SAFETY STOCKS FOR ASSEMBLY SYSTEMS WITH FIXED PRODUCTION INTERVALS

Candace A. Yano

Department of Industrial & Operations Engineering
The University of Michigan
Ann Arbor, Michigan 48109-2117

Robert C. Carlson
Department of Industrial Engineering
and Engineering Management
Stanford University
Stanford, CA 94305

Technical Report 87-7

February 1987

Revised June 1987 and November 1987

SAFETY STOCKS FOR ASSEMBLY SYSTEMS WITH FIXED PRODUCTION INTERVALS

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

Robert C. Carlson
Department of Industrial Engineering
and Engineering Managment
Stanford University
Stanford, CA 94305

February 1987
Revised June 1987 and November 1987

SAFETY STOCKS FOR ASSEMBLY SYSTEMS WITH FIXED PRODUCTION INTERVALS

ABSTRACT

We present a heuristic approach for setting safety stock levels in a simple two-level assembly system facing random demand, where production or assembly of each part is done at fixed intervals. The objective is to minimize expected inventory holding and shortage costs, or alternatively, to minimize expected inventory holding costs subject to a fill-rate constraint. Results from the heuristic and from simulation studies indicate that component safety stock may not be economical in most situations. We also provide some qualitative insights into factors affecting safety stock economics.

SAFETY STOCK FOR ASSEMBLY SYSTEMS WITH FIXED PRODUCTION INTERVALS

INTRODUCTION

Scheduling multi-stage batch manufacturing facilities facing uncertain demand is not an easy task. Recent research has focused on relatively simple systems with the intent of characterizing or developing insight into optimal policies. Schmidt and Nahmias (1985) investigated a periodic review assembly system with one finished product and two components and having no setup costs. The optimal policy was shown to have a single critical number (order-up-to point) at each stage of production. The actual production quantity is modified if there is insufficient input material or if it is known that a component mate will not be available. Lambrecht, et al. (1984) analyze the same system with setup costs. They indicate that an (s,S) policy at each stage is optimal.

Continuous review models have been developed by DeBodt and Graves (1985). Various other approaches have been developed, including those of Meal (1979) and Miller (1979). Finally, simulation studies, too numerous to mention here, have been done addressing the problem (e.g., Whybark and Williams (1976)).

Our objective in this paper is to characterize the safety stock problem and to develop heuristic approaches for safety stock management in the context of an assembly system using fixed production intervals. By this we mean that production is done intermittently, with a fixed time between starts of successive production runs. We sometimes use the term production cycle instead of production interval. The use of fixed production intervals usually gives rise to periodic review inventory policies where the review period is the greatest common divisor of the production intervals. It has been suggested by Elmaghraby (1978), Caie and Maxwell (1981), and others, that using fixed production intervals simplifies both production planning (i.e., dealing with capacity constraints), and detailed scheduling. Even when demand is uncertain and production quantities are permitted to vary from one batch to the next in response to demand fluctuations, having fixed production intervals can still simplify production planning and scheduling. At a minimum, one knows whether or not a setup must be done for an item and has a general idea of how much must be produced (e.g., a two-week supply). Furthermore, since one has an estimate of the quantity to be produced, it is possible to make relatively accurate estimates of the required processing time. It is important to note that the basic time period need not be large. Therefore, if setup costs and setup times are small, the production intervals can be small, keeping both inventory and total production leadtimes low.

The actual lot-timing can be done using demand forecasts and any applicable optimal or heuristic lot-sizing technique which can be constrained to give fixed production intervals. For the case of constant demand and costs, the algorithms of Blackburn and Millen (1982), or Maxwell and Muckstadt (1985) could be used. With time-varying demand or costs, another approach would need to be used (e.g., Afentakis et al. (1984), Afentakis and Gavish (1986)).

In the next section we describe our problem in more detail. We then describe an approximate analytical model which can aid in determining the most advantageous "locations" for safety stock. Finally, we present some computational results from which we draw general conclusions.

PROBLEM DESCRIPTION AND ASSUMPTIONS

We analyze a simple assembly system with one finished product and two components. Demand for the finished product is uncertain but stationary in time, and for simplicity, is assumed to have a normal distribution in each period. (The approach can be extended to non-stationary demand processes). All shortages are backordered and all orders are filled using a first-come, firstserved policy.

Each component is produced or procured at a fixed interval which is an integer multiple of the finished product assembly interval. The production schedule, therefore, is a nested schedule. Each step in the production process is assumed to have a known, deterministic leadtime, which includes time for procurement, fabrication, and delivery to the assembly stage (in the case of components) or assembly and delivery to the warehouse (in the case of the finished product). We assume that production intervals have been specified so that there is sufficient capacity to produce the desired quantities within the leadtime. If the demand process is highly non-stationary, then we assume that sufficient production smoothing has been done to make the fixed production schedule achievable. We also assume that production or procurement of components is done on a just-in-time basis (i.e., as late as possible), that production yields are deterministic, and that supply of materials to the component stage is perfectly reliable.

We also assume that the system uses a (T,S) or base stock policy in which one orders up to S every T periods. In such a system, the production quantity at each stage always replaces whatever has been supplied (either to a successor stage or to customers) since the most recent order was placed. Therefore, demands observed by the component stages are aggregates (over time) of external customer demands.

Since the production intervals are fixed, a modified policy (T,S) policy is clearly optimal. Schmidt and Nahmias have indicated that actual assembly quantities may be limited by availability of components, and production quantities of the components should be modified from the prescribed value if it is known in advance that a mate will not be available. Unfortunately, most manufacturing information systems are not sophisticated enough to check <u>future</u>

availability of component mates before ordering or producing a component. Even if such a capability were to exist, the sequence in which the production decisions are made would affect the ultimate production quantities--and it is highly unlikely that the information system could accommodate joint optimization of these production decisions. Thus, because of these practical considerations, we will sacrifice optimality by implementing a literal (T,S) policy.

