SCHOOL OF OPERATIONS RESEARCH
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK

TECHNICAL REPORT NO. 313

September 1976

SOME APPROXIMATIONS IN MULTI-ITEM, MULTI-ECHELON
INVENTORY SYSTEMS FOR RECOVERABLE ITEMS

by

John A. Muckstadt

This research was supported in part by the Office of Naval Research under contract N00014-75-C-1172, Task NR 042-335 and the RAND Corporation under Project RAND.

TABLE OF CONTENTS

	PAGE
ABSTRACT	i
SECTION	
1. INTRODUCTION	1
2. THE APPROXIMATION PROBLEM	3
3. COMPUTING OPTIMAL SOLUTIONS FOR PROBLEMS 2 and 3	11
4. A COMPARISON OF ALTERNATIVE SOLUTION PROCEDURES FOR SOLVING PROBLEM 1	14
5. A COMPUTATIONAL COMPARISON OF VARIOUS ALGORITHMS	21
REFERENCES	26

ABSTRACT

Almost a decade ago Sherbrooke formulated the well known METRIC model for determining optimal stock levels for recoverable items for two echelon inventory systems [3]. Subsequently Fox and Landi [2] proposed a Lagrangian approach for obtaining item stock levels for each location. In this paper we develop a method for estimating the value of the optimal Lagrangian multiplier used in the Fox-Landi algorithm, present alternative ways for determining system stock levels, and compare these proposed approaches with the Fox-Landi algorithm and other solution techniques. The conclusion of this study is that the proposed approximation methods significantly reduce computation time for determining system stock levels without degrading the quality of the solution.

1. INTRODUCTION

Almost a decade ago Sherbrooke formulated the well-known METRIC model for determining optimal stock levels for recoverable items--items subject to repair when they fail--in a two-echelon setting [3]. In particular, he studied the Air Force's two-echelon supply system. This system consists of a set of bases and their supporting depots. Primary demands occur at the bases while depots are central repair and inventory stocking points which resupply bases when necessary. When a failure occurs at a base, a demand is placed on the base supply organization for a corresponding replacement part. Depending on the nature of the failure, the failed part is then either repaired at that base or is sent to a depot for repair. Resupply of the base supply organization comes from the base maintenance organization if repair is accomplished at the base; otherwise, resupply comes from a depot. In either case, the organization resupplying the base supply activity does so by exchanging on a one-for-one basis a serviceable part for the failed part. Thus the inventory policy for placing orders on the base's maintenance organization or a depot is an $(s-1, s)$ policy.

Sherbrooke presented a model (METRIC) for determining both depot and base stock levels for all items for this system. In particular, the problem he formulated was to minimize the average total number of base backorders existing at an arbitrary point in time subject to a constraint on system investment; that is,

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^m \sum_{i=1}^n \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{io})) \\ \text{subject to} \quad & \sum_{j=0}^m \sum_{i=1}^n c_i s_{ij} \leq C, \end{aligned} \tag{1}$$

where n represents the number of items,
 m represents the number of bases,
 s_{ij} represents the stock level at base j for item i ,
 s_{i0} represents the depot stock level for item i ,
 λ_{ij} represents the expected daily demand rate for item i at base j ,
 c_i represents the unit cost for item i ,
 C represents the budget constraint,
 $T_{ij}(s_{i0})$ represents the average resupply time for base j for item i
 given the depot stock level for item i is s_{i0} ,
 and $p(x|y)$ represents the probability that x units are in the resupply
 system given that the expected number of units in the resupply
 system is y .

Furthermore, Sherbrooke shows that $T_{ij}(s_{i0})$ can be expressed as

$$T_{ij}(s_{i0}) = r_{ij}A_{ij} + (1-r_{ij})(B_{ij} + \delta(s_{i0}) \cdot D_i),$$

where A_{ij} is the average base repair time for item i at base j ,
 r_{ij} is the proportion of demands requiring base repair for item i
 at base j ,

B_{ij} is the average depot to base order-and-ship time at base j for
 item i ,

D_i is the average depot repair cycle time for item i ,

$$\delta(s_{i0}) \cdot D_i = \frac{1}{\lambda_i} \sum_{x > s_{i0}} (x - s_{i0}) p(x | \lambda_i D_i), \text{ the expected delay per depot}$$

demand for item i ,

$$\text{and } \lambda_i = \sum_{j=1}^m (1-r_{ij}) \lambda_{ij}, \text{ the expected daily depot demand rate for item } i.$$

In the remainder of the paper i will refer to an item and j will refer to a base ($j=0$ represents the depot); thus i and j will always be elements of the sets $\{1, \dots, n\}$ and $\{0, \dots, m\}$, respectively. Additionally, an integer k appearing in the text to the right of the statement of a problem or equation will designate for future reference that problem or equation. For a complete description of this problem's background and formulation see reference 3.

