Procurement of Common Components in a Stochastic Environment

Ramesh Srinivasan*
Rangarajan Jayaraman*
James A. Rappold†
Robin O. Roundy‡
Sridhar Tayur§

(November 1992; Revised September 1998)

Abstract

We consider a multi-period, two level product-component configuration where many of the components are common to several products. The demands for the products in any period are stochastic and independent. Our objective is to determine the component quantities that are to be ordered every period that satisfy a pre-specified service level placed jointly on the products. We consider three approximations and compare the quality of solutions obtained. Our focus is on quickly solving medium to large scale problems. Through a large simulation study, we demonstrate the performance and behavior of our algorithms as well as provide insights into the benefits of commonality for complex product structures.

KEY WORDS: COMMONALITY, INVENTORY, SERVICE CONSTRAINT, APPROXIMATIONS.

*IBM T.J. Watson Research Center, Yorktown Heights, NY 10598
†University of Wisconsin, Madison, WI 53706
‡Cornell University, Ithaca, NY 14850
§Carnegie Mellon University, Pittsburgh, PA 15213

1 Introduction

The management of component inventories in assemble-to-order environments poses several challenges. The systems we are concerned with are characterized by long procurement lead times for components which are subsequently used in the assembly of several different finished end products. The time required for assembly is negligible compared to the long procurement lead times.

We assume the system operates in the following manner. On a periodic basis, the procurement of components is decided and orders are placed several months in advance of their consumption based on forecasts of stochastic demand for end products. These components remain in a component inventory until they are required for assembly. Firm customer orders arrive two to five days before their due date and are satisfied in a make-to-order (or assemble-to-order) fashion. It is assumed that no finished goods inventory is maintained. Once the demand is known, the allocation of components to assembled end products is made over a two day planning horizon. The assembly operation has sufficient capacity to fill customer orders, and thus customer service depends solely on the availability of components. We further assume that there are no fixed ordering costs or fixed assembly setup costs which affect inventory and production decisions.

The system described is common in the computer industry and in other lean manufacturing environments where the product assembly times are short and the component procurement times are long. For example, in electronic card assembly and test plants, the time for assembly and test operations is typically one to two weeks, while the time for procuring high-technology integrated circuits is typically two to three months. In computer assembly plants, the time for assembly and test is typically one or two days, while the time for procuring high-technology subassemblies, such as hard disks, may be longer than one month. In these examples, the majority of the value of the product is in its componentry, and thus warrants careful planning.

The demand for the end products is stochastic and highly non-stationary due to the short customer visibility horizons of two to five days. In most cases, these horizons have been shortened as a result of ongoing Just-In-Time (JIT) efforts.

Due to the costs associated with maintaining inventories and to the risk of obsolescence, it is desirable to keep component inventories as low as possible. As depicted in Figure 1, components (shown as circles) are combined in different quantities (shown on arcs) according to a Bill-of-Material (BOM) structure to produce a wide variety of end products (shown as squares). Typically, there is a high degree of commonality of components among the end products. The demand for components may be translated from the end product demand through the BOM structure. Thus, component demand arrives in sets and coordinated planning is essential to achieve a high level of service at the end product level. For an end product, all required components must be ready before assembly begins. Furthermore, because it is difficult to disassemble a product once it is assembled in practice, we will assume that a component is consumed once it has been used in the assembly of an end product.

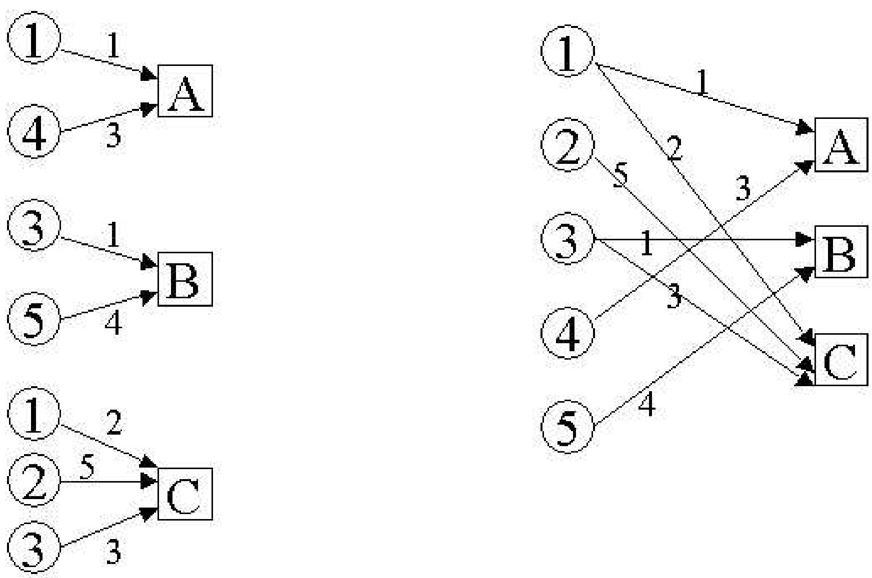

Figure 1: A Three-Product, Five-Component Example BOM Structure

Current Materials Requirements Planning (MRP) approaches are unsuitable for controlling operations of this type since these approaches consider only deterministic end product demand. Ad-hoc procedures to deal with the randomness of demands have lead to poor plant performance since they are unable to effectively model the demand correlation among components. Moreover, the nervousness of the MRP system may lead to excessive component inventories without the benefit of high service levels for the end products. Therefore, an alternative planning method for dependent demand components is necessary which is suitable for very large problems and which considers both the stochastic nature of demand as well as the commonality of the components. The objective of this paper is to provide such a method to determine base stock levels of components in large industrial settings.

The paper is organized as follows. In section 2 we review the relevant literature. In section 3 we formulate the problem of interest, review available methods for exact solutions, and make some observations that lead to tractable approximations. In section 4 we describe three approximations and demonstrate the feasibility of solutions obtained from these approximations to the original problem. In section 5 we provide the solution methods to the approximations. In section 6 some refinements to the solutions are proposed. In section 7 we solve a real-sized problem as well as other large problems to provide insight into the effectiveness of our approach and into the relationships between commonality and inventory costs. In section 8 we offer concluding remarks. The Appendix contains technical details of one of the approximations.

2 Literature Review

Research thus far in this area has dealt with simple product structures, stationary demands and small numbers of products and components, and is restricted to a single period (Baker (1985), Collier (1981), Collier (1982), Gerchak and Henig (1986)). Also see McClain, et al. (1984) and Gerchak and Henig (1989). These papers have provided insight into the trade-offs between standardization and price, proofs of the benefits of coordinated component stocking, and solution methods for special demand distributions and for small sized problems. Baker, Magazine, and Nuttle (1986) consider a two-product, three-component system whose demands are uniformly distributed. Each product requires one unique component and one common component. All three of the components have the same price and one of each is used to make the respective products. The authors show that planning for the common component in a coordinated manner reduces the overall inventory costs. Gerchak, Magazine, and Gamble (1988) extended this work further by considering a more general product structure in which n products have a general joint probability distribution describing demand. Each product requires one unique component with price p_i, and one common component which is common to all n products. The benefit of commonality was demonstrated but no algorithm or procedure to determine the component amounts was derived. The more recent work of Eynan and Rosenblatt (1991) considers two products whose demands are uniformly distributed, with common components costing more than the components they replace. They provide conditions when it is cost effective to replace unique components with more expensive common components. This work is also restricted to the case where one of each component is required for assembly. A different approach is taken in Kannan, Mount, and Tayur (1995) where a randomized algorithm is developed that comes within ϵ of the optimal cost with a high probability in polynomial time for the general problem: n final end products with m components where u_{ij} is the number of component i used in product j, and the end product demands have a joint probability distribution that is log-concave. Problems with up to 10 components can be solved using this methodology. In Swaminathan and Tayur (1998), a multi-period model with backorder costs rather than service level constraints is developed and analyzed. The need for theoretical exactness in the solution in previous papers has precluded developing implementable approximations for industrial size problems (over 50 end products and over 100 components), which is the main motivation for our work.