There is one final assumption which is critical to the analytic model developed in the next section. We assume that shortages only occur in the last period with positive demand in a production cycle (i.e., at the end of the production cycle). In reality, shortages may occur earlier than this. Nevertheless, if production interval of an item is not significantly longer than the production interval of its successor, and shortage costs are sufficiently high to encourage a moderate to high level of customer service, there is a minimal chance of significant shortage quantities earlier than the assumption specifies. In extensive simulation studies, we found that shortages rarely occur earlier than we assume here. Jonsson and Silver (1987), in related work on multi-stage, periodic review distribution systems, show that almost all shortages occur at the end of a cycle.

We are concerned with minimizing expected total inventory and shortage costs. Since the production intervals are fixed, only the safety stock portion of the base stock level (but not cycle stock) is controllable. A cost per unit per period is charged on average safety stock. Inventory holding costs are not charged during production leadtimes since these costs are fixed under the assumptions of complete backordering, constant leadtimes and fixed production intervals. A cost per unit short is charged for each backorder. Because of the shortage timing assumption, for practical purposes, the duration of each backorder can be taken as one period. (For any reasonable level of customer service, backorders are filled by the subsequent production run.) Thus, the analysis can incorporate a shortage cost per unit per unit time quite easily.

Throughout the remainder of the paper, we use the following notation:

D = average demand per period for the finished product;

σ = standard deviation of finished product demand;

T_i = production interval (time between production runs) for item i;

L_i = production leadtime for item i;

h_i = inventory holding cost per period for item i;

k_i = safety stock multiplier for item i.

The base stock level for item i, S_i, equals $T_iD + k_i \sqrt{T_i + L_i}\ \sigma$.

HEURISTIC MODEL

Our objective is to develop an approximate analytic representation of the expected cost of inventory resulting from, and the reduction of shortages achieved by increasing the base stock levels. By comparing the cost with the benefit, we can determine whether increasing the base stock of an item is economically attractive.

This marginal approach to the problem does not fully capture the interactions among the base stock decisions, and is not guaranteed to be optimal because the objective function for the problem may not be convex. Thus, the models are approximate, and one would not be advised to use them alone to optimize base stock quantities. The value of the models is to quantify some critical factors in base stock decisions, and in so doing, provide (or at least confirm) some managerial insights which might be transferred to more realistic settings.

The approach involves posing the following question. Suppose the base stock level (of a specified item) were increased. How much would it ultimately cost us, on average, and what is the expected reduction in shortages? The algorithm, in a greedy manner, adds a unit to the base stock level of the item with greatest net benefit per unit time. The expected costs and savings are recomputed using the new base stock levels and the procedure repeats until there are no further net benefits. If the shortage cost is known, the benefits are computed in terms of shortage cost reductions. If a fill-rate (fraction of demand filled from stock) is specified, the benefit is measured in terms of reduction of shortages. (In this case, however, it is difficult to estimate when the desired fill-rate is achieved.) The details of the approach appear in the appendix. The remainder of this section is devoted to a qualitative explanation of the reasoning underlying the computation of costs and benefits.

An important corollary to the shortage timing assumption is that adding a unit to the base stock of any item can reduce shortages only once during that item's production interval. Thus, for each item, both expected inventory cost increases and expected shortage cost reductions are computed over the item's own production interval. With this in mind, we next develop an approach for computing these values for the finished product and then for the components.

The expected cost of adding a unit to the base stock level of the finished product can be approximated as the product of its holding cost and the expected time (during a production cycle) that it spends in inventory. If it is not used to satisfy a demand, it remains in inventory for the entire cycle. If it is used to satisfy a demand, by the shortage timing assumption, it remains in inventory for one period less. The probability that it is needed to satisfy a demand depends upon the current base stock level. (If the base stock level is already high, the probability that more inventory is needed is small). The expected reduction of shortages (measured in units) is simply the probability the additional unit is needed to satisfy a demand. (These arguments are based upon newsboy-like analyses and are exact for a single-level system, but are approximations for a two-level system because they assume that components will always be available.)

The expected cost of increasing the base stock level of a component is more complicated. If the additional unit is not ordered by the assembly stage, it will remain in component inventory for the entire production cycle. If it is ordered by the assembly stage, it will remain in component inventory for one assembly production interval less. At that point in time, it will be assembled into a unit of finished product if the component mate is available. (Otherwise, it will remain in component inventory.) If assembly of the extra unit occurs, that extra unit may remain in finished product inventory for an entire assembly production cycle if it is not needed to satisfy a demand. Otherwise, it will remain in finished product inventory for one period less (by the shortage timing assumption). In computing the cost incurred from the extra finished product inventory, we include only the holding cost of the value added in the assembly process because the component mate was already in inventory. By computing the costs and probabilities of each of the four possible outcomes (shown in Figure 1), we can obtain the expected cost due to increasing the base stock level of a component.

FIGURE 1

The probability that the additional unit actually reduces finished product shortage is the probability that the fourth path in Figure 1 occurs, i.e., the unit is ordered by the assembly stage, a mate is available, and the assembled finished product is needed to satisfy a demand. The major point to be noted in this analysis is that increasing the base stock level of a component can lead to costs exceeding the component inventory holding cost (because of possible assembly into a finished product), but several events must occur in concert for the unit to provide a reduction of shortages.

The implementation of a literal (T,S) policy sometimes results in ordering up to S even though it is known in advance that some of the units will not have a mate available. Since the procedure for computing expected inventory costs assumes this literal implementation, the expected inventory holding costs due to increasing the component base stock level may be overstated. It should also be pointed out that the potential exists for double-counting the third path (one time for each component) when mating occurs. Both components would be charged with the cost of holding the value added in the assembly process, causing overstatements of the cost of incremental component safety stock. Thus, there are at least two reasons for the approach to underestimate the desirability of component safety stock. However, many approximations are made throughout the analysis, and the assumption that finished product safety stock costs and benefits can be determined assuming 100% service from the components is much more critical.

COMPUTATIONAL RESULTS

Results indicate that for nearly all realistic situations, safety stock level for second-level components should be zero. Initially we designed a complete factorial experiment for which the parameter values are shown in Table 1. We ran the algorithm for a few of these problems (those with relatively low component holding costs) using various shortage cost values which would be expected to yield fill-rates from 90% to 98%. In all of these cases, the algorithm provided solutions with no component safety stock.

