

**SAFETY STOCKS IN MULTI-STAGE
SERIAL PRODUCTION SYSTEMS**

Candace Arai Yano

**Department of Industrial & Operations Engineering
The University of Michigan
Ann Arbor, Michigan 49109-2117**

Technical Report 86-14

April 1986

Revised March 1987 and November 1987

SAFETY STOCKS IN MULTI-STAGE
SERIAL PRODUCTION SYSTEMS

Candace Arai Yano
Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, Michigan 48109-2117

April 1986
Revised March 1987 and November 1987

This work was supported in part by National Science Foundation Grant
DMC 8504644.

SAFETY STOCKS IN MULTI-STAGE SERIAL PRODUCTION SYSTEMS

ABSTRACT

Some recent research on safety stock positioning in two-level manufacturing systems has indicated that in certain types of systems, most or all safety stock should be held in finished product form. Questions have been raised as to whether these results can be generalized to deeper product structures. We investigate the problem of determining safety stocks for serial systems with cyclic production schedules to study the effect of product structure depth on the desirability of component safety stock.

The objective is to minimize the sum of expected shortage costs for the finished product and expected inventory holding costs. A single-pass heuristic algorithm is developed which provides a useful starting point for more detailed studies. In addition, simulation results indicate that a policy with only finished product safety stock performs reasonably well.

SAFETY STOCKS IN MULTI-STAGE SERIAL PRODUCTION SYSTEMS

1.0 INTRODUCTION

Problems associated with uncertain demand continue to plague manufacturing planning and control systems even though many of the issues have been recognized for some time. Practitioners deal with these problems primarily through the use of safety stock, rescheduling, or a combination of the two methods. It is still not well understood when, where, and how much of each should be used, nor how best to use them in combination. However, some recent and ongoing research is beginning to answer some of these questions for safety stock in simple production systems.

We investigate the safety stock positioning problem for serial systems in an attempt to understand of the effects of product structure depth on the desirability of component safety stock. By serial we mean that each intermediate production stage has only one immediate predecessor and one immediate successor. The objective is to determine safety stock levels consistent with a shortage cost per unit (or equivalent fill-rate) criterion where only demand is uncertain and shortages of finished product are considered critical. We analyze the special case where cyclic, nested production schedules are used which, to our knowledge, has not been studied previously.

The most closely related research includes pioneering work by Clark and Scarf (1960), who show that in periodic review serial systems with no setup costs, the optimal policy has a single critical number at each stage. They also show that introducing a setup cost at the final stage of production changes the form of the optimal policy to (s,S) at each stage. DeBodt and Graves (1985) develop approximate order quantity-reorder point policies for a serial system having setup costs at all stages.

Considerable analytical work has also been done on safety stocks in

two-level arborescent assembly and distribution systems (see, for example, Lambrecht, Muckstadt, Luyten (1984), Schmidt and Nahmias (1985), Schwarz (1985)), which generally indicate that most safety stock should be held closest to the source of uncertainty when only demand is uncertain. Heuristic policies have been suggested by Meal (1979) and Miller (1979), among others. There are also numerous simulation-based studies, some of which appear in the references in Maxwell et al. (1984). The reader is referred to Baker (1985), Baker, Magazine and Nuttle (1986), McClelland and Wagner (1985), and Gerchak and Henig (1986) for some recent work on safety stocks in systems with components common to multiple end-items. The issue remains unresolved for general product structures. Some questions also remain as to whether these initial results can be generalized to deeper product structures.

In section 2 we describe the problem and assumptions. Section 3 presents a heuristic algorithm to determine safety stock quantities. Section 4 includes computational results using the algorithm and a simulation study to test the quality of the solutions. We conclude with a summary and discussion in Section 5.

2.0 DESCRIPTION OF THE PROBLEM

We address the problem of determining a safety stock level for each stage of production of a single product with a serial product structure. The finished product is produced to stock and the producer is concerned with minimizing the sum of inventory holding costs and per unit shortage costs for the finished product in conditions of uncertain demand. Excess demands are assumed to be backordered.

We assume that the demand distribution is stationary in time. To simplify the exposition, however, we will use demands which are independent and identically distributed and have a normal distribution. The approach can be

extended fairly easily to any situation with a stationary, infinitely divisible distribution of forecast errors (including environments where the process mean is non-stationary).

At each stage in the process, production is done using a cyclic schedule. The time between production runs of an item is an integral multiple of the time between production runs of its successor. Thus, the policy is, in some sense, nested. However, the presence of positive leadtimes and random demands results in production quantities that deviate a little from those in a literal nested schedule. The frequency of production runs is determined by an optimal or heuristic multi-stage approach (e.g., Crowston, et al. (1973), Blackburn and Millen (1982), Afentakis, Gavish and Karmarkar (1984), Roundy (1986) and Afentakis and Gavish (1986)). A lot-for-lot policy, which is commonly used in serial systems, falls into this framework. We also assume that there is unlimited production or storage capacity, or that capacity constraints have been dealt with in the lot-timing decisions.

For simplicity, we assume that leadtimes are constant and positive, but each stage may have a different leadtime. Since each production run at a stage represents the same time-supply, the assumption of constant leadtimes is not unrealistic. We further assume that supply and production yields are known and deterministic.

It is assumed that production runs start at the beginning of the period and that demand occurs during the period. The system operates using an order-up-to type of policy at each stage, where the policy is to produce enough to last a production cycle plus a production leadtime ahead (on average), plus safety stock, when applicable, whenever a production run is scheduled. We also assume that demand information is made available to all stages at the end of each period and future dependent demands are updated accordingly. This type of updating is typical in many manufacturing information systems such as

Manufacturing Resource Planning. As a result of this updating, component production stages may have partial information about imminent dependent demands, even before they occur.