3 Problem Formulation

In the literature, standard discrete time inventory models use backorder costs to penalize inventory shortages and holding costs to penalize excess inventory. The focus in that methodology is to prove the optimality of and computation of base stock policies, or (s, S) policies in the presence of fixed ordering costs. These policies assume that any desired amount of inventory can be procured within a predetermined lead time. However, the situation we are faced with here is a two-stage decision process.

First, one has to select the amounts of components to be delivered well in advance. Once these quantities have been determined and the orders placed, the amounts that will be delivered are fixed and dynamic ordering after the occurrence of demands is not possible. Second, after the demand is realized, an allocation decision is made in which components are allocated to the assembly of end products. These decisions may be determined with a linear program subject to the following constraints. The assembly operation has a service level constraint in which they must meet all backorders from the previous period plus a pre-specified fraction of the current demand for all products with a given probability. These probabilities may be stipulated by management. This type of aggregate service specification is preferred by the managers we interviewed. Only one probability needs to be specified while the service levels of the different products can be optionally differentiated through pre-specified fractions.

In this setting, the decisions of primary interest are the amounts of each component to be ordered in each time period of the planning horizon. The goal is to meet the service level for the end products at minimum cost. The decision model is re-run in each time period on a rolling horizon basis. In the development of the model, we will use the following notation:

T periods in the planning horizon;

L procurement lead time for all components;

n_p number of products;

n_c number of components;

t index for time periods;

j index for products;

i index for components;

p_i price of one unit of component i;

p^j price of one unit of product j;

h_i holding cost per unit per period of component i;

u_{ij} number of components i needed in one unit of product j;

X_t probability of satisfying outstanding backorders and the desired fraction of the current demand for all products, in period t. The service level target as specified by management;

y^{jt} management-specified fraction of demand of product j in period t that should be met in period t;

q^{jt} cumulative amount of product j satisfying demand through period t;

q_{it} cumulative amount of component i delivered from supplier through period t (decision variable);

d^{jt} cumulative demand for product j through period t (random variable);

d_{it} cumulative demand for component i through time t (random variable);

s^{jt} cumulative amount of product j produced through period t (decision variable);

K_t a variable taking a value of 1 if the aggregate service objective is met in period t, and taking a value less than one otherwise;

$\hat{d}^{jt}$ cumulative required amount of product j necessary to satisfy the pre-determined fraction of end product demand, y^{jt}, through time t;

$\hat{d}_{it}$ cumulative required amount of component i necessary to satisfy the pre-determined fraction of end product demands, y^{jt}, for all j through time t.

We assume that the end product demand in period t, $(d^{jt} - d^{j,t-1})$, is observed at the beginning of time period t. Thereafter, s^{jt} and $q_{i,t+L}$ are chosen for all j and all i, where $L \geq 1$ is the procurement lead time for all components. The case where lead times are component-specific will not be covered in this paper. The variables must be non-anticipatory. That is, q_{it} is a function of d^{js} for all $s \leq t - L$, but not for $s > t - L$. Also, s^{jt} is a function of d^{js} for $s \leq t$, but not for $s > t$.

The cumulative demand for component i through time period t is defined as,

$$d_{it} = \sum_j u_{ij} d^{jt}.$$

Even if the demands for finished products are independent, the component-level demands $(d_{it}, 1 \leq i \leq n_c)$ may be highly correlated. Let $\hat{d}^{jt}$ denote the cumulative total requirement for product j through time t to achieve the service level target, X_t, in period t. That is, we define,

$$\hat{d}^{jt} = d^{j,t-1} + y^{jt} \cdot (d^{jt} - d^{j,t-1}). \tag{1}$$

Similarly, let $\hat{d}_{it}$ denote the cumulative total requirement for component i through time t to achieve the service level constraint in period t. Thus,

$$\hat{d}_{it} = \sum_j u_{ij} \hat{d}^{jt}. \tag{2}$$

The target service level, as specified by management, is expressed as,

$$\Pr\{s^{jt} \geq \hat{d}^{jt} \ \forall j\} \geq X_t. \tag{3}$$

We formulate the component procurement and production problem as follows:

$$(P) \quad \min_{q_{it}} \quad E\left\{\sum_{t=1}^{T}\sum_{i=1}^{n_c} p_i(q_{it} - \sum_{j=1}^{n_p} u_{ij}s^{jt})\right\} \tag{4}$$

$$\text{s.t.} \quad \Pr\{q_{it} \geq \hat{d}_{it} \ \forall i\} \geq X_t \quad \forall t \tag{5}$$

$$s^{jt} \geq (1-K_t)s^{j,t-1} + K_t \cdot \hat{d}^{jt} \quad \forall j \ \forall t \tag{6}$$

$$s^{jt} \leq d^{jt} \quad \forall j \ \forall t \tag{7}$$

$$q_{it} \geq \sum_{j=1}^{n_p} u_{ij} \cdot s^{jt} \quad \forall i \ \forall t \tag{8}$$

$$q_{i,t-1} \leq q_{it} \quad \forall i \ \forall t \tag{9}$$

$$s^{j,t-1} \leq s^{jt} \quad \forall j \ \forall t \tag{10}$$

$$K_t = \min_i \left\{1, \frac{q_{it} - \sum_{j=1}^{n_p} u_{ij}s^{j,t-1}}{\sum_{j=1}^{n_p} u_{ij}(\hat{d}^{jt} - s^{j,t-1})} \ \forall i\right\} \quad \forall t \tag{11}$$

$$q_{it}, s^{jt} \geq 0 \quad \forall i \ \forall j \ \forall t \tag{12}$$

Equation (5) represents the aggregate service level constraint for period t in which sufficient components are required to satisfy the demand backlog and pre-specified fractions (y^{jt}) of the current period's demand for all products with probability X_t. The demand fraction for different products (y^{jt}) may be different depending on marketing or other considerations and may vary with time. Equation (6) states that for each product, if $K_t = 1$, the production of product j in period t should cover the backlog and the management specified fraction of demand. Equation (7) ensures that no end products are assembled in a given period that are not sold immediately. Recall that we assume that in the assembly stage, the capacity is sufficiently large so that there is no need to assemble in advance for future periods. Equation (8) expresses the manner in which component availability constrains production. Equations (9) and (10) state that order quantities and production quantities are non-negative. Equation

(11) requires $K_t = 1$ if possible, but allows K_t to be less than one to make the formulation feasible for any demand realizations and ordering decisions that render $K_t = 1$ infeasible. From equations (7), (8), and (10), both the numerator and denominator of the ratio in (11) are non-negative. In the evaluation of (11), we will adopt the convention $0/0 = \infty$. Under equations (6) and (11), the two versions of the service constraint, (3) and (5), are equivalent. The decision variables q_{it} and s^{jt} are non-negative. The objective function is the expected cumulative inventory investment cost in components over T periods.