Subsequently Fox and Landi suggested a Lagrangian approach for solving Problem 1 [2]. One of the major obstacles to the successful implementation of METRIC using the Fox-Landi algorithm is the requirement of estimating an appropriate value for the Lagrangian multiplier. An important and related problem is the lengthy computer run time required to obtain an optimal solution to Problem 1 when using this algorithm.

The purposes of this paper are to present an approach for obtaining an estimate of the optimal Lagrange multiplier value required in the Fox-Landi algorithm, to present two new methods for determining stock levels, and to compare these methods with the Fox-Landi method and other techniques. The proposed approach eliminates the particularly time consuming portion of the Fox-Landi algorithm devoted to searching for the best Lagrange multiplier value. The conclusion of the study is that the proposed approximation methods significantly reduce computation time for determining stock levels without degrading the quality of the solution.

2. THE APPROXIMATION PROBLEM

We begin this section by constructing a problem that is a continuous approximation to Problem 1. We next state and prove two theorems that are the basis for an algorithm developed in the next section which can be used to solve this approximating problem.

Recall that the total average base backorders existing at any point in time for item i can be expressed as

$$\sum_{j=1}^m \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})).$$

Two useful probability distributions for describing the demand process are the Poisson and negative binomial distributions. As shown in reference 1, this implies that if demand has a Poisson or negative binomial distribution, then for a given $\lambda_{ij} T_{ij}(s_{i0})$, the form of $p(x | \lambda_{ij} T_{ij}(s_{i0}))$, the probability distribution representing the number of units in resupply of item i at base j at any point in time, is a Poisson or negative binomial distribution, respectively.

Experimental data gathered during the conduct of this study indicate that when $p(x | \lambda_{ij} T_{ij}(s_{i0}))$ is either a Poisson or negative binomial distribution, the above total expected backorder expression can be closely approximated by an exponential function. That an exponential function accurately approximates this expression should not be entirely unexpected. First, for budgets of practical interest the item stock levels, s_{ij} , are normally much larger than the average demand during the resupply time. For example, the probability of running out of stock during the resupply time is often much less than .15 in real Air Force applications. Thus, the only probabilities entering the backorder calculation are the tail probabilities. In the tails, the Poisson and negative binomial distributions behave almost like the geometric distribution; that is, each succeeding probability is roughly a constant proportion of its predecessor. Consequently, when s_{ij} is large relative to $\lambda_{ij} T_{ij}(s_{i0})$, the expected number of backorders existing at any point in time at location j for item i is approximately a geometric function of s_{ij} .

Therefore an exponential function is a useful continuous approximation to this relationship between expected backorders at a location and the item's stock level at that location.

Furthermore, total expected base backorders exhibit this same behavior. If demand has either a Poisson or negative binomial distribution (or, for that matter, any other compound Poisson distribution), then the total number of units of an item in resupply across all bases also has a Poisson or negative binomial distribution, respectively, given independence of demand among bases and assuming the "order size" distribution is the same at all bases. Since in most practical situations total system stock substantially exceeds the total expected number of units in resupply, the tail of the distribution describing the total number of units in resupply is the only portion of the distribution of importance. As an approximation, this distribution can be used to determine the nature of the relationship between total expected base backorders and total system stock. For the reasons discussed previously, an exponential function should adequately represent this relationship as well.

Thus we will approximate total system backorders for item i , that is,

$$\sum_{j=1}^m \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})),$$

with an exponential function of the form

$$B_i(N_i) \equiv a_i e^{-b_i N_i}.$$

In this approximation N_i represents total system stock. The parameters $a_i > 0$ and $b_i > 0$ are estimated using regression analysis. The data used in the regression analysis are the backorder data obtained from the solution to

the problem

$$\text{minimize } \sum_{j=1}^m \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0}))$$

subject to

$$\sum_{j=0}^m s_{ij} = N_i, \quad \text{and}$$

$$s_{ij} = 0, 1, \dots, N_i,$$

for several appropriate values of N_i .

