SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK 14853-3801

TECHNICAL REPORT NO. 1284

March 2001

A COMBINATORIAL MULTI-INDENTURE, MULTI-ITEM INVENTORY MODEL FOR NASA'S REUSABLE LAUNCH VEHICLE PROGRAM

By

Kathryn E. Caggiano and John A. Muckstadt

A Combinatorial Multi-Indenture, Multi-Item Inventory Model for NASA's Reusable Launch Vehicle Program*

Kathryn E. Caggiano and John A. Muckstadt †

December 2000

*This work was supported in part by NASA Grant NAG8-1716.

†School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853

Abstract

In this paper we describe a mathematical model and solution approach for determining spare inventory levels for a multi-item, multi-indenture system of recoverable items in which failed parts are removed according to a fixed cyclic schedule. This model was motivated by NASA's desire to better understand and control maintenance costs associated with their future Reusable Launch Vehicle (RLV) programs. Our model differs from existing recoverable part models in two fundamental ways. First, the vast majority of existing recoverable item models assume that part failure processes are Poisson or compound Poisson. Our analyses have shown that this assumption is inappropriate for a system characterized by a small number of flights (over the maximum repair cycle time) that adhere to a fixed mission and maintenance pattern. Instead, we use an explicit combinatorial model to capture part failures over time and develop approximation methods based on this combinatorial model. Second, instead of attempting to characterize stationary distributions and measures for random points in time, the performance measures of interest in our model correspond to *specific* points in time, namely, the points at which vehicles are scheduled to have their maintenance completed (i.e., RLV due dates). By using the proposed combinatorial model and methods in lieu of a Poisson-based model, NASA could reduce inventory investment levels by many millions of dollars.

1 Introduction

Because of the enormous cost associated with space vehicle launch delays, the successful operation of NASA's Reusable Launch Vehicle (RLV) programs depends critically on their ability to maintain an on-time launch schedule. While certain types of launch delays cannot be avoided, such as weather-related delays, those that result from prolonged or delayed vehicle maintenance *are* largely avoidable by investing properly in resources that support the maintenance process.

In this paper, we examine the impact of part shortages on vehicle maintenance schedules and describe a method for investing a given budget in spare parts inventory so as to minimize this impact. We concern ourselves only with *recoverable* or *repairable* parts. That is, parts that are either repaired in-house or sent to a contractor for repair when they fail. Typically these are more complex, expensive parts, and maintaining excessive inventory in these items can be both costly and wasteful.

There is no shortage of literature on repairable parts models. Within the realm of multi-echelon models are numerous extensions and generalizations of the METRIC model Sherbrooke (1968), a multi-item, multi-echelon inventory management system developed for the United States Air Force in the late 1960s, and widely held to be the first practical application of multi-echelon inventory theory. Although the list is far from exhaustive, notable among the extentions of METRIC are Simon (1971), Muckstadt (1973), Shanker (1981), Graves (1985), Moinzadeh and Lee (1986), Sherbrooke (1986), Lee and Moinzadeh (1987), Svoronos and Zipkin (1991), and Wang et al. (2000).

Compared to their multi-echelon cousins, multi-indenture models have received relatively little attention in the literature. Sherbrooke (1971) considers the single-base case and describes a method for evaluating the expected number of vehicles that are not operationally ready due to supply (NORS), but shows that the corresponding optimization model is not tractible since the objective function is not separable. Using the Sherbrooke (1971) model, Silver (1972) shows that the ready rate objective function is separable in a special case, and uses a heuristic technique to develop a set of potential solutions from the special case solution. Muckstadt (1973) presents MOD-METRIC, an extension of the METRIC model that considers multi-indenture items within the multi-echelon framework. Using ideas from Graves (1985), Sherbrooke (1986) shows that the accuracy of multi-indenture models can be improved by assuming that the distribution of the number of top-level items in resupply is negative binomial instead of Poisson. Sherbrooke merges these ideas into the VARI-METRIC model Slay (1980) to arrive at an improved multi-echelon, multi-indenture model.

The model and analysis we present in this paper differ from other multi-indenture models in two fundamental ways. First, all of the aforementioned models assume that part failure processes are Poisson or compound Poisson. This assumption is inappropriate for a system characterized by a small number of flights that adhere to a fixed mission and maintenance pattern, and using a Poisson model to plan spare parts inventory for such a system can have costly consequences. In this paper, we use an explicit combinatorial model to capture part failures over time and develop approximation methods based on this combinatorial model. Second, instead of attempting to characterize stationary distributions and measures for a *random* point in time, the performance measures of interest in our model correspond to specific points in time. That is, the points at which vehicles are scheduled to have their maintenance completed (i.e., RLV due dates).

Like the models of Silver (1972), Muckstadt (1973), and Sherbrooke (1986), our model considers two classes of recoverable parts: line replaceable units (LRUs), and shop replaceable units (SRUs). LRUs are modular assemblies (or subsystems) that are tested in-place on an RLV during the inter-launch maintenance process. If an LRU fails its test, the part is removed from the vehicle and replaced with a spare LRU (if one is available). Meanwhile, the failed LRU enters a repair process in which the cause of its failure is identified and remedied.

SRUs are components, or subassemblies, of LRUs. An LRU failure may be due to a failure in one or more of its SRU components. During the LRU repair process, if an LRU failure is determined to be the result of an SRU failure, then the failed SRU is removed from the LRU and replaced with a spare SRU (if one is available). The failed SRU then enters its own repair process. Components needed to repair SRUs are presumed to be readily available and are not considered in our model.

While the availability of both LRUs and SRUs is vital to the maintenance process, only a lack of LRUs can cause a delay in a vehicle maintenance schedule. A lack of SRU components can delay the *repair* of one or more LRUs, but it does not necessarily cause a schedule delay. For this reason, we define a *backorder* to be an unfulfilled order for an LRU. A backorder is *delay-causing* if it still exists at end of the scheduled maintenance period for the RLV. In these terms, we are interested in measuring the *expected number of delay-causing backorders* at the end of a scheduled maintenance period.

In Section 2, we describe the framework upon which our model is based, including the RLV ground maintenance process, and the repair processes for LRUs and SRUs. In Section 3 we formulate our problem as a mathematical program and describe how an optimal tradeoff curve can be computed in a straightforward manner. In Section 4 we discuss methods for approximating the probabilities that are the key components of the objective function. An example illustrating the differences between the approximation methods is given in Section 5, along with a comparison of the resulting solutions with those from a Poisson model. We briefly summarize our contributions in Section 6.

2 Modeling Framework

Unless otherwise stated, we will assume that all times and interval lengths are integer-valued.

2.1 The RLV Ground Maintenance Process

We consider a scenario in which identical vehicles maintain a fixed mission pattern such that successive vehicle missions are offset by γ time units. Each vehicle returning from a mission must undergo a fixed maintenance schedule before the vehicle can be prepared to launch again. Maintenance cycles are scheduled to be τ time units in length and are also offset by γ time units. See Figure 1.

On each RLV there is exactly one LRU of each type. For each LRU type l, the maintenance schedule dictates a specific time δ_l (relative to the start of the maintenance cycle) at which the LRU of type l is checked for failure. See Figure 2. We let $\tau_l = \tau - \delta_l$

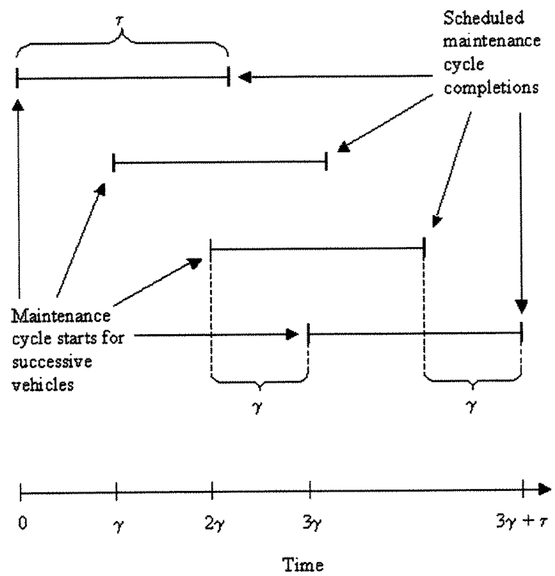

Figure 1: The RLV Maintenance Process

denote the *effective* maximum maintenance schedule length for LRU type l. That is, the time that elapses between the test of an LRU of type l and the corresponding maintenance schedule due date. Regardless of when the LRUs are checked during the maintenance schedule, all failed LRUs must be replaced by the end of the maintenance schedule in order to avoid a delay. It is at these due dates ($i\gamma + \tau$ for $i = 0, 1, 2, \ldots$) that we are interested in measuring backorders, since these are the points in time at which existing backorders can become delay-causing.