The amount of computational effort required to solve stochastic programming problems of this level of complexity – an objective function of the expected value type, linear constraints, correlated demands d_{it} in the chance constraint – is prohibitive. As stated in Wets (1989), the number of points required in a ten dimensional case to compute the objective function value and the probability constraint accurately exceeds 10^{10} points. Even problems with an objective function that is of the type $\sum_{i=1}^{n_c} p^i q^i$ can be solved only for a small number of dimensions. Komaromi (1986) provides a dual method to solve a problem of this type; however, for the case when the random variables are correlated, only a Monte-Carlo type approximation is possible to compute the gradients as well as the probabilistic constraint value. The computational expense is large mainly because the gradient at a point is a an $n - 1$ dimensional conditional distribution. Namely, it is,

$$\frac{dF(y_1, \ldots, y_n)}{dy_i} = F(y_1, \ldots, y_{i-1}, y_{i+1}, \ldots, y_n | y_i) f(y_i).$$

This conditional probability is not separable due to dependence.

Other methods of chance constrained programming that rely on deterministic equivalents require the probabilistic constraints to be separable. Only the normal and multi-variate random variable cases have been successfully solved. See Kolbin (1977) for an extensive survey of methods and applications. For many successful applications of this approach to lower dimensional problems, see Prekopa (1973).

Our setting is such that the demand variables are not necessarily normal or gamma, and the number of dimensions is at least 50. Thus, we turn to approximations to solve (P) which are based on the following useful observations.

1. Setting Production Variables: For every given set of monotone q_{it}, the selection of s^{jt} for a given t is a linear programming problem that can be solved after the demands up to and including period t have been observed.

2. Eliminating Production Variables: If we approximate the objective function by a function of a form that is independent of s^{jt}, then the decisions q_{it} can be made without considering the constraints that contain s^{jt}. Consequently, all constraints in (P) that involve s^{jt} may be dropped.

3. Time Period Separability: If we relax (9), we have separability by time period. The objective function may then be approximated by using component-level demands as,

$$G_t(q_{it}, 1 \leq i \leq n_c) = E\left\{\sum_{i=1}^{n_c} p_i(q_{it} - d_{it})^+\right\}. \tag{13}$$

These observations lead to the following simpler problem:

$$(P_t) \quad \min_{q_{it}} \quad G_t(q_{it}, 1 \leq i \leq n_c) \tag{14}$$

$$\text{s.t.} \quad \Pr\{q_{it} \geq \hat{d}_{it} \ \forall i\} \geq X_t. \tag{15}$$

For each t, (P_t) is solved for the component decisions $(q_{it}, 1 \leq i \leq n_c)$. If (9) fails, then the cumulative component order quantities q_{it} may be adjusted upwards in later time periods. In all of our heuristics, our strategy will be to use the time period separability approximation and to frequently eliminate the production variables.

One standard approach to solving stochastic programming problems such as (P) and (P_t) is to approximate the stochastic demand vector $(\hat{d}_{it}, 1 \leq i \leq n_c)$ by a discrete random variable. As we mentioned earlier, Wets (1989) indicates that 10^{10}

vectors would be required in the case of 10 components. Since our objective is to solve problems having between 50 and 200 components, this approach is impractical for solving (P_t). Our personal experience confirms this conclusion.

On the other hand, any practical solution to (P) or (P_t) will require the evalutation of the left-hand side of the service level constraint (5). Suppose that the stochastic demand vector $(\hat{d}^{jt}, 1 \leq j \leq n_p)$ has a distribution that is uniform within the unit cube. Then the service constraint (3) would be equivalent to finding the volume of a polyhedron. The decision version of this problem (is the volume greater than, say, V) is NP-Hard. See Dyer and Frieze (1988) for details. The best algorithms for this problem are all Monte-Carlo or pseudo-Monte-Carlo algorithms which can be viewed as approximating the stochastic demand vector $(\hat{d}^{jt}, 1 \leq j \leq n_p)$ by a discrete random variable (see Gritzmann and Klee (1994)).

Consequently, whereas Monte-Carlo type algorithms seem hopeless for the direct solution of (P_t), they seem to be a necessary element of any practical solution to (P_t) since they represent the best currently available approach to evaluating the service constraint (3) or (5) accurately. The curse of dimensionality plays a key role here. Since (P_t) is a 50-dimensional optimization problem for a 50 component problem, the problem of evaluating (3) or (5) for a given set of component order quantities can be viewed as the problem of evaluating the expectation of a one-dimensional random variable which takes on values of either 0 or 1.

These observations are crucial in defining the approach that we have taken to the problem. Our strategy is to develop analytical approaches for computing component order quantities and to use Monte-Carlo techniques to adapt the order quantities to satisfy the service constraint.

4 Approximations to (P)

In this section we discuss heuristics and present three tractable approximations for solving (P).

After several discussions, managers and forecasters at an IBM facility proposed the following model as a vehicle for quantifying the variance of the demand for the finished products. Demand is modeled as a symmetric piecewise-linear distribution, called a "trapezoidal distribution", as illustrated in Figure 2. For this distribution, the value of the mean of the density is three times as large as the value at the extremes. Managers specify the mean and the range of the demand distribution. In most of the computations, a Normal approximation with the same mean and variance was used. Based on discussions with the managers, a Normal approximation to the distribution would be acceptable.

(I) THE ORDER BY PRODUCT APPROXIMATION (OBP): In this approach, appropriate stock levels, q^{jt}, for each product j are computed separately. The product stock levels are then translated into component stock levels, q_{it}. This approach ignores the benefits of risk pooling that arise when a component can be used in several different products. Using this approach and by relaxing (10), problem (P) may be expressed as:

$$\begin{aligned} \min_{q^{jt}} \quad & E\left\{\sum_{t=1}^{T}\sum_{j=1}^{n_p} p^j(q^{jt}-s^{jt})\right\} \\ \text{s.t.} \quad & \Pr\{q^{jt} \geq \hat{d}^{jt}\ \forall j\} \geq X_t, \quad \forall t, \\ & s^{jt} = min(q^{jt}, d^{jt}), \end{aligned}$$

solved separately for each period t. This may be rewritten as,

$$\begin{aligned} (P-OBP_t) \quad \min_{q^{jt}} \quad & E\left\{\sum_{t=1}^{T}\sum_{j=1}^{n_p} p^j(q^{jt}-d^{jt})^+\right\} \\ \text{s.t.} \quad & \Pr\{q^{jt} \geq \hat{d}^{jt}\ \forall j\} \geq X_t, \quad \forall t, \end{aligned} \tag{16}$$

where $(a)^+ = max(0, a)$. Since the demands for the final product are assumed to be independent, the joint probability term in (16) separates into the product of individual probability terms. Let $\hat{F}^{jt}$ and $\hat{f}^{jt}$ be the cumulative distribution and density functions of $\hat{d}^{jt}$, respectively. Taking logarithms of (18) we obtain,

$$\sum_{j=1}^{n_p} \log \hat{F}^{jt}(q^{jt}) - \log X_t \geq 0. \tag{17}$$

$(P - OBP_t)$ is a convex program if and only if $\hat{f}^{jt}/\hat{F}^{jt}$ is decreasing. This property is satisfied for the Normal distribution among others. Let F^{jt} be the cumulative distribution function of d^{jt}. The Kuhn-Tucker conditions for $(P - OBP_t)$ are:

$$\begin{aligned} \frac{\hat{f}^{jt}(q^{jt})}{\hat{F}^{jt}(q^{jt})} &= \frac{p^j}{\lambda_t} F^{jt}(q^{jt}) \qquad \forall j = 1, \ldots, n_p, \\ \lambda_t\Big(\sum_{j=1}^{n_p} \log \hat{F}^{jt}(q^{jt}) - \log X_t\Big) &= 0. \end{aligned}$$

These equations are easily solved by searching over λ_t. For a given λ_t, the order quantities are obtained from,

$$q_{it} = \sum_i u_{ij} q^{jt}.$$

The following theorem implies that if commonality is not exploited, the result is a feasible solution that has higher inventory costs. This is the general form of the results obtained by Baker, Magazine, and Nuttle (1986) and Gerchak, Magazine, and Gamble (1988). The proof is straightforward and is omitted.

