parts operational at $t = 0$. It remains to be determined what time should be counted in t : flying time, power-on time, time equivalent of cycles, etc. This problem is discussed further in Sec. IV.

Note that the failure rate λ'_i of the Poisson input of $\{Z_i(t), t \geq 0\}$ is a multiple of the λ_i used here, where the factor is the number of shuttle vehicles in operation adjusted by any other system, such as the Shuttle Avionics Integrated Laboratory (SAIL), that generates part failures.

To obtain the distribution of the $\{U_i\}$, we simply condition on whether a failure has been discovered by time t , where t denotes the time the present vehicle landed, or some time during the refurbishing

process in a turnaround network in preparation for the next launch.

Thus, for $u \geq 0$,

$$\text{Prob}[U_i \leq u] = \text{Prob}[N_i(t) = 0] + \text{Prob}[N_i(t) > 0] \text{Prob}[W_i \leq u],$$

and equals zero otherwise. In giving an overview of the model, we defined

$$U^j = \max\{U_k, k \in M_j\}.$$

It follows that

$$\text{Prob}[U^j \leq u] = \prod_{k \in M_j} \text{Prob}[U_k \leq u].$$

It is clear that this approach can be further generalized to allow each part type k to have a constant component V_k , since in such case,

$$\text{Prob}[U^j \leq u] = \prod_{k \in M_j} \text{Prob}[U_k \leq u - V_k], \quad \text{for } u \geq \max\{V_k, k \in M_j\}$$

and equals zero otherwise.

THE DISTRIBUTION OF DELAY

We discussed above how some STAR and AFTEC networks led us to a representation of the shuttle network with a delay given by

$$D = \max\{K_0, K_1 + U^1, K_2 + U^2\} - K^* + U^3$$

where the constant $K^* = \max\{K_0, K_1, K_2\}$. In the last section, we showed how to obtain the distributions of the $\{U_i\}$. The presence of groups of arcs in parallel is no problem; however, having the resulting distribution of this parallel system connected in series with another group of arcs may cause technical problems unless the parts of the latter group are different from those in the parallel groups. If such is the case, then the resulting distribution is determined by the convolution

$$\text{Prob}[D \leq e] = \int_0^e \text{Prob}[U^{12} \leq K^* + e - y] d(\text{Prob}[U^3 \leq y]),$$

where

$$\text{Prob}[U^{12} \leq K^* + u] = \text{Prob}[U^1 \leq K^* - K_1 + u] \text{Prob}[U^2 \leq K^* - K_2 + u]$$

for $u \geq 0$.

If there are parts in common among the groups of arcs in series, then convolutions are not appropriate, since the independence assumption would no longer hold. However, with or without independence, the expected value of the delay, $E[D]$, can always be obtained as the sum of the expected values of the components in series, i.e.,

$$E[D] = E[U^{12}] + E[U^3]$$

As a practical matter, the independence assumption can be ignored and

the convolution approach can still be used. However, further study is required to evaluate the validity of this approximation and its domain of applicability.

PERFORMANCE MEASURES DERIVED FROM THE DELAY RANDOM VARIABLE

Once the probability distribution of D is found, various performance measures can easily be derived: for example, the expected launch delay, the variance or standard deviation of launch delay, the probability of no launch delay, the expected launch delay time given that there *is* a delay, and the probability that launch delay is less or greater than any specified value. An interesting extension is possible if there exists some delay time, say D_0 , such that, if the launch of some shuttle of interest is delayed longer than D_0 , the launch of the next shuttle will also be delayed; then the probability distribution of D also yields the probability that D will exceed D_0 , that is, the probability that the next shuttle will also be delayed.

Note, too, that if a penalty function, $P(d)$ dollars, due to a shuttle launch delay of d days, can be specified, the expected penalty per shuttle launch due to delays can also be computed directly from the probability distribution of D .

ADAPTING THE MODEL TO DEAL WITH DATA VARIABILITY

In the case of simple Poisson input, the number of parts in the repair pipeline, X_i , has a Poisson distribution; therefore, the variance-to-mean ratio for X_i equals unity. In cases where the the quality of the data is in question, it may be appropriate to model this lack of confidence by using a probability distribution whose variance-to-mean ratio is greater than unity. In fact, attempts have been made to estimate this ratio using Bayesian methods [1].

A suitable model to incorporate this ratio parameter is to assume a compound Poisson process input with a logarithmic compounding distribution (see Feller [16], Sherbrooke [17]). In this case, the number of parts in the repair pipeline has a negative binomial distribution. We will show below that results obtained above easily generalize to accommodate this approach.

Also, when using Bayesian methods to incorporate recent history in the estimates of revised failure rates, it is a common approach to assume a gamma "prior" distribution. This approach is discussed in Sec. IV. Again the resulting distribution is negative binomial.

In preparation for future use of either of the above approaches, we generalize the above results for compound Poisson input. We emphasize that the modeling of this type of input is merely a device that can serve either the variance-to-mean ratio representation or the Bayesian approach. In particular, we do not assume that in the shuttle environment parts failures occur in batches. To be sure, further study and comparisons are needed in this area to evaluate the impact of either approach.

We now generalize the above results. For simplicity, we again drop the subscript i , the part type.

A stochastic process, $\{X(t), t \geq 0\}$, is said to be a compound Poisson process if it can be represented, for $t > 0$, by

$$X(t) = \sum_{k=1}^{N(t)} I_k$$

where $\{N(t), t \geq 0\}$ is a simple Poisson process with mean

$$m(t) = E[N(t)] = (\lambda'p/(1 - p)) t, \quad 0 < p < 1,$$

and $\{I_n, n = 1, 2, \dots\}$ is a family of independent random variables having a common distribution called the compounding distribution. The Poisson process and the compounding distribution are assumed to be independent. $\{N(t), t \geq 0\}$ is the cumulative number of customers that arrived by time t , whereas I_n denotes the demands that the n th customer brings. By a conditioning argument, it follows that $X(t)$ has the compound Poisson distribution

$$\text{Prob}[X(t) = x] = \sum_{n=0}^{\infty} \text{Prob}[N(t) = n] \text{Prob} \left[\sum_{k=0}^n I_k = x \right],$$

where

$$\text{Prob} \left[\sum_{k=0}^n I_k = x \right]$$

is obtained by the n -fold convolution of the distribution of I_k .

If the compounding distribution is logarithmic, that is,

$$\text{Prob}[I_n = y] = (1 - p)^y / (y \log p^{-1}) \quad 0 < p < 1,$$

then the distribution of $X(t)$ can be shown to be negative binomial [16] with parameter $m(t)$ and density function

$$\text{Prob}[X(t) = x] = \binom{m(t)+x-1}{x} p^{m(t)} (1-p)^x,$$

where $m(t) = E[N(t)] = (\lambda'p/(1-p))t$. The mean

$$E[X(t)] = m(t)(1-p)/p = \lambda't,$$

and the variance

$$\text{Var}[X(t)] = m(t)(1-p)/p^2.$$

Therefore, the variance-to-mean ratio equals $1/p$.

Feeney and Sherbrooke [18] proved that, with compound Poisson input and arbitrary repair distribution, the number in the steady state pipeline has a compound Poisson distribution, provided that all the demands that a customer brings complete repair at the same time:

$$\text{Prob}[X = x] = \sum_{n=0}^{\infty} \text{Prob}[X' = n] \text{Prob}\left[\sum_{k=0}^n I_k = x\right],$$

where X' , the number of customers in the pipeline, has a simple Poisson distribution with mean $(\lambda'p/(1-p))/\mu$, and $\{I_n\}$ correspond to the compounding distribution. X is the total number of parts still in repair.