TABLE 1

Before continuing, we present two examples which were typical of the results, and which demonstrate the effect of component holding costs and production intervals. Symmetric product structures were chosen for ease of presentation of the results, but results for asymmetric structures are similar.

We found it helpful to plot average total setup and inventory holding cost (not including shortage costs) versus fill-rate because it reflects the main tradeoff over a range of possible shortage costs. In these figures, the inventory cost includes cycle stock costs and safety stock costs, but recall that only the safety stock costs are controllable.

Figures 2 and 3 show average inventory cost versus fill-rate relationships as finished product safety stock is increased and as component safety stock (for one component, or both, as applicable) is increased. These figures are for two situations in which the algorithm indicates that no component safety stock should be used. The plotted values represent averages obtained from 50 simulation runs, each with a 24-period horizon. The same sets of randomly generated demands (common random numbers) were used for each set of safety stock multipliers, so as to reduce the variance of the differences among the results.

FIGURES 2 AND 3

Note that Figure 2 represents a case in which parameter values were chosen so as to afford advantage to component safety stock relative to finished product safety stock. The components have very low holding costs and the likelihood of reducing shortages is large because the mate is always available. (Here, safety stock of both components was increased simultaneously).

The case illustrated in Figure 3 is more typical. Only 20% of the value of the product is added at the last stage (versus 80% in Figure 2). In this case, component safety stock is extremely costly and provides little increase in the fill rate.

Next we investigated what component inventory holding costs might make component safety stock economical. To do so, we attempted to find component inventory holding costs which (on the basis of the algorithm) would make one indifferent between adding finished product and component safety stock. For the system parameters in Table 1, in no case did we get a cutoff component inventory holding cost of more than 30% of the finished product inventory holding cost, and in only one case did we get component inventory holding costs totalling more than 40% of the finished product inventory holding cost.

Some patterns in these results also suggested that certain combinations of parameters might make component safety stock beneficial. In particular, it appeared that a component having a low inventory holding cost and a short production cycle (e.g., equal to the assembly cycle) in conjunction with a mate having a relatively high holding cost and a long production cycle, is a prime candidate for safety stock. The reasons are quite clearly explained through the cost-benefit computations. The costs incurred by increasing the base stock of such a component are relatively small since (a) the production cycle is short, and (b) the inventory holding costs, both of the component itself and of the value added in the assembly process, are small. The benefit of increasing the base stock of the component is enhanced by the long production cycle of the mate, since, as a result, the mate is available if needed in most instances.

We simulated several systems having parameters as those described above and found that some component safety stock was beneficial, but in most cases, the saving was less than 1% of total costs (see Yano, 1981). Thus, it appeared that other factors would be needed to make component safety stock of significant benefit.

We performed some additional studies for situations with high shortage costs. We found that initially (starting from a point with no safety stock), it was desirable to add finished product safety stock. Eventually, the relative cost of more finished product safety stock became prohibitive (due to decreasing marginal returns), and some component safety stock became desirable. We note, however, that these high shortage costs corresponded to fill-rates well in excess of 99% (usually greater than 99.5%), so some firms would not find these results particularly relevant.

Results from the algorithm and from simulation studies indicate that several factors must be present simultaneously in order for some component safety stock to be cost-effective. They are:

(1) The holding cost of the component must be very small relative to that of the finished product (i.e., 10% or less).

(2) The mates must have long production intervals leading to infrequent stockouts.

(3) The shortage cost (or desired fill-rate) must be very high.

SUMMARY AND DISCUSSION

Through the use of a heuristic analytical procedure and simulation, we have found that component safety stock is normally not economical in two-level assembly systems with two components. The results obtained here may be applicable to other product structures. If there are more than two components, the likelihood decreases that all component mates are available. Therefore, the expected benefit of using safety stock of any particular component decreases. This in turn causes a decrease in the optimal safety stock quantity and the potential savings to be gained from using component stock.

Intuition would also indicate that as one moves deeper into an arborescent product structure, a much larger number of events must occur simultaneously in order for component safety stock to have a beneficial effect on customer service. The joint probability that all the advantageous events occur simultaneously decreases approximately geometrically with the number of levels. The expected holding costs arising from an additional unit of safety stock tends to decline at a slower rate, since inevitably some of the additional safety stock will be incorporated into units of successor items for which inventory holding costs tend to be higher (due to value added in production). Therefore, the expected cost of reducing shortages tends to increase as one moves deeper into the product structure.

On the other hand, one effect of a deeper product structure is a greater degree of supply uncertainty (which is actually the result of demand uncertainty) because one stage of production may not receive what it ordered. In our two-level system, the component production facilities always received the needed inputs; only the assembly facility sometimes did not receive what it ordered. In a system with several levels, most production facilities would not receive what they ordered at least a portion of the time. This could lead to a need for more component safety stock.

Two points of view on safety stock are commonly held. The first view is that safety stock should be placed nearest to the source of uncertainty. The other view is that holding product in a less complete form should provide some protection at lower cost. We had hoped that the latter would be true, making the component safety stock decision an important one. Although the results did not bear out this conjecture, they nevertheless provide some useful insights which might be transferred to more complex systems. In particular, we have found that three factors are needed for component safety stock to be of benefit: (1) very small cost of holding the component relative to holding the finished product; (2) infrequent production of the mate(s) which minimizes mating difficulties; and (3) very high shortage costs (or fill-rate). It is interesting to note that the conclusions of Schwarz (1985) in a very different setting are similar qualitatively.

The model, however, has several limitations which need to be addressed in future research. It deals only with a two-level system where the components are not used in any other finished product, but in real systems there are many stages and many common components. Issues related to component commonality in two-level systems are beginning to be addressed by Axsater and Nuttle (1987), Baker (1985), Baker et al. (1985), and Gerchak and Henig (1986).

