Component procurement and allocation for products assembled to forecast: Risk-pooling effects

by S. J. Grotzinger
R. Srinivasan
R. Akella
S. Bollapragada

This paper considers procurement and allocation policies in a manufacturing environment where common components are assembled into various products that have stochastic demands. The components are allocated to the assembly of a product at a time when product demand is still uncertain (assemble to forecast, ATF). The special case of one component shared by *N* different products is analyzed, and insights into the general problem are obtained for the situation in which the common component can be reallocated to different products as product demands change. An allocation policy is developed for general distributions and prices in an ATF environment. The policy first addresses anomalies in the state of the system and then, for a feasible state, minimizes the expected excess finished-goods inventory. A procurement level that is nearly optimal is obtained from a Monte Carlo simulation in which the probability of satisfying all of the random product demands simultaneously is considered relative to this allocation policy. Numerical studies indicate that the total component and finished-goods inventory is significantly reduced by an allocation policy that incorporates risk pooling while still fulfilling service-level requirements.

Introduction

Multiplant coordination involves correlating needs in a chain of facilities in which upstream facilities supply parts consisting of components and subassemblies to manufacturing facilities downstream. At the final assembly facility, these parts are assigned to particular products and then released into an assembly process to meet a demand for a particular finished product. These releases into the final assembly process are determined to satisfy or *service* the probable demand requirements for a set of products that will exist at the end of this stage. An appropriate *policy* must be devised to determine how parts are allocated toward finished products of a particular type and released into the final assembly process. If an assembly cycle time is long, these product demands are unknown at the time the release is made, but a probability distribution may be known and utilized. Suppose the cycle or lead time is four weeks for the final assembly process and the demand for finished products must be satisfied 90% of the time. The releases must be determined now to meet the currently unknown demand. Sufficient finished goods must be available four weeks from now to satisfy or service the random demand with the given probability. This probability of 0.9 is often referred to as the service level.

To satisfy stochastic demand requirements with a certain probability, it is necessary to maintain inventories at various levels. The overall objective is to minimize the total inventory in the system that is required to achieve a desired service level. The total inventory is the sum of the value of the unsold finished products along with the value of the unutilized components/subassemblies. This inventory is often composed of components that are common to a set of distinct products. For example, the same microprocessor is used in a variety of different personal computer models. Components common to various products and an appropriate allocation policy can be exploited to reduce the total inventory in such systems. A common component procured for the probable demand of one product can be allocated to another distinct product, if a shift in the demand patterns is discerned prior to actual release into the final assembly process. This is referred to as *risk pooling.*

Two different production environments occur in manufacturing. The first is an assembly facility in which the final assembly process is long enough to require the releases to be made before the actual demand for the product is known with certainty. The allocation policy must explicitly address the complexity arising from this uncertainty. This mode of manufacturing is commonly referred to as "assembling to forecast" (ATF). In the other environment, because an assembly operation is short, it is possible to observe the actual demand and assemble to order (ATO). The available parts are allocated and then released to meet these known demands for the various products. The number of components/subassemblies allocated to a particular product equals the number required to exactly meet the known demand for that product. However, if the procurement lead times for parts are significantly long, it will be necessary in both of these production environments to order these components or begin fabrication of the subassemblies well in advance of the time when these parts are to be released into the final assembly process.

The electronics industry provides an example of a multiplant coordination problem. A typical electronics manufacturing process can be roughly divided into three stages of production. To start with, there is semiconductor manufacturing, where chips are fabricated from wafers. Printed circuit boards are then assembled in card assembly plants. At the final stage, the boards are assembled into system units to meet actual customer demands. The manufacturing cycle times vary for the three processes. Chip fabrication is in the range of three to six months. For various manufacturers, the cycle time for assembling printed circuit cards or system units can range from several days to several weeks for each of these processes. In some large electronics firms, the three manufacturing processes are done at different sites, which introduces lead times for transportation. The raw materials or parts at all three stages of production may have to be procured to forecast if the procurement lead time is long. The assembly of system units may be done upon realization of market demands if the cycle time for assembly is short. The chip manufacturing plants have to assemble to forecast to address their significant cycle times.

In these multiplant manufacturing networks where demand is unknown at the time of procurement and release, exploiting commonality, improving serviceability, and reducing inventories while respecting lead times are strategically important goals in many industries, but very difficult problems to address. Current practice ignores commonality and does not explicitly address demand uncertainty, resulting in excessive finished-goods inventory for some products but significant backlogs for others. Even with expensive inventories at both the parts and finished-product level, poor serviceability is often realized.