As before, if the $\{I_n\}$ have a logarithmic compounding distribution, the resulting distribution of the number in the pipeline under steady

state is negative binomial, but this time with parameter $(\lambda'p/(1-p))/\mu$ instead of $m(t)$.

The results for the delay random variable generalize for the compound Poisson process. If we assume the exponential distribution, for $w \geq 0$,

$$\text{Prob}[W \leq w] = \text{Prob}[Y \leq w] \quad \text{if } S = 0.$$

For $S > 0$, we condition on the actual number of customers, X' , in the pipeline, before an arrival that just occurred. This yields

$$\text{Prob}[W \leq w] = \sum_{x=0}^{\infty} \text{Prob}[W \leq w \mid X' = x] \text{Prob}[X' = x].$$

Consider H_x defined as follows:

$H_x(w)$ = Probability that, within w time units, the number of total parts (the sum of all batches) in repair does not exceed S , the stock level, given that there are x parts already in repair, and one more customer just arrived. Also, no more arrivals are assumed to occur.

Clearly,

$$\text{Prob}[W \leq w \mid X' = x] \geq H_x(w).$$

It follows that

$$H_x(w) = \sum_{j=0}^{x+1} \binom{x+1}{j} A^j(w) (1 - A(w))^{x+1-j} \text{ Prob} \left[\sum_{m=0}^{x+1-j} I_m \leq S \right]$$

A comparable expression can be obtained for the case of a general repair distribution.

ADAPTING THE MODEL TO DEAL WITH FINITE SERVERS AND A FINITE SOURCE POPULATION

A major difference between the STS environment and that of the Air Force is the number of sorties and flying hours involved in the flying program. Obviously, this number is relatively small for the STS case. The difference is relevant in that, for the Air Force case, one can be more easily convinced that using the infinite calling population assumption is appropriate. But it may be necessary, in making the infinite population assumption, to restrict the domain of validity of the model for the STS environment to a smaller class of components that share certain characteristics.

A second assumption made in most logistics models is that of slack repair capacity, or, in more precise terms, of an infinite number of repair servers. We believe that for the STS, this assumption may be of no major consequence because of the low expected maintenance demand rates, unless repair times are excessively long. Both of these issues are addressed in the following paragraphs.

First, recall that the results shown previously were derived for failures arriving according to a simple Poisson process, whereas the repair was assumed to have a general distribution with an infinite number of servers. Also, an infinite population source was implicitly

assumed. We now take the initial steps necessary to generalize the above results by discarding some of these restrictions.

We assume the exponential repair distribution to obtain a birth and death process for which many results are available [15,19]. In particular, we will look at the number of parts in the queueing system.

The probability distribution for the number in the queueing system under steady state is known for any number of servers (see Hillier and Lieberman [15], pp. 397-399). For example, the distribution is geometric for the single server case:

$$\text{Prob}[X_i = x] = (1 - \rho'_i) (\rho'_i)^x,$$

where $\rho'_i = \lambda'_i / \mu_i$.

For the single server and finite source population, the distribution of the number in the system is also found in Hillier:

$$\text{Prob}[X_i = 0] = 1 / \left(\sum_{n=0}^M [M! / (M - n)! (\rho_i)^n] \right),$$

and

$$\begin{aligned} \text{Prob}[X_i = k] &= M! / (M - k)! (\rho_i)^k \text{Prob}[X_i = 0], \\ k &= 1, \dots, M, \end{aligned}$$

where M is the size of the source population (in some cases, M will be the number of shuttle vehicles) and ρ_i is the ratio λ_i / μ_i .

For the case of infinite servers and a finite source population, the distribution can be derived from the steady state birth and death equations (see [15], p. 394). In this case, we have

$$\text{Prob}[X_i = 0] = 1/(1 + \rho_i)^M,$$

and

$$\text{Prob}[X_i = k] = M!/(k!(M - k)!) (\rho_i)^k \text{Prob}[X_i = 0],$$
$$k = 1, \dots, M.$$

This finite source model can also be incorporated into an analytical network approach, but additional development and evaluation are required.

SUMMARY

Part of the research problem called for the identification of requirements for an analytical modeling capability that would relate logistics support decisions and resource requirements to the capability to meet launch schedules. Several different approaches were assessed and the most promising has been presented in this section. We have not only shown the theoretical feasibility of the approach, but have also performed the initial steps of the computational model development. In this summary, we review what has been accomplished and what remains to be done.

One of the model requirements is that it incorporate the project network structure information, including slack times. The network must include all the major shuttle components: orbiter, solid rocket boosters, and external tank. The example in Sec. I shows that the network can, at times, have an overriding impact on the spare parts quantity and mix required to meet a particular launch goal. The modeling approach developed in the present section assumes certain

network characteristics and is not generalizable to all networks. It is applicable only to networks of the type shown in Fig. 3, which is similar to the AFTEC LCOM network. That network, however, does not provide sufficient detail about the demand and fill nodes for individual components, and this approach remains to be evaluated using a more detailed network.

Another part of the research problem requires that the model provide a probabilistic launch delay measure along with resulting statistics, since such a measure deals specifically with the ability to launch on time. These measures are needed to estimate the potential costs of shuttle launch delays caused by an imbalanced logistics support policy.

An important characteristic of the STS program, in contrast to Air Force logistics situations, is its small fleet size and flying program. The model presented here has the potential to incorporate the small fleet size. In this section we have laid the groundwork for dealing with this problem. The preliminary results presented here can be used to estimate the sensitivity and impact of small fleet size on the proposed shuttle operations, and can be incorporated in our capability assessment approach as necessary.

Section IV, below, points out some problems with data which we assume will be corrected over time. Meanwhile, the uncertainty in the currently available data should be considered explicitly in the model. The approach presented here can incorporate any variance-to-mean ratio as a distribution parameter.

We believe that we have developed a potentially powerful assessment approach. We trust that after completion of the additional development

and evaluation tasks described above, the resulting model can be used for STS logistics policy studies.

This section derives the approximate probability distribution of the time one must wait for a given part. The section also demonstrates how data uncertainty as well as finite repair capacity and finite source population can be incorporated so as to derive a more appropriate probability distribution of the time one must wait for a given part.

We recommend that NASA undertake a computational development of this distribution. Such a development, together with a representative network corresponding to the prelaunch schedule of operations, would provide an excellent estimate of the delay of a typical shuttle vehicle as a function of spares stockage levels, repair and transportation times, failure rates, and the underlying network.

Of course, we recognize that not all items need necessarily be considered and that a separate study should be made to discover those items that should be included in such models.

III. SPARES STOCKAGE OPTIMIZATION

The problem discussed in this section has two elements. The first is to find a method for estimating the relationship between spares investment level and a direct, meaningful measure of system performance, such as expected launch delay. An explicit representation of that relationship would allow NASA to determine the spares investment level required to support any specified level of system performance, or, conversely, to specify the desired level of performance in full light of its costs.

The second element of the problem is to ensure that, for each level of performance, the required spares investment level is minimal. Each of these components of the problem implies the other. What is needed, then, is not only an explicit representation of the relationship between performance and cost, but one in which each point is an optimum in the sense that it represents the least-cost mix of spares for the specific level of performance, and, conversely, represents the best possible performance for the specific level of investment.

The computation of such a relationship depends on estimates of component characteristics that emerge from definition of the system's maintenance concept and repair level decisions, and the quality of the estimated relationship depends on the quality of the estimates of component characteristics. Issues related to the quality of available component-level data are discussed in Sec. IV, where we recommend several important steps for NASA to take to improve those data.