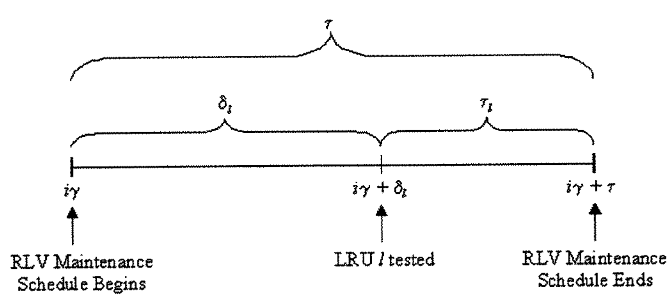

Figure 2: One Maintenance Cycle

2.2 The LRU Repair Process

When an LRU fails, it is either sent away for repair, or it enters an in-house repair process. In this section we describe the process for in-house repairs.

Recall that LRUs of type l are checked for failure at times $i\gamma + \delta_l$, $i = 0, 1, 2, \ldots$. Each LRU l fails with probability p_l, independent of any other LRU failures. If an LRU of type l fails, the failed LRU is removed from the vehicle and transported to its designated repair site where it is queued for service. This transport time is denoted t_l^{in}. Once at the repair facility, the LRU waits in the repair queue for q_l units of time before being serviced. The LRU service process itself has two phases, failure diagnosis and repair.

During the failure diagnosis phase, which takes d_l units of time, tests are performed to determine the cause of the LRU failure. The failure may be the result of a failed SRU component. If the needed SRU component is available in stock (or if no component is needed), the LRU repair phase commences immediately. If the needed component is not available, the LRU must wait until the component becomes available. This waiting time is denoted w_l. In this paper, we explicitly derive the distribution of w_l.

During the repair phase, which takes r_l units of time, any faulty SRU components are replaced, and the LRU is prepared for subsequent use on a vehicle. Once an LRU has been repaired, it is transported to a stock location where it is then available for use. The time required to transport the repaired LRU from the repair facility to the stock location is denoted t_l^{out}.

Figure 3 depicts the entire LRU repair cycle. In our model, the repair cycle time for a failed LRU of type l is denoted $\hat{C}_l$.

2.3 The SRU Repair Process

As described above, SRU failures are detected during the LRU failure diagnosis phase. Specifically, SRUs that are components of LRU type l will have their failures detected at times $i\gamma + \delta_l + C_l^b$, $i = 0, 1, 2, \ldots$. We assume that at most one SRU failure occurs for each LRU failure.

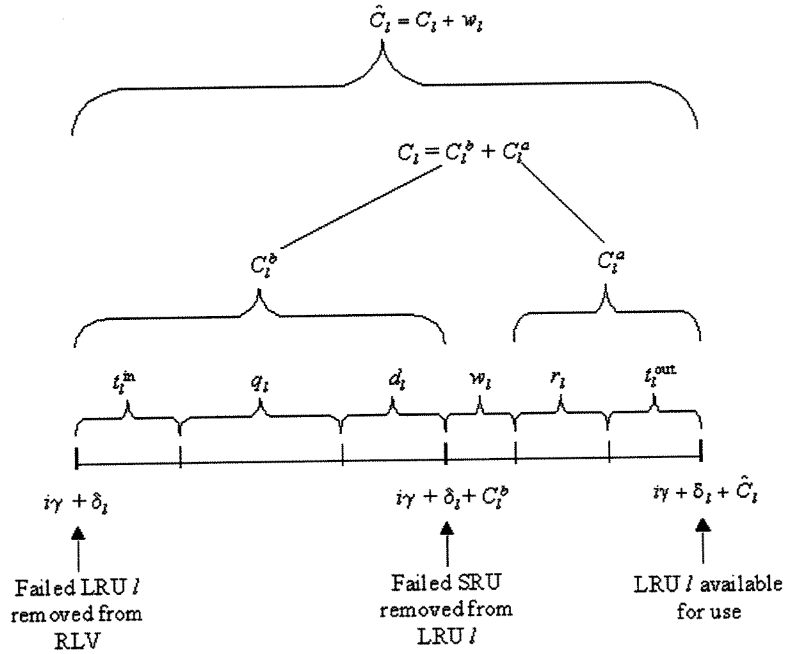

Figure 3: The LRU Repair Cycle

For a given RLV maintenance schedule, an SRU j fails with probability p_j, independent of other SRU failures that are not components of the same LRU. Given that LRU l fails, the probability that a component SRU j is the cause of the failure is p_j/p_l. When an SRU failure is detected, the failed SRU undergoes a repair process similar to the LRU repair process. The repair cycle time for a failed SRU of type j is denoted $\hat{C}_j$ and is assumed to be constant in our model.

2.4 Repair Cycle Times

Figure 4 depicts some of the major factors that influence the various components of the LRU repair cycle times. Note that SRU spare inventory levels and SRU repair cycle times affect the LRU repair cycle times through the amount of waiting time they induce in the process. The LRU repair cycle times, in turn, dictate the LRU spare inventory levels that are necessary to avoid shortages. While our analysis of LRU repair cycle times is focused solely on understanding the nature of the waiting times w_l, it is extremely important to understand how each resource (or lack thereof) contributes to the various components of repair cycle time. We examine our problem under the assumption that *without investing in SRU spare parts inventory, the LRU waiting times will constitute a significant fraction of the LRU repair cycle time.*

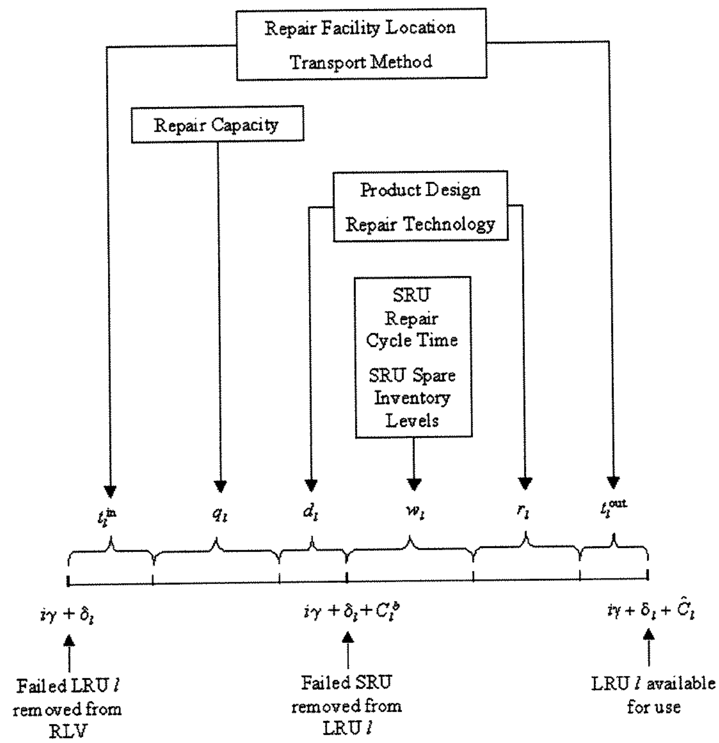

Figure 4: Factors Impacting LRU Repair Cycle Time

2.5 Counting Delay-Causing Backorders

Consider an LRU type l and an arbitrary maintenance schedule due date $T^* = i^*\gamma + \tau$. Suppose that backorders exist for LRUs of type l at time T^*. Not all of these backorders are necessarily delay-causing, since there may be backorders that correspond to vehicles whose due dates are later than T^*. The backorders that are delay-causing at T^* are precisely those that are associated with vehicles whose due dates are at or before T^*. Equivalently, the backorders for LRU type l that are delay-causing at T^* are those that were recognized at or before time $T^* - \tau_l$.

Let N_l denote *the number of LRUs of type l in the repair process at an arbitrary maintenance schedule due date $T^* = i^*\gamma + \tau$ that entered the repair process at or before time $T^* - \tau_l$*. If there are s_l spares of LRU type l, then there will be a delay-causing backorder at maintenance due date T^* if $N_l > s_l$. More specifically, if $N_l = s_l + k$, then k vehicles will be delayed due to LRU type l at time T^*. The k delayed vehicles are *not* necessarily the k vehicles that have most recently landed. Since backorders are filled on a first-come, first-serve basis, and since we assume no cannibalization of parts, the k delayed vehicles will be the ones having the k most recent test failures of LRU l.

Let D_l be *the number of delay-causing backorders for LRU type l at an arbitrary maintenance schedule due date T^**, so that

$$D_l = [N_l - s_l]^+. \tag{2.1}$$

Note that while D_l is a function of s_l, N_l is not. However, N_l *does* depend upon the stock levels s_j of SRUs that are components of LRU type l, since these stock levels affect w_l.

3 Problem Formulation

We begin this section with a summary of modeling assumptions. We then define notation for the parameters and variables that are key to the problem formulation. We conclude the section by formulating our problem as a mathematical program and outlining a practical solution approach.

3.1 Modeling Assumptions

In addition to the basic structure outlined in the preceding sections, we make the following modeling assumptions about the recoverable parts, the failure processes, and the repair processes:

1. Each RLV contains exactly one LRU of each type.

2. Each LRU contains at most one SRU of each type.

3. Each SRU component is found in exactly one LRU type. LRUs do not share common SRU components.

4. Each LRU type is either repaired in-house or sent away for repair.

5. If an LRU type is sent away for repair, its constituent SRU types are also repaired in the process. Hence, these SRU types are not considered in our model.