We assume that the shortage cost per unit is sufficiently high that shortages (at all levels in the product structure) occur only in the last period having positive demand in each production cycle. In most situations, this is a very accurate approximation, as we will demonstrate later. Silver and Jonsson (1987) provide support for this assumption in their analyses of two-level distribution systems.

We wish to determine safety stock quantities at each stage given a shortage cost per unit. Since the timing of orders is fixed and demand is stationary, determining safety stock levels is equivalent to determining an order-up-to point for each stage.

In the next section we describe our approach to the problem and an algorithm which implements that approach.

3.0 APPROACH TO THE PROBLEM

This section is divided into six parts. We first provide a formulation of and motivation for our approach to the problem. In section 3.1 we provide justification for the shortage timing assumption, which is critical to the developments that follow. Then, in sections 3.2 and 3.3 we develop the model details necessary for the algorithm (which is the embodiment of our approach) appearing in section 3.4. Finally section 3.5 describes possible approximation biases.

We use the following notation throughout the paper:

h_i = inventory holding cost per unit per period at stage i

π = shortage cost per unit for the finished product

T_i = production interval at stage i

L_i = leadtime for process at stage i

μ = mean finished product demand during a period

σ = standard deviation of finished product demand during a period

SS_i = safety stock at stage i (in units)

k_i = safety stock multiplier, where $SS_i = k_i \sqrt{n} \sigma$

and n represents an appropriate number of periods

$\Phi(\cdot)$ = standard normal cumulative distribution function

The stages are numbered so that stage N is the first production process to be completed and the last process to be completed is called stage 1. Our assumption about nested cyclic production schedules implies that $T_i \leq T_{i+1}$ and that T_{i+1}/T_i is always a non-negative integer.

We can formulate the problem approximately by assuming that cycle stock carrying costs are fixed for a given vector of production intervals, \mathbf{T} . The problem then becomes:

$$\text{minimize } H(\mathbf{SS}) + \pi G(\mathbf{SS}|\mathbf{T})$$

where \mathbf{SS} is the safety stock vector, $H(\mathbf{SS})$ = expected inventory holding cost per period using \mathbf{SS} , and $G(\mathbf{SS}|\mathbf{T})$ = expected number of shortages per period using \mathbf{SS} given \mathbf{T} . It is difficult even to represent expected inventory holding costs accurately because of interactions between stages. Analytic methods to determine (or even to estimate) shortages for multi-stage systems given safety stock quantities and production intervals do not exist and are not likely to be forthcoming. Thus, other methods must be used to circumvent these difficulties.

We develop a sequential procedure which uses a simple marginal approach to the safety stock problem. In order to retain the sequential nature of the procedure, it is necessary to assume in determining the safety stock at stage i

that stage $i + 1$ is always able to satisfy requirements at stage i . Thus, there is no mechanism in the procedure for stage i to compensate for uncertainty of supply from stage $i + 1$. However, in determining the cost and usefulness of safety stock at stage $i + 1$, its potential inability to supply stage i is considered. The net result of this assumption is that the safety stock values obtained from the algorithm are approximate. They are, however, better estimates than we might expect because of some mitigating factors in the algorithm which we discuss more fully in section 3.5.

The approach is a greedy-type algorithm in which one adds safety stock only up to the point where (on the margin) the benefit (measured by reduction of shortage costs per cycle) equals the expected holding cost per cycle from an additional unit of safety stock. Of course, it is difficult to equate these exactly, but we try to equate them up to a practical tolerance level. Since both the benefit and cost are reflected over one production cycle, and the production cycles differ among items, we use a basic result of renewal processes (Ross, 1970) to obtain the expected net benefit per unit time for each item. The procedure is applied using these expected net benefit values.

We model the expected reduction of shortage costs per cycle in section 3.2. The expected (incremental) holding cost per cycle is derived in section 3.3.

3.1 Justification of the Shortage Timing Assumption

Recall that the shortage timing assumption indicates that shortages only occur in the last period with positive demand in a production cycle. It has been stated in this manner because the cyclic, nested production schedule causes dependent demands to be lumpy (i.e., in some periods demands are zero). Conceptually, however, it is easier to think of the assumption as stating that shortages occur at the end of the cycle. The probability of a shortage occurring earlier than this is very small in most situations. In the special

case where $T_i/T_{i+1} = 1$, for stage i , the "last period with positive demand in a production cycle" is the only period with positive demand. Therefore, the probability of an early stockout occasion is zero. As the T_i/T_{i+1} ratio increases, as the leadtime for stage $i + 1$ increases, and as demand variability increases, the early stockout probability increases (for fixed safety stock at stage $i + 1$). It does not increase quickly, however, because of the offsetting of high and low demands which naturally occur.

Table 1 indicates the probability of an early stockout situation at stage $i+1$ for various values of T_i and T_{i+1}/T_i multiples when $L_{i+1}=T_{i+1}$, $\mu/\sigma = 3$ and $k_{i+1} = 0$ when the demands are distributed $N(\mu, \sigma^2)$. In other words, the table shows the probability that stage $i+1$ stocks out "early" when the demand variance is relatively high, there is no safety stock at stage $i+1$, and the leadtime at stage i is as long as its own production interval. While these probabilities do not represent the "worst case," they do reflect results for a relatively adverse set of parameters. Thus, many real situations would have much lower probabilities. Jonsson and Silver (1987) have obtained qualitatively similar results in a distribution setting. (Details of how the probabilities are computed appear in Appendix A).

TABLE 1

Because the probabilities are so small, we ignore their effects in our approximate model. As a result, certain costs and probability estimates may be slightly biased. We will discuss the potential biases later in the paper.

3.2 Derivation of Expected Reduction of Shortages

We first remark that since demand is stationary, stage i follows an (R_i, T_i) policy (order up to R_i every T_i periods), and each order is equivalent to the sum of demands since the previous order was placed. This is also true for stage 1, whose requirements are external demands. Each requirement at each stage

therefore can be linked directly to finished product demand over some time interval.