Here we address the problem of computing the optimal spares stockage posture for any specified level of spares investment. We develop approaches to two forms of this problem for the specific network structure discussed in Sec. II--one in which we minimize expected delay subject to a budget constraint, and one in which we maximize the probability that delay does not exceed a specified length of time, given a budget constraint. Item resupply times are inputs to these computations; therefore, repair level decisions can be evaluated in terms of spares investment requirements and system performance.

It is important to realize that the computation of a stockage posture is optimal for some specified set of component characteristics. However, those characteristics are largely determined by selection of transportation mode; location, depth, and scope of repair; tool and equipment investment levels; test equipment software capability; and other characteristics of the logistics system. Therefore, each computed estimate of the performance/cost relationship must be viewed as pertinent only to one set of assumptions about the system, its dimensions, and its operating characteristics. The models developed here are useful and powerful, but only when applied to the problem in the perspective of the entire logistics system. For example, one might compute the least-cost stockage posture required to deliver some specified level of expected launch delay based on a set of reasonable assumptions about repair times. Then one could vary the repair times to reflect alternative levels of repair, repair locations, or transportation modes, and observe explicitly how the cost of the stockage posture changed to achieve the same level of performance as

before. Also, if one could estimate reasonably the costs of the alternatives in repair and transportation, a "least-cost" solution could be approximated to the combination of interdependent decision problems. This gives some sense of the integrated view of the problem that is needed to employ these models most constructively.

In the remainder of this section, we develop the spares optimization models. Again, because the discussion is necessarily mathematical, we encourage the general reader to continue reading at the start of Sec. IV.

SPARES OPTIMIZATION WITH THE AFTEC LCOM NETWORK

As described in Sec. II, the delay random variable of a NASA project network can be written as

$$D(S) = \max\{K_0, U^1(S^1) + K_1, U^2(S^2) + K_2\} + U^3(S^3) - K^*$$

where the stock vector is $S = (S^1, S^2, S^3)$ and S^j is the vector of stock corresponding to components in M_j , $j = 1, 2, 3$. We are assuming that M_3 has no items in common with M_1 or M_2 . We have made the dependence on stock, S , explicit, because that is the variable to be optimized. In Sec. II we found that

$$\begin{aligned} \text{Prob}[D(S) \leq e] = & \int_0^e \text{Prob}[U^1(S^1) \leq K^* - K_1 + e - x] \text{Prob}[U^2(S^2) \leq K^* - K_2 + e - x] \\ & d(\text{Prob}[U^3(S^3) \leq x]), \end{aligned}$$

where $K^* = \max\{K_0, K_1, K_2\}$.

Two Measures of Effectiveness

We will examine two measures of effectiveness, ME1 and ME2.

ME1: One approach is to allocate stock according to the expected delay given by

$$E[D(S)] = \int_0^{\infty} \text{Prob}[D(S) > y] dy.$$

Given a budget level $B > 0$ for stock, an appropriate optimal stockage allocation can be determined by solving the stockage optimization problem

$$\min_S E[D(S)],$$

subject to

$$\sum_j C_j S_j \leq B,$$
$$S_j \geq 0, S_j = 0, 1, 2, \dots; j = 1, 2, \dots$$

That is, S is within the budget B . C_j is the unit procurement cost of stock of type j .

ME2: Another standard approach is to choose some time e and allocate stock so as to maximize the probability that project delay does not exceed e . The stockage optimization problem is then

$$\max_S \text{Prob}[D(S) \leq e],$$

subject to S , which is within the budget B .

Solution Methods for ME1 and ME2

Both stockage optimization problems, ME1 and ME2, have nonseparable objective functions and are not amenable to standard solution methods. However, we have developed approaches to these problems that are computationally feasible for problems with a large number of items to be stocked.

In the first place we state a result to be found in [20].

Fact: Let $b(y;S)$ be a positive function for each y and each stockage vector S . If two stockage vectors S' and S'' can be found so that

$$\int b(y;S') \log b(y;S'') dy > \int b(y;S') \log b(y;S') dy,$$

then

$$\int b(y;S'') dy > \int b(y;S') dy .$$

We now show how this fact can be used to develop solution methods for ME1 and ME2.

Solution of ME1: Since

$$E[D(S)] = E[\max\{K_0, U^1(S^1) + K_1, U^2(S^2) + K_2\}] + E[U^3(S^3)] - K^*,$$

we may write

$$\begin{aligned} E[D(S)] = & \int_0^{\infty} (1 - \text{Prob}[U^1(S^1) \leq K^* - K_1 + y] \text{Prob}[U^2(S^2) \leq K^* - K_2 + y]) dy \\ & + \int_0^{\infty} (1 - \text{Prob}[U^3(S^3) \leq y]) dy - K^*, \end{aligned}$$

and, therefore, for a suitably large L we have

$$\begin{aligned} E[D(S)] = & L - \int_0^L \text{Prob}[U^1(S^1) \leq K^* - K_1 + y] \text{Prob}[U_j(S_j) \leq K^* - K_2 + y] dy \\ & + L - \int_0^L \text{Prob}[U^3(S^3) \leq y] dy - K^*. \end{aligned}$$

Therefore, the optimal stockage allocation for problem ME1 may be found by solving the problem

ME1:

$$\begin{aligned} \max_S \quad & \int_0^L \left(\prod_{j \in M_1} \text{Prob}[U_j(S_j) \leq K^* - K_1 + y] \prod_{j \in M_2} \text{Prob}[U_j(S_j) \leq K^* - K_2 + y] \right) dy \\ & + \int_0^L \left(\prod_{j \in M_3} \text{Prob}[U_j(S_j) \leq y] \right) dy, \end{aligned}$$

subject to

$$\sum_{j \in M_1 \cup M_2} C_j S_j + \sum_{j \in M_3} C_j S_j \leq B,$$

$$S_j \geq 0, \text{ integer-valued.}$$

Problem ME1 is a nonseparable problem in the integer stock variables; hence, standard solution methods, such as dynamic programming and marginal allocation, do not directly apply. However, the Fact described above is directly applicable.

For ease of explanation, let the budget $B > 0$ be divided into two parts, $B^1 > 0$ and $B^2 > 0$, so that $B = B^1 + B^2$.

We now define a subproblem, FM1, of the original problem.

FM1:

$$\max_S \int_0^L b(y;S) dy,$$

subject to S^1 , where S^2 is within the budget B^1 , and

$$b(y;S) =$$

$$\prod_{j \in M_1} \text{Prob}[U_j(S_j) \leq K^* - K_1 + y] \prod_{j \in M_2} \text{Prob}[U_j(S_j) \leq K^* - K_2 + y],$$

and, for FM1, it is understood that the stock vector S includes only those items j such that $j \in M_1 \cup M_2$.

We now use the results from [20], stated above, in solving this subproblem.

Stockage Algorithm: Select any stock vector S' that satisfies the budget constraint. Solve

FM1(S'):

$$\max_S \int_0^L b(y;S') \log b(y;S) dy$$

subject to S^1 , S^2 within the budget B^1 . Let S'' be optimal. The hope is that this subproblem is easier to solve. From the Fact above, if

$$\int_0^L b(y; S'') \log b(y; S'') dy > \int_0^L b(y; S') \log b(y; S') dy,$$

then

$$\int_0^L b(y; S'') dy > \int_0^L b(y; S') dy.$$

That is, S'' is better than S' for the problem FM1, the one we are interested in. Set $S' = S''$ and resolve FM1(S').