6. All LRU and SRU failures are repairable.

7. In calculating the distributions of N_l and N_j, we ignore the fact that there are a finite number of physical vehicles in rotation. That is, we do not attempt to capture the fact that if a backorder persists for too long, the vehicle cannot complete a mission before its next scheduled maintenance cycle.

8. Each LRU of type l fails with probability p_l, independent of other LRU failures.

9. Each failed LRU contains at most one failed SRU component.

10. SRU failures are independent of each other across maintenance cycles. Each SRU of type j fails with probability p_j. The probability that SRU $j \in J_l$ has failed given that LRU l has failed is p_j/p_l.

11. SRU repair cycle times $\hat{C}_j$ are known and constant.

12. For LRU types repaired in-house, all repair cycle time components except w_l are known and constant.

13. For LRU types sent away for repair, repair cycle times $\hat{C}_l$ are are i.i.d. bounded random variables with known distributions.

14. Backorders for the same LRU type are filled on a first-come, first-serve basis.

15. LRUs in the repair process that are waiting for the same SRU type are allocated SRUs on a first-come, first-serve basis.

3.2 Notation

Unless otherwise stated, we will use the following notation throughout this paper.

Sets

L - the set of LRU types, indexed by l,

J - the set of SRU types, indexed by j, and

J_l - the set of SRU types that are components of LRU type l.

Decision Variables

s_l - the spare stock level of LRU type l,

s_j - the spare stock level of SRU type j, and

s_{J_l} - the vector of spare stock levels of SRU types $j \in J_l$.

Independent Parameters

RLV Maintenance Process:

τ - the scheduled length of an RLV maintenance cycle, and

γ - the time between the starts of consecutive RLV maintenance cycles.

LRU Repair Process:

δ_l - the time from the start of an RLV maintenance schedule to the point at which LRU type l is checked,

τ_l - the time from when LRU type l is checked until the scheduled end of the corresponding RLV maintenance cycle (i.e., $\tau - \delta_l$),

p_l - the failure probability of LRU type l,

t_l^{in} - the time to transport LRU type l to its designated repair facility,

q_l - the repair queue time of LRU type l,

d_l - the failure diagnosis time for LRU type l,

r_l - the repair time for LRU type l,

t_l^{out} - the time to transport LRU type l from its designated repair facility,

C_l^b - the portion of the repair cycle time up to and including failure diagnosis for an LRU type l that is repaired in-house (i.e., $t_l^{\text{in}} + q_l + d_l$),

C_l^a - the portion of the repair cycle time that occurs after all needed SRU components are available for an LRU type l that is repaired in-house (i.e., $r_l + t_l^{\text{out}}$),

C_l - the minimum repair cycle time for an LRU type l that is repaired in-house (i.e., $t_l^{\text{in}} + q_l + d_l + r_l + t_l^{\text{out}}$), and

c_l - the unit cost of LRU type l.

SRU Repair Process:

p_j - the failure probability of SRU type j,

p_j/p_l - the failure probability of SRU type $j \in J_l$ given that an LRU of type l has failed,

$\hat{C}_j$ - the repair cycle time for SRU type j, and

c_j - the unit cost of SRU type j.

Other:

B - the target budget level.

Dependent Parameters

w_j - the length of time a failed LRU of type l needing SRU type $j \in J_l$ waits for SRU j,

w_l - the length of time a failed LRU of type l waits for needed SRU components,

$\hat{C}_l$ - the total repair cycle time for LRU type l (i.e., $C_l + w_l$ for in-house repairs),

N_l - the number of units of LRU type l in the repair process at an arbitrary maintenance schedule due date $T^* = i^*\gamma + \tau$ that entered the repair process at or before time $T^* - \tau_l$,

N_j - the number of units of SRU type $j \in J_l$ in the repair process at an arbitrary failure diagnosis completion time $T_l' = i'\gamma + \delta_l + C_l^b$ for LRU type l, and

D_l - the number of delay-causing backorders for LRU type l at an arbitrary maintenance schedule due date T^*.

3.3 Mathematical Formulation

In formulating our problem as a mathematical program, we have chosen the objective of *minimizing the total expected number of delay-causing backorders at an arbitrary maintenance schedule due date.* Several alternative objectives are possible, many of which are discussed at length in Sherbrooke (1971). As Sherbrooke points out, the expected backorder criterion is sensitive to the *duration* of a stockout condition, not merely the probability of one. Letting $\mathcal{D}_l(s_l, s_{J_l})$ represent the *expected number of backorders for LRU type l at an arbitrary maintenance schedule due date*, we have:

$$\mathcal{D}_l(s_l, s_{J_l}) \equiv \sum_{k>s_l} (k - s_l) \Pr[N_l = k | s_{J_l}], \tag{3.1}$$

where s_{J_l} denotes the set of stock levels $s_j, j \in J_l$. Given this definition, we state the *spare parts budget allocation problem* (SPBAP) as:

$$\text{(SPBAP) minimize} \quad \sum_{l \in L} \mathcal{D}_l(s_l, s_{J_l}) \tag{3.2}$$

subject to

$$\sum_{l \in L} c_l s_l + \sum_{j \in J} c_j s_j \leq B \tag{3.3}$$

$$s_l \geq 0 \text{ and integer } \forall l \in L, \tag{3.4}$$

$$s_j \geq 0 \text{ and integer } \forall j \in J. \tag{3.5}$$

Note that, for a *fixed* SRU spare parts strategy, 3.2 is discretely convex in s_l for all $l \in L$. Hence, if we fix an SRU spare parts strategy and replace the original target budget B with the residual target budget $B' = B - \sum_{jinJ} c_j s_j$, the problem can be solved for a budget level

$$\bar{B} \in (B' - \max_{l \in L} c_l, B']$$

using straightforward marginal analysis. Unfortunately, 3.2 is *not* discretely convex in s_j, $j \in J$. However, we can still construct a tradeoff curve for the problem by using a semi-enumerative approach, which we describe next.

3.4 Solution Approach

Our solution approach is similar to that of MOD-METRIC Muckstadt (1973). Detailed, step-by-step procedures are given in Caggiano and Muckstadt (2000). We provide a descriptive overview here.

Our approach breaks the problem down by *LRU family*, where an LRU family l consists of the LRU type l and the SRU types $j \in J_l$. Note that LRU families are *disjoint*, since SRU types are not shared across LRU types. Hence, the only spare stock levels impacting the objective function term $\mathcal{D}_l(s_l, s_{J_l})$ are s_l and $s_j, j \in J_l$. Conversely, the spare stock levels s_l and $s_j, j \in J_l$ only impact the term $\mathcal{D}_l(s_l, s_{J_l})$. Because of this property, we can concentrate our efforts on understanding budget allocation tradeoffs *within* each LRU family, and then put the pieces back together to arrive at a final solution.

Specifically, for each LRU family, we construct an *LRU family curve*, $\mathcal{T}_l$, depicting the minimum $\mathcal{D}_l(s_l, s_{J_l})$ that can be achieved at various family investment levels. That is,

$$\mathcal{T}_l = \left\{(B_l^i, \mathcal{D}_l^i, \mathcal{S}_l^i = (s_l^i, s_{J_l}^i)), i = 0, ..., n_{\mathcal{T}_l}\right\}$$

is a set of $n_{\mathcal{T}_l}$ points, ordered by increasing budget level, where each point i corresponds to a spare parts strategy $\mathcal{S}_l^i$ for LRU family l, the budget B_l^i required for this strategy, and the expected delay-causing backorders for LRU type l, $\mathcal{D}_l^i$, achieved by this strategy. Our LRU family curve construction is such that, when plotted, the points $(B_l^i, \mathcal{D}_l^i), i = 0, ..., n_{\mathcal{T}_l}$ describe a discretely convex function.

The key to constructing an LRU family curve lies in understanding how various levels of investment in the LRU family should be allocated among the LRU type l and the SRU types $j \in J_l$. We begin by constructing an *SRU tradeoff curve*, which depicts the best performance that can be achieved at various investment levels if we invest in SRUs only, and then complete the process by using this information to evaluate spare level combinations.

Recall that the sole purpose of holding spare inventory of an SRU type $j \in J_l$ is to decrease the waiting time component w_l of the repair cycle time for LRU type l. In Section 4, for each LRU type l, we formally introduce the variables $w_j, j \in J_l$, where w_j denotes the waiting time that an LRU of type l experiences *given* that it requires an SRU of type j to complete repair. $E[w_j|s_j]$, the expected waiting time for SRU type j, is a decreasing, discretely convex function in the spare level s_j. Because this is true for all $j \in J_l$, we can construct an SRU tradeoff curve using marginal analysis that depicts, for a range of budget levels, the lowest possible value of $E[w_l|s_{J_l}]$ that can be achieved, where $E[w_l|s_{J_l}]$ is the expected waiting time for SRU components experienced by a failed LRU of type l. The constructed SRU tradeoff curve is discretely convex in the investment level.

Once the SRU tradeoff curve is constructed, we use a semi-enumerative method to construct the LRU family curve. This method considers a set of budget levels between 0 and B that are dense enough to ensure that all undominated points will be found. For each budget level, the procedure evaluates all possible combinations of the LRU spare level s_l and the best complementary point on the SRU tradeoff curve to arrive at a single undominated point on the LRU family curve. Once the set of undominated points has been constructed, the procedure pares the curve points further by selecting those points that are on the convex minorant of the set of undominated points. See Figure 5.