Suppose that we were to increase the safety stock level of item i by one unit. A series of events must occur if that additional unit is to be instrumental in preventing a finished product shortage that would have occurred otherwise. First, the "order" from stage $i-1$ must be large enough so that this extra unit is "needed." An order at stage i being "large enough" is the same as the sum of demands over some appropriate interval being "large enough." Now, even if that unit of item i is "needed" at stage $i-1$, the unit of item $i-1$ that now exists (by virtue of the unit of item i being available), must also be needed by stage $i-2$ if that unit of item i is to be instrumental in preventing a shortage. This cascades all the way to stage 1. The set of all events which may occur in an N -stage system are diagrammed in Figure 1. Each path corresponds to one possible sequence of events.

FIGURE 1

We use a simple example of a two-stage system to explain how each of these events corresponds to the demand for the finished product. The interested reader is referred to Appendix B for details of the general case. Figure 2 depicts the flow of goods between stages and over time in a system with $T_1 = 2$, $L_1 = 1$, $T_2 = 4$, and $L_2 = 3$. One batch at stage 2 and two batches at stage 1 are supposed to supply (directly or indirectly) customer demands over 4 time periods. Stage 1 places two orders on stage 2, one in period $t + T_1 - T_2$ and one in period t . Since it is using an order-up-to policy, the first order is equal to the sum of demands in the periods from $t - T_2$ through $t + T_1 - T_2 - 1$ (which happens to be only two periods in this case). Similarly, its next order is the sum of the demands in the two succeeding periods.

Stage 2 already had forewarning of the demands (at least forecasts), so it started the corresponding production run in period $t + T_1 - T_2 - L_2$. Under the shortage timing assumption, stage 2 should have no difficulty satisfying the order of stage 1 in period $t + T_1 - T_2$, but it may have difficulty satisfying the second of the two orders.

FIGURE 2

Referring again to the two-stage example and Figure 2, stage 2 is responsible for the demand uncertainty from period $t + T_1 - T_2 - L_2$ through period $t + 2T_1 - T_2 - 1$. (Recall that the second of the two orders from stage 1 is equal to demand in periods $t + T_1 - T_2$ and $t + 2T_1 - T_2 - 1$). At that point, future demand uncertainty becomes the responsibility of stage 1.

The main point to be noted in this special case (which is also true for the general case) is that each stage faces uncertainty during its own production leadtime and the time from the first to the last withdrawal by the successor stage and the various time intervals do not overlap. Therefore, the probabilities of events on various branches in a path are independent since the forecast errors are assumed to be i.i.d.. In each case, the probability can be expressed as a function of a standard normal distribution.

We can now express the probabilities for each outcome for an event tree with the last path having index m as

$$\begin{aligned}
 p_{0m} &= \prod_{i=1}^m (1 - \Phi(k_i)) \\
 p_{jm} &= \Phi(k_j) \prod_{i=j+1}^m [1 - \Phi(k_i)]
 \end{aligned} \tag{3}$$

where we use the convention $\prod_a^b = 1$ if $b < a$.

3.3 Derivation of Expected Cost of Additional Unit

Recall that the paths indexed 0 through i in Figure 2 indicate all possible outcomes for an incremental unit of safety stock at stage i . We have already determined the probability of each path in section 3.2. We now need to determine the cost of each path in the tree. Since we have assumed that shortages occur only in the last period having positive demand prior to the end of a production cycle, for each path, we know exactly whether or not the incremental unit is used by the subsequent stage, and if so, when.

We must compute expected inventory holding costs incurred during an entire typical cycle of length T_i . The reason for this is that there is (only) one opportunity during a cycle for item i safety stock to prevent a stockout situation at stage $i - 1$, and similarly up to stage 1. Thus, during any T_i periods, stage i has only one opportunity to reduce finished product shortages.

Now consider the calculation of the cost of path $j \geq 1$ in a tree with last path having index m . Suppose that an additional unit of safety stock is added at stage m . Observe that adding safety stock at stage m does not affect the sequence of events (or the costs associated with those events) on paths indexed $m+1, \dots, N$. Thus, we only need to consider the effect of stage m safety stock upon events (and their costs) in paths $1, \dots, m$.

The j th path has events "item i needed" for $i = j + 1$ to m . By our shortage timing assumption, the time spent by the extra unit at these stages will be $T_i - T_{i-1}$, respectively. The j th path also has event "item j not needed" as its final branch. Thus, the incremental unit of j remains in inventory during the entire cycle, or T_j periods.

Let c_{jm} = cost of path j in a tree having last path with index m . We can now write

$$c_{jm} = \sum_{i=j+1}^m h_i(T_i - T_{i-1}) + h_j T_j \quad (4)$$

For c_{0m} , the only difference is that the last branch is "item 1 needed." Thus, the incremental unit of item 1 remains in inventory for $T_1 - 1$ periods, and

$$c_{0m} = \sum_{i=2}^m h_i(T_i - T_{i-1}) + h_1(T_1 - 1) \quad (5)$$

The expected cost term for item i is

$$\sum_{j=0}^i c_{jm} p_{jm} \quad (6)$$

This is simply the expected inventory holding cost which will be incurred if one adds an incremental unit of safety stock at stage i .

Recall that we would like to compare πp_{0m} (the expected benefit during a cycle) with the expression in (6), which is the expected cost during a cycle. Because we are using a sequential procedure, however, we do not need to compare the net benefit per unit time for the various items. Thus, for each item we only need to ascertain whether the net benefit is positive by comparing π to C_i , where

$$C_i = \left[\sum_{j=0}^i c_{jm} p_{jm} \right] / p_{0m} \quad (7)$$

The numerator reflects the expected inventory holding cost incurred over the production cycle of the item, while the denominator reflects the expected reduction of shortages (fraction of a unit) achieved over the same interval.