It remains only to determine whether FM1(S') can easily be solved. By the definition of this problem, we may rewrite FM1(S') as

FM1(S'):

$$\max_S \sum_{j \in M_1 \cup M_2} u_j(S_j; S')$$

subject to S^1, S^2 within the budget B^1 , where

$$u_j(S_j; S') = \int_0^L b(y; S') \log \text{Prob}[U_j(S_j) \leq K^* - K + y] dy,$$

and K equals K_1 or K_2 , depending on j being in M_1 or M_2 . Therefore, the objective function for FM1(S') is a separable function in the stockage decision variables S_j , $j \in M_1 \cup M_2$, and standard methods, such as dynamic programming, or perhaps marginal allocation, can be used to solve FM1(S').

The point is that the difficult stockage problem FM1 can be approached by solving a sequence of easier problems of the form FM1(S'). This same idea can then be applied to the second part of ME1, written as

SM1:

$$\max \int_0^L \prod \text{Prob}[U_j(S_j) \leq y] dy,$$

where S^3 is within the budget B^2 . Therefore, optimal solutions for ME1 can be found relatively easily. Of course, it is not at all obvious how a given budget, B, should be optimally partitioned into B^1 and B^2 . That is, to determine how much of the budget should be allocated to stock in M_1 and M_2 , and how much should be allocated to stock in M_3 , one must begin by essentially solving the overall problem. This situation is not uncommon, however, and various "resource directive" strategies can be employed to yield insights into a nearly optimal decomposition of the budget into B^{1*} and B^{2*} . In particular, an approximate Lagrange multiplier for FM1, with respect to B^1 , can be found that represents the gradient or rate of change of the optimal value of FM1 with respect to changes in B^1 . Similarly, an approximate Lagrange multiplier for SM1, with respect to B^2 , can also be found. One then compares these two multipliers to determine the direction of change for B^1 and B^2 . Since $B^1 + B^2 = B$ must hold, one budget will be increased and the other decreased. We are hopeful that this procedure will quickly provide an optimal, or nearly optimal, decomposition of the budget B into B^{1*} and B^{2*} .

Solution of ME2

Problem ME2 is also not separable in the stockage decision variables. However, the method discussed above for expected delay minimization is applicable.

Briefly, let S' be a stockage vector that satisfies the budget constraint. Define

$$b(x; S') =$$

$$\text{Prob}[U^1(S'^1) \leq K^* - K_1 + e - x] \text{ Prob}[U^2(S'^2) \leq K^* - K_2 + e - x]$$

$$\cdot d\{\text{Prob}[U^3(S'^3) \leq x]\},$$

where $d\{\}$ is the probability density function of U^3 . We then solve

FM2(S'):

$$\max_S \int_0^L b(x; S') \log b(x; S) dx,$$

where S is within the budget B , to find an improved stockage allocation S'' , as in the solution procedure for ME1. The objective function for problem FM2(S') can then be rewritten as

$$\begin{aligned}
 & \sum_{j \in M_1} \int_0^L b(x; S') \log \text{Prob}[U_j(S_j) \leq K^* - K_1 + e - x] dx \\
 & + \sum_{j \in M_2} \int_0^L b(x; S') \log \text{Prob}[U_j(S_j) \leq K^* - K_2 + e - x] dx \\
 & + \int_0^L b(x; S') \log d\{\text{Prob}[U^3(S^3) \leq x]\} dx.
 \end{aligned}$$

The first and second sums are separable in the stockage variables S_j , $j \in M_1 \cup M_2$. Therefore, standard procedures such as dynamic programming or marginal allocation may be employed with respect to these variables. However,

$$d\{\text{Prob}[U^3(S^3) \leq x]\} = \sum_{j \in M_3} r_j(x; S_j) \prod_{j \in M_3} \text{Prob}[U_j(S_j) \leq x],$$

where

$$r_j(x; S_j) = d\{\text{Prob}[U_j(S_j) \leq x]\} / \text{Prob}[U_j(S_j) \leq x].$$

This implies that

$$\begin{aligned} & \int b(x; S') \log d\{\text{Prob}[U^3(S^3) \leq x]\} dx \\ &= \sum_{j \in M_3} \int b(x; S') \log \text{Prob}[U_j(S_j) \leq x] dx \\ &+ \int b(x; S') \log \sum_{j \in M_3} r_j(x; S_j) dx \quad . \end{aligned}$$

The first term after the equals sign is separable in the stockage decision variables S_j , $j \in M_3$, but the second one is not. However, we are confident that approximations can be made so that the ideas or strategies presented for problem ME1 will remain valid for ME2.

AN APPROACH TO MORE GENERAL NETWORKS

If the project network corresponding to the prelaunch schedule of operations cannot be well represented by a project network of the type presented in Sec. II, then the probability distribution of project network delay cannot easily be written in terms of the individual probability distributions of delay corresponding to individual line replaceable units. Moreover, even when the probability distribution of delay can be written for more general networks, the stockage optimization problems, in terms of this distribution, are likely to prove intractable. We therefore present the following stockage optimization procedure, in terms of arbitrary project networks, to demonstrate that certain optimization methods can overcome the above-mentioned difficulties.

For instance, Fig. 4 represents a project network, where each X_i is a random variable representing delay of the corresponding arc. There are three paths, the length of each being a random variable. The length of the first path is

$$P_1 = X_1 + X_4;$$

that of the second path is

$$P_2 = X_2 + X_5;$$

and that of the third path is

$$P_3 = X_1 + X_3 + X_4.$$




Fig. 4 — A network with correlated path lengths

The time random variable for project duration is then

$$\begin{aligned} T &= \max \{P_1, P_2, P_3\} \\ &= \max \{X_1 + X_4, X_2 + X_5, X_2 + X_3 + X_4\}. \end{aligned}$$

Here, P_1 and P_3 are correlated and P_2 and P_3 are also correlated. Therefore, the probability distribution of T is not the product of the probability distributions of each path length.

The example above demonstrates the need for alternative stockage optimization procedures, given the results of repair level analyses, when the project network corresponding to the prelaunch schedule of operations is of a more general nature than that described in Sec. II.

The standard deviation of delay of an individual LRU is relatively large compared with its mean. The cumulative probability distribution of delay of an individual LRU is typified in Fig. 5. Specifically, the height at the origin represents the probability of no delay, and the cumulative probability distribution then climbs rather slowly toward unity. Thus, even though the probability of no delay can be relatively large, the expected delay can also be relatively large. Note that the probability distribution depends upon the stock level of the item as well as upon failure rate and repair time.

Even though we may not be able to easily compute the probability distribution of delay for the project network corresponding to the prelaunch schedule of operations, Fig. 5 represents the general nature of this distribution.




Fig. 5 — Typical delay distribution of an item

Therefore, a stockage model capable of determining good stockage allocations in all project network situations is needed. To motivate such a model, we consider the delay caused by a single LRU_i , denoting this delay random variable by $U_i(S_i)$, where S_i is the spares stock level of LRU_i . Let $E[U_i(S_i)]$ denote mean delay and let $\sigma(S_i)$ denote the standard deviation of the delay caused by LRU_i . By Chebyshev's inequality [21], we have

$$\text{Prob}[U_i(S_i) \geq E[U_i(S_i)] + k\sigma_i(S_i)] \leq 1/k^2.$$

That is, the probability that delay, as a function of spares level, will exceed mean delay plus a multiple, $k > 0$, of the standard deviation is less than $1/k^2$. For instance, if $k = 3$, then the above probability is less than $1/9$. Moreover, it is known that the probability is usually lower than that suggested by Chebyshev's inequality.

The above discussion suggests that with each LRU_i we associate a delay function, a function of the stock level, denoted by

$$D_i(S_i) = E[U_i(S_i)] + k\sigma_i(S_i).$$

This function will serve as an approximation for the delay associated with LRU_i , and $D_i(S_i)$ can be computed by using the distribution of $U_i(S_i)$ developed in Sec. II.