Given these convex LRU family curves as input, we use another marginal analysis procedure to construct a final tradeoff curve for the entire problem. The main drawback to this construction approach is that the final tradeoff curve does not contain a point for every possible budget level, and the difference in budget level from one point to the next on the final curve will be equal to the difference (in budget level) between two consecutive points on one of the LRU family curves. Hence, if the LRU family curves have large differences in budget level from point to point, the final tradeoff curve will have large differences as well. Realistically, we expect the final tradeoff curves to be reasonably dense relative to any practical target budget level B.

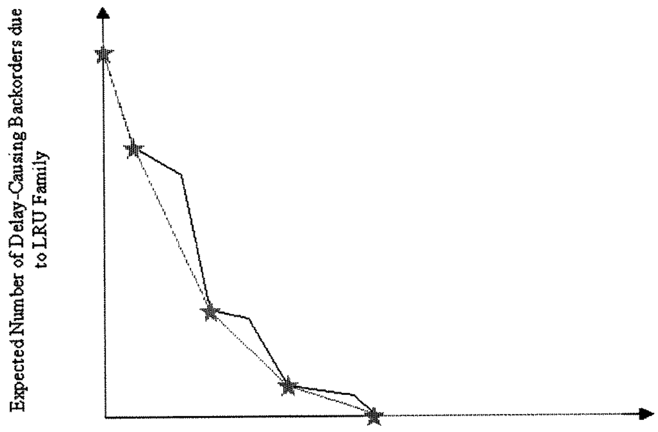

Figure 5: An LRU Family Curve.

4 The Distribution of N_l

Recall that N_l is the number of units of LRU type l in the repair process at an arbitrary maintenance schedule due date $T^* = i^*\gamma + \tau$ that entered the repair process at or before time $T^* - \tau_l$.

We begin by reviewing a simple proof that N_l has a binomial distribution when the LRU repair cycle time $\hat{C}_l$ is a known constant. We then characterize N_l as the sum of two independent random variables when $\hat{C}_l$ is a bounded random variable. Using this characterization, we derive the exact distribution of N_l for those LRU types l that are sent away for repair, and we outline three approximation methods for those LRU types l that are repaired in-house. All of the approximation methods for in-house repairs depend on the waiting time w_l, and we derive the exact distributions of w_l and $\hat{C}_l$ as part of our analysis.

4.1 The Distribution of N_l When $\hat{C}_l$ is Fixed

Suppose that for some LRU type l, $\hat{C}_l$ is a known, fixed length of time. Then N_l is exactly the number of LRUs of type l that entered the repair process in the time interval $(T^* - \hat{C}_l, T^* - \tau_l]$. Any LRUs of type l that entered the repair process at or before $T^* - \hat{C}_l$

will have completed the process by T^*, and any LRUs of type l that entered the repair process after $T^* - \hat{C}_l$ will still be in the repair process at T^*.

If $\hat{C}_l \le \tau_l$, then $\Pr[N_l = 0] = 1$, since every failure can be repaired before the end of the same maintenance cycle. (Hence, it is not necessary to carry any spare inventory for LRU l.) If $\hat{C}_l > \tau_l$, we can write

$$\hat{C}_l = \tau_l + \hat{m}_l\gamma + \hat{\varepsilon}_l,$$

where $\hat{m}_l$ and $\hat{\varepsilon}_l$ are unique nonnegative integers with $\hat{\varepsilon}_l < \gamma$. If $\hat{\varepsilon}_l > 0$, then exactly $\hat{m}_l + 1$ tests of LRUs of type l are conducted in the interval $(T^* - \hat{C}_l, T^* - \tau_l]$. If $\hat{\varepsilon}_l = 0$, then exactly $\hat{m}_l$ tests of LRUs of type l are conducted in this interval. Since LRUs of type l fail independently of one another with probability p_l, we have that $N_l \sim \mathbf{Bin}(\hat{m}_l + 1, p_l)$ if $\hat{\varepsilon}_l > 0$ and $N_l \sim \mathbf{Bin}(\hat{m}_l, p_l)$ if $\hat{\varepsilon}_l = 0$. That is:

$$\Pr[N_l = k | \hat{C}_l = \tau_l + \hat{m}_l\gamma + \hat{\varepsilon}_l] = \begin{cases} \binom{\hat{m}_l + 1}{k} p_l^k (1-p_l)^{\hat{m}_l + 1 - k} & \text{, if } \hat{\varepsilon}_l > 0. \\ \binom{\hat{m}_l}{k} p_l^k (1-p_l)^{\hat{m}_l - k} & \text{, if } \hat{\varepsilon}_l = 0. \end{cases} \tag{4.1}$$

4.2 Characterizing N_l When $\hat{C}_l$ is a Random Variable

Now suppose that $\hat{C}_l$ is a random variable whose distribution is known and bounded. In this case, we can write:

$$\Pr[\hat{C}_l \le \tau_l] + \sum_{k=0}^{\overline{m_l}} \Pr[\tau_l + k\gamma < \hat{C}_l \le \tau_l + (k+1)\gamma] = 1,$$

where $\overline{m_l}$ is *the smallest nonnegative integer such that the sum is equal to 1.* For purposes of exposition, we will assume that $\Pr[\hat{C}_l \le \tau_l] = 0$. (All of our analysis is easily extended if $\Pr[\hat{C}_l \le \tau_l] > 0$. See Caggiano and Muckstadt (2000) for details.) In this case, there also exists a largest nonnegative integer $\underline{m_l}$ (where $\underline{m_l} \le \overline{m_l}$) such that

$$\sum_{k=\underline{m_l}}^{\overline{m_l}} \Pr[\tau_l + k\gamma < \hat{C}_l \le \tau_l + (k+1)\gamma] = 1. \tag{4.2}$$

Given these integers, we can express N_l as the sum of two random variables:

$$N_l = N_l^{\le \underline{m_l}} + N_l^{> \underline{m_l}},$$

where $N_l^{\leq \underline{m_l}}$ denotes the number of units of LRU l in the repair process at T^* that entered the repair process in the time interval $[T^* - (\tau_l + \underline{m_l}\gamma), T^* - \tau_l]$, and $N_l^{> \underline{m_l}}$ denotes the number of units of LRU l in the repair process at T^* that entered the repair process in the time interval $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - (\tau_l + \underline{m_l}\gamma))$. Note that these intervals are disjoint. Since $(\tau_l + \underline{m_l}\gamma)$ is a strict lower bound on $\hat{C}_l$, $N_l^{\leq \underline{m_l}}$ is simply the number of LRUs of type l that entered the repair process in the time interval $[T^* - (\tau_l + \underline{m_l}\gamma), T^* - \tau_l]$. Thus, since the LRU failure process is independent of the repair process, $N_l^{\leq \underline{m_l}}$ and $N_l^{> \underline{m_l}}$ are independent, and

$$\Pr[N_l = k] = \sum_{n=0}^{k} \Pr[N_l^{\leq \underline{m_l}} = n] \Pr[N_l^{> \underline{m_l}} = (k - n)]. \tag{4.3}$$

So, to capture the distribution of N_l, it suffices to capture the distributions of $N_l^{\leq \underline{m_l}}$ and $N_l^{> \underline{m_l}}$. The former is easy, while the latter is difficult.

There are $\underline{m_l} + 1$ points in the time interval $[T^* - (\tau_l + \underline{m_l}\gamma), T^* - \tau_l]$ at which LRUs of type l are checked for failure. Thus, $N_l^{\leq \underline{m_l}} \sim \mathbf{Bin}(\underline{m_l} + 1, p_l)$, or:

$$\Pr[N_l^{\leq \underline{m_l}} = k] = \binom{\underline{m_l} + 1}{k} p_l^k (1 - p_l)^{\underline{m_l}+1-k}. \tag{4.4}$$

Let $I = \{\underline{m_l} + 1, \ldots, \overline{m_l}\}$ denote the index set of checkpoints in the time interval $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - (\tau_l + \underline{m_l}\gamma))$. There are $\overline{m_l} - \underline{m_l}$ checkpoints in this interval. At any *one* checkpoint $u \in I$, an LRU of type l fails with probability p_l, and the probability that the repair cycle time will extend beyond T^* is given by $\Pr[\hat{C}_l > \tau_l + u\gamma]$. While this information is sufficient to compute the expected value of N_l (see below), we need more information to compute the distribution of N_l since the latter depends upon the *joint* distribution of the repair cycle times associated with the $\overline{m_l} - \underline{m_l}$ checkpoints in $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - (\tau_l + \underline{m_l}\gamma))$.

The expected value of N_l is given by:

$$\begin{aligned} E[N_l] &= E[N_l^{\leq \underline{m_l}}] + E[N_l^{> \underline{m_l}}] \\ &= (m_l + 1)p_l + \sum_{k=\underline{m_l}+1}^{\overline{m_l}} p_l \Pr[\hat{C}_l > \tau + k\gamma] \\ &= p_l \left[(m_l + 1) + \sum_{k=\underline{m_l}+1}^{\overline{m_l}} \Pr[\hat{C}_l > \tau + k\gamma] \right]. \end{aligned} \tag{4.5}$$

Note that 4.5 follows directly from Little's Law Buzacott and Shanthikumar (1993). Arrivals of LRU type l to the repair process occur at a rate of p_l per γ-window, and the term in large brackets is the expected number of γ-windows beyond τ_l that an LRU of type l remains in the repair process.