Much insight into these general questions can be gained by restricting consideration to a single manufacturing facility. The process of procuring parts and their subsequent allocation to different products, where both decisions have significant lead times associated with them, also arises in this context. Henceforth these questions will be pursued from the perspective of a single manufacturing facility. In such a plant, the components arrive after a certain lead time and they are then released to be assembled into various products. Because the cycle time for product assembly may be significant, products must be assembled so as to meet a currently unknown demand. The procurement and allocation decisions must be made so that the demand for finished products is met with a certain service level and the total cost of the inventory in the system is simultaneously reduced. The demand for the products is random, but its marginal probability distribution is known. There is a high degree of commonality of components among the products. Once again the risk-pooling effects of component commonality and an appropriate allocation policy are to be exploited to reduce the total inventory.

An ATO environment having N products with one common component and special pricing and distribution restrictions has been considered previously [1–4]. Recently the ATO problem has been addressed in the presence of general product structure, prices, multiple time periods, and independent product demands [5]. In this paper, the setting of N products with one common component, general pricing structures, and general distributions is addressed in an ATF environment in order to obtain further insights into risk-pooling effects. An allocation policy is developed which assesses current inventory levels and respects the cycle time required for the assembly process. The allocation policy is then coupled with a simulation technique to implicitly address serviceability. Numerical implementation and integration into a simulation experiment suggest the potential for significant inventory reductions from risk pooling in this ATF environment.

The remaining sections are organized as follows. The next section considers the overall planning problem for the ATF environment. The formulation considers multiple products, components, and time periods relative to an *aggregate service measure*. An aggregate service measure reflects the probability of meeting the random demands for all products simultaneously. To assess the significance of inventory reductions resulting from commonality of components and an appropriate allocation policy, a special case of the general formulation is considered in detail. This special case, which considers a single component to be used in several products, is addressed in the section on the single-component problem. An allocation policy is developed which first examines the state of the system for any anomalies, then assesses feasibility concerns, and subsequently makes an optimal allocation when a feasible state exists. When a feasible solution exists, an iterative algorithm based on nonlinear optimality criteria is used to determine how to release assemblies into the manufacturing process relative to some as yet unknown demand. Expected excess finished-goods inventory is the objective minimized. The algorithm is applicable to a single-stage model with N products and one common component, general pricing structure, and arbitrary distributions of finite support or truncated distributions. Relative to this allocation policy, a Monte Carlo simulation procedure to determine a near-optimal *order-up-to level* for the common component is then developed in the section entitled "Order-up-to by simulation." The order-up-to level is the quantity of parts on hand at the beginning of each period. All of these parts need not be allocated to specific products, but they are all available for allocation. Enough parts are then procured to bring the amount on hand back up to this level for availability in the next period. A procedure to obtain the order-up-to level when N unique components are used instead of a common component is also discussed in this section. In the section on computational study, the numerical results of experiments to assess the reduction in total inventory levels obtainable by exploiting component commonality and an appropriate allocation policy are presented for a varying number of products, demand distributions, service levels, and prices. Insights into risk-pooling effects suggested by these numerical results are also discussed. A summary discussion is provided in the section on conclusions, and the proof of optimality of the solution to the allocation problem is presented in the Appendix.

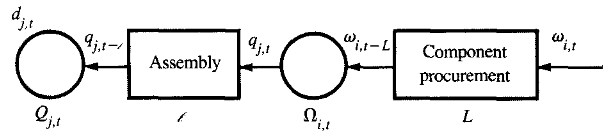

Figure 1

Schematic representation of a two-stage production system.

A general problem formulation

This section formulates the problem of determining component procurement and product release levels in the presence of significant lead times and demand uncertainty in a two-stage production system (**Figure 1**). The formulation presented considers several products, components, and time periods. This general formulation provides the overall context for the multiple-products, single-component, and single-period problem that is actually solved in the next section. Further, the general formulation presented here is new and has not been presented before. The two stages shown in Figure 1 involve component procurement/fabrication and release into the assembly process. Components are procured from vendors or fabricated in internal plants. The lead time for component procurement is assumed to be the same for all components. When the components arrive, they are stored in the warehouse, and quantities required subsequently are released into the assembly process. This release is based on allocation decisions made by considering the current state of the system, the service level, and the inventory cost objectives. The assembly cycle time is assumed to be the same for all of the products. The finished products arrive at the output buffer from which demand is serviced.