3.4 Algorithm

For any given value of k , we can now compute C_i if m (i.e., the stage for which safety stock is being considered) is specified. The role of the index m is explained next. Observe that for any i , each term in the summation in (7) is increasing in k_j , $j = 1, \dots, i$. This is true by virtue of the definitions of p_{jm}

in (2) and (3). (Details of the proof appear in Appendix C). Therefore, C_i is increasing in each k_j , $j = 1, \dots, i$, or in words, C_i increases if safety stock at any succeeding stage is increased, while safety stock at preceding stages does not affect it. This is consistent with the assumption that stage $i+1$ can always satisfy stage i , which permits us to retain the sequential nature of the procedure, and justifies the use of the proposed marginal approach under this assumption. Solving for k_j , $j = 1, \dots, n$ in sequence, we can obtain a solution, since only successor stages affect the decision at stage i . In solving for the safety stock level at stage i , we therefore need to use a tree with last path having index i .

An algorithm follows:

Step 1. $i = m = 1$.

Step 2. Compute p_{jm} for $j \leq m$ using (2) and (3).

Step 3. Compute c_{jm} for $j \leq m$ using (5) and (6).

Step 4. Calculate C_i using (7). If $C_i < \pi$, increment k_i , and return to step 2. Otherwise increment i and m and return to step 2 if $i \leq N$.

Note that in the algorithm, we compare C_i with π . This is equivalent to comparing the expected benefit of the unit of safety stock, namely πp_{0m} , with the expected cost in (6). Thus, safety stock is added only if the expected benefit exceeds the expected cost, as estimated using the procedure described above.

Because of the monotonicity of the C_i values, the algorithm has a unique solution (within the limits of the step size). It is not, however, guaranteed to be optimal because it is a greedy procedure, and the objective function has not been shown to be convex.

Observe that the safety stock multipliers are independent of the leadtimes, but the safety stock quantity at stage i (denoted as SS_i) does depend upon the

leadtimes. We have

$$SS_i = k_i \sqrt{T_i + L_i - T_{i-1}} \sigma, \quad i = 2, \dots, N$$

$$SS_1 = k_1 \sqrt{T_1 + L_1} \sigma$$

3.5 Approximation Bias

There are two factors which contribute to an approximation bias. The first is the shortage timing assumption. The second is the assumption of perfectly reliable supply from stage $i + 1$ to stage i in determining safety stock at stage i . We discuss the effects of the two assumptions in turn.

At each stage, some shortages may occur earlier than our shortage timing assumption specifies. Therefore, items may be held for shorter periods of time than we have assumed. Yet, if shortages occur early at intermediate stages, an incremental unit will be transformed into its successor earlier than anticipated. It thereby faces the possibility of being held for a longer period of time at a (usually) higher inventory holding cost rate. Thus, it is difficult to determine whether the c_{jm} values under- or over-estimate the true inventory holding costs.

Each term in the products of equations (2) and (3) represents the probability of shortage at some stage i given the shortage timing assumption. But the actual probability is slightly larger since shortages may actually occur earlier, and therefore affect supply to subsequent stages. The same is true for the probability of shortage at stage 1 in the denominator of the righthand side of equation (2). Thus, p_{0m} understates the true probability. On the other hand, $\Phi(k_j)$ in equation (3) overestimates the probability that "item j is not needed." In other words, the probability of an item j shortage is understated because the shortage timing assumption is not exact. Since this term is overstated, but the product term is understated, the net effect (of the product of the two terms) is indeterminate.

We now know that the approximation understates p_{0m} , but the effect on the c_{jm} s and p_{jm} s, $j > 0$ is indeterminate. Therefore, it is difficult to predict whether the algorithm will be biased due to these effects, and if so, in what direction. We investigate this more fully in the next section.

We now turn to the effects of our assumption about perfectly reliable supply. The first order effects are clear: since each stage does not compensate for possible uncertainty of supply, the safety stock values from the algorithm will be too low. On the other hand, the second order effects are more subtle. Suppose that the algorithm gives safety stock multipliers for stages $i-1$, i and $i + 1$ equal to k_{i-1}' , k_i' and k_{i+1}' , respectively. Stage i should compensate for uncertain supply by using $k_i > k_i'$ since $k_{i+1}' < \infty$. Also, when considered alone, k_{i-1} should be greater than k_{i-1}' for the same reason. But with k_{i-1} increased from k_{i-1}' , there is less motivation for holding safety stock at stage i because it becomes relatively more expensive to do so.

To illustrate this effect with an example, consider a three stage system. Stage 1 "thought" that stage 2 would be a perfectly reliable supplier, as did stage 2 of stage 3. Stage 2 turns out to be unreliable occasionally. So stage 1 increases its safety stock. This causes stage 2 to reduce its safety stock. But now stage 2 finds out that stage 3 is sometimes unreliable, so stage 2 increases its safety stock from the reduced level. The net effect at stage 2 is difficult to predict. If stage 2 has a net increase in its safety stock, stage 1 now wants to decrease its safety stock, and vice versa. The net overall effect at each stage is unknown. What becomes clear is that decisions at one stage may not be nicely monotonic in the decisions at other stages, and that even sophisticated multi-pass or iterative approaches (if they could be developed) may not converge (at all or to the correct solution). Indeed, there is no guarantee that C is a nicely behaved function of \mathbf{k} , although it is evident

that C_i is convex increasing in k_i if k_j , $j \neq i$ are held fixed. It is not, obvious, however, how C_i changes with k_j in general if one wishes to incorporate the effects of unreliable supply.

We hypothesize that the net overall effect of the approximation is that the algorithm underestimates the optimal safety stock multipliers. We test this hypothesis in the following section.

