Onderzoeksrapport Nr. 8215

PROTECTIVE STOCKS IN MULTI-STAGE
PRODUCTION SYSTEMS

by

Marc R. LAMBRECHT^{*}
John A. MUCKSTADT[°]
Robert LUYTEN^{*}

Wettelijk Depot : D/1982/2376/24

^{*} Katholieke Universiteit Leuven, Dep. of Applied Economics

[°] School of Operations Research and Industrial Engineering, Cornell University.

This report was published as Technical Report No. 562, School of Operations Research and Industrial Engineering, Cornell University.

This research was supported in part by the National Science Foundation under grant MSF-81-21429 and by the Onderzoeksfonds K.U. Leuven under grant OT/IX/7.

THE UNIVERSITY OF CHICAGO

 

BY

ROBERT L. HICKS
AND
 

RESEARCH REPORT - 1962

I. INTRODUCTION

Material Requirements Planning (MRP) systems have been developed during the past two decades to provide a mechanism for coordinating the replenishment of components, subassemblies, and final products in multi-stage production and distribution systems. The MRP systems used in practice are based on the assumption that all factors affecting production are known with certainty over a product's total manufacturing lead time. However, uncertainty does exist in several forms--variability in demand from period to period, uncertainty in the supply from stage to stage due to the variability in the yields from each production batch, and uncertainty in the lead times for both production and procurement. Protection against the disruptions in production caused by these sources of uncertainty is required, however, ~~and~~ often results in buffer stocks. These stocks can exist as safety stock or as inventory resulting from the existence of safety time. Safety stock is inventory that, on the average, always exists. It is the expected amount of on-hand inventory at a stage just before production of a lot is available or a procurement order arrives at that stage. Roughly speaking, safety time occurs when a production batch or procurement lot becomes available before it is needed to produce some subsequent assembly or finished product.

Safety stock is needed to protect against uncertainties in the demand for the final product over the manufacturing lead time. It is also required for components having uncertain yields; in this case, safety stock preserves lot size integrity and keeps the number of setups from growing, particularly at bottleneck work centers. Finally, safety stock may be desirable for components with multiple uses. In case the

components and assemblies are produced at regular intervals (cyclical production schedules), safety stock may be desirable so that the manufacturing lead times for the assemblies using the components are maintained at or near their nominal values.

Safety time, on the other hand, usually exists when production and/or procurement lead times are uncertain; but, it may exist to protect against demand variability, as we will subsequently observe.

In a practical situation, however, the origin of protective stocks may be difficult to determine. We will examine some relatively simple problems in this paper for which we can measure safety stock and safety time. As we will see, it is impossible to determine their values by simply examining the form of the policy being followed.

Academic researchers have ~~recognized~~ the need to create buffer stocks of the types we have mentioned while many practitioners do not admit that such stocks are necessary at intermediate stages of production. Research to date has produced no results that provide a fundamental understanding of the nature of protective stocks.

In fact, the behavior of MRP derived policies is largely not understood much less the form and characteristics of optimal policies. Research on protective stocks in multi-stage systems has primarily been experimental. Simulation models have been used to evaluate the effects of various policies on service levels and inventory investment in very special environments. The work to BANERJEE [2] and WHYBARK and WILLIAMS [13] are illustrations of this approach. MRP folklore normally suggests protecting against uncertainty by creating safety stock of final products. MILLER [10] suggests an alternative to this policy. He

recommends the utilization of a "pipeline" safety stock strategy that is based on the recognition of the dependent demand relationships between final products and intermediate stages and the value added cost structure. His hedging strategy is designed to keep safety stocks in the supply pipeline rather than in the stockroom in finished product form. CARLSON and YANO [3] [4] also developed a methodology to tackle the safety stock problem. Their approach is primarily experimental. They conclude for a special situation that finished product safety stocks are in general more appropriate than intermediate level safety stocks. NEW [11] critically reviews the existing uncertainty protection mechanisms in available MRP systems. MEAL [9] developed a statistical method for determining the proper time to produce at each stage.

As we have said, it is our opinion that these research efforts have provided limited insight into when, why, or how much safety stock and safety time should exist in multi-stage production systems. Furthermore, this research has produced no computationally feasible approaches for computing operating policies for even slightly complex systems. In this paper it is our goal to address these issues for situations in which demand variability is the sole source of uncertainty. To achieve this goal, we have organized the remainder of the paper as follows.

In Section II, we develop a mathematical model--a dynamic programming model--that can be used to find an optimal operating policy for multi-stage systems when demand is the only source of uncertainty, the planning horizon is finite and a periodic review policy is followed. We next show how this problem can be viewed as a Markovian Decision Problem (MDP). In Section III we show that the exact methods developed in

Section II are of limited practical value and review the well known results of CLARK and SCARF [5,6] for analyzing serial systems. We indicate how their method can be modified to obtain solutions for assembly type systems.

In Section IV we present results of an experiment. The form and nature of the operating policies are discussed. In particular, we show that generalized (s,S) policies appear to be optimal. The experiment shows that the modified Clark-Scarf approach, which also yields policies of the (s,S) type, provides excellent results. The computational results indicate that both safety time and stock can exist at intermediate stages. Examples are presented illustrating the presence of safety time and stock.

In Section V we analyze the form ~~of~~ the policies obtained when following MRP logic. Specifically, we observe that the policy followed using MRP logic is an (s,S) policy. As such, both safety stock and time can exist when the demand for the final products is uncertain. This occurs, as we show with an example, even though their existence is not anticipated when establishing the production schedule. We conclude that the (s,S) policies generated using MRP logic will perform well, in an expected cost sense, when lead times and demand variability per period are relatively low.

II. MODEL FORMULATION

In this section we present two modelling approaches. First, we show how to formulate the multi-stage production planning problem as a dynamic programming problem when the time horizon is finite. Second, we show that the problem can be modelled as a Markovian Decision Problem. Attention in this case is limited to the situation where the demand process is stationary and the planning horizon is infinite; however, the same observation can be made for a finite horizon as well. The models we present are based on the following assumptions:

1. The product structures are known;
2. A periodic review policy is followed;
3. The production and procurement lead times, which are measured in periods, are known constants;
4. The demand distributions are known; demand is independent from period-to-period for each final product;
5. All excess demand for final products is backordered. We also assume that the relevant costs consist of production costs, at each stage, which are composed of a fixed setup cost per production run and a constant cost per unit produced, holding costs, which are charged proportional to the echelon stock at the end of a period, and linear shortage costs, which are charged proportional to the number of units of final product backordered at the end of each period.

II.A. A Dynamic Programming Formulation

The multi-stage production planning problem we have discussed can be formulated as a dynamic programming problem. The formulation we give is appropriate when the planning horizon is finite and when each item is either a final product or component, but not both. The formulation can be easily modified to accommodate situations in which items can be both final products, say spare parts, as well as a component for a higher level assembly. The model also allows production capacity to be taken into account. As we will show, the formulation is quite complicated and hence, except for relatively simple problems is computationally intractable. This is due to the size of the state space. For this problem, the state of the system reflects the number of units on-hand for each item and the number of units that will become available for use in each time period as a consequence of previous production decisions. Thus, if lead times are long and/or the number of values that can be assumed by the state variables is large, the state space set can consist of an enormous number of elements.