All procurement and allocation decisions are made at the beginning of time periods. Demand for a period is realized during that time period and is completely known at the end of that time period, i.e., the beginning of the next time period.

The following notation is utilized in the formulation:

- J index set for products
- I index set for components
- ℓ manufacturing cycle time for all products
- L procurement lead time for all components
- u_{ij} usage count of component i in product j
- Q_{jt} net inventory/backlog of product j at the beginning of period t
- Ω_{it} net inventory of component i at the beginning of period t
- c_i cost of component i
- h_j holding cost of product j
- d_{jt} random variable representing the demand for product j in period t realized during period t, known completely at the beginning of period $t+1$.
- β_t service level for period t
- S_t state of the system known at time t before allocation and procurement decisions are made;

$$\equiv \{Q_{jt}, (q_{j\tau}, \tau = t-1, \cdots, t-\ell+1), \Omega_{i,t-1}, (\omega_{i\tau}, \tau = t-1, \cdots, t-L)\}$$

The decision variables for allocation and procurement are represented by

- q_{jt} allocation decision for product j to be released at the beginning of period t toward finished goods in period $t+\ell$
- ω_{it} procurement decision for component i made at the beginning of period t for delivery at the beginning of period $t+L$

In the formulation given below, the objective function minimizes the new inventory investment to be made at the beginning of period t so that the demands for products in period $t+\ell+L$, including backlogs, are met with a probability of $\beta_{t+\ell+L}$. The minimization is done over the procurement decisions, ω_{it}, made at t relative to an allocation policy, $g(\cdot)$, to be specified. This allocation policy will be employed in periods $t, \cdots, t+L$. The result of the allocation done at $t=\tau$ is the set of releases $q_{j\tau}$, for all j. Equation (3) and Equation (4) are inventory balance equations.

$$\min_{\omega_{it}} \sum_i c_i \omega_{it}, \tag{1}$$

such that

$$\Pr\{Q_{j,t+\ell+L} \geq d_{j,t+\ell+L}, \ \forall\ j\} \geq \beta_{t+\ell+L}, \tag{2}$$

$$Q_{j\tau} = Q_{j,\tau-1} + q_{j,\tau-\ell} - d_{j,\tau-1},$$
$$\forall\ j,\ \tau = t, \cdots, t+\ell+L, \tag{3}$$

$$\Omega_{i\tau} = \Omega_{j,\tau-1} + \omega_{i,\tau-L} - \sum_j u_{ij} q_{j\tau},$$
$$\forall\ j,\ \tau = t, \cdots, t+\ell+L, \tag{4}$$

$$q_{j\tau} = g(S_\tau, j),\ \tau = t, \cdots, t+L, \tag{5}$$

where $g(\cdot)$ is an *allocation policy* that is specified.

$$\Omega_{i\tau}, \omega_{i\tau}, q_{j\tau} \geq 0,\ \forall\ i,\ \forall\ j,\ \forall\ \tau. \tag{6}$$

In this formulation, S_t is known, and one set of procurement decisions, ω_{it}, is made at t. Also, the allocation policy, $g(\cdot)$, to be employed in periods $t, \cdots, t+L$ is specified at t. However, only the current allocation decisions, q_{jt}, are made at t; the remaining allocation decisions, $q_{j\tau}$, $\tau = t+1, \cdots, t+L$, are made at τ. This problem may be reformulated as a stochastic dynamic program with $L+1$ stages. However, the problem is inherently difficult to solve as a consequence of three features being present together: (a) the stochastic nature of the demands for products, (b) allocations being made at multiple stages (time periods), and (c) component commonality.

Before attempting to solve this problem, it is desirable to understand how significant these features of the problem are. This facilitates the choice of features to retain and to omit from the formulation. The stochastic nature of the product demands cannot be ignored. The allocation decisions must be made dynamically, i.e., in multiple stages, because the backlogs or excess inventories accrued up to a certain point in time are an important factor to be considered in determining what to release into the assembly process at that time. Furthermore, these allocation decisions must be made when the product demands are still unknown, as contrasted with an ATO situation where allocation decisions are made after demands are known with certainty. This aspect of allocation needs full consideration. The effect of component commonality on safety stock reductions in ATO systems has been studied [1–5]. Empirical studies in [5] show a significant reduction in the required procurement levels compared to the levels required when commonality is ignored. Such studies have not been done in the case of ATF systems in the presence of service-level constraints. The objective in the following sections is to determine the reduction in total inventory obtainable from component commonality and an appropriate allocation policy in the ATF environment. The methodology developed there is applicable to general pricing structures and arbitrary distributions for product demands. On the basis of the significance of the reductions that may be achieved, recommendations on whether component commonality should be explicitly considered in the general formulation presented in this section will be made.