4.0 EXPERIMENTAL RESULTS

We selected problems representing three and four stage systems with a wide range of parameter values for study. For the three-stage systems, we let $T_1 = 1, 2, \text{ and } 4$, with $T_i/T_{i-1} = 1, 2, \text{ or } 4$, with a maximum value of $T_i = 8$. This resulted in seventeed combinations of T_i values. We normalized $h_1 = 1.0$ and let $h_i/h_{i-1} = 0.2, 0.4, 0.6, \text{ and } 0.8$. The shortage cost, π , may take values 4, 36, 100, and 196. For the four-stage systems, we used similar parameter values but let $h_i/h_{i-1} = 0.2, 0.5, \text{ or } 0.8$. Since the safety stock multiplier is independent of the leadtimes for $L_i > 0$, we used $L_i = 1$ for all i .

Using the algorithm described in section 3, we determined safety stock multipliers for these problems to investigate the effects of the various parameters on the safety stock levels. The algorithm gave positive safety stock multipliers only at stage 1.

We then repeated the exercise using $h_i/h_{i-1} = 0.05, 0.10, \text{ and } 0.15$ with the belief that smaller component holding costs would make intermediate safety stocks more desirable. The results were the same as discussed above. These results indicate that if intermediate safety stocks are desirable, it is because of the one major effect which was not considered in our model--the need to compensate for the predecessor's unreliability.

To investigate the magnitude of biases of the algorithm, we selected a small sample from among the parameter sets (shown in Table 2) for which positive

component safety stock might be economical. We used a search procedure which involved evaluating (via simulation) all points in a three-dimensional cube around the best solution obtained so far. If an improved solution was found, the procedure was repeated. The process continued until the solution could not be improved. We used a step size for k_1 of 0.10. Recall that we are implementing order-up-to policies. For these policies, it is well-known that total cost is a convex function of a single order-up-to point if all other order-up-to points are held fixed. Thus, the total cost function is, at least in some sense, well behaved. We used this information to judiciously select other solutions for evaluation. It should also be noted that the objective function is very flat in some dimensions because of the relatively small holding costs for components. The search procedure was necessarily limited by the number of safety stock parameters to be varied, and the fact that simulation had to be used to evaluate the performance of each safety stock policy. Nevertheless, we are confident that the best solution from the grid search gives a cost which is very close to the minimum cost for the reasons cited above.

TABLE 2

Each replication of a given system was run for eight times the production cycle of the component with the longest cycle, and was initialized in steady state. Twenty-five replications were run for each applicable system and set of safety stock parameters. The average costs from 25 simulation runs are listed in Table 3 for the best solution from the grid search, the solution from the algorithm, and the "best" solution with no component safety stock. We found the cost functions to be relatively well-behaved despite the possible interactions discussed earlier.

TABLE 3

It appears that the algorithm provides good "ballpark" safety stock multipliers. However, there seems to be considerable latitude in "trading" finished product safety stock for second-level component safety stock without a significant change in total cost. Nevertheless, costs were reduced by decreasing finished product safety stock in favor of second-level component safety stock. Observe, however, that the problems have $h_1 = 1.0$ and $h_2 = 0.4$, which is rather uncommon in a serial production system (60% of the value added in the last stage of production). One might expect considerably less latitude in this tradeoff for more typical cost structures.

We note that the shortage costs used here are high relative to the inventory holding costs. Moreover, the differences between the cost of the solution from the algorithm and the cost of the "best" solution with no component safety stock are small. This suggests that there is relatively little benefit from increasing finished product safety stock above the value given by the algorithm. Indeed, the characteristics of the solutions from the search procedure suggest that the benefits from additional finished product safety stock are so small, that it would be better to increase second-level safety stock at the expense of reducing finished product safety stock.

The "best" solution with no component safety stock would be an acceptable solution in most instances, even for these situations in which component safety stock would appear to be economically advantageous. In all five problems, the relative cost penalty from using either the solution from the algorithm or the best solution with no component safety stock is relatively small (11% or less).

These results corroborate earlier work on two-stage assembly systems which found that in most circumstances, only finished product safety stock should be held, or that a policy using only finished product safety stock performs well. The results from our experiments indicate that even in the absence of component mating problems (which tend to make component safety stock less desirable),

safety stock of components is warranted only in special situations.

Of course, we probably could have concluded this by comparing a more extensive grid search with the "finished-product-only" policy, but the algorithm significantly reduced the grid search by providing a reasonable starting point. The algorithm also provides fairly good solutions with very little effort when compared with simulation.

All of the trends from the numerical results described above may not generalize to other distributions, since right hand tail probabilities are critical. Most good forecasting models for real systems, however, are able to provide forecasts with near-normal forecast errors. If necessary, the approach developed here can be adapted to a wide range of forecast error distributions.

Some recent research by Baker, Guerrero, and Southard (1986) suggests that more careful control of component safety stock through the use of Miller's (1979) idea of hedging and slower (less than 100%) safety stock replenishment rates could make component safety stock economically advantageous. It may well be that, in addition to production cycle combinations and cost structures with special characteristics, control policies much more sophisticated than those available in most MRP systems are needed to make component safety stock a viable substitute for a portion of the finished goods inventories.

Although we have investigated the safety stock positioning problem in serial systems, the results also have implications for a variety of related issues and decisions. We discuss some of the insights that we have gained, not with the intent of providing conclusive results, but rather possible directions for future research.

In general, we know that there is a tradeoff between cycle stock and safety stock. Thus, there is a tradeoff between cycle stock and safety stock. This was recognized long ago for continuous review inventory models but has not been addressed fully for multi-stage production systems in a periodic review

environment. There are important implications for implementation of just-in-time systems in situations where uncertainty cannot be eliminated. The price for reducing cycle stock may be a lot of additional safety stock, primarily in the form of finished product.

The results raise the question of when component safety stock should be used. Consider the role of safety stock in smoothing flow. Except in situations with highly non-stationary demand, component safety stock may only help to smooth the flow if it is important that large batches of specified size be retained intact. Otherwise, a cyclic production timing policy can provide much of the smoothing.