Leadtimes were assumed to be deterministic in our model, but would be more realistically modeled as random variables because the quantity produced in each production run varies along with demand, and because of the effects of capacity constraints on leadtimes. Random leadtimes will affect the magnitude of optimal safety stock quantities, but it is not yet clear how they will affect its allocation among stages. Some recent work dealing with random leadtimes in systems with multiple stages includes Karmarkar (1987) and Yano (1987).

We have assumed that production yields are perfect, or at least deterministic, but this is rarely ever true. Issues of random yields in systems with multiple stages are beginning to be investigated (e.g., Lee and Yano, 1985). Finally, we have assumed that production of components is done on a just-in-time basis, but this is not always possible because of capacity and scheduling considerations. Graves recently (1987) proposed a method to model the interplay between safety stock and capacity flexibility, an important first step toward resolving this issue. Much more research is needed to deal with scheduling decisions when demand is uncertain.

TABLE 1

Problem Parameters

Parameter	Values
T_1	2,4
T_2, T_3	T_1, $2T_1$, $3T_1$
h_1	1.0
h_2, h_3	.10, .25, .40
D	200 per period
σ	10, 30, 50 per period
L_1, L_2, L_3	1,5

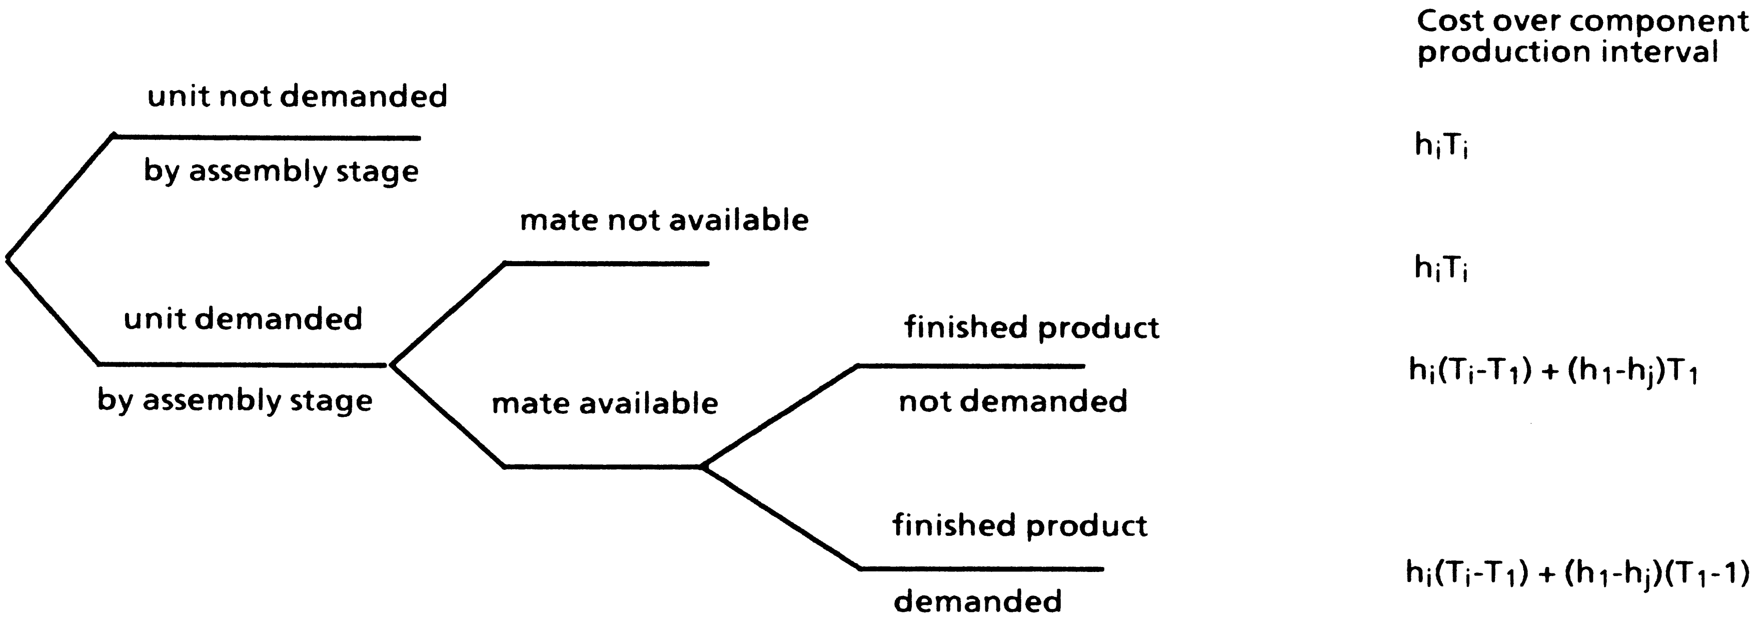

Figure 1
Uses and Costs of Component Safety Stock

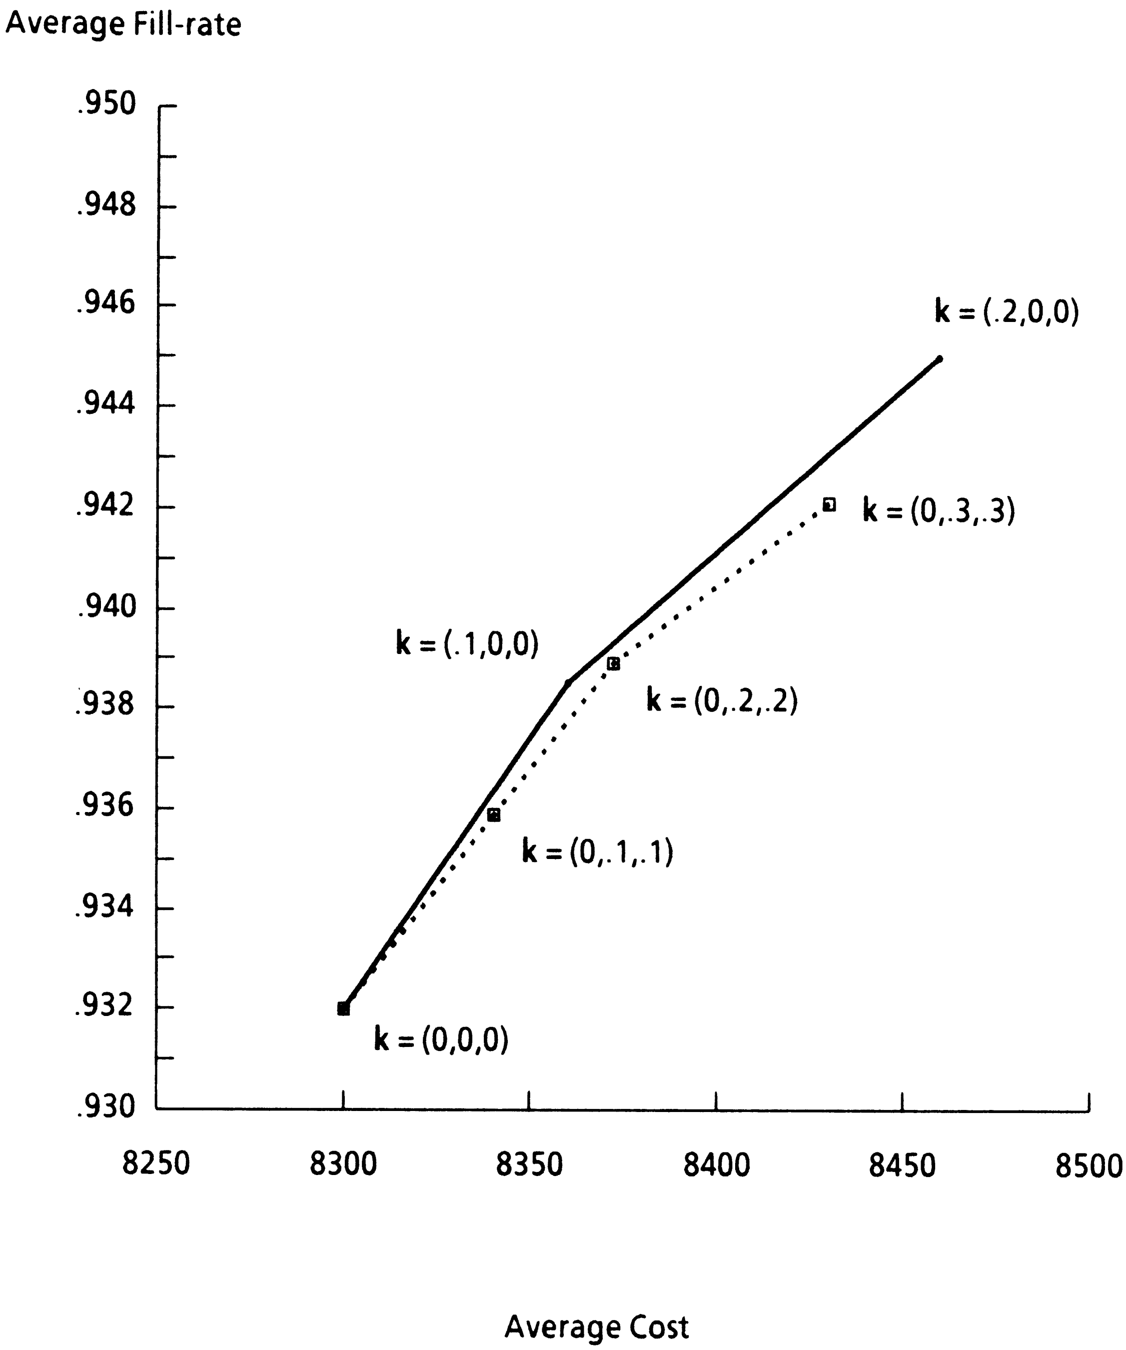

Figure 2
Comparison of Finished Product
and Component Safety Stock
for **T** = (2,2,2), **h** = (1,0.1,0.1),
L = (1,1,1), $\sigma = 30$