We now formulate a continuous approximation to Problem 1 in which the exponential representation of total system backorders for an item is used. In this approximation problem the decision variables are the total system stock, N_i , rather than the stock levels for each location, s_{ij} . As we shall see, the main reason for studying this approximation problem is that it is a vehicle for obtaining an estimate of the optimal Lagrangian multiplier value used in the Fox-Landi algorithm. The approximation problem is formulated as

$$\text{minimize } \sum_{i=1}^n B_i(N_i)$$

subject to

$$\sum_{i=1}^n c_i N_i \leq C,$$

$$N_i \geq 0.$$

(2)

Note that N_i is a continuous variable in this approximation. The optimality conditions (Kuhn-Tucker conditions) for this problem are as follows:

Find $\theta_1 \geq 0$ such that

$$a) \quad \frac{dB_i}{dN_i} + \theta_1 c_i \geq 0,$$

$$b) \quad \sum_{i=1}^n c_i N_i \leq C,$$

$$N_i \geq 0,$$

$$c) \quad \theta_1 \left(\sum_{i=1}^n c_i N_i - C \right) = 0,$$

and

$$d) \quad N_i \left(\frac{dB_i}{dN_i} + \theta_1 c_i \right) = 0.$$

A relaxed version of Problem 2 in which the non-negativity constraint on the item stock level is removed is

$$\text{minimize} \quad \sum_{i=1}^n B_i(N_i)$$

subject to

$$\sum_{i=1}^n c_i N_i \leq C.$$

(3)

The optimality conditions for this problem are:

Find $\theta_2 \geq 0$ such that

$$a) \quad \frac{dB_i}{dN_i} + \theta_2 c_i = 0,$$

$$b) \quad \sum_{i=1}^n c_i N_i \leq C,$$

$$c) \quad \theta_2 \left(\sum_{i=1}^n c_i N_i - C \right) = 0,$$

and

$$d) \quad N_i \left(\frac{dB_i}{dN_i} + \theta_2 c_i \right) = 0.$$

We now explore the relationship between Problems 2 and 3 in detail.

Suppose we obtained a solution to Problem 3 (we'll show how to find its solution in the next section). Let N_i^1 represent the optimal solution to Problem 2, and N_i^2 represent the optimal solution to Problem 3. If $N_i^2 \geq 0$ for all i , then $N_i^1 = N_i^2$ and the objective function values are equal.

Suppose, however, that $N_i^2 < 0$ for at least one value of i . Let

$$\bar{N}_i = \max(0, N_i^2)$$

and

$$\bar{C} \equiv \sum_{i=1}^n c_i \bar{N}_i.$$

Since $\bar{N}_i \geq N_i^2$ for all i and $\bar{N}_i > N_i^2$ for at least one i , $\bar{C} > C$.

Suppose Problem 2 is modified slightly so that the right-hand side value C is replaced by $\bar{C}$. This modified problem is

$$\begin{aligned} & \text{minimize} \quad \sum_{i=1}^n B_i(N_i) \\ & \text{subject to} \quad \sum_{i=1}^n c_i N_i \leq \bar{C}, \end{aligned} \tag{4}$$

$$N_i \geq 0.$$

The optimality conditions for this problem are the same as those given for Problem 2 after substituting $\bar{C}$ for C . Also, let $\bar{\theta}$ represent the optimal value of the Lagrangian multiplier for Problem 4.

In solving Problem 3, we will obtain a value for θ_2 . We now show that $\bar{\theta} = \theta_2$, and that $\bar{N}_i = \max(0, N_i^2)$ is the optimal solution to Problem 4 by demonstrating that these values satisfy the Kuhn-Tucker conditions corresponding to Problem 4.

By construction,

$$\sum_{i=1}^n c_i \bar{N}_i = \bar{C}, \quad \bar{N}_i \geq 0, \quad \text{and} \quad \bar{\theta} \left(\sum_{i=1}^n c_i \bar{N}_i - \bar{C} \right) = 0.$$

If $\bar{\theta} = \theta_2$, $\bar{\theta} \geq 0$ since $\theta_2 \geq 0$. Suppose $\bar{N}_i = N_i^2$; that is, $N_i^2 \geq 0$. Then

$$\left. \frac{dB_i}{dN_i} \right|_{N_i = \bar{N}_i} = \left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^2} \quad \text{and}$$

$$0 = \left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^2} + \theta_2 c_i = \left. \frac{dB_i}{dN_i} \right|_{N_i = \bar{N}_i} + \bar{\theta} c_i.$$

By assumption there exists at least one value of i for which $\bar{N}_i > N_i^2$; that is, $\bar{N}_i = 0$ while $N_i^2 < 0$. Since

$$\left. \frac{dB_i}{dN_i} \right|_{N_i = 0} > \left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^2},$$

due to the exponential form of $B_i(N_i)$, and

$$\left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^2} + \theta_2 c_i = 0, \quad \text{we know that} \quad \left. \frac{dB_i}{dN_i} \right|_{N_i = 0} + \bar{\theta} c_i > 0.$$

Consequently, the optimal solution to Problem 4 is $N_i = \bar{N}_i = \max\{0, N_i^2\}$. Furthermore, the optimality conditions are satisfied when $\bar{\theta}$ is equal to θ_2 .