Theorem 1 *If (9) and (10) are relaxed, then any solution that is feasible for $(P - OBP_t)$ is feasible for (P).*

(II) The Order by Component Approximation (OBC): In this approximation, the product structure is suppressed; however, its effects are partially captured in the probability term and in the objective function. The objective function is approximated using component-level demands as,

$$G_t(q_{it}, 1 \leq i \leq n_c) = \; E\left\{\sum_{i=1}^{n_c} p_i (q_{it} - d_{it})^+\right\}. \tag{18}$$

For each time period t, we formulate the problem,

$$(P - OBC_t) \quad \min_{q_{it}} \quad G_t(q_{it}, 1 \leq i \leq n_c) \tag{19}$$

$$\text{s.t.} \quad \Pr\{q_{it} \geq \hat{d}_{it} \ \forall i\} \geq X_t. \tag{20}$$

Recall that the joint probability term in (5) and its gradient are hard to obtain for dependent random variables. Thus, we will approximate the probability term by a product of terms (each term from a different component) which results in a tractable problem. Fortunately, the solution of this approximation is feasible for $(P - OBC_t)$. This can be shown using the following result with $Y_i = -\hat{d}_{it}$.

Theorem 2 *(Tong (1980)) If random variables Y_i, $i = 1, \ldots, N$ are associated, then* $\Pr\{Y_i \geq x_i, \forall i\} \geq \prod_{i=1}^{N} \Pr\{Y_i \geq x_i\}$.

After taking logarithms, (20) becomes,

$$\sum_{i=1}^{n_c} \log \hat{F}_{it}(q_{it}) - \log X_t \quad \geq \quad 0, \tag{21}$$

where $\hat{F}_{it}$ is the marginal distribution function of $\hat{d}_{it}$.

The LHS of (21) is concave if $\log \bar{\hat{F}}_i(q_{it})$ is concave, *i.e.,* if $\bar{\hat{F}}_i(q_{it})$ is log-concave, where $\bar{\hat{F}}_i(q_{it}) = 1 - \hat{F}_i(q_{it})$. A class of distributions that satisfy this is the PF_2 class, the Polya-frequency functions of order two. Example distributions in the PF_2 class are the Normal distribution and the Erlang distribution. Thus, the lower bound approximation to $(P - OBC_t)$ is a convex program if we assume log-concave distributions of end product demands, since positive linear combinations of PF_2 random variables are in PF_2 (Keilson and Sumita (1982)).

The Kuhn-Tucker conditions for $(P - OBC_t)$ are similar to those in the Order by Product approximation, as is the solution method.

The Order by Product (OBP) and Order by Component (OBC) strategies can be thought of as general approaches that are relatively easy to adopt. Both of the algorithms described above require optimization across families of components or products, which may or may not be done properly in practice. Both are based on conservative approximations so both approaches will systematically use too much inventory and are likely to exceed the service level target.

(III) THE CLARK APPROXIMATION (CL): The third approximation to (P) approximates the probability in (5) and its derivatives with respect to q_{it} without assuming independence among the component demands. This is done by assuming the demands are Normally distributed and by using results from Clark (1960). The required computations are messy but efficient, and the details are given in the Appendix. We then minimize (19) subject to our approximation to (5). This is done by dualizing (5) and using subgradient optimization to minimize the Lagrangian.

Unfortunately, no analogue to Theorems 1 and 2 exists for the Clark approximation. However in all of our computational tests, the Clark approximation has been overly optimistic in its expected service levels for a given stocking decision.

5 Extending the Approximations

We have discussed three tractable approximations to (P), all of which allow us to compute the component order quantities, q_{it}. The actual service levels attained by these orders are typically very different from the target service level, X_t. Both the OBP and OBC approximations under-estimate the probability of achieving the service level target, so the constrained optimization algorithms create an excess of inventory. The situation is reversed with the CL approximation, which over-estimates expected service levels and consequently results in inadequate inventory levels.

Our approach is based on seeking opportunities for adjusting inventory levels, computed either by OBC or CL, while meeting the target service level. The first step simultaneously scales q_{it}, for all i. It is accomplished using either of two methods, *direct scaling* or λ-*scaling*. After this reduction, there is a possibility that certain component inventory levels are in excess with respect to the inventories of other components. These component inventories are reduced further while still satisfying the service requirement. We refer to this second step as *slack reduction.*

To evaluate the service level in (5) we resort to simulation. We generate a random sample of N product demand vectors $(\hat{d}^{jt}, 1 \leq j \leq n_p)$, and use (2) to compute component demand vectors $(\hat{d}_{it}, 1 \leq i \leq n_c)$. Given a set of component order quantities, we perform *direct scaling* by multiplying the solution $(q_{it}, 1 \leq i \leq n_c)$ by a scalar α. This α is selected so that $\alpha q_{it} \geq \hat{d}_{it}\ \forall i$ for $\lceil X_t N \rceil$ of the N demand vectors in the random sample, where $\lceil x \rceil$ is the ceiling of x. This can be done very efficiently.

The λ-*scaling* is done as follows. Note that both the OBC approximation and the CL approximation use Lagrangian relaxation to solve an optimization problem that has a single constraint. This is done by writing the constraint in the form $f(q_{it}, 1 \leq i \leq n_c) \leq 0$, and by forming the Lagrangian by adding $\lambda \cdot f(q_{it}, 1 \leq i \leq n_c)$ to the cost. For each non-negative value of λ, the Lagrangian can be minimized, producing a solution $(q_{it}, 1 \leq i \leq n_c)$. The goal is to select λ so that $\lambda \cdot f(q_{it}, 1 \leq i \leq n_c) = 0$. To perform λ-*scaling*, we alter this approach by selecting λ so that the corresponding solution satisfies $q_{it} \geq \hat{d}_{it} \forall i$, for $\lceil X_t N \rceil$ of the N demand vectors in the random sample.

After performing either *direct scaling* or λ-*scaling*, the *slack reduction* step is performed as follows. Returning from component space (n_c dimensions) to product space (n_p dimensions), we let,

$$D = \left\{ (\hat{d}^{jt}, 1 \leq j \leq n_p) : q_{it} \geq \sum_{j=1}^{n_p} u_{ij} \hat{d}^{jt} \quad \forall i \right\},$$

be the set of product demand vectors that can be met from component stock. Component stock levels computed by the methods described above usually have components i for which $\{(\hat{d}^{jt}, 1 \leq j \leq n_p) : q_{it} = \sum_{j=1}^{n_p} u_{ij} \hat{d}^{jt}\} \cap D = \phi$. This implies that q_{it} can be further reduced without altering D or the service level. Although this may be accomplished using linear programming, we adopt the simpler approach of reducing

q_{it} to the quantity,

$$\max\left\{\sum_{j=1}^{n_p} u_{ij}\hat{d}^{jt} : (\hat{d}^{jt}, 1 \leq j \leq n_p) \in D\right\}.$$

We recognize that there is an inherent statistical bias in the scaling and slack reduction techniques that we have described. In our tests, a sample size of $N = 2,500$ consistently proved to be large enough to make this bias inconsequential.