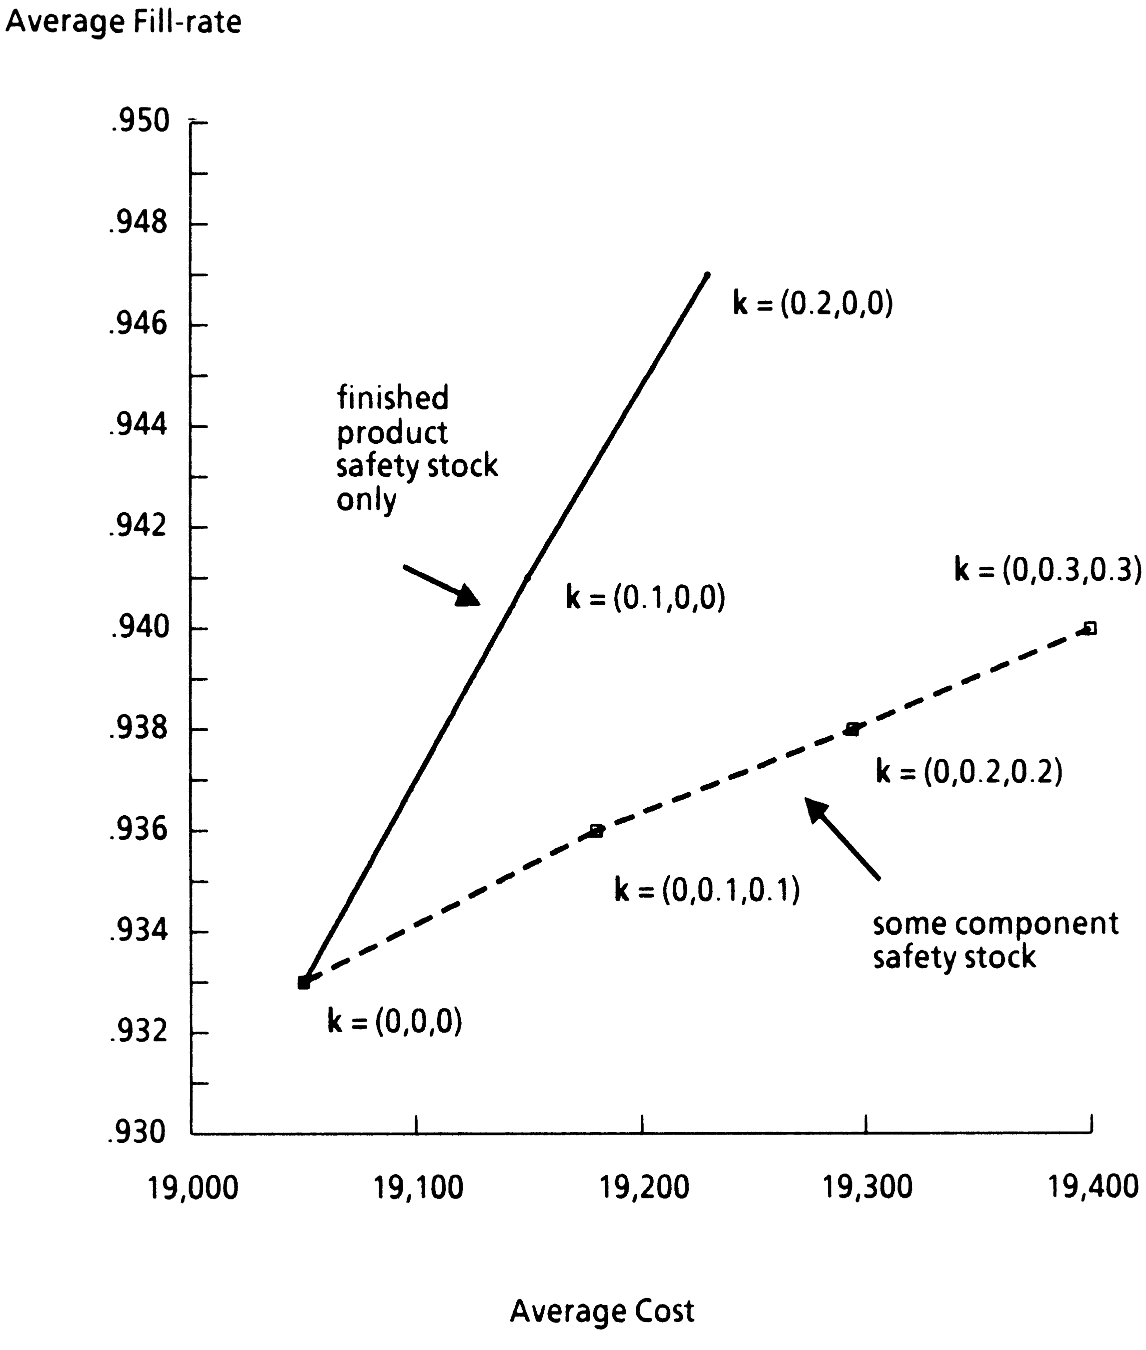

Figure 3

Efficiency of Finished Product and Component

Safety Stock for **T** = (2,4,4), **h** = (1,0.4,0.4), **L** = (1,1,1)

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(2), 222-239.

Axsater, S. and H.L.W. Nuttle (1987), "Combining Items for Lot Sizing in Multi-Level Assembly Systems," forthcoming in Int. J. Prod. Res. 25(6).

Baker, K.R. (1985), "Safety Stocks and Component Commonality," Journal of Operations Management 6(1) 13-22.

Baker, K.R., M.J. Magazine, H.L.W. Nuttle (1986), "The Effect of Commonality on Safety Stock 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.

Caie, J. P. and W. L. Maxwell (1981), "Hierarchical Machine Load Planning". In Multi-Level Production/Inventory Control Systems: Theory and Practice, L.B. Schwarz (ed.), New York: North Holland Publishing Co.

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.

Elmaghraby, S.E. (1978), "The Economic Lot Scheduling Problem (ELSP): Review and Extensions," Management Science 24(6), 587-598.

Gerchak, Y. and M. Henig (1986), "An Inventory Model with Component Commonality," Operations Research Letters 5(3), 157-160.

Graves, S.C. (1987), "Safety Stocks in Manufacturing Systems," Journal of Manufacturing and Operations Management (to appear).

Jonsson, H. and E.A. Silver (1987), "Analysis of a Two-Echelon Inventory Control System with Complete Redistribution," Management Science 33(2), 215-227.

Karmarkar, U.S. (1987), "Lot Sizes, Lead Times, and In-Process Inventories," Management Science 33(3), 409-418.

Lambrecht, M. R., J. A. Muckstadt, and R. Luyten (1984), "Protective Stocks in Multi-Stage Productions Systems." Int. J. Prod. Res. 22, 1001-1027.

Lee, H.L. and C.A. Yano, (1985), "Production Control in Multi-Stage Systems with Variable Yield Losses," Operations Research (to appear).

Maxwell, W.L. and J. A. Muckstadt (1985), "Establishing Consistent and Realistic Reorder Intervals in Production-Distribution Systems." Operations Research 33(6), 1316-1341.

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.

Schmidt, C. and S. Nahmias (1985), "Optimal Policy for a Two Stage Assembly System Under Random Demand," Operations Research 33(5) , 1130-1145.

Schwarz, L.B. (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.

Whybark, D. C. and J. G. Williams (1976), "Material Requirements Planning Under Uncertainty," Decision Science 7, 595-606.

Yano, C.A. (1981), "Safety Stocks in Material Requirements Planning Systems," Unpublished Ph.D. Dissertation, Department of Industrial Engineering and Engineering Management, Stanford University.

Yano, C.A. (1987), "Setting Planned Leadtimes in Serial Production Systems with Tardiness Costs," Management Science 33(1), 95-106.

APPENDIX A

In this appendix we describe the algorithm in further detail.