Given a spares vector, S , that satisfies the budget constraint

$$\sum_j C_j S_j \leq B,$$

we wish to find a spares stockage allocation S^* that is optimal in the sense that it minimizes an approximate project network delay, where the delay due to LRU_i is given by the function

$$D_i(S_i) = E[U_i(S_i)] + k\sigma_i(S_i),$$

and the approximate project network delay is the length of the critical path resulting from $\{D_i(S_i)\}$.

Let $L_1(S)$, $L_2(S)$, ..., $L_N(S)$ be the lengths, given a spares stockage allocation, of the N paths of the project network corresponding to the prelaunch schedule of operations. Then,

$$L_j(S) = k_j + \sum_{i \in P_j} D_i(S_i),$$

where P_j denotes those arcs of path j that correspond to LRU delays, and k_j is the sum of the times of the remaining arcs of path j . Figure 6 clarifies the notation. Here, $k_j = 16$ and two arcs on this path correspond to LRUs. The length of path j , as a function of S_{11} and S_5 , is

$$L_j(S) = 16 + D_{11}(S_{11}) + D_5(S_5).$$




Fig. 6 — Typical path in a project network

The spares stockage optimization problem is

ME3:

$$\min \max \{L_1(S), L_2(S), \dots, L_N(S)\},$$

subject to S within the budget B .

Even though the project network may contain a very large number of paths, N , problem ME3 can be relatively easily solved. Assume that we have found $n < N$ paths of the project network. We then solve the optimization problem

ME3(n):

$$\begin{aligned} \min y \\ y \geq k_j + \sum_{i \in P_j} D_i(S_i), \quad j = 1, \dots, n, \end{aligned}$$

subject to S within the budget, B , where $S_r \geq 0$, integer-valued, for all r , and we let y^* , S^* be an optimal solution. Then S^* is an optimal allocation for ME3, and y^* is the project duration, if, and only if,

$$y^* \geq k_j + \sum_{i \in P_j} D_i(S_i^*), \quad j = 1, \dots, N.$$

Of course, we know the above inequality to be true for all the paths of ME3(n); however, we do not know whether this inequality holds for all the paths not contained in ME3(n). However, by solving a single longest-path problem for the project network with arc lengths corresponding to

LRUs given by $D_i(S_i^*)$, we may determine whether these inequalities hold.

Let j^* index such a longest path. If

$$y^* \geq k_{j^*} + \sum_{i \in P_{j^*}} D_i(S_i)$$

then S^* is an optimal allocation for ME3. If this constraint is not satisfied, we replace j^* by $n + 1$ and $S_{j^*}^*$ by S_{j^*} , and introduce the resulting inequality as the constraint, corresponding to path $n + 1$, into problem ME3(n) to form a new problem with one more path constraint, ME3(n + 1).

The above process must terminate in a finite number of steps with an optimal spares allocation for problem ME3. Moreover, the delay function $D_i(S_i)$ is such that each ME3(n) problem can be solved relatively efficiently. A preliminary computer program has been developed at Rand for this purpose. The potential value of this approach lies not only in computing spares requirements for general networks, but also in providing additional computational capability for networks of the AFTEC LCOM type.

SUMMARY

We have demonstrated the feasibility of developing stockage optimization techniques in terms of measures of system performance (e.g., launch delay) that incorporate the project network corresponding to the prelaunch schedule of operations. The problems are mathematically tractable and we are confident that the techniques developed above will also apply to the model versions that will incorporate aspects of a finite source population and limited repair capacity.

The spares optimization methods of Sec. III, and their corresponding computer programs, need to be fully developed to handle the large number of different items typical of the Space Transportation System. Moreover, these methods must be modified so that (1) shop replaceable units can be explicitly taken into account, (2) the finite source population and limited repair capacity can be incorporated, and (3) the inherent uncertainty in the input data is not ignored.

Additionally, the underlying "NASA NET," the most appropriate project network of the prelaunch schedule of operations, needs to be developed. Modifications of our basic methods can then be fully tested and evaluated in that context, which is the one most suitable for assessing spares requirements and repair level decisions. For instance, there may be environments in which the results of current spares and repair level decisions are quite adequate in terms of our measures of system performance. But we cannot identify those environments until the above-mentioned developments have been carried out. We therefore recommend that they be undertaken.

IV. DATA ISSUES IN SPARES PROVISIONING

The data available to support provisioning decisions for the orbiter were based on failure rates observed on components of heavy aircraft [22, p. 148]. Among the available data are estimates of maintenance demand rates (MDRs), i.e., rates of component removals that generate demands for spares. There are two fundamental problems with these MDRs. The first is that they implicitly assume knowledge of component operating time, and operating time is not now routinely recorded; therefore, the MDRs are not appropriate for use in spares requirements computations without at least estimates of operating time. The second problem is not unique to the orbiter provisioning problem, or even to NASA operations in general: initial estimates of component characteristics are notoriously unreliable. Thus, the orbiter initial provisioning problem is complicated by a severe paucity of useful, reliable information. In addition, because the maintenance concept has not achieved final form nor all of the level-of-repair decisions been reached, initial estimates of other component characteristics, such as repair and transportation times, which depend on the level-of-repair decision and maintenance concept, cannot be viewed as reliable either.

Another very important characteristic of any collection of initial MDRs is that, in the aggregate, they should be consistent with reasonable estimates of the mean time between failure (MTBF) rates of the orbiter as a whole. In other words, the inverse of the sum of the initial MDRs of all the LRUs on the vehicle ought to be consistent with the expected MTBF of the vehicle. If this is not the case, there may be

bias in the MDRs that could lead to a decision on initial spares investment that would be inconsistent with the desired level of system performance; i.e., serious over- or underinvestment could result.

Discussions with NASA have revealed that there is no plan to use component removal or failure data from the first six flights to revise the initial MDRs. Furthermore, no systematized collection of component failure or removal data or operating time has been implemented. We infer from this that NASA intends to make decisions on both investment level and spares mix on the basis of the initial estimates alone. Our experience with initial provisioning problems suggests that, if the initial estimates alone constitute the basis for those decisions, the performance of the resulting stockage posture could probably be achieved at dramatically less cost (or, conversely, given a specified investment level, the performance of the stockage posture could be dramatically improved) if an alternative strategy were used that took advantage of the body of theory that has emerged from the initial provisioning scenario.

NASA could take several steps that would substantially improve the cost-effectiveness of its initial provisioning strategy. In the remainder of this section, we discuss these steps and offer observations and analyses directly applicable to the orbiter provisioning problem.

THE F-16 CASE

The discussion that follows draws heavily on a recent study [23] that examined the initial provisioning of the U.S. Air Force's F-16 aircraft program. Initial provisioning was applied to the first two years of scheduled production, 150 aircraft. The scope of the study was limited to 810 recoverable (reparable) line replaceable units (LRUs) and

shop replaceable units (SRUs) unique to the F-16 aircraft. Common spares and consumables were excluded. The study specifically addressed the issue of the usefulness of early operational data in revising initial estimates of component maintenance factors (component removals per 100 flying hours). The distinction between maintenance factors and MDRs is that MDRs estimate or measure component removals per 1000 hours of operating time. In the F-16 program, maintenance factors were used because of their consistency with Air Force data systems. The study examined the accuracy of initial estimates of maintenance factors and unit prices. It concluded that unit price estimates were quite accurate, at least for the sample chosen, but that initial estimates of component maintenance factors were heavily, systematically, and positively biased. The accuracy of other initial estimates of component characteristics, such as not-reparable-this-station (NRTS) rates and repair times, was not examined.