We will use the following nomenclature in the dynamic programming model. Let

N = the number of finished goods in the system,

$A(i)$ = the set of items that are immediate successors of item i ,

Z_{it} = the production level for item i in period t ,

τ_i = the lead time for item i ,

Y_{it} = the on-hand inventory of item i at the end of period t ,

$C_{it}(Z_{it})$ = the production cost function for item i in period t ,

h_i = the echelon holding cost for item i ,

p_i = the shortage cost for item i ($p_i = 0$ if item i is not a final product),

$S(i)$ = the set of all successor items (both immediate and indirect) of item i ,

E = the set of final items,

$F_{i,t}(\cdot)$ = marginal demand distribution function for final product $i \in E$ in period t , and

$F_t(\cdot, \cdot, \dots, \cdot)$ = the joint demand distribution for all final products $i \in E$ in period t .

Let $f_n(\cdot)$ represent the minimum expected cost of operating the system assuming n periods remain in the planning horizon. Time periods are numbered to reflect the number of periods remaining in the planning horizon, as illustrated as follows.

To reduce the difficulty of writing the argument for the function f , we simply use $(Y_{i,n}, Z_{i,n}, \dots, Z_{i,n+\tau_i-1})$ to represent the state of the system for all items at the beginning of period n following the production decision for that period. Then

$$f_n(Y_{i,n+1}, Z_{i,n+1}, \dots, Z_{i,n+\tau_i}) =$$

$$\min_{Z_{in} \geq 0} \left\{ \sum_i C_{in}(Z_{in}) \right.$$

$$Z_{in} \in P \quad (\text{production costs})$$

$$+ \left\{ \sum_{i \notin E} \left(h_i (Y_{i,n+1} + Z_{i,n+\tau_i} - \sum_{j \in A(i)} Z_{jn}) \right) \right.$$

(holding costs for units of component i on-hand)

$$+ \sum_{j \in S(i) - E \cap S(i)} h_i (Y_{j,n+1} + \sum_{\ell=0}^{\tau_j} Z_{j,n+\ell} - \sum_{k \in A(j)} Z_{kn})$$

(holding costs for units of i in units of successor items that are not finished goods)

$$+ \sum_{j \in S(i) \cap E} \left\{ h_i \left(\max(0, Y_{j,n+1} + Z_{j,n+\tau_j}) \int_0^{\infty} (Y_{j,n+1} + Z_{j,n+\tau_j} - u) dF_{jn}(u) + \sum_{\ell=0}^{\tau_j-1} Z_{j,n+\ell} \right) \right\}$$

(holding costs for units of i in finished items in production or on-hand)

$$+ \sum_{i \in E} \left\{ p_i \int_0^{\infty} \max(0, Y_{i,n+1} + Z_{i,n+\tau_i}) (u - (Y_{i,n+1} + Z_{i,n+\tau_i})) dF_{in}(u) + \right. \\ \left. h_i \int_0^{\max(0, Y_{i,n+1} + Z_{i,n+\tau_i})} (Y_{i,n+1} + Z_{i,n+\tau_i} - u) dF_{in}(u) \right\}$$

(penalty costs for shortage + holding costs for finished items)

$$+ \int \dots \int f_{n-1} \left((Y_{i,n+1} + Z_{i,n+\tau_i} - \sum_{j \in A(i)} Z_{jn}), Z_{in}, \dots, Z_{i,n+\tau_i-1} \right) dF_n(u_1, \dots, u_N) \\ \text{or} \\ u_i \text{ if } i \in E$$

where P is the set of constraints reflecting lead time and production restrictions for each item in each period.

The problem is solved beginning at the end of the horizon and working backwards until the beginning of the horizon is reached. The dynamic programming formulation is based on the echelon-stock concept. The echelon stock for facility i equals the installation stock at that facility plus the on-hand stock position at all succeeding facilities. As a consequence, the echelon holding costs reflect the value added structure. The characterization of optimal solutions is given in terms of echelon stock.

II.B. Markovian Decision Processes (MDP)

When the demand process is stationary, the costs are time independent and the planning horizon is infinite, the dynamic programming problem introduced in Section II.A can be viewed and modeled as a Markovian Decision Process problem, as we will now see. The system is observed at time $t = 0, 1, 2, \dots$ and classified into one of a finite number of states, where, as before, the states represent on-hand plus on-order inventory for each stage. For every state i a number of decision alternatives, k , is possible, labeled $1, 2, \dots, K(i)$. Whenever the system is in state i and decision k is made, the system moves to a new state j , with known transition probability p_{ij}^k , which is determined based on the demand distribution. Let q_i^k represent the expected cost incurred when the system is in state i and decision k is made.

The transition probabilities are time invariant since the demand distribution is stationary. Thus the process of moving from

state-to-state is a Markov process. Furthermore, since the expected costs, q_i^k , and the transition probabilities, p_{ij}^k , depend on the decision that is made, the model has the form of a Markovian Decision Process. Consequently, the solution can be obtained using the successive approximation procedure developed by HOWARD [7]. This is clearly preferable to using a standard dynamic programming approach for solving the problem. However, due to the size of the state space, even this method is of little practical value.

II.C. Comments Concerning the Preceding Formulations

In the preceding subsections we showed that the problem of determining the optimal production plan for each stage can be formulated as a dynamic program, and when the demand process and costs are time independent, as a MDP. We repeat that the state space in either formulation grows rapidly as the lead time and number of demand possibilities in each period increase. Thus these models are of little practical use; however, they are of considerable value for testing the quality of the solutions obtained using heuristics, which is how we shall use them.

III. A HEURISTIC: A MODIFIED CLARK AND SCARF APPROXIMATION PROCEDURE

Since the exact methods for finding an optimal solution to the multi-stage production planning problem presented in Section II are computationally intractable for most realistic problems, it is necessary to develop an alternative method for finding good, if not optimal, solutions. In this section we suggest such a procedure. However, we limit our attention to assembly or serial type systems, such as displayed in Figure 1. By an assembly system, we mean a product structure in which each item has at most one successor item, but any number of predecessors. We will use the numbering of stages as given in Figure 1; namely, stage 1 is the final stage and stage M is the beginning stage.

FIGURE 1a: A Three Stage Assembly System

FIGURE 1b: An M Stage Serial System

The procedure we suggest is a modification of one developed by CLARK and SCARF [6]. We adapt their approach for solving serial system problems to the more general assembly system problems. Since our heuristic is based on the Clark and Scarf approximation, we begin this section by reviewing it.

III.A. The Clark and Scarf Approximation

The methods described in Section II are computationally intractable because the state space becomes too large to solve problems of a practical size. The Clark and Scarf approximation circumvents this problem. Rather than solving the problem for all stages simultaneously, Clark and Scarf suggest solving a series of simpler problems, one for each stage in the serial system. The replenishment policy is computed separately for each stage, starting with stage 1 and ending with stage M . Given the cost structure, which includes a fixed setup cost, a variable production cost, a holding and a shortage cost, the optimal policy at each stage is of the (s, S) type [1]. The optimal policy is always stated in terms of echelon stock rather than installation stock.