6 Computational Experiments

In this section, we present our results based on a large simulation study designed to examine the performance and behavior of the different algorithms in a wide variety of scenarios. The algorithms that will be examined are abbreviated and listed as:

- OBP is the ORDER BY PRODUCT APPROXIMATION in which component commonality is ignored;
- OBC is the ORDER BY COMPONENT APPROXIMATION;
- OBC-λ is the ORDER BY COMPONENT APPROXIMATION with λ-*scaling* and *slack reduction*;
- OBC-D is the ORDER BY COMPONENT APPROXIMATION with *direct scaling* and *slack reduction*;
- CL is the CLARK APPROXIMATION;
- CL-λ is the CLARK APPROXIMATION with λ-*scaling* and *slack reduction*;
- CL-D is the CLARK APPROXIMATION with *direct scaling* and *slack reduction.*

Based on management preference, a "trapezoidal" probability distribution is used for the simulation of end product demands. For this distribution, the relative height of the density at the mean is three times the height at the minimum and maximum values. We define the *radius* of this distribution to be the difference between the maximum and minimum values. An example of this distribution is shown in Figure 2.

This section is divided into four subsections as follows. First, we describe the experiment factors of interest which will define a particular system scenario. Second, the experiments are described. Third, the results of the study are presented. Fourth, insights are presented into the tradeoff between component commonality and product simplification.

6.1 Factors Under Study

In attempting to understand the performance and behavior of the different algorithms, we will alter several system attributes. Descriptions of the the factors of interest are listed below and the specific factor levels under study are shown in Table 1.

- **Number of Components** is the number of different components in the system;
- **Number of Products** is the number of different end products which are demanded;
- **Product-to-Component Ratio** measures the ratio of the number of products to the number of components. While the problem size may vary, a Product-to-Component Ratio of 60/42 refers to a system with 60 products and 42 components for a total of $60 \times 42 = 2520$ possible arcs. When varying the Product-to-Component Ratio (Experiment 3), the total number of possible arcs was fixed at approximately 2500;
- **BOM ratio** measures the level of component commonality in the system. It is defined as the average number of different components per end product divided

by the total number of components. For example, for a system with 100 products and 50 components, a BOM ratio of $8/50 = 0.16$ implies that, on average, an end product is assembled from 8 different components;

- **Demand Mean** is an interval representing the possible mean demand values for individual end products. All mean demand values for end products were generated randomly from this interval. For example, an interval of (900, 1100) implies that the mean demand value for each end product is generated randomly from a uniform distribution between 900 and 1100;

- **Demand Variability** is the radius of an interval with center 1000 representing the trapezoidal demand distribution for all end products, as shown in Figure 2. An interval of (800,1200) implies that the demand distribution radius is 400 and has minimum 800, mean 1000, and maximum 1200 for all end products;

- **Target Service Level**, X_t, is the overall target service level across all products, as specified by management;

- **Product Service Levels**, y^{jt}, are the product-specific target service levels.

6.2 Experiments

We first compare the performance of the different algorithms on a base case (**Experiment 1**) that is representative in size and in structure of a current bill of materials (BOM) structure in an IBM plant. In all of the experiments, a single time period is considered and the usage quantity for all components is precisely one unit for all end products. The factor levels used in each experiment are summarized in Table 2.

Experiment 1 considers a system in which 50 components are used in the assembly of 50 products. The BOM ratio is set to 8/50 =0.16. While the average number of components used per end product is 8, the actual number of components used for each

Factor Under Study	**Levels**
Number of Components	25, 50, 75, 100, 125
Number of Products	25, 50, 75, 100, 125
Product-to-Component Ratio	100/25, 60/42, 50/50, 42/60
BOM Ratio	8/50, 12/50, 16/50, 20/50, 24/50
Demand Mean	$1000, \pm 100, \pm 200, \pm 300, \pm 400$
Demand Variability	$1000, \pm 170, \pm 340, \pm 500, \pm 670, \pm 840$
Target Overall Service X_t	82%, 86%, 90%, 94%, 98%
Product Service Level y^{jt}	95% for all j

Table 1: Factor Levels Under Study

end product is generated randomly from a discrete uniform distribution between 1 and 15, so that the average number of components per end product is 8. The product demand is generated from a trapezoidal distribution. For this experiment, for all end products, the mean demand is 1000 and the radius of the distribution is 500. Component costs are randomly generated from a uniform distribution between $0.00001 and $1,000. The overall target service level, X_t, is set at 90% and the product-specific service levels, y^{jt} are set at 95% for all products.

In **Experiments 2** through **7**, all data other than the factor being varied (see Table 2) is the same as in **Experiment 1**. In **Experiment 2**, we vary the problem size by altering the number of components and products in the systems. The various problem sizes are represented as (products $\times$ components). In **Experiment 3**, the product-to-component ratio is varied. In **Experiment 4**, the BOM ratio is varied. In **Experiment 5**, the end product demand means are varied across end products. An interval such as (1000, 1000) implies that all end product mean demand are equal to 1000. An interval such as (900, 1100) implies that end product demand means lie between 900 and 1100, and are not equal across end products. In **Experiment 6**,

	Experiment Number						
Factor	Exp 1	Exp 2	Exp 3	Exp 4	Exp 5	Exp 6	Exp 7
Components, n_c	50	*	*	50	50	50	50
Products, n_p	50	*	*	50	50	50	50
Number of arcs in BOM	2500	*	2500	2500	2500	2500	2500
Product-to-Component Ratio	50/50	Constant	*	50/50	50/50	50/50	50/50
BOM Ratio	8/50	8/50	8/50	*	8/50	8/50	8/50
Demand Mean	1000	1000	1000	1000	*	1000	1000
Demand Min/Max Range	±500	±500	±500	±500	±500	*	±500
Target Overall Service, X_t	90%	90%	90%	90%	90%	90%	*
Target Product Service, y^{jt}	95%	95%	95%	95%	95%	95%	95%

Table 2: Factor levels used for each experiment. An asterisk indicates which factors were varied according to the levels shown in Table 1.

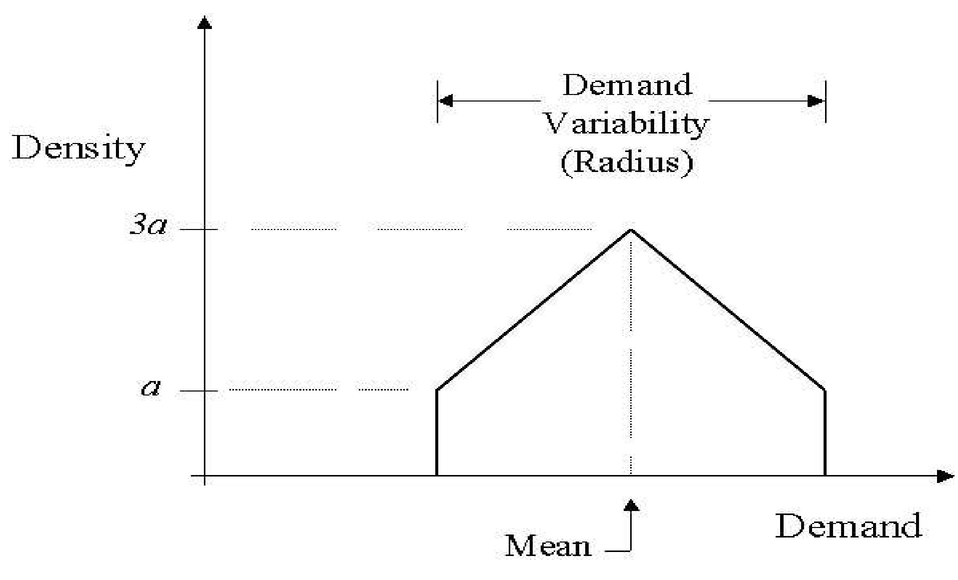

Figure 2: Example of a trapezoidal distribution.

the coefficient of variation of demand is varied. Note that the coefficient of variation is not directly changed. Instead, the radius of the trapezoidal distribution is changed, while fixing the mean. In **Experiment 7**, the overall target service level, X_t, is varied while the product-specific service levels, y^{jt} are fixed.