Initial Estimates

In the F-16 program, the prime contractor, General Dynamics, recommended that a mean flying time between failure (MFTBF) rate of 2.9 be included in the weapon system's specifications. Data collected during the first two and a half years of the weapon system's life on the 810 recoverable LRUs and SRUs peculiar to the F-16 yielded an MFTBF of 5.82. In addition to the 810 peculiar recoverables, the F-16 consists of approximately 1200 common recoverables and many consumables as well. An allocation of roughly half the total number of failures on the aircraft to the 810 peculiar recoverables seems reasonable to us, since they tend to be the most complex and most prone to failure. Therefore, the MFTBF of 2.9 recommended by General Dynamics seems consistent with

the observed MFTBF on the 810 items of 5.82. However, a calculation of the MFTBF using the initial estimates of the 810 items yields a result of 1.45. If, in addition, common recoverables and consumables were accounted for, the calculated MFTBF would be well below 1.0.

This simple arithmetic applied to the initial estimates of maintenance factors would have suggested, as early as two years before the operational deployment of the first aircraft, that they were very heavily, positively biased. Such bias results in underestimating the performance expected from alternative spares investment levels. It can significantly diminish the cost-effectiveness of a program's initial stockage posture.

The lesson here for NASA in provisioning spares for STS is clear: Develop a set of initial MDRs that are well founded in engineering judgment, tempered with data from other programs, and consistent in the aggregate with the vehicle's system-level reliability characteristics. Such a set of initial estimates would be of great utility, not only in determining an appropriate investment level for initial spares, but also in computing the most effective mix of spares for that investment level. As we will discuss, such estimates, when modified by early operational data, however sparse, are powerful aids in computing spares stockage postures that are robust in the face of the uncertainties that pervade initial provisioning decisions. The estimates should be developed without using data from early missions because techniques are available for modifying the estimates with the observed data in a way more nearly optimal than using human judgment alone. Furthermore, MDRs should be redefined as component removals per flight hour or mission, or NASA should implement systematic collection of operating times. In any

event, the use of the MDR in requirements computations must be borne in mind. It should not be viewed, for example, as an estimate or measure of "true" failures, but rather as an estimate or measure of component removals that induce the need for spares.

Explicating Uncertainty in Initial Estimates

Initial estimates of component characteristics are matters of substantial uncertainty. A well-known, well-developed body of theory known as Bayesian learning suggests that it is constructive to characterize one's uncertainty about such estimates by modeling them as random variables using probability distributions that best characterize the uncertainty. Probability distributions used in this way are called *a priori* distributions, or simply *priors*. The selection of priors for the shuttle vehicle provisioning problem will depend on the specific form of the spares requirements model used and the objective function incorporated in it. It is the prior probability distribution that is revised by the observed data, the shape and breadth of the prior actually specifying how much weight is given to the initial estimate and how much to the observations. The theory of Bayesian learning is discussed in [18]. Its application to inventory systems is well known; some applications are discussed in [25-27]. Additional explanation and applications may be found in [28].

Revision of Initial Estimates

Based on data from the F-16 program, a method was developed for revising initial estimates of maintenance factors with early operational data that dramatically improved their accuracy, thereby improving the accuracy of the computed relationship between aircraft availability and

the cost and mix of spares. It also dramatically improved the cost-effectiveness of the computed mix of spares for any specified investment level. The technique applied Bayesian learning, coupled with a linear correction factor derived from the operational data, to the problem of deciding how to modify the initial estimates with the observed data. This method is called the Bayes-Lin technique. The fact that early observations are useful in revising initial estimates was demonstrated clearly and powerfully.

Although development of the Bayes-Lin technique was based on a fighter aircraft acquisition program of very different dimensions from those of the STS, its fundamental logic applies to virtually any initial spares provisioning problem. In its implementation, however, it might look quite different from the technique in its application to the F-16. For example, the maintenance factors for the shuttle might be adjusted to reflect ground operating time, since that represents a much larger proportion of total operating time than in the aircraft case. As discussed earlier, the expected availability of the shuttle is a measure of secondary interest; the primary focus should be on the expected launch delay due to parts shortages. Thus, while the best method of formulating the mathematics for the shuttle problem may differ in detail from the F-16 case, its fundamental, underlying logic is essentially the same. Appropriate models exist that can be used in the shuttle application and that are logically consistent with the spares requirements methodology selected.

Interviews suggest that NASA and contractor personnel place little or no value on component removal data generated in the first six space flights. What we know about the power of the Bayes-Lin technique

suggests strongly that ignoring those early data would be a serious mistake. The view that there are too few data to be useful is clearly refuted by the F-16 analysis. It may be helpful here to relate a story told by Professor Howard Raiffa in [29, pp. 20, 21]:

Professor Ward Edwards, a psychologist at the University of Michigan, has investigated the intuitive reactions of many subjects to experimental, probabilistic evidence. In one of his experiments he poses the following problem.

"I have two canvas book bags filled with poker chips. The first bag contains 70 green chips and 30 white chips, and I shall refer to this as the predominantly green bag. The second bag contains 70 white chips and 30 green chips, and I shall refer to this as the predominantly white bag. The chips are all identical except for color. I now mix up the two bags so that you don't know which is which, and put one of them aside. I shall be concerned with your judgments about whether the remaining bag is predominantly green or not. Now suppose you choose 12 chips at random with replacement from this remaining bag and it turns out that you draw eight green chips and four white chips, in some particular order. What do you think the odds are that the bag you have sampled from is predominantly green?"

At a cocktail party a few years ago I asked a group of lawyers, who were discussing the interpretation of probabilistic evidence, what they would answer as subjects in Edwards' experiment. First of all, they wanted to know whether there was any malice aforethought in the actions of the experimenter. I assured them of the neutrality of the experimenter, and told them that it would be appropriate to assign a .5 chance to "predominantly green" before any sampling took place.

"In this case," one lawyer exclaimed after thinking awhile, "I would bet the unknown bag is predominantly white."

"No, you don't understand," one of his colleagues retorted, "you have drawn eight greens and four whites from this bag. Not the other way around."

"Yes, I understand, but in my experience at the bar, life is just plain perverse, and I would still bet on predominantly white! But I really am not a betting man."

The other lawyers all agreed that this was not a very rational thing to do - that the evidence was in favor of the bag's being predominantly green.

"But by how much?" I persisted. After a while a consensus emerged: The evidence is meager; the odds might go up from 50-50 to 55-45; but "...as lawyers we are trained to be skeptical, so we would slant our best judgments downward and act as if the odds were still roughly 50-50."

The answer to the question "By how much?" can be computed in a straightforward fashion ..., and there is no controversy about the answer. The probability that the bag is predominantly green, given a sample of eight green and four white chips, is .964. Yes, .964. This bag is predominantly green "beyond a reasonable doubt." This story points out the fact that most subjects vastly underestimate the power of a small sample. The lawyers described above had an extreme reaction, but even my statistics students clustered their guesses around .70.

The logic of this story is especially pertinent to the initial provisioning problem and the reluctance of many logisticians to use early observations in revising initial estimates of component characteristics. Nevertheless, as the graphs in Fig. 7 show, such data may be dramatically more useful than intuition might suggest. We show several curves drawn from the F-16 experience that explicate the relationship between aircraft availability and initial spares investment level. One of the solid curves is computed from initial estimates alone, another from the initial estimates revised with one month of operational data representing only 49.2 flying hours of experience, and the other from the initial estimates revised with six months of operational data representing 963.5 flying hours. The dashed curves are computed based on two years of operational data collected *after* the initial six months of operation that generated the data used to revise the estimates. The dashed curves approximate what would have happened for various investment levels based on the stockage postures computed with the estimates of maintenance factors used with the respective solid curves. The very large difference between "computed" and "actual" where




Fig. 7 — Improvements in availability versus cost using the Bayes-Lin technique

only the initial estimates were used is due to the powerful, systematic, positive bias in the initial estimates of the maintenance factors. Note the dramatic improvement in the availability-vs.-cost curve in the case where only one month of operational experience is used to revise the estimates, and the remarkable improvement in predictive accuracy.