The shortage costs used when finding the myopic policy for stages 2 through M are not specified explicitly in advance. Rather, they measure the additional costs at subsequent stages that are incurred due to the policy followed at a given stage. For example, the imputed cost at stage 2 is determined using the following logic.

Suppose the optimal policy has been calculated for stage 1. The solution states that whenever the echelon 1 stock is less than s_1 , production is started for an amount sufficient to bring the echelon 1

stock up to S_1 , where s_j and S_j are the policy variables for echelon j . However, the echelon stock for stage 2 may not always be of sufficient size to raise stage 1's echelon stock to S_1 . When it isn't, an additional cost is incurred at stage 1, since stage 1 is no longer following its optimal policy. The expectation of these resulting incremental costs for stage 1 is the imputed shortage or penalty cost function for stage 2. The penalty functions for the remaining stages can be calculated in a similar fashion. We subsequently present the mathematical details that implement these notions.

The Clark and Scarf concept for decomposing the complex production problem does not always find the optimal policy. When the fixed setup cost is zero for stages 1 through $M-1$, they show that the solution is optimal. However, when the fixed setup cost is positive for any of the stages 1 through $M-1$, then the computed policy may not be optimal [5,6]. In the latter case, they suggest an approximation method that yields upper and lower bound estimates on total cost.

We present the details for the Clark and Scarf approximation procedure for a two stage serial system. We show how calculations are made for stage 1 followed by those for stage 2. The procedure can easily be extended to any number of stages.

a. Computation for Stage 1

The following steps are used to determine the locally optimal (s, S) policy for stage 1. As we did earlier, time periods are numbered to reflect the number of periods remaining in the horizon. The algorithm for this stage consists of three steps, which we now present.

Step 1: Calculate $g_{1,n}(y)$ for all y and let $S_{1,n}$ be the value of y that solves the problem

$$\min_y [g_{1,n}(y)] = \min_y [c_{1,n}(y) + L_{1,n}(y) + \sum_{q=0}^{\infty} f_{1,n-1}(y-q)p_n(q)], \text{ where}$$

$g_{1,n}(y)$ = the minimum expected cost for stage 1 for period n through the end of the horizon (period 1),

$c_{1,n}(y)$ = the cost of producing y units for arrival at time $n-\tau_1$,

$p_n(q)$ = the probability distribution (discrete distribution) of demand for period n , and

$$L_{1,n}(y) = h_1 \sum_{Q=0}^y (y-Q)p_{n,n+\tau_1}(Q) + p_1 \sum_{Q=y+1}^{\infty} (Q-y)p_{n,n+\tau_1}(Q),$$

where Q is the demand for periods n through $n-\tau_1$ ($Q = q_n + \dots + q_{n-\tau_1}$).

Step 2: Let $s_{1,n}$ be the smallest value for which

$$g_{1,n}(s_{1,n}) \leq g_{1,n}(S_{1,n}) + K_1,$$

where K_1 is the setup cost for stage 1.

Step 3: Calculate $f_{1,n}(x_{1,n})$ for all $x_{1,n}$, where $x_{1,n}$ is the inventory on-hand and on-order at the beginning of period n .

$$f_{1,n}(x_{1,n}) = \begin{cases} L_{1,n}(x_{1,n}) + \sum_q f_{1,n-1}(x_{1,n}-q)p_n(q), & x_{1,n} \geq s_{1,n} \\ K_1 + c_{1,n}(S_{1,n}-x_{1,n}) + L_{1,n}(S_{1,n}) + \sum_q f_{1,n-1}(S_{1,n}-q)p_n(q), & x_{1,n} < s_{1,n} \end{cases}$$

$$f_{1,0}(\cdot) = 0,$$

The procedure obviously starts at $n = 1$ and proceeds backwards. Given the fixed charge cost structure, there is an optimal policy of the $(s_{1,n}, S_{1,n})$ type; furthermore, if $K_1 = 0$, $s_{1,n} = S_{1,n}$.

b. Computation for Stage 2

In this subsection we show how to compute stage 2's locally optimal (s, S) policy. To minimize notational difficulties, we assume a zero lead time at stage 2. The three step procedure is as follows.

Step 1: Calculate $\bar{g}_{2,n}(y)$ for all y and let $S_{2,n}$ be the value of y that solves the problem

$$\min_y [\bar{g}_{2,n}(y)] = \min_y [c_{2,n}(y) + \bar{L}_{2,n}(y) + \sum_{q=0}^{\infty} f_{2,n-1}(y-q)p_n(q)].$$

The expected holding and shortage cost function $\bar{L}_{2,n}(y)$ includes a penalty function $\Lambda(y)$. Rather than using the exact function $\Lambda(y)$, Clark and Scarf propose bounding it. The upper bound function $\bar{\Lambda}_n(y)$ is obtained by charging an additional shortage cost whenever the echelon stock $(x_{2,n})$ for stage 2, is less than $S_{1,n}$ (even in cases where the stage 1 order point is not reached). The lower bound, $\underline{\Lambda}_n(y)$ is defined to be different from zero only when $x_{2,n} \leq s_{1,n}$.

We have chosen to use $\bar{\Lambda}_n(y)$ in our calculations of $\bar{L}_{2,n}(y)$. $\bar{\Lambda}_n(y)$ is obtained as follows. When the on-hand and on-order inventory at stage 1, $x_{1,n}$, is less than $s_{1,n}$, an order is placed of size $(S_{1,n} - x_{1,n})$. Then the expected minimum cost would be

$$K_1 + c_{1,n}(S_{1,n} - x_{1,n}) + L_{1,n}(S_{1,n}) + \sum_q f_{1,n-1}(S_{1,n} - q)p_n(q). \quad (1)$$

But, if $x_{2,n} < S_{1,n}$ we will not be able to order up to $S_{1,n}$, in which case we are restricted to order up to $x_{2,n}$. Then the stage 1 expected cost is

$$K_1 + c_{1,n}(y - x_{1,n}) + L_{1,n}(y) + \sum_q f_{1,n-1}(y-q)p_n(q), \text{ where we let } y = x_{2,n}. \quad (2)$$

$\bar{\Lambda}_n(y)$ is defined as the difference between (2) and (1) when $y < S_{1,n}$ and 0 otherwise; that is,

$$\bar{\Lambda}_n(y) = \begin{cases} 0 & , y \geq S_{1,n}, \\ c_{1,n}(y - S_{1,n}) + L_{1,n}(y) - L_{1,n}(S_{1,n}) + \sum_q [f_{1,n-1}(y-q) - f_{1,n-1}(S_{1,n}-q)]p_n(q), & y < S_{1,n}, \end{cases} \quad (3)$$

and, therefore,

$$\bar{L}_{2,n}(y) = h_2 \sum_{Q=0}^y (y-Q)p_{n,n+\tau_1}(Q) + \bar{\Lambda}_n(y).$$

The final two steps of the procedure are as follows.