In the sections that follow, we will derive exact and approximate distributions for $N_l^{>\underline{m_l}}$ in the context of our multi-indenture problem.

4.3 The Distribution of N_l for LRU Types Sent Away for Repair

Recall that for LRU types sent away for repair, repair cycle times $\hat{C}_l$ are assumed to be i.i.d. bounded random variables with known distributions. Assuming that $\Pr[\hat{C}_l \leq \tau] = 0$ and that $\underline{m_l}$ and $\overline{m_l}$ are defined as in 4.2, we have that:

$$\Pr[N_l^{>\underline{m_l}} = k] = \sum_{S \subseteq I : |S| = k} \left[p_l^k \prod_{u \in S} \Pr[\hat{C}_l > \tau_l + u\gamma] \right] \left[\prod_{v \in I \setminus S} \left(1 - p_l \Pr[\hat{C}_l > \tau_l + v\gamma]\right) \right], \tag{4.6}$$

where $I = \{\underline{m_l}+1, \ldots, \overline{m_l}\}$ denotes the index set of checkpoints in the time interval $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - (\tau_l + \underline{m_l}\gamma))$. Using Equations 4.4 and 4.6 for all relevant k, Equation 4.3 can be used to derive the exact distribution of N_l.

Note that the computational time required to compute the probabilities in 4.6 for all relevant k is $O(2^{\overline{m_l}})$. Clearly this approach is intractable for large values of $\overline{m_l}$. However, recall that this model is *specifically intended* for situations where $\overline{m_l}$ is small. Otherwise, a model assuming a Poisson failure process is more appropriate. In NASA's case, the number of vehicle landings over the maximum repair cycle time for an LRU is currently small and is expected to remain small for several years.

4.4 The Distribution of N_l for LRU Types Repaired In-House

Recall that for LRU types l repaired in-house, the repair cycle time $\hat{C}_l = C_l + w_l$, where C_l is a known constant and w_l is the waiting time for SRU components. Since w_l is a function of the SRU spare levels s_{J_l}, the distribution of N_l will also be a function of these spare levels.

First, we derive the distributions of w_l and $\hat{C}_l$. We then outline three methods for approximating the distribution of N_l and discuss the advantages and disadvantages of each method.

4.4.1 The Distributions of w_l and $\hat{C}_l$

Recall that LRUs of type l are tested at times $i\gamma + \delta_l$, for $i = 0, 1, 2, \ldots$. If an LRU failure occurs in a given cycle i', and the failure entails a failed SRU $j \in J_l$, the failed SRU will be removed from the LRU at time $i'\gamma + \delta_l + C_l^b$ and replaced with a spare *if one is available.*

We define N_j to be *the number of units of SRU type $j \in J_l$ in the SRU repair process at an arbitrary failure diagnosis completion time $T_l' = i'\gamma + \delta_l + C_l^b$ for LRU l* (not including the failed SRU that enters at time T_l'). If $N_j < s_j$, then at least one spare SRU of type j is available at this time and $w_l = 0$ for the failed LRU needing the part. However, if $N_j \geq s_j$, then no spares are available, and the failed LRU must wait to complete its repair until an SRU j becomes available. Note from its definition that N_j is exactly the number of SRUs of type j that entered the repair process in the time interval $(T_l' - \hat{C}_j, T_l')$. If $\hat{C}_j \leq \gamma$, then $\Pr[N_j = 0] = 1$. Otherwise, we can write

$$\hat{C}_j = \hat{m}_j\gamma + \hat{\varepsilon}_j,$$

where $\hat{m}_j$ and $\hat{\varepsilon}_j$ are unique nonnegative integers with $\hat{\varepsilon}_j < \gamma$. From our discussion in Section 4.1, it is immediate that $N_j \sim \mathbf{Bin}(\hat{m}_j, p_j)$ if $\hat{\varepsilon}_j > 0$ and $N_j \sim \mathbf{Bin}(\hat{m}_j - 1, p_j)$ if $\hat{\varepsilon}_j = 0$. That is:

$$\Pr[N_j = k | \hat{C}_j = \hat{m}_j\gamma + \hat{\varepsilon}_j] = \begin{cases} \binom{\hat{m}_j}{k} p_j^k (1 - p_j)^{\hat{m}_j - k} & \text{, if } \hat{\varepsilon}_j > 0. \\ \binom{\hat{m}_j - 1}{k} p_j^k (1 - p_j)^{\hat{m}_j - 1 - k} & \text{, if } \hat{\varepsilon}_j = 0. \end{cases} \tag{4.7}$$

For $j \in J_l$, let w_j denote *the time that a failed LRU diagnosed at T_l' must wait for an SRU component, given that the diagnosis revealed a faulty SRU of type j.* If $N_j = s_j + r - 1$, then the LRU will have to wait for the rth SRU of type j to complete its repair, since $r - 1$ LRUs needing SRU type j are in line ahead of it at time T_l'. Thus, our LRU will receive the SRU of type j that has been in the repair process the rth longest at time T_l'. If the rth oldest outstanding failure of SRU type j was diagnosed at time $T_l' - (\hat{m}_j - k)\gamma$ for some $k \in [0, 1, ..., \hat{m}_j]$, then the waiting time $w_j = k\gamma + \hat{\varepsilon}_j$. Note that when $s_j = 0$, the LRU will have to wait until its own faulty SRU can be repaired. Hence, $\Pr[w_j = \hat{m}_j\gamma + \hat{\varepsilon}_j | s_j = 0] = 1$. When $s_j > \hat{m}_j$, the LRU will never have to wait, so $\Pr[w_j = 0 | s_j > \hat{m}_j] = 1$. The following theorem describes the distribution of w_j in the case that $0 < s_j \leq \hat{m}_j$:

Theorem 1 *Given a constant SRU repair cycle time $\hat{C}_j > \gamma$ with the unique representation*

$$\hat{C}_j = \hat{m}_j\gamma + \hat{\varepsilon}_j, \;\; \hat{m}_j > 0 \;\; and \;\; \gamma > \hat{\varepsilon}_j \geq 0,$$

if $\hat{\varepsilon}_j > 0$, then the distribution of w_j, for any $0 < s_j \leq \hat{m}_j$, is given by:

$$\Pr[w_j = 0 | s_j] = \sum_{k=0}^{s_j - 1} \binom{\hat{m}_j}{k} p_j^k (1 - p_j)^{\hat{m}_j - k}, \tag{4.8}$$

$$\Pr[w_j = k\gamma + \hat{\varepsilon}_j | s_j] = \begin{cases} \binom{\hat{m}_j - k - 1}{s_j - 1} p_j^{s_j} (1 - p_j)^{\hat{m}_j - k - s_j} & , \; for \;\; k = 0, ..., (\hat{m}_j - s_j), \\ 0 & , \; for \;\; k > (\hat{m}_j - s_j). \end{cases} \tag{4.9}$$

If $\hat{\varepsilon}_j = 0$, then the distribution of w_j, for any $0 < s_j \leq \hat{m}_j$, is given by:

$$\Pr[w_j = 0 | s_j] = \sum_{k=0}^{s_j - 1} \binom{\hat{m}_j - 1}{k} p_j^k (1 - p_j)^{\hat{m}_j - 1 - k}, \tag{4.10}$$

$$\Pr[w_j = k\gamma | s_j] = \begin{cases} \binom{\hat{m}_j - k - 1}{s_j - 1} p_j^{s_j} (1 - p_j)^{\hat{m}_j - k - s_j} & , \; for \;\; k = 1, ..., (\hat{m}_j - s_j), \\ 0 & , \; for \;\; k > (\hat{m}_j - s_j). \end{cases} \tag{4.11}$$

Proof. Since $\Pr[w_j = 0] = \Pr[N_j < s_j]$, 4.8 and 4.10 are immediate from 4.7. To see that 4.9 and 4.11 are true, note that when $N_j = s_j + r - 1$, the rth oldest outstanding failure of SRU type j at T_l' is the same as *the s_jth most recent failure of SRU type j*, not including the failure at T_l'. Thus, the probability that $w_j = k\gamma + \hat{\varepsilon}_j$ for any $k \in [1, ..., \hat{m}_j]$ is precisely the probability that the s_jth most recent failure of SRU type j was diagnosed at time $T_l' - (\hat{m}_j - k)\gamma$. (If $\hat{\varepsilon}_j > 0$, this is true for $k = 0$ as well.) For this to happen, we must have exactly $s_j - 1$ failures within the $(\hat{m}_j - k - 1)$ most recent checkpoints (not including T_l'), and a failure on the $(\hat{m}_j - k)$th most recent checkpoint. There are $\binom{\hat{m}_j - k - 1}{s_j - 1}$ ways that the former can happen (each occurring with probability $p_j^{s_j-1}(1-p_j)^{\hat{m}_j-k-s_j}$), and only one way for the latter to happen (which occurs with probability p_j). For $k > (\hat{m}_j - s_j)$, there are fewer than s_j checkpoints in the range $[T_l' - (\hat{m}_j - k)\gamma, T_l')$, so it is impossible for the s_jth most recent failure to have occurred at $T_l' - (\hat{m}_j - k)\gamma$. □