Single-component problem

Risk-pooling effects are assessed in an ATF environment from studies of a single common component and N products. In this regard, the optimization in (1) and the service requirement (2) are addressed by use of a Monte Carlo simulation relative to a specific allocation policy $g(\cdot)$ to be developed in the section which follows. This single-component problem requires the determination of the procurement and subsequent allocation to the N products. The allocation problem is addressed first and is followed by a determination of an order-up-to level.

The allocation policy $g(\cdot)$ uniquely determines the releases into the final assembly process to satisfy an as-yet-unrealized demand. To determine this policy, four distinct cases must be considered. Case one addresses an anomaly that may arise from the current state of the system; cases two and three assess whether or not the state of the system is feasible. With *a priori* knowledge that a feasible state exists, case four subsequently determines an optimal allocation relative to expected excess finished-goods inventory.

- *Allocation problem*

First, the *policy* of optimally allocating a known order-up-to quantity Ω to the N different products is considered. The cases assessing anomalies in the state of the system are discussed after an optimal allocation is formulated. The component index i is not needed for the single-common-component problem. The component usage rates u_j are assumed equal to unity, and the time index is dropped. The component index i and time index t are also dropped from all the variables. With these simplifying assumptions, the notation is reduced to the following:

Ω	order-up-to quantity or amount of the common component available
d_j	unknown demand for product j
$\bar{d}_j$	maximum realizable demand for product j
$f_j(\cdot)$	marginal density function for demand for product j
$F_j(\cdot)$	marginal cumulative distribution function for demand for product j

The procedure is to determine q_j, the amount of the common component released into the manufacturing process toward demand for product j; $\hat{q}_j$ is used in place of Q_{jt} to emphasize the fact that this is a known quantity as far as this allocation procedure is concerned.

Only finished-goods inventories are considered. That is, this part of the allocation policy assumes that all of the available common components, Ω, will be released into the system and allocated to some product j. The formulation permits arbitrary price relationships and unrelated distributions for the different products. However, it is assumed that the density functions, representing product demands, are positive on the domain of interest, i.e.,

$$f_j(z) > 0 \;\; \forall \;\; z \in (0, \bar{d}_j]. \tag{7}$$

The objective is to determine allocations q_j for each assembly j so that the expected excess finished-goods inventory costs are minimized. Since the expectation of a sum of random variables is equal to the sum of the expectations of the individual random variables, we have

$$\min_{q_j} E\left\{\sum_{j=1}^{N} h_j \cdot [q_j + \hat{q}_j - d_j]^+\right\}$$

$$= \min_{q_j} \sum_{j=1}^{N} h_j \cdot E([q_j + \hat{q}_j - d_j]^+).$$

As discussed in the Appendix, the expectation of excess inventory for product j can be computed by using the marginal density function. Thus, the optimization problem to be solved is as follows:

$$\min_{q_j} \sum_{j=1}^{N} h_j \cdot \int_0^{\hat{q}_j+q_j} [\hat{q}_j + q_j - d_j] \cdot f_j(d_j) d\, d_j, \tag{8}$$

subject to

$$\sum_{j=1}^{N} q_j = \Omega, \tag{9}$$

$$[-\hat{q}_j]^+ \le q_j \le \bar{d}_j - \hat{q}_j \qquad j = 1, 2, \cdots, N, \tag{10}$$

where $E(x)$ is the expected value of x and $[x]^+ \equiv \max\{0, x\}$.

Before proceeding, the cases addressing anomalies in the state of the system and feasibility concerns are motivated and then specified.