Step 2: Let $s_{2,n}$ be the smallest value for which

$$\bar{g}_{2,n}(s_{2,n}) \leq \bar{g}_{2,n}(S_{2,n}) + K_2.$$

Step 3: Calculate $f_{2,n}(x_{2,n})$, for all $x_{2,n}$, where $x_{2,n}$ is the echelon stock for stage 2, and

$$f_{2,n}(x_{2,n}) = \begin{cases} \bar{L}_{2,n}(x_{2,n}) + \sum_q f_{2,n-1}(x_{2,n}-q)p_n(q), & x_{2,n} \geq s_{2,n} \\ K_2 + c_{2,n}(s_{2,n}-x_{2,n}) + \bar{L}_{2,n}(s_{2,n}) + \\ \sum_q f_{2,n-1}(s_{2,n}-q)p_n(q), & x_{2,n} < s_{2,n}. \end{cases}$$

c. Observations Concerning the Clark and Scarf Approximation Procedure

There are several important observations we have made concerning the behavior of the Clark and Scarf approximation procedure.

As we have seen, the policy values are calculated independently from stage-to-stage. This is the method's ~~major~~ computational advantage over the algorithms we presented in Section II. Unfortunately, the policy variables must sometimes be adjusted to reflect the interactions among the stages. For example, in a two stage system, the policy at stage 2 is determined independently of that found for stage 1; however, the replenishment policy at stage 1 depends on the echelon stock for stage 2. In other words, $(s_{1,n}, S_{1,n})$ depends on $x_{2,n}$, as we now illustrate.

Suppose in period n that the control parameters for stage 1 are given by $s_{1,n} = 3$ and $S_{1,n} = 5$ and the graph of $g_{1,n}(y)$ is as shown in Figure 2. Suppose $x_{2,n} \geq S_{1,n}$. Then the control parameters shown in Figure 2 remain unchanged ($s_{1,n} = 3$ and $S_{1,n} = 5$). If, however, $x_{2,n} < S_{1,n}$, we recompute the control parameters in the following way. Set the new order-up-to level, say $S'_{1,n}$, equal to $x_{2,n}$ and determine the

FIGURE 2. Graph of $g_{1,n}(y)$

corresponding $s'_{1,n}$ to be the smallest value for which

$$g_{1,n}(s'_{1,n}) \leq g_{1,n}(S'_{1,n}) + K_1, \text{ where}$$

$$S'_{1,n} = x_{2,n}.$$

In Figure 2, we see that $S'_{1,n} = 2$ and $s'_{1,n} = 2$, if $x_{2,n} = 2$. In general, we can conclude that the control parameters at stage 1 are functions of $x_{2,n}$. We call policies in which such modifications are made generalized (s, S) policies. Thus the solutions generated by the Clark and Scarf procedure may need to be modified in certain cases if an optimal policy is to be followed.

Our computational experiments have also identified a major weakness in the Clark and Scarf algorithm when the cost structure is such that the setup-to-holding cost ratios for successive stages satisfy $\frac{K_j}{h_j} < \frac{K_{s(j)}}{h_{s(j)}}$, where $s(j)$ represents the successor stage to stage j . We illustrate the potential problems using the following examples.

Suppose demand in each period is for either 0, 1, or 2 units of the final product with probabilities of .25, .50, and .25, respectively. Furthermore, assume the production lead time for stage 1 is one period and for stage 2 is 0 periods. We study three examples, whose cost data are given in Table I.

Example	Item	K_j	h_j	p_j
1	1 2	50 1	1 10	50/100/200
2	1 2	4 40	4 5	100/200/400
3	1 2 3	50 1 10	1 10 5	50/100/200

TABLE I. Data for Example Problems

Examples 1 and 2 are each two stage serial systems whereas Example 3 is a two stage assembly system whose structure is shown in Figure 3.

FIGURE 3. Assembly Structure for Example 3.

Let's examine the setup to holding cost structure for these three examples. For Example 1,

$$\frac{K_2}{h_2} = \frac{1}{10} < \frac{50}{1} = \frac{K_1}{h_1};$$

for Example 2,

$$\frac{K_2}{h_2} = \frac{40}{5} > \frac{4}{4} = \frac{K_1}{h_1};$$

and for Example 3,

$$\left. \begin{array}{l} \frac{K_2}{h_2} = \frac{1}{10} \\ \frac{K_3}{h_3} = \frac{10}{5} \end{array} \right\} < \frac{50}{1} = \frac{K_1}{h_1}.$$

In a deterministic situation, for cost structures that correspond to Examples 1 and 3, the lot sizes for all items should be identical, as shown by Maxwell

and Muckstadt [8] and others. Thus the production for these items should be planned to occur simultaneously when demand occurs continuously and at a constant rate. For Example 2, the lot size for stage 2 will be greater than that for stage 1, since $\frac{K_2}{h_2} > \frac{K_1}{h_1}$. These observations are essentially the result of using the economic lot size formula to calculate production quantities. This observation should be a guide to identifying reasonable policies for probabilistic demand problems. Namely, in serial or assembly type systems we should generally have $S_j - s_j \geq S_{s(j)} - s_{s(j)}$. Otherwise, several production lots may be produced for item j per lot of $s(j)$. In cases where the production lot for item j exceeds $S_j - s_j$, it is possible to have $S_{s(j)} - s_{s(j)} \leq S_j - s_j$. However, the actual production lot sizes will normally correspond to the relationship suggested by the solution to the related deterministic problem.

The solutions found using the MDP and the Clark and Scarf methods for these three examples are given in Tables II-IV. We see that the Clark and Scarf method obtains the wrong solution for Examples 1 and 3. Furthermore, the annual costs obtained when following the Clark and Scarf approach are substantially greater than the optimal cost for Examples 1 and 3. This occurs because the policies found using the Clark and Scarf method are not coordinated when they are determined. This occurred for all other examples we evaluated where the natural lot size (EOQ lot size found using echelon holding costs) for item 1 is greater than that for item 2. We also found that the Clark and Scarf solutions are normally optimal or extremely close to being optimal (well within 1% of the optimal average annual cost) when the natural lot sizes are in the right order, as they are in Example 2. Observe that the optimal value of $S_1 - s_1$ is smaller than

Item	Shortage Costs					
	50		100		200	
	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N
2	(3,3)	(2,5)	(4,4)	(3,5)	(4,4)	(3,5)
1	(3,12)	(2,5)	(3,12)	(3,5)	(4,13)	(3,5)

TABLE II. Policies for Example 1

Item	Shortage Costs					
	100		200		400	
	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N
2	(3,6)	(3,6)	(3,6)	(3,6)	(4,7)	(4,7)
1	(3,4)	(3,4)	(3,4)	(3,4)	(4,4)	(4,4)

TABLE III. Policies for Example 2

Item	Shortage Costs					
	50		100		200	
	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N	CS SOL'N	OPT SOL'N
3	(4,9)	(2,4)	(5,9)	(2,5)	(6,9)	(3,5)
2	(3,3)	(2,4)	(4,4)	(2,5)	(4,4)	(3,5)
1	(3,12)	(2,4)	(4,13)	(2,5)	(4,13)	(3,5)

TABLE IV. Policies for Example 3

the Clark and Scarf method for Examples 1 and 3 while $S_2 - s_2$ is greater. The reason that the optimal policies are identical for each stage is that the lead times are 0 for stage 2 items. This need not be the case when $\tau_2 > 0$.

Our observations concerning the performance of the Clark and Scarf approximation procedure suggests that some modifications should be made to improve its performance.