Since at most one SRU failure occurs per LRU failure, we have as an immediate corollary to Theorem 1:

Corollary 1 *For any* $0 \leq \varepsilon < \gamma$*:*

$$
\begin{aligned}
\Pr[w_l = 0 | s_{J_l}] &= 1 - \frac{1}{p_l} \sum_{j \in J_l} p_j \left(1 - \Pr[w_j = 0 | s_j]\right), \\
\Pr[(k-1)\gamma + \varepsilon < w_l \leq k\gamma + \varepsilon | s_{J_l}] &= \begin{cases} \frac{1}{p_l} \sum_{j \in J_l} p_j \Pr[w_j = k_{j,\varepsilon}\gamma + \hat{\varepsilon}_j | s_j], \\ \qquad \textit{for} \quad k = 1, \ldots, (\max_{j \in J_l}(\hat{m}_j - s_j) + 1), \\ 0, \quad \textit{for} \quad k > \max_{j \in J_l}(\hat{m}_j - s_j) + 1, \end{cases} \\
\textit{where} \quad k_{j,\varepsilon} &= \begin{cases} k & \textit{if} \quad \hat{\varepsilon}_j \leq \varepsilon. \\ k-1 & \textit{if} \quad \hat{\varepsilon}_j > \varepsilon. \end{cases}
\end{aligned}
\tag{4.12}
$$

Also,

$$
E[w_l | s_{J_l}] = \sum_{j \in J_l} E[w_j | s_j] \frac{p_j}{p_l}. \tag{4.13}
$$

The expected values of $w_j, j \in J_l$ can be computed directly from their respective distributions. However, the following representation, which can be derived using Little's Law, is in many instances computationally preferable:

$$
E[w_j | s_j] = \begin{cases} \hat{\varepsilon}_j \Pr[N_j \geq s_j] + \frac{\gamma}{p_j} \left[\sum_{k=s_j+1}^{\hat{m}_j} (k - s_j) \Pr[N_j = k] \right] & \text{, if } \hat{\varepsilon}_j > 0. \\ \gamma \Pr[N_j \geq s_j] + \frac{\gamma}{p_j} \left[\sum_{k=s_j+1}^{\hat{m}_j - 1} (k - s_j) \Pr[N_j = k] \right] & \text{, if } \hat{\varepsilon}_j = 0. \end{cases} \tag{4.14}
$$

Now that we have the distribution of w_l, the distribution of $\hat{C}_l = C_l + w_l$ is immediate since C_l is a constant. However, we wish to represent the distribution of $\hat{C}_l$ in the γ-window form of 4.2. We show how to do this for the case that $C_l > \tau_l$ and $\varepsilon_l = (C_l - \tau_l)\text{div}\gamma > 0$. (The other cases are similar).

Writing $C_l = \tau_l + m_l\gamma + \varepsilon_l$, where m_l and ε_l are unique nonnegative integers and $0 < \varepsilon_l < \gamma$, we can set $\underline{m_l} = m_l$ and $\overline{m_l} = m_l + (\max_{j \in J_l}(\hat{m}_j - s_j) + 1)$ and conclude that:

$$\sum_{k=\underline{m_l}}^{\overline{m_l}} \Pr[\tau_l + k\gamma < \hat{C}_l \leq \tau_l + (k+1)\gamma | s_{J_l}] = 1.$$

or, equivalently,

$$\sum_{k=0}^{\overline{m_l}-\underline{m_l}} \Pr[\tau_l + (\underline{m_l} + k)\gamma < \hat{C}_l \leq \tau_l + (\underline{m_l} + k + 1)\gamma | s_{J_l}] = 1.$$

Using the fact that $w_l = \hat{C}_l - C_l$, we have:

$$\begin{aligned}
\Pr[\tau_l + \underline{m_l}\gamma < \ & \hat{C}_l \leq \tau_l + (\underline{m_l} + 1)\gamma | s_{J_l}] \\
&= \Pr[w_l \leq (\gamma - \varepsilon_l) | s_{J_l}] \\
&= \frac{1}{p_l} \sum_{j \in J_l} p_j \Pr[w_j \leq (\gamma - \varepsilon_l) | s_j] \\
&= \frac{1}{p_l} \left[\sum_{j \in J_l} p_j \Pr[w_j = 0 | s_j] + \sum_{j \in J_l : 0 < \hat{\varepsilon_j} \leq (\gamma - \varepsilon_l)} p_j \Pr[w_j = \hat{\varepsilon_j} | s_j] \right]. \\
\Pr[\tau_l + (\underline{m_l} + k)\gamma < \ & \hat{C}_l \leq \tau_l + (\underline{m_l} + k + 1)\gamma | s_{J_l}] \\
&= \Pr[(k-1)\gamma + (\gamma - \varepsilon_l) < w_l \leq k\gamma + (\gamma - \varepsilon_l) | s_{J_l}], \\
&\quad \text{for} \quad k = 1, ..., (\overline{m_l} - \underline{m_l}).
\end{aligned} \tag{4.15}$$

Since $(\gamma - \varepsilon_l) < \gamma$, the probabilities $\Pr[(k-1)\gamma + (\gamma - \varepsilon_l) < w_l \leq k\gamma + (\gamma - \varepsilon_l) | s_{J_l}]$ in the last expression are given by 4.12 with ε set to $(\gamma - \varepsilon_l)$.

For ease of exposition in the three methods described next, we will continue to assume that $\tau_l < C_l = \tau_l + m_l\gamma + \varepsilon_l$, where m_l and ε_l are unique nonnegative integers and $0 < \varepsilon_l < \gamma$.

4.4.2 Approximating the Distribution of N_l - Method 1

In this simple method, we minimize computation by assuming that $w_l = \lceil E[w_l | s_{J_l}] \rceil$ for all failed LRUs of type l, so that $\hat{C}_l = C_l + \lceil E[w_l | s_{J_l}] \rceil$ is a constant. In terms of the characterization of N_l given in Section 4.2, this method is equivalent to $N_l^{>m_l} \sim$ **Bin**(n_l, p_l), where

$$n_l = \begin{cases} (\varepsilon_l + \lceil E[w_l | s_{J_l}] \rceil)\text{div}\gamma & , \text{if } (\varepsilon_l + \lceil E[w_l | s_{J_l}] \rceil)\text{mod}\gamma > 0. \\ (\varepsilon_l + \lceil E[w_l | s_{J_l}] \rceil)\text{div}\gamma - 1 & , \text{if } (\varepsilon_l + \lceil E[w_l | s_{J_l}] \rceil)\text{mod}\gamma = 0. \end{cases} \tag{4.16}$$

Hence, this method assumes that $N_l \sim \mathbf{Bin}(m_l + 1 + n_l, p_l)$, where n_l is given by 4.16.

The primary advantage of Method 1 is that $E[w_l|s_{J_l}]$ can be computed efficiently, and hence, so can the distribution of N_l. The obvious drawback is that Method 1 does not account for the variability of w_l at all. This can easily result in LRU family curve points i with under- and overestimated values of $\mathcal{D}_l^i$. We will see an example of this in Section 5.

4.4.3 Approximating the Distribution of N_l - Method 2

In Method 2, we compute the distribution of $\hat{C}_l$ as in 4.15 and then use 4.6 to approximate the distribution of $N_l^{>\underline{m_l}}$. This method is only approximate since 4.6 assumes independence of repair cycle times. For in-house repairs, LRU repair cycle times (specifically, the waiting times w_l) are *not* independent since LRUs of type l that fail in close proximity to one another will compete for the same set of spare SRUs. Computationally, Method 2 is much less efficient than Method 1 since the distributions for w_l, $\hat{C}_l$, and N_l must be recomputed each time any $s_j, j \in J_l$ is changed. As we shall see, however, the computed distribution of N_l under Method 2 incorporates the variability of w_l and, in general, results in backorder estimates and stock level solutions that are very close to optimal.

4.4.4 Approximating the Distribution of N_l - Method 3

In Method 3, we explicitly derive the distribution of $N_l^{>\underline{m_l}}$ described in Section 4.2 for the special case in which the SRU repair cycle times $\hat{C}_j = \hat{m_j}\gamma + \hat{\varepsilon_j}, j \in J_l$ are all the same for LRU type l. For ease of exposition, we will assume that $\hat{m_{J_l}} > 0$ and $\hat{\varepsilon_{J_l}} > 0$, where $\hat{m_{J_l}}$ and $\hat{\varepsilon_{J_l}}$ denote the common parameters. We will also assume that $\hat{\varepsilon_{J_l}} + \varepsilon_l > \gamma$. In this case, $\underline{m_l} = m_l$ and $\overline{m_l} = m_l + (\max_{j \in J_l}(\hat{m_{J_l}} - s_j) + 1)$.