For all experiments, 10 random instances of the problem were generated. Tables 4 through 10 show the average cost and service level results for each experiment averaged over the instances. In all of the experiments, the end product demand rates were adjusted to keep the component demand rates constant. For example, as the problem size was varied from 50×50 to 100×100 with a fixed BOM ratio, the number of products doubled. In order to keep the component demand rates identical across problem sizes, the product demand rates were adjusted. The cost figure represents the expected excess inventory cost associated with meeting a single period's demand, as computed in (18). That is, the expected excess inventory is the expected amount by which the component inventory exceeds the component demands.

6.3 Experiment Results

SERVICE. Although a target service level is specified in all algorithms, different methods under-achieved or over-achieved this target service level. As expected, the OBC approximation over-achieved the target service level. The service level achieved by the Clark approximation was consistently lower than the target service level. Recall that by design, the algorithms based on λ-*scaling* and *direct scaling* will achieve the overall target service level. For example, the OBC-λ approximation will achieve the specified service level provided that a large number of iterations are done in the Monte-Carlo scaling procedures. As shown in Table 3, the computational time taken by all these methods is very reasonable even for large size problems. Therefore, achieving the specified service level with a high level of confidence is computationally viable.

COST. The expected excess inventory is computed analytically assuming that the component demands are Normally distributed. The OBC method appears to be a good and fast initial solution. OBC-λ and OBC-D are improvements over OBC and were generally found to yield close results. CL-D was found to be the most reasonable subgradient-based optimization method among the set of methods CL, CL-D, and CL-λ. The method based on OBP, which ignores component commonality, resulted in higher inventory costs and a higher than specified service level due to increased inventory levels. This is due to the fact that by ignoring commonality, whenever there is at least one unique component in each end product, the OBP method will result in higher inventory levels for the same service level.

As expected, the expected excess inventory increases as the variability in end product demand increases. As the overall target service level increases, the expected excess inventory appears to increase at a faster rate. As the level of component commonality decreases, the expected excess inventory increases.

Overall, the OBC-λ method consistently provided the lowest cost answer among the methods in all of the experiments for a specific level of service. We found the Clark approximation to be sensitive to starting solutions. We also believe that repeated invocations of Clark's approximation on a large number of random variables used in our subgradient method degrades the solution.

Examples of the computation time required for the approaches to solve the problems on an Sparc 10 workstation are listed in Table 3. These time estimates are for single period problems with 2,500 iterations in the Monte-Carlo scaling procedures employed in the OBC-λ and CL-D methods. For multi-period problems, the computation times should be multiplied by the number of time periods, since the multi-period problem is decomposed into several single period problems. These time estimates are extremely reasonable considering the fact that we are solving very large planning problems.

Method	Computation Time (Minutes)
OBC	0.05
OBC-λ	0.62
OBC-D	0.48
CL	1.20
CL-λ	11.70
CL-D	2.30
OBP	0.13

Table 3: Computation Times on a Sparc 10 Workstation

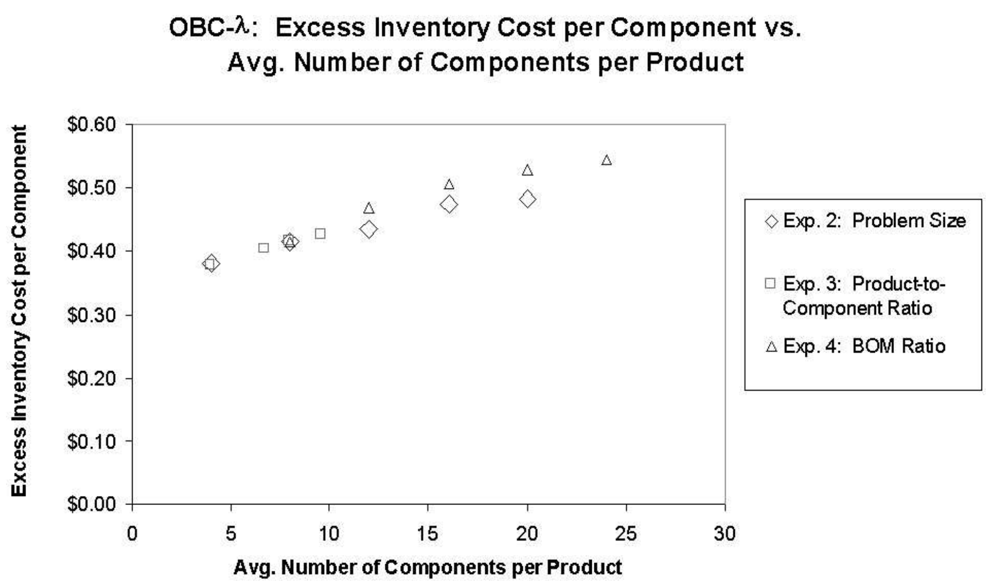

Figure 3: Impact of Factors on Normalized Inventory Cost

6.4 Experiment Insights

In the component planning problem, the imbedded BOM structure dictates the following tradeoff in the control of the system. At the product level, all components in their proper quantities are required before beginning the assembly of an end product. Therefore, for each product, it is desirable to require as few components as possible. We call this the *need for assembly* which is measured as the average number of components per product. At the component level, as the number of common components increases, so do the advantages of *component risk pooling*. Therefore, for each component, it is desirable to be used in as many products as possible. We measure this as the average number of products per component. The central question is how the tradeoff between *need for assembly* and *component risk pooling* is affected by various factors, and which has a significant impact on cost and service.

In **Experiments 2** and **4**, varying the factor of interest creates opposing forces on the system cost. As the problem size increases from 25×25 to 125×125, the average number of components per product increases from 4 to 20 in order to keep the BOM ratio constant. The increased number of components per product requires additional component stock and hence increases the expected cost. On the other hand, at the component level, as the problem size increases, the number of products per component increases from 4 to 20 (since the number of products equals the number of components) in order to keep the BOM ratio constant. Although this creates an opportunity for risk pooling, the benefits are slight and are insufficient to offset the additional cost associated with the *need for assembly*. As a result, the system cost rises. Similarly in Experiment 4, varying the BOM ratio causes the same effect.

In **Experiment 3**, varying the product-to-component ratio while fixing the BOM ratio and the number of arcs in the BOM structure benefits both the *need for assembly* and *component risk pooling*. At the product level, as the number of components decreases, so does the average number of components per product. Thus, the expected excess cost decreases. At the component level, as the number of components decreases, the number of products increases to maintain the same number of arcs. Consequently, the average number of products per component increases and the benefits of risk pooling further lower the expected cost.

In order to make the costs comparable, we normalize the costs in terms of cost per component. These normalized costs are plotted in Figure 3.

Based on these observations, the cost impact from the *need for assembly* is substantially greater than the benefits of risk pooling. Alternatively, having fewer components per product seems to have more impact on cost than having more products per component, for a given level of service.

7 Conclusion

The assembly of components into products is prevalent across a wide spectrum of industries, such as in computer manufacturing and in automobile assembly. Furthermore, component commonality is a growing trend as design for manufacture and assembly philosophies are combined with product variety proliferation.

In this research, we have provided very effective methods for solving medium to large scale component procurement planning problems both in terms of quality of solution and in terms of the computational time required. We have demonstrated that an approach ignoring commonality is vastly inferior to the other methods developed. The quantitative effects of degree of commonality, target service level, and the degree of variability of demand on inventory levels are amply demonstrated in this work through the methods and computational experiments presented. Not only can these methods be used to make procurement decisions and obtain trade-offs between inventory and service, these methods may also be effectively used in evaluating the expected impact of component commonality on inventory investments during early product design stages as well as of ongoing re-engineering efforts. Further research is needed in procurement planning problems in multi-stage production systems where procurement lead times are distinct and random, and where end product demands are correlated among themselves as well as across time periods.