III.B. A Modified Clark and Scarf Algorithm for Assembly Systems

Modifications to the Clark and Scarf algorithm that rectify some of the problems we observed are presented in this section. We first discuss how the algorithm can be extended to obtain solutions to problems having assembly type product structures.

That the Clark and Scarf approximation algorithm can be easily extended to assembly type systems should not be surprising. Let us illustrate how this can be done for a two stage assembly system, such as the one for Example 3 shown in Figure 3. First, obtain a locally optimal solution for stage 1, using the method discussed in Section III.A.a. Once this solution is obtained, calculate the penalty function and separately apply the method described in Section III.A.b to components 2 and 3.

Obviously this approach can be extended to any assembly system. We note that this extension possesses the shortcomings we discussed for the Clark and Scarf algorithm plus an additional one. We pointed out that by separately calculating policies for each item we can obtain policies, in certain cases, that are not coordinated. This can obviously happen for the extension we have described. However, it can occur in an additional way as well.

Again consider Example 3. If the control policies calculated for components two and three are not coordinated, then it may be the case that

there are a large number of units on hand for component 2 but none for component 3 thereby making it impossible to produce any units of item 1. Thus the implementation of the (s,S) policies at stage 2 must be modified to account for what we call horizontal interaction. The following example illustrates this need for coordination. The data for the example are given in Table V. The demand distribution is the same as for Examples 1-3. The optimal solution is given in Table VI. As can be seen from Table VI, the replenishment control parameters for item 1 depend on the minimum of x_2 and x_3 , the echelon stock for components 2 and 3, respectively. This obviously results in the desired coordination.

Item	K_i	h_i	p_i
1	10	2.7	500
2	25	1.26	
3	25	2.52	

TABLE V: Data for Example 4

Component 3	Component 2	Assembly 1				
(4,7)	(4,9)	$\text{Min}(x_2, x_3)$				
		4 5 6 7				
		$s_1(x_2, x_3)$				
		4 4 4 4				
		$S_1(x_2, x_3)$				
		4 5 6 5				

TABLE VI: Optimal Solution for Example 4.

The modification we propose to the Clark and Scarf algorithm pertains to the coordination problem we observed in serial structures when $\frac{K_j}{h_j} < \frac{K_{s(j)}}{h_{s(j)}}$. The modification will be explained for a two stage serial system. Extensions to other serial or assembly systems can be accomplished in a straightforward manner. We call this heuristic, the modified Clark and Scarf (MCS) procedure.

MCS Procedure

Step 1: If $\frac{K_j}{h_j} > \frac{K_{s(j)}}{h_{s(j)}}$, then compute the operating policies as described in Section III.A, that is, using the method suggested originally by Clark and Scarf. Otherwise, go to Step 2.

Step 2: If $\tau_2 \leq 1/2 \tau_1$, go to Step 3; otherwise, go to Step 4.

Step 3: Collapse the problem into a single stage problem and apply Clark and Scarf's procedure for this resulting single stage problem. As a consequence both stages will follow the same replenishment policy.

(The set-up cost and the holding cost for the collapsed problem are respectively given by $K_1 + K_2$ and $h_1 + h_2$. The lead time is given by $\tau_1 + \tau_2$.)

Step 4: Compute the replenishment policy for each stage separately.

- a. Stage 2 policy: Collapse the two stage problem into a single stage problem and obtain the common replenishment policy. Use the resulting policy for stage 2. (The set-up cost and holding costs used when finding the policy are $K_1 + K_2$ and $h_1 + h_2$, respectively.)

- b. Stage 1 policy: Use Clark and Scarf's procedure with parameters K_1 , h_1 and p_1 to obtain the policy at stage 1.

Our computational experiments suggest that when τ_2 is low relative to τ_1 , it is generally better, in terms of expected costs, to implement the collapsed policy instead of computing the policies separately for each stage. We therefore have suggested that the collapsed policy be used when $\tau_2 \leq 1/2 \tau_1$. We point out that additional experimentation is needed to identify the conditions under which Step 3 or Step 4 should be selected. Our limited experience suggests that $\tau_2 \leq 1/2 \tau_1$ works well; however, we have no proof that this policy is the best possible one.

We comment that the optimal solutions for Examples 1 and 3 are found when using the proposed MCS procedure.

IV. EXPERIMENTAL RESULTS

We conducted an experiment to achieve two goals. Our first objective was to compare the solutions obtained using the MCS procedure with optimal solutions. Thus our goal is to identify when the MCS works well. The second goal was to find conditions under which safety stock and safety time exist. As we have pointed out previously, the amount of safety stock and time that exist are a function of the policy that is followed. Consequently, it is not surprising that our experiments show that under certain circumstances both safety time and stock should exist.

IV.A. The Experiment

Optimal policies and those found using the MCS procedure were calculated for a number of lead time, demand distribution, and cost combinations. In all cases, a two stage serial production system was examined. Thus this experiment obviously does not allow us to make general statements concerning the potential usefulness of the MCS algorithm for solving more complicated problems. However, the results indicate that such an approach, or some further modification of it, may very well be a good heuristic to use to obtain cost effective operating policies for more complex product structures.

The cost structures we used in the experiment are given in Table VII. The fixed production costs (K_i), the echelon holding costs (h_i), and the finished product shortage cost (p_1) are given for two structures. In the first case, $\frac{K_2}{h_2} = 1 < 6 = \frac{K_1}{h_1}$, and in the second case $\frac{K_2}{h_2} = 6 > 1 = \frac{K_1}{h_1}$. Thus the modification we made to the Clark and Scarf algorithm will be important in all cases where the first cost structure is used.

Two demand distributions for the final product are used in the experiment. The expected demand is the same for both distributions; however,

the variance is higher for the second distribution. The specification of these distributions is given in Table VIII.

Stage	Structure 1			Structure 2		
	K_i	h_i	p_i	K_i	h_i	p_i
1	30	5	150	9	9	150
2	9	9	--	30	5	--

TABLE VII. Cost Data for the Experiment

Demand Distribution A		Demand Distribution B	
Demand Per Period	Probability	Demand Per Period	Probability
0	0	0	.4
1	.4	1	0
2	.3	2	.1
3	.2	3	.2
4	.1	4	.3
Mean = 2 Variance = 1		Mean = 2 Variance = 3	

TABLE VIII. Demand Distributions for the Experiment

One of the most important factors affecting the values of the policy parameters, and hence the amount of safety stock and time, is the lead time. The lead time combinations given in Table IX were used in the experiment. These values, measured in periods, were selected to determine the effect on the selection of an operating policy when the relative lead time lengths differ between the two stages. As can be expected, these differences significantly change the values of the policy parameters.

Lead Time Combination	Lead Time For Stage 1 (τ_1)	Lead Time For Stage 2 (τ_2)
1	1	0
2	1	1
3	2	1
4	1	5
5	5	1

TABLE IX. Lead Time Combinations Used in the Experiment

Twenty problems were constructed using various combinations of the cost structures, demand distributions and lead time values given in Tables VII-IX. Table X gives the exact specification of each of these problems.