Theorem 1. $\theta_1 \geq \theta_2$.

Proof: The optimal objective function value for Problem 2 is a convex, differentiable, strictly decreasing function of the available budget, C . Since the slope of this function is equal to the negative of the Lagrangian multiplier value, $\theta_1 \geq \bar{\theta}$ since $C \leq \bar{C}$. But $\theta_2 = \bar{\theta}$, so $\theta_1 \geq \theta_2$. |

Corollary. $\theta_1 > \theta_2$ when $\bar{C} > C$.

Next we compare N_i^1 with $\bar{N}_i$. If $C = \bar{C}$, then $N_i^1 = \bar{N}_i$ for all i . Now let us suppose $\bar{C} > C$ so that $\theta_1 > \theta_2 = \bar{\theta}$. Let us examine the two cases $\bar{N}_i > 0$ and $\bar{N}_i = 0$ separately.

First assume $\bar{N}_i > 0$. Then

$$\left. \frac{dB_i}{dN_i} \right|_{N_i = \bar{N}_i} + \bar{\theta} c_i = 0.$$

Furthermore, if $N_i^1 > 0$, then $\left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^1} + \theta_1 c_i = 0$.

Since $\theta_1 c_i > \bar{\theta} c_i = - \left. \frac{dB_i}{dN_i} \right|_{N_i = \bar{N}_i}$, $\left. \frac{dB_i}{dN_i} \right|_{N_i = \bar{N}_i} > \left. \frac{dB_i}{dN_i} \right|_{N_i = N_i^1}$, and $N_i^1 < \bar{N}_i$.

If $N_i^1 = 0$, then $\bar{N}_i > N_i^1$.

Next assume $\bar{N}_i = 0$. Since

$$\left. \frac{dB_i}{dN_i} \right|_{N_i=0} + \theta_1 c_i > \left. \frac{dB_i}{dN_i} \right|_{N_i=0} + \bar{\theta} c_i \geq 0, \text{ it follows that } N_i^1 = 0$$

by complementary slackness. Thus we have proven the following theorem.

Theorem 2. $\bar{N}_i \geq N_i^1$; additionally, $\bar{N}_i > N_i^1$ whenever $\bar{N}_i > 0$.

In this section we established several important relationships among Problems 2, 3, and 4. In the next section we develop a simple algorithm for solving Problem 2 based on these relationships. We will begin the next section by showing how to find the solution to Problem 3. As we have just demonstrated, once we have the solution to Problem 3 we also have the solution to Problem 4. From Theorem 2, we then have an upper bound on the value of N_i^1 . In particular, if $\bar{N}_i = 0$, then $N_i^1 = 0$. Combining this observation with the implications of Theorem 1 and its corollary provides the bases for the proposed algorithm for solving Problem 2.

3. COMPUTING OPTIMAL SOLUTIONS FOR PROBLEMS 2 AND 3

We begin this section by developing a method for determining the optimal solution to Problem 3. Observe that the optimal solution must satisfy the following two conditions:

$$\frac{dB_i}{dN_i} + \theta_2 c_i = 0$$

and

$$\sum_{i=1}^n c_i N_i = C.$$

The second condition must hold since each $B_i(N_i)$ is a strictly decreasing function of N_i .

Since

$$B_i(N_i) = a_i e^{-b_i N_i},$$

where $a_i, b_i > 0$, the first condition states that

$$\theta_2 = \frac{a_i b_i e^{-b_i N_i}}{c_i} > 0,$$

or

$$\hat{\theta} \equiv \ln \theta_2 = \ln \left\{ \frac{a_i b_i}{c_i} \right\} - b_i N_i.$$

Letting

$$d_i = \ln \left\{ \frac{a_i b_i}{c_i} \right\},$$

we see that

$$N_i = \frac{d_i - \hat{\theta}}{b_i}.$$

From the second condition we know that

$$\sum_{i=1}^n c_i \left\{ \frac{d_i - \hat{\theta}}{b_i} \right\} = C.$$

Thus

$$\hat{\theta} = \left\{ \sum_{i=1}^n \frac{c_i d_i}{b_i} - C \right\} / \left\{ \sum_{i=1}^n \frac{c_i}{b_i} \right\}.$$