Recall that $N_l^{>\underline{m_l}}$ denotes the number of units of LRU type l in the repair process at T^* that entered the repair process in the time interval $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - (\tau_l + \underline{m_l}\gamma))$. In our case, this is equivalent to the number of units of LRU type l in the repair process at T^* that entered the repair process in the time interval $[T^* - (\tau_l + \overline{m_l}\gamma), T^* - C_l)$.

Now, recall that for LRU type l, $C_l = C_l^b + C_l^a$, where C_l^b denotes the portion of the repair cycle time that occurs up through the failure diagnosis completion, and C_l^a

denotes the portion of the repair cycle time that occurs after all SRU components have been received through the end of the cycle. (See Figure 3.) Thus, we can redefine $N_l^{\underline{>m_l}}$ as *the number of LRUs of type l that are awaiting SRU components at time $T^* - C_l^a$*.

Since at most one SRU failure occurs per LRU failure, we have that:

$$\begin{aligned} N_l^{\underline{>m_l}} &= \sum_{j \in J_l} [\text{Number of LRUs of type } l \text{ awaiting SRU } j \text{ at } T^* - C_l^a\ ] \\ &= \sum_{j \in J_l} [\text{Number of SRUs of type } j \text{ in repair at } T^* - C_l^a \text{ in excess of } s_j] \\ &= \sum_{j \in J_l} [\tilde{N}_j - s_j]^+, \end{aligned} \tag{4.17}$$

where $\tilde{N}_j$ denotes the number of SRUs of type j in repair at time $T^* - C_l^a$. However, $\tilde{N}_j$ is also the number of SRUs of type j that entered the repair process in the time interval $(T^* - C_l^a - \hat{C}_j, T^* - C_l^a]$. Since $\hat{C}_j = \hat{m}_j \gamma + \hat{\varepsilon}_j$, and since SRUs of type $j \in J_l$ potentially enter the repair process at times $i\gamma + \delta_l + C_l^b$, for $i = 0, 1, \ldots$, we must determine the number of potential entrances in the time interval $(T^* - C_l^a - \hat{C}_j, T^* - C_l^a]$. A general method for doing this is given in Caggiano and Muckstadt (2000). For the case we are considering (i.e., $C_l > \tau_l$, $\varepsilon_l > 0$, and $\hat{\varepsilon}_j + \varepsilon_l > \gamma$), there are $(\hat{m}_j + 1)$ potential entrances in the time interval $(T^* - C_l^a - \hat{C}_j, T^* - C_l^a]$.

Thus, for all $j \in J_l$, we have that $\tilde{N}_j \sim \mathbf{Bin}(\hat{m}_{J_l} + 1, p_j)$, and the vector $\tilde{N}_{J_l}$ of the number of each type of SRU $j \in J_l$ in the repair process at time $T^* - C_l^a$ has a multinomial distribution with $\hat{m}_{J_l} + 1$ trials and $|J_l| + 1$ categories (one for each SRU $j \in J_l$ and one for the case that no SRU fails). Indexing the SRU types $j \in J_l$ with (1, 2, ..., $|J_l|$), define the sets

$$A_k = \left\{ (a_1, a_2, \ldots, a_{|J_l|}) \in Z_+^{|J_l|} : (\sum_{j \in J_l} a_j) \le \hat{m}_{J_l} + 1 \text{ and } \sum_{j \in J_l} [a_j - s_j]^+ = k \right\}.$$

Using these sets, we can characterize the distribution of $N_l^{\underline{>m_l}}$ as follows:

$$\begin{aligned} &\Pr[N_l^{\underline{>m_l}} = k] = \Pr\left[\sum_{j \in J_l} [\tilde{N}_j - s_j]^+ = k \right] \\ &= \sum_{(a_1, \ldots, a_{|J_l|}) \in A_k} \frac{(\hat{m}_{J_l} + 1)!}{a_1! \ldots a_{|J_l|}! ((\hat{m}_{J_l} + 1) - \sum_{j \in J_l} a_j)!} p_1^{a_1} \ldots p_{|J_l|}^{a_{|J_l|}} (1 - \sum_{j \in J_l} p_j)^{((\hat{m}_{J_l}+1) - \sum_{j \in J_l} a_j)}. \end{aligned} \tag{4.18}$$

The advantage of Method 3 over the others is that it captures the exact distribution of N_l (i.e., without assuming independence of the repair cycle times $\hat{C}_l$ from cycle to cycle). However, the sets A_k are difficult to characterize, and computing the distribution of $N_l^{\underline{>m_l}}$ using 4.18 is computationally undesirable if either $|J_l|$ or $\hat{m}_{J_l}$ is large. While we have assumed throughout this paper that $\hat{m}_{J_l}$ is reasonably small, we have *not* made any assumptions about the size of the sets J_l.

5 Comparison of Models

In this section we compare the LRU family curves obtained using the three methods outlined in the previous section with the LRU family curve obtained using a comparative Poisson model.

There are many ways that one could construct a comparative model based on a Poisson failure process. In order to have a fair comparison, we felt that the model should have at least the following properties:

- Failures of LRU type l occur according to a Poisson process with mean $\lambda_l = p_l/\gamma$.
- Failures of SRU type j occur according to a Poisson process with mean $\lambda_j = p_j/\gamma$.
- An LRU of type l that fails at time t must be replaced by time $t + \tau_l$ in order to avoid a delay-causing backorder.

A commonly used model assumes that the distribution of N_l is Poisson, and Sherbrooke has given an expression for the expected value of N_l in the case that backorders cause vehicle delays immediately Sherbrooke (1986). Using the above properties in conjunction with this expression, it is not hard to derive the following results:

- For each LRU type l with outsourced repair and a repair cycle time distribution such that $\Pr[\hat{C}_l \leq \tau_l] = 0$, the distribution of N_l is Poisson with mean $E[N_l] = \lambda_l(E[\hat{C}_l] - \tau_l)$ Feeney and Sherbrooke (1966).

- For each LRU type l repaired in-house having $C_l \geq \tau_l$, the distribution of N_l has mean

$$\lambda_l(C_l - \tau_l) + \sum_{j \in J_l} [\sum_{k > s_j} (k - s_j) \Pr[N_j = k]], \tag{5.1}$$

where N_j has a Poisson distribution with mean $\lambda_j \hat{C}_j$ Sherbrooke (1986).

Hence, for our comparative model, we assume that N_l has a Poisson distribution with a mean given by 5.1.

Table 1: Example Problem Data

	RLV		LRU		SRU1	SRU2
γ	10	p_l	0.5	p_j	0.1	0.4
τ	30	τ_l	25	$\hat{C}_j$	56	56
		C_l	42	c_j	300	400
		c_l	1000			

For ease of exposition we consider a single LRU family with two SRU components. The problem data are given in Table 1. (Examples containing multiple LRUs are illustrated in Caggiano and Muckstadt (2000)). Figures 7 and 8 show different portions of the LRU family curves resulting from Methods 1, 2, 3, and the Poisson model. Figure 6 gives the corresponding curve point information for each method. Method 3 gives exact values for the expected number of delay-causing backorders (denoted by E[D] in the tables), and hence provides a basis of comparison for the other methods. While we were able to supply the optimal tradeoff curve for this small example, Method 3's complicated characterization of N_l makes computation of the distribution extremely difficult in general.

There are several important observations. First, the Method 1 curve illustrates the danger of ignoring the variability of w_l in the N_l computation. For instance, the solution at the 3900 level claims that no delays will occur if the (2, 1, 4) stock level strategy is followed. Clearly this is not the case.

Method 1 - Final Curve Points

		Stock Levels		
Investment	E[D]	LRU	SRU1	SRU2
0	4	0	0	0
400	3	0	0	1
800	2	0	0	2
1800	1.0625	1	0	2
2900	0.25	1	1	4
3900	0	2	1	4

Method 2 - Final Curve Points

		Stock Levels		
Investment	E[D]	LRU	SRU1	SRU2
0	4	0	0	0
400	3.046656	0	0	1
800	2.279936	0	0	2
1100	1.811377	0	1	2
1500	1.355697	0	1	3
2500	0.526776	1	1	3
3500	0.109398	2	1	3
3900	0.049602	2	1	4
4200	0.016002	2	2	4
4600	0.005385	2	2	5
4900	0.001357	2	3	5
5200	0.000223	3	2	4
5600	1.4E-05	2	4	6
5900	2.5E-07	2	5	6
6200	0	2	6	6

Method 3 (Exact) - Final Curve Points

		Stock Levels		
Investment	E[D]	LRU	SRU1	SRU2
0	4	0	0	0
400	3.046656	0	0	1
800	2.279936	0	0	2
1100	1.811377	0	1	2
1500	1.355697	0	1	3
2500	0.533291	1	1	3
3500	0.123316	2	1	3
3900	0.055092	2	1	4
4200	0.018283	2	2	4
4600	0.005995	2	2	5
4900	0.001384	2	3	5
5300	0.00036	2	3	6
5600	1.45E-05	2	4	6
5900	2.5E-07	2	5	6
6200	0	2	6	6