There are several situations, however, where component safety stock may be much more important. In each of these situations, however, the tradeoffs are not simply a question of safety stock at stage n versus safety stock at stage m . Bottleneck production processes require safety stock to ensure that the process can remain busy. Here, the tradeoff is between inventory and capacity. Components which are inputs to or outputs of processes with variable yields also need some provision for safety stock. Again, the tradeoff is inventory vs. capacity, and in some cases, customer service. Components ordered from unreliable suppliers, which do not fill orders as specified, may require safety stock. However, minor quantity variations probably can be tolerated without much safety stock, depending upon the size of the orders. Timing variations call for safety stock, particularly if facilities would be stopped because a shortage of the part. Again, this is a question of safety stock and capacity. In our scenario, there are no explicit capacity constraints, but in reality the duration of the leadtimes and the time between production runs are affected by capacity limitations, which, in turn, affect safety stock levels. Some of these issues are beginning to be addressed by Graves (1987). Finally, safety stock may

reduce the need to re-schedule.

Thus, in conclusion, we continue to believe that the safety stock positioning problem is an important one, but as a result of this research we also believe that uncertainty of demand may not be the motivating force for component safety stock. Certainly money can be saved through the appropriate placement and quantity of safety stocks and for this, simple heuristic procedures such as the one discussed here may serve as a more than adequate guideline. The role of safety stock relative to rescheduling and capacity considerations probably will provide much more economic motivation for safety stock as well as much greater opportunity for cost reduction.*

* The author wishes to express appreciation to Karla Bourland for her assistance in developing the simulation model, to the National Science Foundation for support under grant DMC 8504644, and to two anonymous referees for their helpful comments on earlier drafts of this paper.

TABLE 1

Probability of "Early" Shortage Occurrence

$$\mu/\sigma = 3, k_{i+1} = 0, L_{i+1} = T_{i+1}$$

		T_i				
		1	2	3	4	6
T_{i+1}/T_i	2	.01700	.00135	.00012	$<10^{-5}$	$<10^{-5}$
	3	.06681	.01700	.00466	.00122	.00012
	4	.11123	.04181	.01700	.00714	.00135

TABLE 2

Parameters for Simulated Systems

 $(h = (1.0, 0.4, 0.24), \sigma = 30)$

PROBLEM	T	π
1	(1,2,2)	36
2	(1,4,4)	36
3	(2,4,4)	100
4	(2,8,8)	100
5	(4,8,8)	100

TABLE 3

Results of Simulation Study

Problem	Grid Search		Algorithm		Best Solution with $k_2 = k_3 = 0$	
	Solution	Cost	Solution	Cost	Solution	Cost
1	(1.43,1.0,0)	2298	(1.93,0,0)	2558	(2.13,0,0)	2536
2	(1.43,.8,0)	6365	(1.93,0,0)	6975	(2.13,0,0)	6882
3	(1.86,.6,.8)	8655	(2.06,0,0)	9307	(2.26,0,0)	9242
4	(1.86,.8,0)	24502	(2.06,0,0)	27111	(2.46,0,0)	26262
5	(1.56,.8,0)	30433	(1.76,0,0)	31751	(2.16,0,0)	31256

REFERENCES

- Afentakis, P. and B. Gavish (1986), "Optimal Lot Sizing Algorithms for Complex Product Structures," Operations Research 34(2), 237-249.
- Afentakis, P., B. Gavish, and U. Karmarkar (1984), "Computationally Efficient Optimal Solutions to the Lot-Sizing Problem in Multistage Assembly Systems," Management Science 30 (1), 222-239.
- Baker, K.R. (1986), "Safety Stocks and Component Commonality," Journal of Operations Management 6(1), 13-22.
- Baker, K.R., H.H. Guerrero, and M.H. Southard (1985), "Dynamics of Hedging the Master Schedule," Int. J. Prod. Res. 24(6), 1475-1484.
- Baker, K.R., M.J. Magazine and H.L.W. Nuttle (1986), "The Effect of Commonality in a Simple Inventory Model," Management Science 32(8), 982-988.
- Blackburn, J.D. and R.A. Millen (1982), "Improved Heuristics for Multi-Stage Requirements Planning Systems," Management Science 28 (1), 44-56.
- Clark, A.J. and H. Scarf (1960), "Optimal Policies for a Multi-Echelon Inventory Problem," Management Science 6(4), 475-490.
- Crowston, W.B., M. H. Wagner, and J.F. Williams (1973), "Economic Lot Determination in Multi-Stage Assembly Systems," Management Science 19(5), 517-527.
- DeBodt, M. and S.C. Graves (1985), "Continuous Review Policies for a Multi-echelon Inventory Problem with Stochastic Demand," Management Science 31(10), 1286-1299.
- Gerchak, Y. and M. Henig (1986), "An Inventory Model with Component Commonality," O.R. Letters 5(3), 157-160.
- Graves, S.C. (1987), "Safety Stocks in Manufacturing Systems," forthcoming in Journal of Manufacturing and Operations Management.
- Jonsson, H. and E.A. Silver (1987), "Analysis of a Two-Echelon Inventory Control System with Complete Redistribution," Management Science 33(2), 215-227.
- Lambrecht, M.R., J.A. Muckstadt, and R. Luyten, (1984), "Protective Stocks in Multi-Stage Production Systems," Int. J. Prod. Res. 22(6), 1001-1025.
- Maxwell, W.L., J.A. Muckstadt, L.J. Thomas, and J. Vander Eecken (1983), "A Modeling Framework for Planning and Control of Production in Discrete Parts Manufacturing and Assembly Systems," Interfaces 13(1), 92-104.
- McClelland, M.K. and H.M. Wagner (1985), "The Impact of the Product Structure and Increased Part Commonality on Safety Stock Effectiveness," Working Paper, School of Business Administration, University of North Carolina - Chapel Hill.