Letting

$$\alpha = \sum_{i=1}^n \frac{c_i d_i}{b_i} \text{ and } \beta = \sum_{i=1}^n \frac{c_i}{b_i},$$

we can express $\hat{\theta}$ as

$$\hat{\theta} = \frac{\alpha - C}{\beta}$$

Thus

$$\theta_2 = e^{(\alpha - C)/\beta} \quad (5)$$

and

$$N_i = \frac{d_i - \frac{\alpha - C}{\beta}}{b_i} = \frac{g_i + C}{f_i}, \quad (6)$$

where $g_i = \beta d_i - \alpha$ and $f_i = \beta b_i$. Consequently N_i is a linear function of C . If the budget is incremented by an amount ΔC , then the new value of the stock level for item i , N'_i , satisfies

$$N'_i = N_i + \frac{\Delta C}{f_i}.$$

The optimal solution to Problem 2 has been found if each of the N_i found using Equation 6 is non-negative. If there exists an i for which $N_i < 0$, then we may employ the following algorithm to find the optimal solution to Problem 2. Let $I = \{1, \dots, n\}$ and N_i^1 represent the optimal solution to Problem 2.

Step 0. Solve Problem 3 as described above thereby obtaining an initial value for N_i , $i \in I$.

Step 1. Set $N_i^1 = 0$ for all $N_i < 0$ during the last iteration and delete the corresponding i from I . Recompute α and β , where

$$\alpha = \sum_{i \in I} \left\{ \frac{c_i d_i}{b_i} \right\}$$

and

$$\beta = \sum_{i \in I} \{c_i / b_i\}.$$

Step 2. Using Equation 6, obtain new estimates of N_i for each $i \in I$. If $N_i \geq 0$ for all $i \in I$, then the optimal solution has been found, and $N_i^1 = N_i$ for all $i \in I$ and $N_i^1 = 0$ for all $i=1, \dots, n$ for which $i \notin I$. If there exists some i for which $N_i < 0$, return to Step 1.

It is clear that our solution satisfies all the optimality conditions for Problem 2 except possibly condition (a) for $i \in I$. However, at an earlier iteration (when i was deleted from I) we had

$$\left. \frac{dB_i}{dN_i} \right|_{N_i < 0} + \bar{\theta}_2 c_i = 0,$$

where $\bar{\theta}_2$ is the earlier value of θ_2 . Since $\frac{dB_i}{dN_i}$ is clearly increasing in N_i , and θ_2 increases at each iteration (Theorem 1 and its corollary), condition (a) must hold. Convergence is guaranteed since n is finite.

4. A COMPARISON OF ALTERNATIVE SOLUTION PROCEDURES FOR SOLVING PROBLEM 1

In this section we briefly review three algorithms for solving Problem 1 and compare them to two algorithms designed to obtain a solution for Problem 1 based on the solution to the approximating problem, Problem 2.

The first algorithm we will discuss is the procedure originally proposed by Sherbrooke [3]. It is a marginal analysis algorithm consisting of two phases. In the first phase, each item is examined independently. The optimization problem solved for item i in the first phase has the form

$$\text{minimize} \quad \sum_{j=1}^m \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0}))$$

subject to

$$\sum_{j=0}^m s_{ij} = N_i,$$

$$s_{ij} = 0, 1, \dots, \quad (7)$$

to this knapsack problem. Clearly other procedures could be employed to obtain an optimal solution. In any case, this approach requires a substantial amount of storage to save all the $Z_i(N_i)$ values. For moderately sized problems--several thousand items--a storage requirement of 10^6 or more words may be needed to save these values. Furthermore, the computation time required to obtain these $Z_i(N_i)$ values for such problems is very large.