First, it is possible that the net inventory in the system for product j already exceeds the maximum demand which can be realized. Thus, one would not want to release additional assemblies into the pipeline toward demand for product j. In this instance, product j would be dropped from the formulation, and the reformulation for the remaining $N - 1$ products would be addressed:

$$\text{If } \bar{d}_j < \hat{q}_j, \qquad \text{then } q_j^* \equiv 0 \text{ and reformulate.} \tag{11}$$

Second, it is possible that the existing backlogs cannot be met with available resources. In this instance, the policy would be to reduce all existing backlogs by an equal percentage. That is,

$$\text{if } \sum_{j=1}^{N} [-\hat{q}_j]^+ \ge \Omega,$$

$$\text{then } q_j^* = \Omega \cdot \left[\sum_{j=1}^{N} [-\hat{q}_j]^+ \right]^{-1} \cdot [-\hat{q}_j]^+. \tag{12}$$

Third, it may arise that an excess supply of components accumulates. In this instance, the policy is to bring the finished-goods inventory up to the maximum demand levels for all products. That is,

$$\text{if } \sum_{j=1}^{N} (\bar{d}_j - \hat{q}_j) < \Omega, \qquad \text{then } q_j^* = \bar{d}_j - \hat{q}_j \ \ \forall\ j. \tag{13}$$

The remaining components are then held as component inventory, and the procurement for the next period is appropriately reduced in order once again to have the order-up-to quantity Ω available for release in the next period.

Finally, if none of these instances arises, the policy is determined by the optimal solution to the feasible nonlinear optimization problem (8)–(10), which is closely related to the following general form:

$$\max_{y_i} \sum_{i=1}^{N} \Phi_i(y_i), \tag{14}$$

subject to

$$\sum_{j=1}^{N} y_i = B, \tag{15}$$

$$0 \le y_i \qquad \text{for } i = 1, \cdots, N. \tag{16}$$

As noted by Zipkin [6], this form arises in a variety of applied contexts, including the distribution of search effort, marketing, portfolio selection, capital budgeting, reliability, and health care. Luss and Gupta [7] developed an algorithm which is applicable to four specific forms of return functions $\Phi_i(\cdot): \mathbb{R} \rightarrow \mathbb{R}$. This approach subsumed several previous results. Zipkin presented these ranking methods in a unified framework applicable whenever $\Phi_i(\cdot): \mathbb{R} \rightarrow \mathbb{R}$ are continuously differentiable and strictly concave. In addition to enhancing the breadth of applicability, Zipkin's approach in [6] also achieves a significant computational advantage. It requires only a single numerical inversion after termination of the iterative part of the procedure, as opposed to one being required at each iteration. In extending the methodology to the case where upper bounds are also imposed on the variables,

$$y_i \le b_i \qquad \text{for } i = 1, \cdots, N, \tag{17}$$

their technique requires that several related problems of the form (14)–(16) must be solved. The justification is based on a special type of relaxation method. However, by using a sequence of unbounded-variables problems to solve the bounded-variables problem, one is again faced with the possibility of N numerical inversions being required.

The method to be developed here directly addresses the Karush–Kuhn–Tucker conditions, with the upper bound constraints (17) present. As a consequence, it is able to extend to the bounded case the property that only a single numerical inversion need be performed after the iterative portion of the procedure has terminated. This is the calculation required by the step

$$\text{find } \lambda^* \in [-\gamma_{i+1}, -\gamma_i) \ni Z_i(\lambda^*) = \Omega \tag{18}$$

in the procedure specified below.

This allocation policy is employed repeatedly in the determination of the order-up-to level in the next section. For each estimate of the order-up-to level Ω, the allocation subroutine must be invoked 2500 times. Suppose the search to determine Ω requires ten distinct estimates. Then, for $N = 50$ products, the potential of 1 250 000 numerical inversions (18) is reduced to 25 000 by the enhancements developed here. The allocation policy must also be applicable to general distributions—for example, distributions of finite support and truncated distributions. Therefore, the assumptions that $\Phi_i(\cdot): \mathbb{R} \rightarrow \mathbb{R}$ are $\mathscr{C}^1$ and strictly concave everywhere will not be satisfied. This additional complexity is also addressed by the methodology presented here.

The fact that the Hessian of the objective function (8) is a positive semidefinite matrix is verified in the Appendix, implying that the objective function is a convex function. Thus, the Karush–Kuhn–Tucker conditions are sufficient to determine globally optimal allocations. Consequently, a solution to the following set of equations and inequalities is to be determined:

$$h_j \cdot F_j(\hat{q}_j + q_j) + \lambda - \lambda_j + \beta_j = 0 \qquad j = 1, 2, \cdots, N, \tag{19}$$

$$\Omega - \sum_{j=1}^{N} q_j = 0, \tag{20}$$

$$q_j - [-\hat{q}_j]^+ \ge 0 \qquad j = 1, 2, \cdots, N, \tag{21}$$

$$\bar{d}_j - (q_j + \hat{q}_j) \ge 0 \qquad j = 1, 2, \cdots, N, \tag{22}$$

$$\lambda_j \cdot (q_j - [-\hat{q}_j]^+) = 0 \qquad j = 1, 2, \cdots, N, \tag{23}$$