- Meal, H.C. (1979), "Safety Stocks in MRP Systems," Technical Report #166, MIT Operations Research Center.
- Miller, J. (1979), "Hedging the Master Schedule." Chapter 15 in Disaggregation: Problems in Manufacturing and Service Operations. L. Ritzman et al. (eds.), Hingham, MA: Martinus Nijhoff Publishers.
- Roundy, R., (1986), "98% Effective Lot-Sizing for a Multi-Product, Multi-Stage Production/Inventory Systems," Math. of O.R. 11(4), 699-727.
- Ross, S. (1970), Applied Probability Models with Optimization Applications. San Francisco: Holden-Day.
- Schmidt, C.P. and S. Nahmias (1985), "Optimal Policy for a Two Stage Assembly System Under Random Demand," Operations Research 33 (5), 1130-1145.
- Schwarz, L. (1985), "The Partitioning of On-Hand Inventory in a Two-Level Manufacturing or Distribution System: A Newsboy Model Analysis," Working Paper, Krannert Graduate School of Management, Purdue University.
- Yano, C.A. and R.C. Carlson (1987), "Safety Stocks for Assembly Systems with Fixed Production Intervals," forthcoming in Journal of Manufacturing and Operations Management.

APPENDIX A

Derivation of probabilities for Table 1 can be done as explained below. It is assumed that production runs are started at the beginning of a period and that the goods are delivered to the successor at the beginning of some subsequent time period. Demand is assumed to occur during a period. Every time stage $i+1$ produces, it produces up to a quantity which is enough to last $T_{i+1} + L_{i+1}$ periods on average (if there is no safety stock). Suppose it does so at some point in time (call it time zero). The production run is completed at time L_{i+1} and will satisfy orders from stage i in periods L_{i+1} through $L_{i+1} + (T_{i+1}/T_i - 1)T_i$ (at intervals of T_i periods).

Information about finished product demand up to time zero is taken into account in the stage $i+1$ order quantity. Thus, relative to the stage $i+1$ batch in question, stage $i+1$ is uncertain only about demands through $L_{i+1} + (T_{i+1}/T_i - 1)T_i - 1$ (since the last stage i order from the stage $i+1$ batch occurs at the beginning of the following period). Consequently, stage $i+1$ effectively faces a set of equivalent finished product demands over a horizon of $L_{i+1} + (T_{i+1}/T_i - 1)T_i$ periods. An early shortage occurs if it cannot even cover demand over $L_{i+1} + (T_{i+1}/T_i - 2)T_i$ periods, since the second-to-last stage i order occurs T_i periods before the last.

To compute the probability of an early stock, we only need to find the probability that the second-to-last order cannot be satisfied. (If an earlier order from the same batch cannot be satisfied, then it must also be true that the second-to-last order cannot be satisfied). Since stage $i+1$ ordered enough to last T_i periods longer (on average), there would be $T_i D$ units of "safety stock" available. Hence, in computing the probability of an early stockout, it is as if we had $T_i D$ units of safety stock to cover uncertainty over $L_i + (T_{i+1}/T_i - 2)T_i$ periods.

Since we have assumed that demands are i.i.d., for $T_i/T_{i+1} \geq 2$ we can write the probability of early stockout as

$$1 - \Phi\{T_i D / [L_{i+1} + (T_{i+1}/T_i - 2) T_i]\}^{0.5\sigma}.$$

In the special case where $T_i/T_{i+1} = 1$, the early stockout probability is zero since for each stage $i+1$ batch, there is only one period in which a stockout can occur.

APPENDIX B

We develop relationships between the events in Figure 1 and demand for the finished product for the general case in this Appendix. Suppose a production run is begun at period t and is completed at period $t + L_1$. That production run will satisfy demand up to (but not including) period $t + L_1 + T_1$ (at which time the subsequent production run is complete). Therefore, for this production lot, stage 1 faces demand uncertainty from period t through period $t + L_1 + T_1 - 1$, and the probability that a stockout occurs in the last of these periods depends upon the level of item 1 (end-item) safety stock.

Now, continuing to stage 2, we observe that the (input) material for the stage 1 batch represents one withdrawal from the stage 2 batch. If the stage 1 batch does not correspond to the last withdrawal from that stage 2 batch, then we assume that the desired quantity is obtained from stage 2 (by the shortage timing assumption). If, however, it is the last withdrawal from that batch, there may be insufficient stock to cover the stage 1 "demand" depending on the level of item 2 safety stock. The corresponding stage 2 batch began production in period $t + T_1 - T_2 - L_2$. The batch was completed in period $t + T_1 - T_2$ and provides for stage 1 withdrawals in periods $t + T_1 - T_2$, $t + 2T_1 - T_2$, ..., t . Thus, stage 2 faces uncertainty from period $t + T_1 - T_2 - L_2$ through $t - 1$, or a total of $T_2 - T_1 + L_2$ consecutive periods. Note that these periods are disjoint from the periods in which stage 1 faces uncertainty, but they both relate to satisfaction of finished product demand in the same period.

Continuing to stages 3 through n , we would find similarly that (1) the probability that an additional unit of safety stock is "needed" by the succeeding stage depends upon demand in a set of $T_i + L_i - T_{i-1}$ periods at stage i , $i = 2, \dots, n$, and $T_1 + L_1$ periods at stage 1, and (2) that these time

intervals are disjoint. An analogy to a relay race is helpful in understanding these relationships. Each racer is to run a distance $T_i + L_i - T_{i-1}$, where $T_0 = 0$, and the racer with highest index starts the race. Each racer faces uncertainty only during his portion of the race. As soon as the baton is passed, the next runner is responsible. In our system, the baton passing corresponds to delivery of the last withdrawal from a larger batch to the successor stage.