The solutions for these problems are given in Table XI. In each case, the operating policies are stated in terms of echelon stock. For problems 1-12 we give both the solutions obtained using the MCS procedure and the optimal solution. We did not determine the optimal policy for problems 13-20 because the state space was too large for each of the corresponding MDPs. We note that the time required to find the solutions differed substantially. The solutions obtained using the MCS procedure were calculated in a few seconds using an Apple III computer while the optimal solutions were obtained in several minutes on an IBM 370-168.

The solutions found using the MCS procedure were virtually the same as the optimal solution for 8 of the 12 problems. The expected cost per period for the MCS procedure was within .5% of the optimum for 11 of the 12 problems and was within 1.8% of the optimum on the remaining problem. Thus it appears that the MCS procedure provides a reasonable method for calculating operating

Problem Number	Cost Structure	Demand Distribution	Lead Time Combination
1	2	A	1
2	2	B	1
3	1	A	1
4	1	B	1
5	2	A	2
6	2	B	2
7	1	A	2
8	1	B	2
9	2	A	3
10	2	B	3
11	1	A	3
12	1	B	3
13	2	A	4
14	2	B	4
15	1	A	4
16	1	B	4
17	2	A	5
18	2	B	5
19	1	A	5
20	1	B	5

TABLE X: Specification of the Problems Used in the Experiment

Problem Number	Policy Found Using MCS			Optimal Policy		
	First Stage Policy	Second Stage Policy	Expected Cost Per Period	First Stage Policy	Second Stage Policy	Expected Cost Per Period
1	(6,7)	(5,9)	74.06	(6,7)	(5,9)	74.06
2	(7,8)	(7,11)	88.288	(7,8)	If $x_1=6$, (6,11) If $x_1 \neq 6$, (7,11)	88.282
3	(5,7)	(5,7)	82.55	(5,7)	(5,7)	82.55
4	(6,8)	(6,8)	97.09	(6,8)	(6,8)	97.09
5	(6,7)	(8,12)	79.07	(6,7)	(7,12)	79.01
6	(7,8)	(10,14)	95.47	(8,8)	(10,14)	95.45
7	(6,10)	(7,10)	90.124	(6,9)	(7,10)	90.122
8	(7,8)	(9,11)	113.09	(7,11)	(9,11)	111.09
9	(8,9)	(10,14)	96.22	$x_2 \neq 10$, (8,9) $x_2 = 10$, (9,10)	(10,14)	95.75
10	(10,11)	(12,16)	121.55	$x_2 \neq 12$, (10,11) $x_2 = 12$, (10,12)	(12,16)	121.33
11	(10,12)	(10,12)	114.80	(10,12)	(10,12)	114.80
12	(11,14)	(11,14)	143.60	(11,12)	$x_1 \neq 10$, (11,14) $x_1 = 10$, (11,15)	143.12
13	(6,7)	(17,21)	87.317	--	--	--
14	(7,8)	(20,24)	113.095	--	--	--
15	(6,10)	(16,19)	106.35	--	--	--
16	(7,8)	(18,22)	144.40	--	--	--
17	(15,17)	(17,21)	141.80	--	--	--
18	(18,19)	(19,23)	179.40	--	--	--
19	(16,19)	(16,19)	181.10	--	--	--
20	(18,22)	(18,22)	222.90	--	--	--

TABLE XI. Operating Policies for Each Problem

policies for two stage serial systems. Furthermore, it is our opinion that this method will find very good operating policies for more complicated situations. Although we have not proven this, it appears that the optimal policy is always a generalized (s,S) policy and is often a simple (s,S) policy at each stage. Thus the policy produced by the MCS method is the correct or close to the correct form.

The operating policies found depend on several factors. First, observe that in all cases as the variability in the demand process increases, the reorder point at each stage is strictly increasing. Thus additional protection is needed as the demand variability increases. Furthermore, the expected cost per period increases as well.

Second, we see that the expected cost is always lower for the second cost structure when comparing the solutions obtained based on cost structure 2 with those found using cost structure 1. Recall that cost structure 2's costs differ from those in structure 1 only in that the setup and holding costs are reversed. The result is due in part to the fact that holding costs are higher for structure 1. It is also attributable to the fact that lot sizes at stage 1 are higher on the average in cost structure 1 than they are in structure 2.

Third, we see that the lead time structure has a substantial affect on the values of the policy variables. Clearly as the lead times increase both the values of the policy variables and the expected costs increase. An interesting result is obtained when comparing lead time structures 4 and 5. Note that the total lead time is 6 periods in each of these lead time structures. However, when the longer lead time is in the second stage, the expected costs are always correspondingly higher.

IV.B. Safety Time and Safety Stock

Our final observations are related to the existence and amount of safety time and safety stock. It is impossible to determine the exact amount of safety stock or time by knowing only the values of the policy variables. We estimated these quantities for each problem using the policies given in Table XI.

Recall that safety stock is the expected amount of on hand inventory at the time an order or lot arrives. Note that safety stock can exist at either stage. Safety time exists only at stage 2. It measures the expected time between the completion of a lot at stage 2 until any units from that lot are used in the production of a lot at stage 1. When computing safety time, we assume stage 2 production lots are used in a FIFO manner by stage 1.

Using the operating policy computed using the MCS procedure for problem 6, we illustrate how the safety stock and time can be calculated. The policies are (7,8) and (10,14) for stages 1 and 2, respectively. Recall that these policies are stated in terms of echelon stock. The lead time is one period for each stage. The behavior of echelon and installation stock at each stage over a 52 period horizon is given in Figure 4. The beginning on hand inventory at stage 1 is 8 and at stage 2 is 6. The demands for the horizon, which were generated randomly from distribution B, are given in Table XII.

Period	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	
Demand	0	4	0	0	4	2	2	0	3	3	0	0	4	4	0	0	3	0	3	0	0	4	3	4	0		
Period	27	28	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	
Demand	3	4	4	3	4	2	0	0	0	0	4	4	4	3	0	0	3	4	4	4	2	4	4	0	4		

TABLE XII. Demand Data Generated From Distribution B for 52 Periods

(a) On-Hand
Stock at
Stage 2

(b) Echelon
Stock for
Stage 2

(c) On-Hand
Stock for
Stage 1

FIGURE 4. Graphs of Inventory Levels for 52 Periods for Problem 6

The dots in Figure 4a represent the on-hand installation stock for stage 2. Note that the graph indicates both beginning and ending inventory for each period. Asterisks indicate the stage's inventory position--on-hand plus on-order stock. The same information is displayed for stage 1 in Figure 4c. In addition, the order quantities, desired and actual, are given for each order. As can be seen, although stage 1's policy is a (7,8) policy, it cannot always be implemented because the required amount of stock from stage 2 is not available. The graph in Figure 4b shows the history of stage 2's echelon stock.

We can make several observations from this sample realization.

First, the average safety stock at stage 1 is 1.8 units. As can be seen from Figure 4c, there is almost always on-hand inventory when a production lot is completed at stage 1. The average safety stock at stage 2 is 0. The on-hand inventory is zero before each order arrives. Thus there is no safety stock at stage 2. However, the safety time at stage 2 is 0.6 periods. This means that an order at stage 2 is available on the average 0.6 periods before it is used in a production at stage 1. As can be seen from Figure 4a, safety time exists during periods 18, 25, 26, 33, 34, 35, 36, 41 and 42. Thus 9 periods of safety time exist for the 15 orders, or there are 0.6 periods of safety time per order on the average.