There are important lessons for the STS program in this analysis. Given a reasonable set of initial estimates of component removal rates, one should base the initial spares investment level decision and the computed mix of spares on revisions to those estimates that incorporate all available operational (component removal) data and the essential logic of the Bayes-Lin revision technique. This approach can be expected to provide substantial improvement over current NASA plans both in the predictive accuracy of the spares requirements model and in the performance of the resulting stockage posture.

SUMMARY AND RECOMMENDATIONS

NASA needs to implement a data collection system that will routinely provide observed values of component characteristics, logistics system performance, and appropriate operating times. Without such data the ability to revise initial estimates of these values is seriously impaired.

The quality of the initial MDR estimates on shuttle components is questionable. As a first step, they should at least be compared, in the aggregate, with expectations about shuttle MTBF to obtain some sense of their bias.

A second step would be to estimate other component characteristics that affect the expected number of each type of component in resupply,

e.g., repair time, order-and-ship time, repair level distribution, condemnation rate, and procurement leadtime. Reliable unit price estimates are also needed. One problem is that many of these estimates cannot be made until the maintenance concept and levels of repair are determined. Therefore, these decisions should be made as soon as practicable.

A third step would be to revise the estimates of component removal rates using all available component removal data, however sparse, from the first several shuttle missions. The Bayes-Lin technique should be used. (An alternative to this step would involve adaptation of the revision technique to a network-analytic spares requirements model. Such an adaptation might involve different probability models from those used in the F-16 case, but the underlying logic of Bayesian revision and linear correction should still be used.)

A fourth step would be to determine the appropriate level of investment in initial spares using an explicit model of the relationship between performance and cost. This would enable the investment decision to be made in full light of the performance it will yield, and the converse.

Finally, just before any subsequent computation of spares requirements, component removal rates should be revised with *all* available data.

V. SUMMARY AND RECOMMENDATIONS

Several important observations and conclusions emerge from this work. NASA needs to develop a well-formulated logistics support strategy; however, it does not have the tools to evaluate alternatives in a way that takes explicit account of the interdependencies of the various components of such a strategy. The maintenance concept, levels of repair, spares investment levels, spares stockage postures, and their interrelationships have a significant impact on the effectiveness and cost of the long-term logistics support of the Space Transportation System.

THE NEED FOR CAPABILITY ASSESSMENT

NASA's decisionmaking about logistics support planning depends in large part on analyses done by contractors using, at best, questionable data. These important policy decisions need to be made with full recognition of their impact on system performance and on the costs of achieving specified levels of system performance. This implies the need for reasonable estimates of a range of alternative maintenance concepts, repair levels, modes of transportation, and spares investment levels, and the capability to compute the least-cost spares stockage posture for any specified level of investment.

The methods described in Sec. II, when fully developed, would enable the estimation of direct and meaningful measures of launch capability as a function of component and logistics support system characteristics. Those characteristics depend, in turn, on the specification of maintenance concept, repair levels, and other policy

decisions. Thus, we have provided an approach that can potentially support an integrated view of the several fundamentally important policy decisions that NASA faces, and we believe that these methods could be helpful in developing a cost-effective logistics support strategy.

The capability assessment model presented here goes far in demonstrating the feasibility of an analytic approach appropriate for the STS environment. Initial steps have also been taken to deal with the small STS fleet size and the uncertainty in the data estimates. Networks similar to the AFTEC LCOM net and further evaluation with a more detailed network that represents the demand and fill nodes are required.

THE SPARES STOCKAGE OPTIMIZATION PROBLEM

All that we have learned about logistics support strategies from other programs suggests that NASA needs to formulate a strategy for the STS that may differ significantly from that in its current plans. It seems clear that NASA cannot simply use a spares requirements computational model from, say, a military aircraft program and apply it to the STS. The distinctly different dimensions of the STS program, i.e., a small number of vehicles and relatively long periods of ground time between missions, prohibit this. Spares requirements models that have been more or less successfully applied to initial provisioning in military aircraft programs depend on certain steady-state assumptions and are oriented toward maximizing aircraft availability, maximizing fill rate, or minimizing expected backorders, at a randomly chosen point in time, given some budget constraint. Such models do not take advantage of the special structure of the STS program, in which a particular vehicle can be grounded for long periods between missions.

The approaches to stockage optimization described in Sec. III take advantage of the capability assessment methods and project network structure of Sec. II. We believe that our approaches can provide a reasonable basis for developing a sound, coherent, integrated spares acquisition and logistics support strategy for the STS.

Spares optimization methods, and the corresponding computer programs, must be fully developed to handle the large number of different items typical of those in the STS. These methods must be modified so that (1) SRUs can be explicitly taken into account, (2) the finite source population and limited repair capacity can be incorporated, and (3) the inherent uncertainty in the input data is not ignored.

As with the capability assessment approach, the most appropriate project network of the prelaunch schedule of operations needs to be developed, so that the modifications can be fully tested, evaluated, and compared in the context of the project network most suitable for spares requirements and repair level decisions evaluation. Although each of the approaches presented will do better than the typical "ready-rate" models, how much better will not be clear until such an evaluation is undertaken.

DATA ISSUES

The several steps discussed in Sec. IV are vital to NASA if it is to make its spares investment decisions with a thorough understanding of their implications for system performance. The ability to compute the least-cost mix of spares and perform credible logistics system capability assessments depends on those same steps. Initial estimates

alone do not constitute a sound basis for determining a cost-effective STS logistics support strategy. Data from the first few STS missions would be of dramatically greater worth, especially when used with the Bayes-Lin method (described in Sec. IV) to revise the initial estimates. It is important that the uncertainty surrounding estimates of maintenance demand rates and other logistics system performance measures be reflected in logistics policy development and resource requirements determination. Failure to achieve this would be to overestimate the ability of the logistics system to support the shuttle schedule.

MANAGEMENT ISSUES

NASA does not now have the data collection systems, the analytic modeling capability, or the management controls it needs for effective logistics management. NASA's current data systems, for example, do not allow systematic recording or estimation of subsystem or component operating times. Yet the only available estimates of component demand rates require at least estimated operating times to be useful. Furthermore, the data systems currently in use in the STS program are largely contractor-operated and lack interface. Spares requirements and logistic support policy recommendations are made by each major contractor independently, using models that may not be appropriate to the unique logistics support problem that the STS launch and recovery environment presents.

It seems to us important that NASA and the Air Force continue to develop, implement, and use the kinds of decision aids discussed in this report within their management framework. The following paragraphs offer some specific recommendations.

RECOMMENDATIONS

The models presented here demonstrate that it is feasible to develop improved analytical logistics modeling capabilities that will relate logistics support decisions and resource requirements to the capability to meet STS launch schedules. Evaluation and demonstration of the payoff to be gained from using these techniques remain to be done. It is recommended that these actions be accomplished in two phases because of the large amount of data that would be required for a full-scale evaluation. Recommendations concerning improved data collection and parameter estimation are detailed below. These are important regardless of the logistics modeling methodologies that NASA may ultimately choose to use.

Improved Data Collection and Parameter Estimation

It is recommended that NASA:

1. Modify or design and implement an integrated data collection system that would routinely provide up-to-date component removal data, repair times, repair level distributions, retrograde shipment times, order-and-ship times, condemnation rates, procurement and repair costs, procurement lead times, and operating times or usage estimates.
2. Assemble whatever data are available (from either formal or informal systems) from previous STS flights and compare these data with those parameters currently being used for logistics planning and resource requirements computations. From this, judge whether revisions to initial estimates are required. If

such is the case, as is likely, revise the initial estimates using the Bayes-Lin technique suggested in this report, or a similar Bayesian technique. Continue this revision process as more flight experience is gained.