Appendix

In this section, we provide the computational details required by the Clark approximation. Specifically, we describe the details for the approximation of the service level constraint (5) and the gradient of the approximate service level with respect to q_i, $i = 1, \ldots, n_c$. This computation is done in two passes - a forward pass and a backward pass.

In the forward pass of the computation, let $Y_i = \hat{d}_i - q_i$, $i = 1, \ldots, n_c$. We will define and compute,

$$\bar{F}(q_1, \ldots, q_{n_c}) \cong 1 - \Pr\{Y_i \leq 0 \ \ \forall i\} = \Pr\left\{\bigvee_{i=1}^{n_c} Y_i \geq 0\right\}.$$

Note that the Y_i's are correlated. To compute the maximum of these n_c correlated random variables, we let $Z_1 = Y_1 \vee Y_2$. Let μ_1 and σ_1^2 be the first two moments of Z_1, respectively. Also, let $\sigma_{1,m}$ for $m = 3, ..., n_c$ be the correlation between Z_1 and $Y_3, \ldots, Y_{n_c}$. We then assume that $(Z_1, Y_3, ..., Y_{n_c})$ has a multivariate normal distribution with the above parameters. We then find the first two moments of $Z_2 = Z_1 \vee Y_3$, and the correlation between Z_2 and Y_k, $k \geq 4$. This procedure is repeated until the first two moments of Z_{n_c-1}, namely μ_{n_c-1} and $\sigma^2_{n_c-1}$, are obtained. The intermediate steps yield μ_i and σ_i^2 for $i = 2, \ldots, n_c-2$. $\bar{F}(q_1, \ldots, q_{n_c}) = \Phi(\frac{-\mu_{n_c-1}}{\sigma_{n_c-1}})$ provides the service level approximation obtained for this set of $(q_1, \ldots, q_{n_c})$. To complete the forward pass, the following formulas are used (see Clark (1960)):

$$\begin{aligned}
(a^i)^2 &= \sigma_{i-1}^2 + \sigma_{i+1}^2 - 2\sigma_{i-1,i+1}, \\
\alpha^i &= \frac{(\mu_{i-1} - \mu_{i+1})}{a^i}, \\
\mu_i &= \mu_{i-1}\Phi(\alpha^i) + \mu_{i+1}\Phi(-\alpha^i) + a^i\phi(\alpha^i),
\end{aligned}$$

$$
\begin{aligned}
\sigma_i^2 + \mu_i^2 &= (\sigma_{i-1}^2 + \mu_{i-1}^2)\Phi(\alpha^i) + (\sigma_{i+1}^2 + \mu_{i+1}^2)\Phi(-\alpha^i) \\
&\quad +(\mu_{i-1} + \mu_{i+1})\phi(\alpha^i), \\
\sigma_{i,m} &= \sigma_{i-1,m}\Phi(\alpha^i) + \sigma_{i+1,m}\Phi(\alpha^i), \quad \text{for} \;\; m = i+2, i+3, ..., n_c.
\end{aligned}
$$

During the backward pass, the gradients are computed with respect to q_i using $\Phi(\alpha^i)$, $\Phi(-\alpha^i)$, $\phi(\alpha^i)$ and its derivative $\phi(\alpha^i)$ for $i = 1, \ldots, n_c-1$, which are computed and stored in the forward pass.

Since $\bar{F}(q_1, ..., q_{n_c}) = \Phi(\frac{-\mu_{n_c-1}}{\sigma_{n_c-1}})$, we have,

$$
\begin{aligned}
\frac{\partial \bar{F}}{\partial \mu_{n_c-1}} &= -\phi\left(\frac{-\mu_{n_c-1}}{\sigma_{n_c-1}}\right)\left(\frac{1}{\sigma_{n_c-1}}\right), \text{ and,} \\
\frac{\partial \bar{F}}{\partial \sigma_{n_c-1}} &= \phi\left(\frac{-\mu_{n_c-1}}{\sigma_{n_c-1}}\right)\left(\frac{\mu_{n_c-1}}{\sigma_{n_c-1}^2}\right).
\end{aligned}
$$

This is the base step for the backward pass. Starting from $i = n_c - 1$, for each $i \leq n_c-1$, we assume that we know $\frac{\partial \bar{F}}{\partial \mu_i}$ and $\frac{\partial \bar{F}}{\partial \sigma_i}$ and $\frac{\partial \bar{F}}{\partial \sigma_{i,i+2}}$ and compute the corresponding quantities for $i-1$, using the equations above. Although many partial derivatives need to be computed, because of the specific construction of the forward pass, stage i depends on the previous stage only through μ_{i-1}, σ_{i-1}^2, μ_{i+1}, σ_{i+1}^2, and $\sigma_{i-1,i+1}$. The procedure finally yields $\frac{\partial \bar{F}}{\partial \mu_i} = -\frac{\partial \bar{F}}{\partial q_i}$ and $\frac{\partial \bar{F}}{\partial \sigma_i}$ for $i = 1, \ldots, n_c$.

Method	Expected Excess Inventory Cost ($ in millions)	Service Levels Achieved (% of demand)
OBC	25.7	98.1%
OBC-λ	20.7	90.0%
OBC-D	22.3	90.0%
CL	19.0	76.8%
CL-λ	21.0	90.0%
CL-D	21.0	90.0%
OBP	45.7	98.4%

Table 4: Simulation results for Experiment 1: The Base Case. Note that for the λ and direct scaling methods, the service level of 90% is achieved.

Costs Incurred ($ in millions)	Problem Size (Products $\times$ Components)				
Method	25×25	$50 \times 50^*$	75×75	100×100	125×125
OBC	11.7	25.7	40.4	58.1	73.5
OBC-λ	9.5	20.7	32.6	47.4	60.2
OBC-D	9.9	22.3	36.3	52.8	67.6
CL	7.9	19.0	30.5	44.0	55.3
CL-λ	9.5	21.0	37.1	49.3	61.0
CL-D	9.7	21.0	33.2	48.4	61.2
OBP	22.4	45.7	68.4	96.3	119.7

Service Levels Achieved	Problem Size (Products $\times$ Components)				
Method	25×25	$50 \times 50^*$	75×75	100×100	125×125
OBC	98.1%	98.1%	98.2%	98.3%	98.2%
CL	76.7%	76.8%	74.5%	73.7%	71.5%
OBP	98.7%	98.4%	97.9%	98.9%	99.0%

Table 5: Simulation results for Experiment 2: Varying Problem Sizes. An asterisk indicates the results of Experiment 1, the base case.

Costs Incurred ($ in millions)	Product-to-Component Ratio			
Method	100/25	60/42	50/50*	42/60
OBC	11.7	21.2	25.7	32.3
OBC-λ	9.4	16.9	20.7	25.5
OBC-D	9.9	18.3	22.3	27.8
CL	7.9	15.3	19.0	23.3
CL-λ	9.4	19.0	21.0	29.5
CL-D	9.7	17.3	21.0	26.3
OBP	22.4	38.6	45.7	54.9

Service Levels Achieved (% of demand)	Product-to-Component Ratio			
Method	100/25	60/42	50/50*	42/60
OBC	97.7%	98.3%	98.1%	98.6%
CL	76.3%	77.4%	76.8%	72.9%
OBP	98.7%	98.3%	98.4%	98.8%

Table 6: Simulation results for Experiment 3: Different Product-to-Component Ratios, with a constant 2500 possible arcs. An asterisk indicates the results of Experiment 1, the base case.