Subsequently Fox and Landi proposed a Lagrangian algorithm for solving Problem 1 [2]. In particular, they formulated the relaxed version of Problem 1 as

$$\min \sum_{j=1}^m \sum_{i=1}^n \sum_{x \geq s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})) + \theta \sum_{j=0}^m \sum_{i=1}^n c_i s_{ij} \quad (9)$$

$$s_{ij} = 0, 1, \dots,$$

where θ is the Lagrangian multiplier. Since Problem 9 is separable by item, its optimal solution can be found by solving the n individual item problems

$$\text{minimize } \sum_{j=1}^m \sum_{x \geq s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})) + \theta \sum_{j=0}^m c_i s_{ij}$$

subject to

$$s_{ij} = 0, 1, \dots$$

This problem, like Problem 8 in Sherbrooke's two-phase method, is solved using a partitioning procedure; that is, it is reformulated as

start (17) here

$$\min_{s_{i0} = 0, 1, \dots} \left\{ \theta c_i s_{i0} + \sum_{j=1}^m \min_{s_{ij} = 0, 1, \dots} \left\{ \sum_{x \geq s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})) + \theta c_i s_{ij} : s_{i0} \text{ fixed} \right\} \right\}, \quad (10)$$

or equivalently as

$$\begin{aligned} &\text{minimize } Z(s_{i0}; \theta) \\ &\quad s_{i0} \\ &\quad s_{i0} = 0, 1, \dots \end{aligned} \tag{11}$$

where

$$\begin{aligned} Z(s_{i0}; \theta) = & \theta c_i s_{i0} \\ & + \sum_{j=1}^m \min_{s_{ij}} \left\{ \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})) + \theta c_i s_{ij} : s_{ij} = 0, 1, \dots; s_{i0} \text{ fixed} \right\}. \end{aligned}$$

To determine $Z(s_{i0}; \theta)$, solve the m base problems

$$\text{minimize } \sum_{s_{ij}} \sum_{x > s_{ij}} (x - s_{ij}) p(x | \lambda_{ij} T_{ij}(s_{i0})) + \theta c_i s_{ij}.$$

The optimal s_{ij} is the smallest non-negative integer for which

$$\sum_{x > s_{ij}} p(x | \lambda_{ij} T_{ij}(s_{i0})) \leq \theta c_i.$$

Problem 10 (or Problem 11) is solved for each item for a given value of θ . This yields a total investment cost corresponding to θ . In the Fox-Landi approach, the "optimal" value of θ is selected from a grid of M equally spaced values

$$\theta_0 > \theta_1 > \dots > \theta_M > 0.$$

18 The optimal value of θ is the θ_K , $K \in \{0, \dots, M\}$, whose corresponding total investment cost is closest to C .

Fox and Landi suggest that their method is a single pass method, that is, only one pass through the item data base is necessary to obtain the optimal solution. The storage requirement to effect this one pass approach is potentially enormous. For a moderately sized problem having 3000 items, 20 bases and $M = 63$, almost 4 million item stock levels must be saved plus possibly millions of additional item data elements reflecting fill rates, probability of no stockout at an arbitrary point in time, expected base backorders, etc. Furthermore, there may be no simple method for estimating suitable bounds on the values of the multipliers thereby requiring much larger values of M to insure adequate approximation of the budget.

In the author's experience, the ability of Air Force personnel to estimate a reasonable range for θ for large problems is not good. It is not surprising that it is difficult for someone to estimate the optimal value of the multiplier. The data used in the model frequently change in real situations thereby causing the optimal value of the multiplier to change. Furthermore, changing the multiplier's magnitude by 10^{-6} or less often causes the corresponding total cost to change by many millions of dollars. Consequently, 2^{10} values of θ have been used in some Air Force applications to make the system "fool proof." In these cases 60 million or more item stock levels would be needed to be stored explicitly--plus a considerable amount of other item and base data--to make the Fox-Landi algorithm a truly one pass method.

On the other hand, if their method is altered so that the item data are examined a second time, it is possible to eliminate virtually all the requirements for secondary storage. In the first pass, only the running total cost corresponding to each θ_K , $K \in \{0, \dots, M\}$, is saved. At the end of this phase the "optimal" multiplier value, θ^* , is established. The second phase of the algorithm requires a second pass through the data base. In the second pass, the optimal stock levels for each location are found for all items by resolving Problem 10 with $\theta = \theta^*$.

In some applications the Fox-Landi one-pass method is clearly infeasible, that is, there may not be enough peripheral storage capacity to save all the data. If storage capacity is available, there is a tradeoff between the time and cost required to store and access the data in the secondary memory using the one-pass method and the time and cost to recompute the stock levels using the second method. For realistic Air Force problems, the two-pass method appears to be the only feasible approach given current hardware constraints if M is large enough to guarantee that a solution can be found that closely approximates the target budget.