The following statistics can also be computed:

Average beginning on-hand inventory at stage 1	is 5.557 units
Average beginning on-hand inventory at stage 2	is 4.634 units
Average ending on-hand inventory at stage 1	is 3.538 units
Average ending on-hand inventory at stage 2	is 2.615 units
Average number of backorders (stage 1)	is 0.3846 units
Number of orders at stage 1	is 31 orders
Number of orders at stage 2	is 15 orders
Average order quantity at 1	is 3.38709 units
Average order quantity at 2	is 7.1333 units
Fill rate	is 0.98130
Service level (Probability of not running out of stock)	is 0.9615

To obtain reliable estimates of the above statistics for the problems in the experiment, we simulated the behavior of some of these systems for 7500 periods. Some of the results of the simulation runs are presented in Table XIII. We can make some observations from these data.

PROBLEM NO.	5	13	14	17	18	19
Policy Stage 1	(6,7)	(6,7)	(7,8)	(15,17)	(18,19)	(16,19)
Policy Stage 2	(8,12)	(17,21)	(20,24)	(17,21)	(19,23)	(16,19)
Safety Stock at 1	2.085	2.0286	2.3126	3.118	5.041	2.305
Safety Stock at 2	0.155	1.278	2.0706	0.109	0	0
Safety Time	0.135	0.2586	0.6050	0.1527	0.647	0
Average order quantity at 1	2.597	2.8303	3.4227	3.301	3.679	4.724
Average order quantity at 2	5.763	5.7541	6.8079	5.759	6.77	4.724
Fill rate	0.977	0.967	0.9569	0.960	0.944	0.935
Service level	0.964	0.9562	0.9512	0.955	0.914	0.932

TABLE XIII. Results of Simulation Runs for Example Problems

If a collapsed policy is followed, there is never any safety stock or safety time for stage 2 no matter what lead time structure exists. The results for problem 19 demonstrate this fact. The reason for this is that the policy parameters are the same at both stages. Consequently, as soon as the units become available from stage 2, the stage 1 policy indicates that production should commence immediately. We also observe that the only situations for which neither safety time nor safety stock exists at stage 2 is for problems in which a collapsed policy was followed.

The lead time structure is also very important. Compare, for example, the results for problems 13 and 14 with those obtained for problems 17 and 18, respectively. Recall that the lead times are $\tau_1 = 1$ and $\tau_2 = 5$ for problems 13 and 14 and $\tau_1 = 5$ and $\tau_2 = 1$ for problems 17 and 18. Observe that the safety stock is much higher when the lead times are longer.

We note that demand variability affects the amount of desired safety stock. This can be seen by comparing the results obtained for problem 13 with those found for problem 14, and those found for problem 17 with those of problem 18. The combined affect of long lead times and high demand variability will result in safety stock at a given stage.

Finally, we note that safety time exists at stage 2 for all cases except when a collapsed policy is followed. The amount of safety time depends on the length of the stage 2 lead time. →

V. MRP POLICIES ARE (s,S) POLICIES

Our final comments are concerned with the form of the inventory policy corresponding to standard MRP logic. It is not hard to see that MRP policies are nothing more than (s,S) policies. This observation follows immediately from the manner in which planned production orders are triggered. Once echelon stock at a stage reaches a point corresponding to a predicted deficit at that stage at the end of the period a lead time in the future, an order is triggered. Thus the reorder point s equals projected demand over a lead time plus one period, as given by the forecast, plus any buffer stocks required for that and succeeding stages. The order-up-to level S depends on the lot sizing rule being followed. Our observation concerning the form of MRP policies does not depend on the lot sizing rule that is followed.

We will illustrate these observations using an example. The cost and lead time data are given in Table XIV. The demand distribution used is Demand Distribution B in Table VIII. The optimal operating policy for this problem is a $(7,8)$ policy at stage 1 and a $(9,13)$ policy at stage 2. The probability of running out of finished product during a period is .06 when following these policies. As a result, we define the service level for this system to be .94.

Stage	K_i	h_i	p_i	τ_i
1	9	9	100	1
2	30	5	--	1

TABLE XIV. Cost and Lead Time Data for the Example Problem

Suppose MRP logic is used to determine when orders are placed at each stage. Furthermore, assume a lot-for-lot reorder size policy is followed.

Recall that the forecasted demand for the finished product is 2 units in each period. Suppose that a .94 service level is desirable. Then a safety stock, or buffer stock, of 4 units is required at stage 1. This quantity is computed based on the assumption that the production lead time is one period for stage 1 and no buffer stock is planned for stage 2.

To illustrate the behavior of the stock levels at each stage, we assume the demand pattern given in Table XII is experienced. We also assume the beginning on-hand inventory at stage 1 is 8 units and at stage 2 is 6 units. No units are in production at either stage initially. We chose these starting values to illustrate that the long term system behavior when following a MRP policy is not affected by initial conditions. As we show, these starting values cannot be achieved if MRP ordering logic is followed.

The results of following MRP ordering logic is shown in Table XV. For each stage and period we give forecasted demand, projected stock on-hand at the end of a period, net requirements, scheduled receipts, and planned orders. We also give the actual on-hand stock at each stage at the beginning of each period both before and after the receipt of scheduled orders. The MRP system operates on a rolling horizon basis. Data for the next 5 periods are given for each category for each period. The calculations are assumed to be made at the beginning of a period. After each period the appropriate adjustments are made to reflect the actual demand that has occurred. The numbers given in Table XV correspond to the calculations normally given when describing the behavior of a MRP system. They do not directly indicate the behavior of echelon stock, however. Graphs of echelon and stage stock levels for 8 periods are given in Figure 5. These graphs show that once the effects of the initial conditions are removed, the policy at stage 1 is a (8,8) policy and at stage 2 is a (10,10) policy, where the policy parameters are

	Period 1					Period 2					Period 3					Period 4					Period 5					Period 6					Period 7				
*B1	8					8					4					4					8					4					2				
BS1	8					8					4					8					8					4					4				
B2	6					6					6					2					2					2					0				
BS2	6					6					6					2					2					2					4				
BSD2	6					6					2					2					2					0					0				
Period	1	2	3	4	5	2	3	4	5	6	3	4	5	6	7	4	5	6	7	8	5	6	7	8	9	6	7	8	9	10	7	8	9	10	11
Stage 1																																			
Forecast of Demand	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	
Projected Stock on Hand at Period's End	6	4	4	4	4	6	4	4	4	4	2	4	4	4	4	6	4	4	4	4	6	4	4	4	4	2	4	4	4	4	2	4	4	4	4
Net. Requir.	0	0	2	2	2	0	0	2	2	2	0	4	2	2	2	0	0	2	2	2	0	0	2	2	2	0	4	2	2	2	0	4	2	2	2
Sched. Receipt	0	0	2	2	2	0	0	2	2	2	0	4	2	2	2	4	0	2	2	2	0	0	2	2	2	0	4	2	2	2	2	4	2	2	2
Planned Orders	0	2	2	2	2	0	2	2	2	2	4	2	2	2	2	0	2	2	2	2	0	2	2	2	2	4	2	2	2	2	4	2	2	2	2
Stage 2																																			
Gr. Requir.	0	2	2	2	2	0	2	2	2	2	4	2	2	2	2	0	2	2	2	2	0	2	2	2	2	4	2	2	2	2	4	2	2	2	2
Projected Stock on Hand at Period's End	6	4	2	0	0	6	4	2	0	0	2	0	0	0	0	2	0	0	0	0	2	0	0	0	0	-2	0	0	0	0	0	0	0	0	0
Net. Requir.	0	0	0	0	2	0	0	0	0	2	0	0	2	2	2	0	0	2	2	2	0	0	2	2	2	0	4	2	2	2	0	2	2	2	2
Sched. Receipt	0	0	0	0	2	0	0	0	0	2	0	0	2	2	2	0	0	2	2	2	0	0	2	2	2	0	4	2	2	2	4	2	2	2	2
Planned Orders	0	0	0	2	2	0	0	0	2	2	0	2	2	2	2	0	2	2	2	2	0	2	2	2	2	4	2	2	2	2	2	2	2	2	2
	Actual Demand = 0					Actual Demand = 4					Actual Demand = 0					Actual Demand = 0					Actual Demand = 4					Actual Demand = 2 Planned Prod. at 1 = 4 Actual Prod. at 1 = 2 Actual Prod. at 2 = 4					Actual Demand = 2 Planned Prod. at 1 = 4 Actual Prod. at 1 = 4 Actual Prod. at 2 = 2				