Costs Incurred ($ in millions)	BOM Ratio				
Method	8/50*	12/50	16/50	20/50	24/50
OBC	25.7	30.1	33.7	36.3	38.3
OBC-λ	20.7	23.4	25.3	26.4	27.2
OBC-D	22.3	25.4	26.4	27.6	28.2
CL	19.0	19.9	19.8	19.2	18.2
CL-λ	21.0	24.1	27.5	29.0	30.1
CL-D	21.0	24.4	27.1	29.1	30.4
OBP	45.7	66.5	87.5	108.0	128.5

Service Levels Achieved	BOM Ratio				
Method	8/50*	12/50	16/50	20/50	24/50
OBC	98.1%	98.4%	98.3%	98.2%	98.5%
CL	76.8%	70.3%	62.2%	55.9%	50.7%
OBP	98.4%	99.1%	99.9%	99.9%	99.9%

Table 7: Simulation results for Experiment 4: Different BOM Ratios. An asterisk indicates the results of Experiment 1, the base case.

Costs Incurred ($ in millions)	Mean Demand (Min, Max)				
Method	(1000, 1000)*	(900, 1100)	(800, 1200)	(700, 1300)	(600, 1400)
OBC	25.7	25.8	25.9	25.9	26.0
OBC-λ	20.7	20.7	20.8	20.9	20.9
OBC-D	22.3	22.4	22.4	22.5	22.6
CL	19.0	19.1	19.1	19.1	19.1
CL-λ	21.0	21.1	21.8	21.6	22.9
CL-D	21.0	21.1	21.2	21.2	21.3
OBP	45.7	45.8	45.8	45.9	46.0

Service Levels Achieved	Mean Demand (Min, Max)				
Method	(1000, 1000)*	(900, 1100)	(800, 1200)	(700, 1300)	(600, 1400)
OBC	98.1%	98.1%	98.1%	98.1%	98.1%
CL	76.8%	76.8%	76.9%	76.2%	76.2%
OBP	98.4%	98.4%	98.4%	98.4%	98.5%

Table 8: Simulation results for Experiment 5: Differences in Mean Demand. An asterisk indicates the results of Experiment 1, the base case.

Costs Incurred ($ in millions)	Demand Distribution (Min, Mean, Max)				
Method	(830, 1000, 1170)	(660, 1000, 1340)	(500, 1000, 1500)*	(330, 1000, 1670)	(160, 1000, 1840)
OBC	5.3	15.8	25.7	36.3	46.9
OBC-λ	3.9	12.4	20.7	29.5	38.3
OBC-D	4.4	13.5	22.3	31.7	41.0
CL	4.4	12.3	19.0	25.5	31.9
CL-λ	4.2	12.9	21.0	31.1	42.0
CL-D	3.9	12.5	21.0	30.4	40.0
OBP	11.9	29.3	45.7	63.1	80.6

Service Levels Achieved	Demand Distribution (Min, Mean, Max)				
Method	(830, 1000, 1170)	(660, 1000, 1340)	(500, 1000, 1500)*	(330, 1000, 1670)	(160, 1000, 1840)
OBC	98.3%	98.1%	98.1%	98.1%	98.1%
CL	89.3%	84.2%	76.8%	69.0%	62.5%
OBP	98.3%	98.4%	98.4%	98.4%	98.4%

Table 9: Simulation results for Experiment 6: Impact of Demand Variability. An asterisk indicates the results of Experiment 1, the base case.

Costs Incurred ($ in millions)	Target Service Level, X_t				
Method	82%	86%	90%*	94%	98%
OBC	23.6	24.5	25.7	27.4	30.7
OBC-λ	19.0	19.9	20.7	22.2	25.3
OBC-D	19.9	21.1	22.3	24.9	29.5
CL	18.3	18.6	19.0	19.2	19.5
CL-λ	19.4	20.7	21.0	22.6	29.5
CL-D	19.3	20.2	21.0	23.0	26.7
OBP	45.7	45.7	45.7	45.7	45.7

Service Levels Achieved	Target Service Level, X_t				
Method	82%	86%	90%*	94%	98%
OBC	95.6%	96.9%	98.1%	99.1%	99.8%
CL	72.0%	73.6%	76.8%	77.1%	78.9%
OBP	98.4%	98.4%	98.4%	98.4%	98.4%

Table 10: Simulation results for Experiment 7: Different Service Levels Targets, X_t. An asterisk indicates the results of Experiment 1, the base case.

References

BAKER, K. R. *Safety Stocks and Component Commonality*, Journal of Operations Management, vol. 6, no. 1, pp. 13-22, 1985.

BAKER, K. R., M. J. MAGAZINE AND H. L. W. NUTTLE. *The Effect of Commonality on Safety Stock in a Simple Inventory Model*, Management Science, vol. 32, no. 8, pp. 982-988, 1986.

CLARK, C. E. *Normal Approximation to a Maximum of Normal Random Variables*, Operations Research, vol 9., pp. 145-155, 1961.

COLLIER, D. A. *The Measurements and Operating Benefits of Component Part Commonality*, Decision Sciences, vol. 12, pp. 85-96, 1981.

COLLIER, D. A. *Aggregate Safety Stock Levels and Component Part Commonality*, Management Science, vol. 28, no. 11, pp. 1296-1303, 1982.

DYER, M. AND FRIEZE, A. M. *On the Complexity of Computing the Volume of a Polyhedron*, SIAM Journal on Computing, vol. 17, no. 5, pp. 967-974, 1988.

EYNAN, A. AND M. J. ROSENBLATT. *The Impact of Increasing Component Commonality in a Two-Level System*, Working Paper, June 1991.

GERCHAK, Y. AND M. HENIG. *An Inventory Model with Component Commonality*, Operations Research Letters, vol. 5, no. 3, pp. 157-160, 1986.

GERCHAK, Y. AND M. HENIG. *Component Commonality in Assemble to Order System: Models and Properties*, Naval Research Logistics, vol. 36, pp. 61-78, 1989.

GERCHAK, Y, M. J. MAGAZINE AND A. B. GAMBLE. *Component Commonality with Service Level Requirements*, Management Science, vol. 34, no. 6, pp. 753-760, 1988.

Kannan, R., J. Mount and S. Tayur. *A Randomized Algorithm for the Component Commonality Problem*, Mathematics of Operations Research, vol. 20, no. 3, pp. 529-549, 1995.

Keilson, J. and U. Sumita. *Uniform Stochastic Ordering and Related Inequalities*, Canadian Journal of Statistics, vol 10, no. 3, pp. 181-198, 1982.

Kolbin, V. V. *Stochastic Programming*, D. Reidel Publishing Company, Boston, 1977.

Komaromi, E. *A Dual Method for Probabilistic Constrained Problems*, Mathematical Programming Study, vol. 28, pp. 94-112, 1986.

McClain, J.O., W.L. Maxwell, J.A. Muckstadt, L.J. Thomas and E.N. Weiss. *Comment on Aggregate Safety Stock Levels and Component Part Commonality*, Management Science, vol. 30, pp. 772-773, 1984.

Prekopa, A. *Contributions to the Theory of Stochastic Programming*, Mathematical Programming, vol. 4, pp. 202-221, 1973.

Rockafellar, T. *Convex Analysis*, University Press, Princeton, NJ, 1970.

Swaminthan, J. and S. Tayur. *Stochastic Programming Models for Managing Product Variety*, in S. Tayur, R. Ganeshan and M. Magazine (eds.) Quantitative Models for Supply Chain Management, Kluwer Academic Press, 1998.

Tong, Y. *Probability Inequalities in Multivariate Distributions*, Academic Press, 1980.

Wets, R. *Stochastic Programming* in Handbooks in Operations Research and Management Science, vol. 1, North Holland, 1989.