3. Estimate the uncertainty surrounding component removal rates and other logistics system performance parameters, and explicitly consider them in making logistics policy decisions (e.g., level of repair decisions) and in determining spares requirements.

Phase I: Initial Prototype Development and Evaluation

It is recommended that the evaluation and full-scale development and implementation of the logistics system capability assessment and spares optimization methodologies be carried out in two phases because of the difficulty in obtaining the necessary, detailed STS recovery task network and component data. The first phase would focus on the evaluation of these methodologies, using a limited set of representative components. If the outcome of the first phase is positive, the second phase would refine the techniques and implementation. For Phase I the following steps are recommended:

1. To the extent technically feasible, extend the methodologies presented in Secs. II and III to include non-Poisson processes with finite populations; multiple stockage points, including Vandenberg Air Force Base and other possible stockage sites; and variance-to-mean ratios other than unity.
2. Identify a subset of components that, to the extent possible, represent the population of all components from each of the

projects (orbiters, boosters, tank and main engine). At the same time, develop criteria to determine the range of components that should generally be considered in such models.

3. For that select subset, develop the detailed network data corresponding to the prelaunch schedule of operations, including the demand and fill nodes for each component, and collect the most up-to-date component data.
4. Determine whether the network representation for these components is compatible with the assumptions inherent in these methodologies. If there are compatibility problems, develop and evaluate network editing techniques that could allow these methods to be used.
5. Evaluate and demonstrate the use of the capability assessment methodology, using simulation for comparisons as appropriate.
6. Evaluate and compare the two stockage optimization techniques, in terms of launch delay or stockage costs, presented in this report (a) with each other and (b) with those techniques currently in use by NASA, using the capability assessment model or simulation as appropriate.

Phase II: Prototype Improvement and Implementation

If the Phase I evaluation results are positive, the following steps for Phase II are recommended:

1. Modify and improve the methodologies based on the Phase I results. In addition, improve them, to the extent feasible, so that they will be suitable for individual projects as well as for overall system assessment, will consider the availability

of manufacturing assets, and will include indentured components (SRUs, etc.).

2. Define and assess their potential use for integrated logistics management of logistics operations, level of repair analyses, development of out year requirements, and procurement and budget decisions.
3. Develop and implement a full-scale system.

Implementation of these several steps can be expected in the longer run to deliver significantly more cost-effective logistics support to the STS program than NASA's current plans.

REFERENCES

1. Sherbrooke, Craig C., *METRIC: A Multi-Echelon Technique for Recoverable Item Control*, The Rand Corporation, RM-5078-PR, November 1966.
2. Muckstadt, John A., "A Model For a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System," *Management Science*, Vol. 20, No. 4, Part I, December 1973.
3. Hillestad, R. J., and Manuel J. Carrillo, *Models and Techniques for Recoverable Item Stockage When Demand and the Repair Process Are Nonstationary--Part I: Performance Measurement*, The Rand Corporation, N-1482-AF, May 1980.
4. Hillestad, R. J., *Dyna-METRIC: Dynamic Multi-Echelon Technique for Recoverable Item Control*, The Rand Corporation, R-2785-AF, July 1982.
5. Fox, B., "Discrete Optimization Via Marginal Analysis," *Management Science*, Vol. 13, No. 3, November 1966.
6. Fox, B. L., and D. M. Landi, *Optimization Problems with One Constraint*, The Rand Corporation, RM-5791-PR, October 1968.
7. Fisher, R. R., Captain, et al., *The Logistics Composite Model: An Overall View*, The Rand Corporation, RM-5544-PR, May 1968.
8. Brooks, R. B., and R. H. Haase, *A Network Approach to Parts Provisioning for Apollo Prelaunch Operations*, The Rand Corporation, RM-4581-NASA, August 1965.
9. Brooks, Robin B., Colleen A. Gillen, John Y. Lu, *Alternative Measures of Supply Performance: Fills, Backorders, Operational Rate, and NORS*, The Rand Corporation, RM-6094-PR, August 1969.
10. Palm, C., "Analysis of Erlang Traffic Formula for Busy Signal Arrangements," *Ericsson Techniques*, No. 5, 1938, pp. 39, 58.
11. Takacs, Lajos, *Introduction to the Theory of Queues*, Oxford University Press, New York, 1962.
12. Kennedy Space Center (KSC) Space Transportation System (STS) Logistics Composite Model (LCOM) Simulation Development: Phase II, Network Diagrams, The BDM Corporation, A-82-075-TR, March 1982.
13. Kennedy Space Center (KSC) Space Transportation System (STS) Logistics Composite Model (LCOM) Simulation Development: Phase II, Final Report, Vol. I, The BDM Corporation, A-82-075-TR, April 1982.

14. *Shuttle Turnaround Analysis Report (STAR 23)*, National Aeronautics and Space Administration, John F. Kennedy Space Center, May 12, 1982.
15. Hillier, Frederick S., and Gerald J. Lieberman, *Operations Research*, Holden Day, San Francisco, 1974.
16. Feller, William, *An Introduction to Probability Theory and Its Applications*, Vol. I, Wiley, New York, 1968.
17. Sherbrooke, Craig C., *Discrete Compound Poisson Processes and Tables of the Geometric Poisson Distribution*, The Rand Corporation, RM-4831-PR, July 1966.
18. Feeney, G. J. and C. C. Sherbrooke, "The (s-1,s) Inventory Policy Under Compound Poisson Demand: A Theory of Recoverable Item Stockage," *Management Science*, Vol. 12, No. 5, January 1966, pp. 391-411.
19. Ross, Sheldon M., *Applied Probability Models with Optimization Applications*, Holden Day, San Francisco, 1970.
20. Jacobsen, S. E., and S. Arunkumar, "Investment in Series and Parallel Systems to Maximize Expected Life," *Management Science*, Vol. 19, No. 9, May 1973.
21. Hogg, Robert V., and Allen T. Craig, *Introduction to Mathematical Statistics*, Macmillan, New York, New York, 1970.
22. *Orbiter Project and Other Selected Aspects of the Space Shuttle Program*, Vol. III, NASA, Office of the Inspector General, A-HQ-81-003, May 1982.
23. Abell, John B., Joan E. Lengel, and F. Michael Slay, *Toward Improved Initial Provisioning Strategies: The F-16 Case*, Logistics Management Institute, Report #ML108, Washington, D.C., April 1982.
24. Raiffa, H., and R. Schlaifer, *Applied Statistical Decision Theory*, Graduate School of Business Administration, Harvard University, 1961.
25. Brown, G. F., Jr., and W. F. Rogers, "A Bayesian Approach to Demand Estimation and Inventory Provisioning," *Naval Research Logistics Quarterly*, 20:607-624, 1973.
26. Haber, Sheldon E., and Rosedith Sitgreaves, "A Methodology for Estimating Expected Usage of Repair Parts with Application to Parts with No Usage History," *Naval Research Logistics Quarterly*, 17:535-546, 1970.
27. McGlothlin, W. H., and R. Radner, *The Use of Bayesian Techniques for Predicting Spare-Parts Demand*, The Rand Corporation, RM-2536, March 1960.

28. Morris, W. T., *Management Science: A Bayesian Introduction*, Prentice-Hall, Englewood Cliffs, N.J., 1968.
29. Raiffa, Howard, *Decision Analysis*, Addison-Wesley, Reading, Massachusetts, 1968.

RAND/R-3083-NASA