APPENDIX C

In the following we show that C_i is monotonically increasing in k_n , $n = 1, \dots, i$. We can rewrite (7) as

$$C_i = \left\{ \sum_{j=1}^i c_{jm} \phi(k_j) \prod_{s=j+1}^m [1 - \phi(k_s)] + c_{0m} \prod_{s=1}^m [1 - \phi(k_s)] \right\} / \prod_{s=1}^m [1 - \phi(k_s)]$$

Simplifying, we get

$$C_i = \sum_{j=1}^i c_{jm} \phi(k_j) / \prod_{s=1}^j [1 - \phi(k_s)] + c_{0m} \tag{C-1}$$

If k_n is increased by Δk for any $n \leq i$, the n th term in the summation becomes

$$c_{nm} \phi(k_n + \Delta k) / [1 - \phi(k_n + \Delta k)] \prod_{s=1}^{n-1} [1 - \phi(k_s)]$$

which is clearly greater than the n th term in (C-1).

Now k_n also appears in other terms where $j > n$. For any such j , the corresponding term becomes

$$c_{jm} \phi(k_j) / [1 - \phi(k_n + \Delta k)] \prod_{\substack{s=1 \\ s \neq n}}^j [1 - \phi(k_s)]$$

which is also larger than the corresponding term in (C-1). If $j < n$, the corresponding term in (A-1) is unaffected by an increase in k_n . Thus, C_i is monotonically increasing in k_n , $n = 1, \dots, i$.

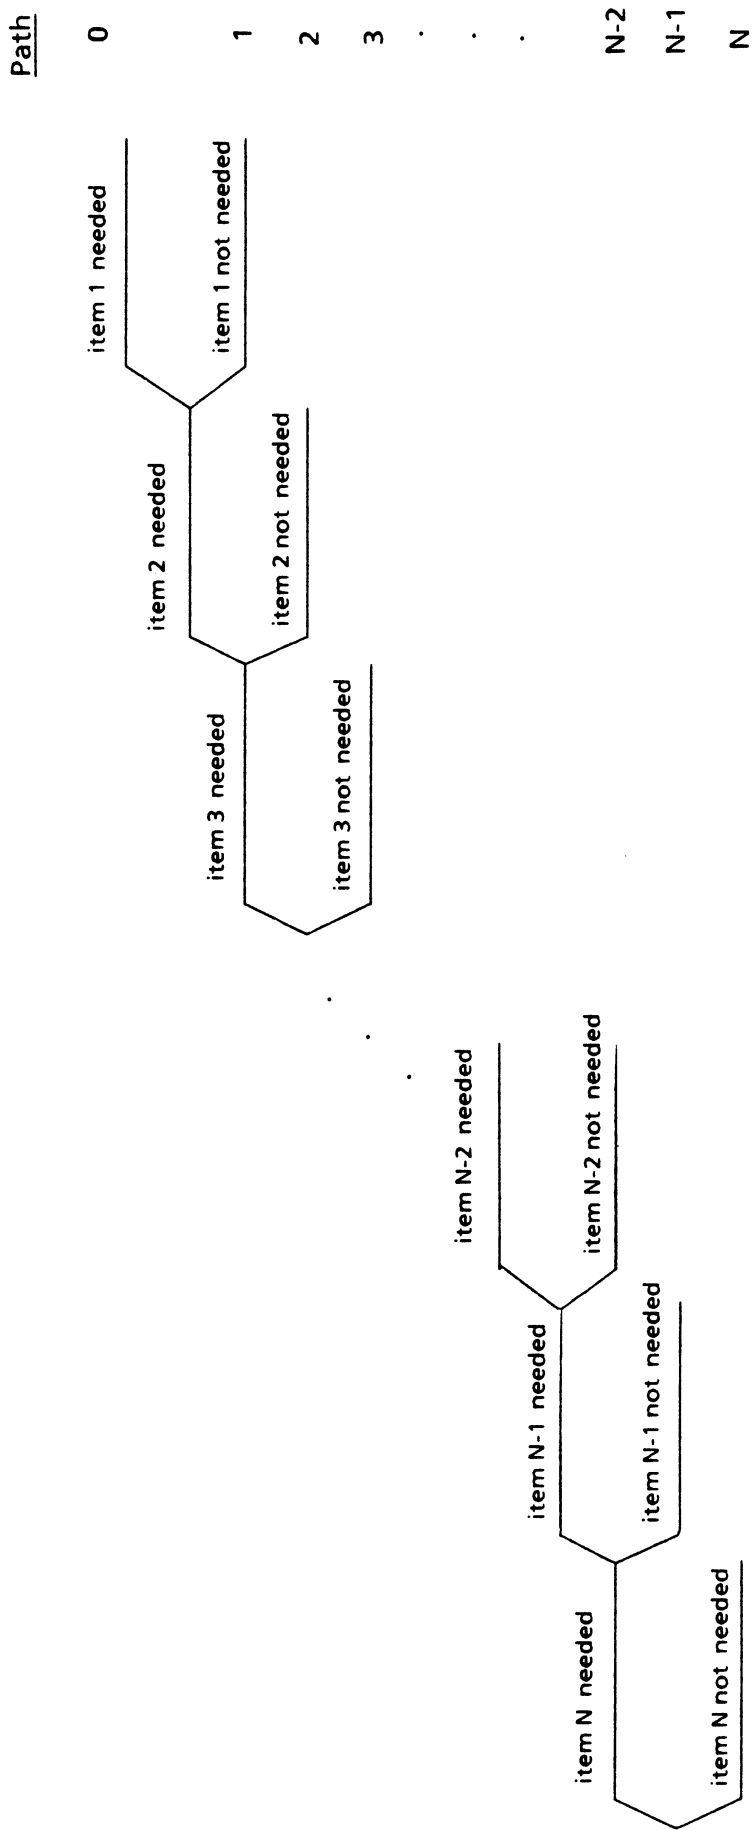


Figure 1
Event Tree for N-Stage System




Figure 2
 Diagram of Inter-stage Flows for
 $T_1 = 2, L_1 = 1, T_2 = 4, L_2 = 3$