*B1: Actual on-hand beginning of the period before scheduled receipt
 BS1: Actual on-hand beginning of the period after scheduled receipt
 B2: Actual on-hand beginning of the period before scheduled receipt
 BS2: Actual on-hand beginning of the period after scheduled receipt
 BSD2: Actual on-hand beginning of the period after scheduled receipt and after delivery to 1

TABLE XV. MRP Schedules for The Example Problem

FIGURE 5.

based on echelon stock. Thus the operating policy dictated by the MRP system is a (s,S) policy at each stage.

Recall that the optimal policy is also of the (s,S) type. Thus the MRP policy is of the correct form, but does not have the policy variables set at their optimal levels. In this example, the average cost per period increases by about 7.4% when following the MRP policy rather than the optimal one. Some other statistical comparisons between the optimal and MRP policies are given in Table XVI. Observe that some safety time occurs as a result of the MRP logic even though it was not planned to exist in advance. Furthermore, observe that the average safety stock is considerably lower than the target value of 4 units. This happens because the desired order quantity and the actual production level at stage 1 are not the same and because there is variability in the demand process. Consequently, the average on-hand inventory does reach the desired level.

	Optimal Policy (7,8), (9,13)	MRP Policy (8,8), (10,10)
Average safety stock at 1	2.029	2.2894
Average safety stock at 2	0	0
Safety time	0.266	0.16
Average order quantity 1	3.058	2.763
Average order quantity 2	7.133	3.433

TABLE XVI. Comparison Between Optimal and MRP Policies

Recall that the calculations made for the preceding example were based on the use of a lot-for-lot scheduling rule. We emphasize that the conclusion we have reached concerning the form of MRP operating policies does not depend on this supposition. For example, if the Silver-Meal [12]

heuristic is used to generate lot sizes, it is easy to show that the MRP policy is a (8,8) policy at stage 1 and a (10,12) policy at stage 2. This policy was simulated for 9000 periods. Some comparisons between this and the optimal policy are given in Table XVII. Note that both safety time and stock exist at stage 2 when the Silver-Meal heuristic is used to determine lot sizes. We also observe that actual lot sizes and safety stock at stage 1 differ substantially from those that were planned.

	Optimal Policy (7,8), (9,13)	MRP (using Silver-Meal Method for Lot Sizing) (8,8), (10,12)
Average Costs	100%	102.57%
Average safety stock at 1	2.2048	2.5929
Average safety stock at 2	0	0.3295
Safety time	0.24	0.57
Average order quantity 1	3.000	3.2144
Average order quantity 2	6.811	3.8755

TABLE XVII. Comparisons Between the Optimal and MRP Policies

In summary, MRP policies are (s,S) policies. The values of the policy parameters depend on the buffer stocks and lot sizing rules that are specified. Consequently, both safety stock and safety time can exist at intermediate stages, as we have shown with our example. The MRP policies generally appear to be of the optimal or nearly optimal form, but normally will not have the policy variables set at optimal values. As long as demand variability is relatively low over the lead time, $\tau_1 + \tau_2$, the MRP policies will be close to optimal, in a cost sense. However, when lead times are long and demand variability is high, then MRP policies can perform very poorly. Thus it is important in these cases to use better heuristics, such as the MCS heuristic, to determine a good operating policy.

REFERENCES

- [1] ARROW, K.J., S. KARLIN, and H. SCARF, Studies in the Mathematical Theory of Inventory and Production, Stanford University Press, Stanford California, 1958.
- [2] BANERJEE, A., "Buffer Inventories for Material Requirements Planning Systems," Working Paper University of Delaware, May 1978.
- [3] CARLSON, R. and A. YANO, "Buffering Against Demand Uncertainty in Material Requirements Planning Systems - Systems with no Emergency Setups," Technical Report #81-2, Department of Industrial Engineering and Engineering Management, Stanford University, April 1981, p. 29.
- [4] CARLSON, R. and A. YANO, "Buffering Against Demand Uncertainty in Material Requirements Planning Systems - Systems with Emergency Setups for Some Components," Technical Report #81-3, Department of Industrial Engineering and Engineering Management, Stanford University, April 1981, p. 27.
- [5] CLARK, A. and H. SCARF, "Optimal Policies for a Multi-Echelon Inventory Problem," Management Science, Vol. 6, No. 4, July 1960, 474-490.
- [6] CLARK, A. and H. SCARF, "Approximate Solutions to a Simple Multi-Echelon Inventory Problem," Chapter 5 in K. Arrow, S. Karlin and H. Scarf (Eds.), Studies in Applied Probability and Management Science, Stanford University Press, Stanford, California, 1962.
- [7] HOWARD, R., Dynamic Programming and Markov Processes, M.I.T. Press, Cambridge, Massachusetts, 1960.
- [8] MAXWELL, W. and MUCKSTADT, J., "Establishing Reorder Intervals in Multi-Stage Production-Distribution Systems," Technical Report #561, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY, 1982.
- [9] MEAL, H.C., "Safety Stocks in MRP Systems," Technical Report #166, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- [10] MILLER, J., "Hedging the Master Schedule," in Disaggregation Problems in Manufacturing and Service Organizations, Eds. L. Ritzman et al., Martinus Nijhoff Publ. Company, 1979, 237-256.
- [11] NEW, C., "Safety Stocks for Requirements Planning," Production and Inventory Management, 2nd. Qtr. 1975, 1-18.
- [12] SILVER, E. and MEAL, H., "A Heuristic for Selecting Lot-size Quantities for the Case of a Deterministic Time-varying Demand Rate and Discrete Opportunities for Replenishment," Production and Inventory Management, Vol. 14, No. 2, 1973, pp. 64-74.
- [13] WHYBARK, C. and J. WILLIAMS, "Material Requirements Planning under Uncertainty," Decision Sciences, Vol. 7, No. 4, 1976, 595-606.