Notation:

T_i = production interval of item i (in periods)

h_i = inventory holding cost of item i per period

L_i = leadtime to produce or procure a batch of item i

π = shortage cost per unit

D = average finished product demand per period

σ = standard deviation of demand during one period

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

$\phi(\cdot)$ = standard normal density

The base stock quantity is expressed as $T_i D + k_i\sqrt{T_i+L_i}\ \sigma$, where the first term is cycle stock and the second term is safety stock. In the safety stock term, k_i is the usual safety stock multiplier and we use T_i+L_i under the square root to reflect the fact that the safety stock quantity should depend upon both the production interval and the leadtime. We use this representation principally for notational simplicity. One could use a quantity, say s_i, instead, but this makes conversion to standard normal notation more difficult. Note that since the cycle stock is constant, the base stock quantity can be adjusted by changing the safety stock quantity.

Finished Product Safety Stock

Suppose we were to add one unit to the finished product base stock level. This unit might be held through the entire finished product production interval if demand is "small." On the other hand, that additional unit could satisfy a demand that otherwise would have been backordered, if the demand during the production interval is "large." By small (large), we mean that demand during the production interval is less than (greater than) cycle stock plus the current safety stock provision. In the case of the assembly stage, $T_1 + L_1$ is precisely the period of time over which the safety stock must provide protection (i.e, from the time a production run begins until the next production batch is available for use). The reason that we must use the interval $T_1 + L_1$ is that the safety stock is incorporated in the order-up-to point, not in a reorder point. Thus, the probability that demand is less than cycle stock plus the current safety stock provision is $\Phi(k_1)$.

By our shortage timing assumption, if a shortage occurs, it occurs in the T_1th period of the production cycle, so the additional unit of safety stock would have been in the system for $T_1 - 1$ periods. Thus, the expected cost due to this unit of safety stock is the expected time that an additional unit of finished product safety stock spends in the system, multiplied by the cost per unit time, which is

$$h_1\{ T_1\Phi(k_1) + (T_1-1)[1 - \Phi(k_1)]\}. \tag{1}$$

The expected decrease in shortages is $[1 - \Phi(k_1)]$ (i.e., one unit multiplied by the probability demand is "high"). So the expected benefit is $[1 - \Phi(k_1)]\ \pi$. The expressions above lead to a form of the newsboy formula, $\Phi(k_1) = a/(a+b)$, where $a = \pi - h_1(T_1 - 1)$ (i.e., the incremental shortage cost per unit), and $b = h_1T_1$ (the incremental overage cost per unit).

<u>Component Safety Stock</u>

To simplify the discussion in this section, let us index the two components by i and j, with i being the component for which we are considering addition of safety stock. Before discussing the details of the model, we mention some basic facts about the operation of the system which may not be readily apparent. The first point is that periodically, production runs of the two components will be delivered to the assembly stage at the same time. Therefore, this represents a possible simultaneous stockout occasion. At other times, one component may have a shortage but the other component will have adequate supply (by the shortage timing assumption). The frequency with which possible simultaneous stockout occasions occur is the least common multiple of T_i and T_j.

A second, related point is that both components observe the same dependent demand, so the stockout probabilities are highly correlated. The degree of correlation depends upon the production intervals and leadtimes.

The third point is that for each component production run, there may be several production runs at the assembly stage, but (by the shortage timing assumption) it is only the last such assembly run corresponding to the component production run for which there may be insufficient component parts.

The foregoing points provide the background for analyzing component safety stock. Consider the addition of a unit of item i safety stock. We might summarize these possible streams of events as in Figure 1. Also indicated in Figure 1 are the marginal costs which would be incurred for each possible stream of events. These costs are explained next.

The first two streams of events in Figure 1 are situations in which nothing happens to the unit of component safety stock. It therefore remains in inventory for the entire production interval at a cost of h_i per period. For the third stream of events, the unit remains in inventory until it is needed for assembly, or for $T_i - T_1$ periods. If the mate is available, it is then assembled into a unit of finished product which is not sold during the production interval, and is thus held for another T_1 periods. (The mating issue arises only immediately before both of the components are about to be delivered to the assembly stage simultaneously). The marginal cost per unit is only $h_1 - h_j$, however, because the unit of component j already existed. In the last stream of events, the costs are the same as for the third stream except that the unit of finished product is sold and is therefore held in inventory for a shorter period of time.

The main difficulty now is to determine the probability that each stream of events occurs. As mentioned earlier, the components observe aggregates of finished product demand, so the probabilities of event occurrences are correlated, making it necessary to find the joint probabilities that the set of events in each stream occurs. This is tedious, since one needs to consider the precise timing of orders and production runs, but it is straightforward conceptually. Details appear in Appendix B.

Let p_n denote the probability that the nth stream occurs. Observe that p_4 is precisely the expected shortage reduction (fraction of a unit). The expected cost incurred by the additional unit of safety stock is the appropriate weighted average of the costs described above. We need to compare this with the expected reduction of shortage costs to determine whether additional safety stock would be economical. More precisely, we need to compare

$$\{h_1(T_i-T_1) + T_1[h_i(p_1+p_2) + (h_1-h_j)(p_3+p_4)] - p_4(h_1-h_j)\} \tag{2}$$

to $\pi\ p_4$. If it is smaller, additional safety stock is economical.

Using the marginal approach described earlier, the policy would be to add a unit (or small increments) of safety stock for the item with the greatest net benefit per unit time until the net benefit is zero or less. In order to capture the interaction of the safety stock decisions using the type of approach described above, it would be necessary to consider simultaneous changes of the safety stock of two or more items. This would involve modeling much more complex cost functions. We viewed such an approach as being too computationally burdensome for the purpose of finding "ballpark" solutions.

Let C_i represent the expected inventory cost during a production cyle of i due to increasing the base stock level of item i. A statement of the algorithm follows.

ALGORITHM

1. Choose a step size, Δk, which is the incremental step size for safety stock and initialize the safety stock multipliers.
2. Calculate C_i for all i using (1) and (2).
3. Find $z = \max_i [\pi\, p_4(i) - C_i]/T_i$ where $p_4(i)$ is the value of p_4 for item i.

   If $z < 0$, stop.