A third way to solve Problem 1 is a slight modification of the Fox-Landi algorithm. This third method, called the bisection method, employs a bisection search to find the optimal value for θ . This procedure requires initial upper and lower bounds on the optimal value of θ . Call these θ_U and θ_L , respectively. The bisection method is as follows:

1. Set $\bar{\theta} = (\theta_U + \theta_L)/2$.
2. Solve Problem 10 with $\theta = \bar{\theta}$ for each item.
3. If the total cost of the solution obtained in Step 2 exceeds C , then replace θ_L with $\bar{\theta}$; otherwise, replace θ_U with $\bar{\theta}$.
4. If a stopping criteria has not been met (such as a fixed number of iterations or an error tolerance), return to Step 1; otherwise, stop.

The major drawback to the bisection approach is that a separate pass through the item data base is required at each iteration of the algorithm.

This algorithm performs very well in terms of convergence and in our experience virtually always produces solutions that are within 1/2 percent of the target budget using 10 bisections.

The closeness of the solutions to the target budget generated by either the Fox-Landi method or the bisection algorithm depends on how broad a range of multiplier values must be searched for a fixed value of M or a fixed number of

bisections. It should be pointed out that both of these methods only yield an approximation to the optimal multiplier value (assuming one exists).

Of the methods discussed thus far, it has been the experience of both the author and Fox and Landi [2] that the latter two algorithms dominate Sherbrooke's algorithm in run times by an order of magnitude or more on real problems given reasonable estimates of upper and lower bounds for the Lagrangian multiplier. Thus in the comparisons we will report, only these two Lagrangian methods will be discussed.

Earlier we described an approximation method for estimating the optimal values of θ and each N_i . Several options are open for implementing this approximation method. One way to implement it is to use a two-phase approach. Call this approach the First Approximation Method. The values of a_i and b_i are computed in the first phase of this method during which the optimal value of θ is also estimated using Equation 5. In the second phase, we solve Problem 10 for each item using the estimate of the optimal θ . This approach has two major advantages over the Fox-Landi method:

- (a) The estimate of the optimal multiplier can be obtained without prespecifying a range of values, and computation time to obtain the estimate does not depend on the uncertainty of the multiplier value.
- (b) The computation time to find an estimate of the optimal multiplier is much smaller.

If the two-pass version of the Fox-Landi algorithm is used, the second phase of that method and the second phase of the approximation method are identical. The one-pass version of the Fox-Landi algorithm requires considerably more storage, and also requires more computer time to determine the optimal stock levels than this approximation method requires.

The First Approximation Method also has the following advantages over the bisection method:

- (a) Only two passes through the data base are required as opposed to seven or more required for the bisection method in practice.
- (b) No stock levels need to be saved; in the bisection method it is necessary to save all stock levels and other data for three multiplier values.

Another algorithm can be employed that directly uses the results of the approximation problem, that is, Problem 2. Call this approach the Second Approximation Method. This algorithm is of interest in situations in which we only want to compute total system stock for each item and are not particularly interested in computing the optimal distribution of the assets. Determining the optimal allocation of a budget among items is of primary importance when purchasing inventory or making budgetary projections for spares for different systems. In these cases, distribution decisions are usually not that critical.

This Second Approximation algorithm also consists of two phases; in the first phase we estimate the values of the a_i and b_i parameters, and in the second phase we determine total system stock for each item using the algorithm described in Section 3 and rounding N_i to the nearest integer. The algorithm requires one pass through the item data base and one pass through an item file consisting of a_i, b_i , and c_i . The major advantage of this approach is that it eliminates the stock allocation phase of both the Fox-Landi algorithm and the First Approximation Method.

5. A COMPUTATIONAL COMPARISON OF VARIOUS ALGORITHMS

The Fox-Landi algorithm, bisection algorithm, and the two approximation methods have been coded and tested on several sample sets of data for the

Air Force's new F-15 fighter. Since all of the tests yielded the same general results, we will discuss only two of them. The first test consisted of a 75 item sample and had 3 operating bases. The flying programs were very different at each base. In the second test, 125 items were included in the sample with demands occurring at 5 bases. In the second test, only the Fox-Landi and the two approximation methods were compared. In all Fox-Landi calculations, a maximum of 128 multiplier values were examined; ten bisections were used in all applications of the bisection method. The run times stated for both approximation algorithms include the time required to estimate the values of a_i and b_i . Furthermore, in both test cases all stock levels for all relevant multiplier values were stored in main memory. Thus the reported computation times, which include compile times which are roughly equal for all the algorithms, are biased in favor of the Fox-Landi method since for larger problems this type of storage would be impossible. Additionally, the range of multiplier values considered in the test of the Fox-Landi and bisection methods was selected after estimating the optimal multiplier value using the First Approximation Method. Thus the test results are biased in favor of them, since the range of multiplier values was much smaller than would normally be the case.