Poisson Model - Final Curve Points

		Stock Levels			
Investment	E[D]	LRU	SRU1	SRU2	Actual E[D]
0	3.63361	0	0	0	4
400	2.754605	0	0	1	3.046656
800	2.101023	0	0	2	2.279936
1100	1.672361	0	1	2	1.811377
1500	1.284409	0	1	3	1.355697
2500	0.561223	1	1	3	0.5332907
3500	0.193579	2	1	3	0.1233158
3900	0.130659	2	1	4	0.0550917
4500	0.054267	3	1	3	0.019243
4900	0.031906	3	1	4	0.005995
5200	0.022305	3	2	4	0.0013838
5500	0.012711	4	1	3	0.0013838
5900	0.006469	4	1	4	0.0003598
6200	0.004106	4	2	4	0.0000145
6500	0.002546	5	1	3	0.0000145

Figure 6: LRU Family Curve Points

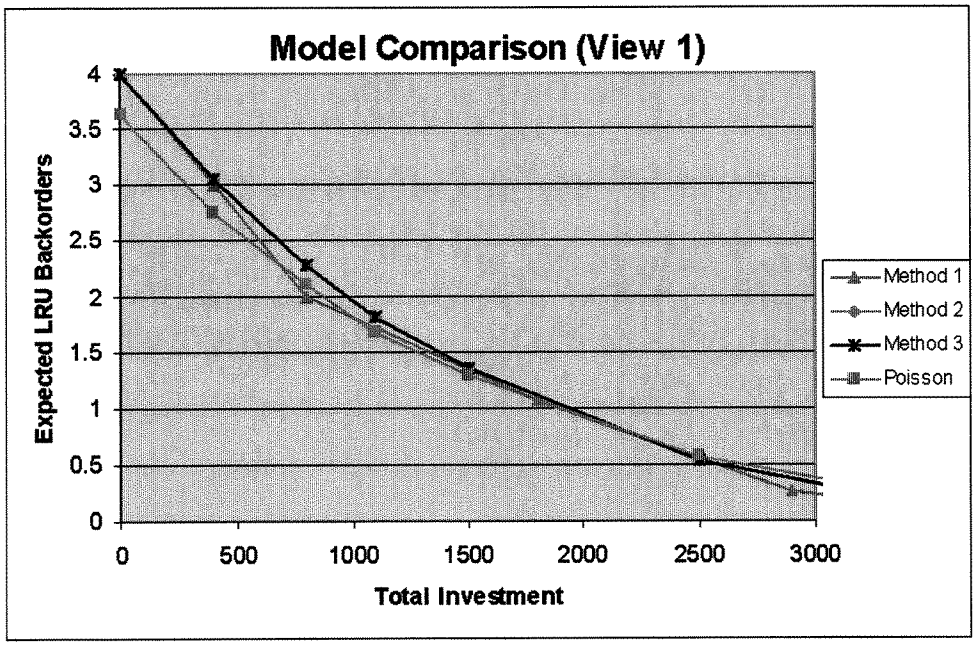

Figure 7: LRU Family Curve Comparison - View 1

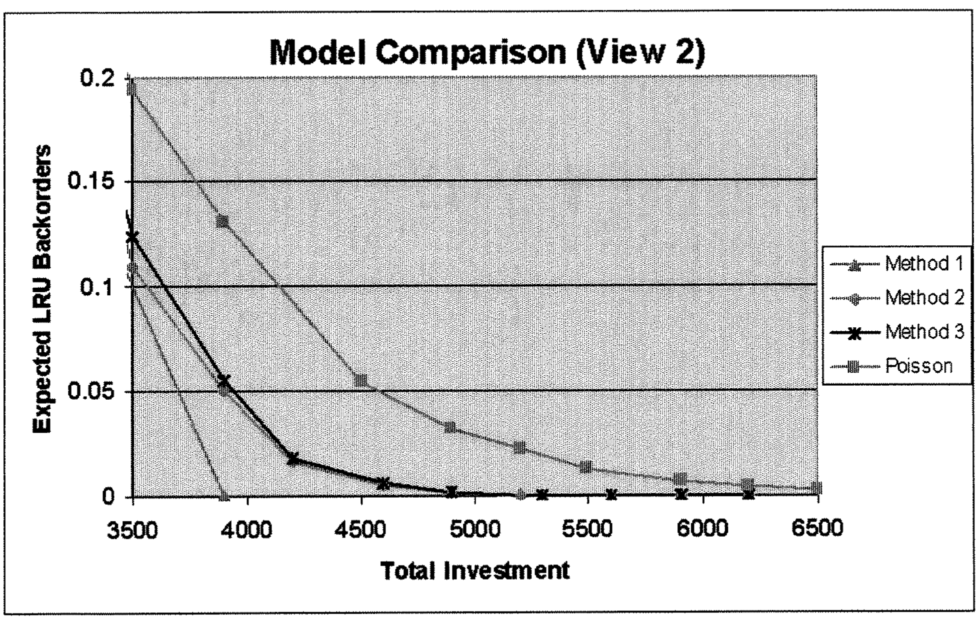

Figure 8: LRU Family Curve Comparison - View 2

Second, the tail points of the curves are particularly important in NASA's case since these are the points that correspond to the low E[D] values that will be desirable for each LRU family. Not only does the Poisson model predict E[D] values that are much higher than the true E[D] values for its curve point stock levels, but the stock levels themselves differ dramatically from the optimal levels. The implication of using the Poisson model is that NASA's inventory investment could be substantially higher than necessary as well as allocated in a suboptimal way.

Finally, we note that the curves in this example are representative of what we have typically encountered during our experimentation. Collectively, our analyses suggest that Method 2 is far superior to Method 1 and the comparative Poisson model, both in terms of estimating E[D] and in terms of allocation.

6 Conclusions

In this paper, we developed a mathematical model and a solution approach for determining and evaluating spare LRU and SRU stock levels when part failures occur according to a fixed cyclic schedule. We presented three methods for approximating the distribution of the number of relevant LRUs in repair at an RLV due date and compared the solution curves obtained using these methods with the solution curve from a comparative Poisson model. Our analyses show that there is ample justification for using a combinatorial model in lieu of a Poisson model when the number of maintenance cycle starts over the LRU repair cycle time is small. Specifically, the Poisson model underestimates the expected shortages for low levels of investment while dramatically overestimating expected shortages in the range of performance that would be desired in practice. For the NASA environment in particular, this could result in expending many millions of dollars in spare parts inventory needlessly. Our proposed Method 2 produces solutions that are far superior without excessive computational effort.

References

Buzacott, J. A. and Shanthikumar, J. G. (1993), *Stochastic Models of Manufacturing Systems*, Englewood Cliffs, NJ: Prentice-Hall.

Caggiano, K. and Muckstadt, J. (2000), "Maintenance Support for the Reusable Launch Vehicle Program: Determining and Evaluating Spare Stock Levels for Recoverable Parts," *Technical Report No. 1269, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853.*

Feeney, G. J. and Sherbrooke, C. C. (1966), "The (S-1,S) Inventory Policy under Compound Poisson Demand," *Management Science*, 12, 391–411.

Graves, S. C. (1985), "A Multi-Echelon Inventory Model for a Repairable Item with One-for-One Replenishment," *Management Science*, 31, 1247–1256.

Lee, H. and Moinzadeh, K. (1987), "Operating Characteristics of a Two-Echelon Inventory System for Repairable and Consumable Items under Batch Ordering and Shipment Policy," *Naval Research Logistics*, 34, 365–380.

Moinzadeh and Lee, H. L. (1986), "Optimal Batch Size and Stocking Levels in Multi-Echelon Repairable Systems," *Management Science*, 32, 1567–1581.

Muckstadt, J. A. (1973), "A Model for a Multi-Item, Multi-Echelon, Multi-Indenture Inventory System," *Management Science*, 20, 472–481.

Shanker, K. (1981), "Exact Analysis of a Two-Echelon Inventory System for Recoverable Items Under Batch Inspection Policy," *Naval Research Logistics Quarterly*, 28, 579–601.

Sherbrooke, C. C. (1968), "METRIC: A Multi-Echelon Technique for Recoverable Item Control," *Operations Research*, 16, 122–141.

— (1971), "An Evaluator for the Number of Operationally Ready Aircraft in a Multilevel Supply System," *Operations Research*, 19, 618–635.

— (1986), "VARI-METRIC: Improved Approximations for Multi-Indenture, Multi-Echelon Availability Models," *Operations Research*, 34, 311–319.

Silver, E. A. (1972), "Inventory Allocation Among an Assembly and its Repairable Subassemblies," *Naval Research Logistics Quarterly*, 19, 261–280.

Simon, R. (1971), "Stationarity Properties of a Two-Echelon Inventory Model for Low Demand Items," *Operations Research*, 19, 761–773.

Slay, M. (1980), "VARI-METRIC - An Approach to Modeling Multi-echelon Resupply when the Demand Process is Poisson with a Gamma Prior," *Working Paper, Logistics Management Institute, Washington, D.C.*

Svoronos, A. and Zipkin, P. (1991), "Evaluation of One-for-One Replenishment Policies for Multiechelon Inventory Systems," *Management Science*, 37, 68–83.

Wang, Y., Cohen, M., and Zheng, Y.-S. (2000), "A Two-Echelon Repairable Inventory System with Stocking-Center-Dependent Depot Replenishment Lead Times," *Technical Memorandum Number 730, Weatherhead School of Management, Case Western Reserve University, Cleveland, OH 44106-7235.*