   Otherwise:

   Set $k_i = k_i + \Delta k$, where $i = \text{argmax}\, [\pi\, p_4(i) - C_i]/T_i$.

   Return to Step 2.

The choice of the initial values of k_i will depend upon the application, but in most cases a safe starting point would be $k_i = 0$ for all i.

We comment that if $T_i = T_j$ <u>and</u> $L_i = L_j$, we can collapse the product structure into a serial system by using an "aggregate" component with holding cost equal to $h_i + h_j$. This forces equal safety stock quantities for the two components.

APPENDIX B

In this appendix we present an example of the procedure for computing the probabilities (p_n values) needed for the component safety stock analysis.

Throughout the analysis, we assume that finished product demand occurs during the period, whereas dependent demand occurs instantaneously at the beginning of a period. We also assume that production runs start at the beginning of a period and only those that are fully complete can be used to satisfy demand. These timing conventions were chosen because they reflect what typically happens in material control systems. One important corollary results from these assumptions. The first is that the component facilities can observe their dependent demand (to be withdrawn from a component batch produced earlier) before production begins. However, the assembly facility only has information up to and including the previous period's demand when the production run is begun.

To simplify the computation, we assume that when each component production run is complete, on-hand inventory is equal to cycle stock plus safety stock. This is always true when $L_i \leq T_1$; it is also true that average on-hand inventory is equal to cycle stock plus safety stock even when $L_i > T_1$.

For assembly production runs prior to simultaneous completions of component production runs, the probability that stream 4 in Figure 1 occurs is (conceptually) the probability that the dependent demands in certain periods are small enough so that component 3 (the mate) is available, yet large enough so that an extra unit of component 2 is ordered for the assembly run, and demand is large enough to require the "extra" unit of finished product. For the remaining situations, p_4 is the probability that dependent demands are large enough so the extra component 2 safety stock is needed and demands are large enough to require the resulting additional unit of finished product.

Probabilities of other event streams can be described in a similar fashion. By translating the dependent demands into their equivalents in terms of finished product demands, we can make the following generalizations for a production run of item i which is begun in period t:

(1) The probability that an additional unit of item i safety stock is actually ordered by the assembly stage depends upon demand in periods

$$t,\ldots,t+L_i+T_i-T_1-1;$$

(2) For situations prior to simultaneous completions of production runs of the components, the probability that item j is available depends upon demand in periods

$$t+T_i+L_i-T_j-L_j,\ldots,t+L_i+T_i-T_1-1; \text{ and}$$

(3) The probability that additional finished product (assembled from the additional components) is actually needed to satisfy a demand is a function of demand in periods

$$t+L_i,\ldots,t+L_i+T_i+L_1-1.$$

Since these time intervals overlap, in general the joint probabilities $p_1,\ldots,p_4$ may have double, triple, and quadruple integral terms.

Consider a system with the following parameters:

$\mathbf{T} = (2,4,6)$

$\mathbf{L} = (1,4,1)$

Observe that items 2 and 3 are delivered to the assembly stage simultaneously every 12 periods. Thus, in 2/3 of the item 2 production runs, we would expect an additional unit of item 2 safety stock to find a mate available with certainty. In the remaining 1/3 of the production runs, availability of the mate depends upon the safety stock provision for item 3 and the demands which actually occur. Let us consider adding to the base stock level of component 2. We will examine an instance in which components both components are delivered to the assembly stage simultaneously (i.e., where component mating is an issue). Suppose t = 1. Then the time intervals of concern are (i) 1 through 6, (ii) 2 through 6, and (iii) 5 through 9. Let W be the random variable for demand in period 1, X for demand in periods 2 through 4, Y for demand in periods 5 and 6, and Z for demand in periods 7 through 9. Thus, we need to make some statements about (1) W + X + Y, (2) X + Y, and (3) Y + Z. We can express the desired probability as

$$P(W + X + Y > \mu_W + \mu_X + \mu_Y + k_j \sqrt{T_j + L_j}\ \sigma$$

$$\text{and}$$

$$X + Y < \mu_X + \mu_Y + k_i \sqrt{T_i + L_i}\ \sigma$$

$$\text{and}$$

$$Y + Z > \mu_Y + \mu_Z + k_1 \sqrt{T_1 + L_1}\ \sigma + \min (k_i \sqrt{T_i + L_i}\ \sigma,\ k_j \sqrt{T_j + L_j}\ \sigma))$$

where μ represents the mean of the designated random variable. The right hand side of the last inequality arises because a finished product shortage is the result of a component shortage only when demand over the appropriate time interval exceeds finished product safety stock plus the lesser of the two component safety stock quantities.

Transforming W, X, Y, and Z into standard normal variates W', X', Y', and Z', respectively, after some algebra this becomes

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta(\sqrt{n_2}\, x + \sqrt{n_3}\, y < k_i \sqrt{T_i + L_i})$$
$$\cdot\ P[W' > (k_j \sqrt{T_j + L_j} - \sqrt{n_2}\, x - \sqrt{n_3}\, y)/\sqrt{n_1}]$$
$$\cdot\ P[Z' > (k_1 \sqrt{T_1 + L_1} + \min (k_i \sqrt{T_i + L_i},\ k_j \sqrt{T_j + L_j} - \sqrt{n_3}\, y))/\sqrt{n_4}]$$
$$\cdot\ \phi(y)\ \phi(x)\ dy\ dx$$

where $\delta = 1$ if the condition is true and 0 otherwise, n_1 denotes the number of periods represented by W (= 1), n_2 the number of periods represented by X (= 3),

n_3 the number of periods represented by Y (= 2), n_4 the number of periods represented by Z (= 3), and $\phi(\cdot)$ is the standard normal density. This expression covers the situations where mating of components is an issue (which occurs 1/3 of the time). In the remaining situations, mating is not an issue. Thus, component j would not need to be considered and the expression is much simpler. To find p_4, we take the the expression above, multiplied by 1/3, and add the appropriate probability when mating is not of concern, multiplied by 2/3.

Each of the p_ns can similarly be represented in closed form (for details see Yano, 1981) and can be computed using standard numerical integration techniques.