The data displayed in Tables I and II indicate how well each approach approximates a given target budget for the two test data sets. Without a doubt the bisection method produced solutions that best matched the target budgets followed in order by the Second Approximation Method, the Fox-Landi method, and the First Approximation Method. As mentioned before, the results are biased in favor of both the Fox-Landi and bisection methods due to the initialization of the range of multiplier values. From a practical viewpoint,

all approaches worked acceptably well in meeting the target budgets. Furthermore, the stock levels generated by the various approaches were virtually the same for similar budgets. Consequently, total system expected backorders, for all practical purposes, are indistinguishable; that is, the backorder versus investment curves virtually coincide among these various approaches. Exact comparison of computed stock levels and expected backorders cannot be made among the competing methods since the allocation of the available budget in each case depends on the way each algorithm estimates the Lagrangian multiplier.

mining
The area in which the methods clearly differ is in computation time. The approximation methods require substantially less time than either the Fox-Landi method or the time consuming bisection method. Other experimentation has shown that the percentage difference in computation times tends to be even more substantial as the number of items considered increases.

p. 24
Thus the approximation methods produce answers that are as good as those produced by either the Fox-Landi method or the bisection method, but with less computational effort. The bisection method did match target budgets slightly better than the approximation methods. However, the approximation algorithms are virtually fool-proof. This is perhaps the greatest advantage of the approximation algorithms. The user does not have to specify the range of multiplier values or the number of bisections in advance. This eliminates one of the main difficulties associated with implementing either the Fox-Landi or bisection algorithms. In view of these observations, the approximation procedures developed here appear to be superior for use on real problems.

Table 1

75-ITEM, 3-BASE TEST CASE

Target Budget	Total Cost (millions of \$)			
	Bisection	Fox-Landi	Approx. I	Approx. II
3.68	3.67	3.68	3.63	3.63
3.97	3.99	3.92	3.82	4.03
4.27	4.27	4.27	4.30	4.18
4.57	4.57	4.57	4.62	4.61
4.87	4.87	4.85	4.87	4.78
5.16	5.16	5.18	5.09	5.17
5.46	5.46	5.42	5.38	5.49
5.76	5.76	5.76	5.75	5.79
6.05	6.06	6.05	6.06	6.08
6.35	6.34	6.38	6.28	6.33
6.65	6.65	6.63	6.63	6.73
6.94	6.89	6.80	6.87	6.92
7.24	7.24	7.19	7.27	7.24
7.54	7.54	7.57	7.68	7.51
7.83	7.84	7.77	7.80	7.83
8.13	8.14	8.24	8.20	8.05
8.43	8.42	8.50	8.42	8.42
8.73	8.73	8.50	8.74	8.77
9.02	9.02	9.04	9.11	9.00
Execution Time (Seconds)	92.57	19.57	11.59	4.57

Table 2

125-ITEM, 5-BASE TEST CASE

Target Budget	Total Cost (millions of \$)		
	Fox-Landi	Approx. I	Approx. II
26.4	26.7	24.8	26.6
27.6	27.6	26.2	27.9
28.7	28.7	27.6	28.9
29.8	30.0	29.5	29.8
31.0	31.2	30.7	30.8
32.1	32.1	32.0	32.2
33.2	33.3	33.1	33.1
34.4	34.4	34.3	34.2
35.4	35.5	35.9	35.7
36.6	36.8	37.0	36.7
37.8	38.0	38.1	37.7
38.9	38.6	39.3	39.2
40.0	39.9	40.6	40.0
41.2	41.1	42.1	41.3
42.3	42.5	43.9	42.4
43.4	43.3	44.7	43.7
44.6	44.5	45.6	44.2
45.7	46.3	46.1	45.9
46.8	47.2	47.3	46.7
Execution Times (Seconds)	36.98	16.28	4.74

NOTE: All programs were run on an IBM 370/168
using the WATFIV compiler.

REFERENCES

1. Feeney, George J. and Sherbrooke, Craig C., "The $(s-1,s)$ Inventory Policy Under Compound Poisson Demand," Management Science, Vol. 12, No. 5, January 1966, pp. 391-411.
2. Fox, Bennett and Landi, M., "Searching for the Multiplier in One-Constraint Optimization Problems," Operations Research, Vol. 18, 1970, pp. 253-262.
3. Sherbrooke, Craig C., "METRIC: A Multi-Echelon Technique for Recoverable Item Control," Operations Research, Vol. 16, 1968, pp. 122-141.