$$\beta_j \cdot (\bar{d}_j - q_j - \hat{q}_j) = 0 \qquad j = 1, 2, \cdots, N, \tag{24}$$

$$\lambda \text{ unrestricted in sign,} \tag{25}$$

INITIALIZE

$i = \eta$, TERM = 0

WHILE TERM $\neq 1$

IF

$$\Omega \leq \sum_{j=1}^{i} \zeta_j(-\gamma_{i+1}) = Z_i(-\gamma_{i+1})$$

THEN find $\lambda^* \in [-\gamma_{i+1}, -\gamma_i) \ni Z_i(\lambda^*) = \Omega$ and set TERM = 1

FOR $j = 1, \cdots, i$

IF $\lambda^* > -h_j$

THEN

$$q_j^* = F_j^{-1}\left(-\frac{\lambda^*}{h_j}\right) - \hat{q}_j = \zeta_j(\lambda^*) \text{ and } \lambda_j^* = 0 \text{ and } \beta_j^* = 0$$

ELSEIF $\lambda^* \leq -h_j$

THEN

$$q_j^* = \bar{d}_j - \hat{q}_j = \zeta_j(\lambda^*) \text{ and } \lambda_j^* = 0 \text{ and } \beta_j^* = -\lambda^* - h_j$$

ENDIF

ENDFOR

FOR $j = i + 1, \cdots, N$

$$q_j^* = 0 \text{ and } \lambda_j^* = \lambda^* + h_j \cdot F_j(\hat{q}_j) \text{ and } \beta_j^* = 0$$

ENDFOR

ELSE

$i \leftarrow i + 1$

ENDIF

ENDWHILE

Figure 2

Iterative procedure for generating solution to Equations (19)–(26).

$$\lambda_j, \beta_j \geq 0 \qquad j = 1, 2, \cdots, N. \tag{26}$$

Before the iterative procedure to generate a solution to conditions (19)–(26) is presented, several additional definitions are developed to facilitate the exposition. The set $\mathcal{N} \equiv \{1, 2, \cdots, N\}$ is partitioned into two sets,

$$\Gamma^- \equiv \{j \in \mathcal{N}: \hat{q}_j \leq 0\},$$

$$\Gamma^+ \equiv \{j \in \mathcal{N}: \hat{q}_j > 0\}. \tag{27}$$

The quantities $\eta \equiv |\Gamma^-|$ and $\nu \equiv -\max\{h_j: j \in \mathcal{N}\}$ are also required, as are the following parameters:

for $j \in \Gamma^-$ define $\gamma_j \equiv 0$,

and

for $j \in \Gamma^+$ define $\gamma_j \equiv h_j \cdot F_j(\hat{q}_j)$. (28)

For ease of exposition and without loss of generality, assume that these quantities occur in the natural order,

$$0 \equiv \gamma_0 = \gamma_1 = \cdots = \gamma_\eta < \gamma_{\eta+1} \cdots \leq \gamma_N \leq -\nu \equiv \gamma_{N+1}. \tag{29}$$

The procedure iteratively considers the subintervals

$$[\nu, -\gamma_N] \cdots [-\gamma_i, -\gamma_{i-1}] \cdots [-\gamma_{\eta+1}, 0]$$

of the corresponding partition of the interval $[\nu, 0]$, and makes use of the following functional forms:

```
Estimate of Ω
WHILE service level is unsatisfied
      Initialize procurement to Ω
      Set q̂_j = 0 ∀ j
      Initialize component inventory and aggregate service to 0
      Initialize all individual product service levels to 0
      FOR periods i = 1, 2, ···, 2500
            Procured components arrive bringing available components up to
            level Ω
            Determine allocations q_j^* with respect to
            Ω and the current state of the system
            Release q_j^*'s
            Update component inventory
            Calculate procurement for next period
            Generate random demands d_j
            Ascertain whether demand was satisfied for each product
            Update q̂_j's
      ENDFOR
      Calculate average component inventory and aggregate service
      Calculate average product inventory and individual service level
      for all products
      Modify order-up-to level Ω
ENDWHILE
```

Figure 3

Experimental framework for Monte Carlo simulation determining the order-up-to level Ω that satisfies a given aggregate service level β.

$$\zeta_j(\lambda) \equiv F_j^{-1}\left(\min\left\{1, \frac{-\lambda}{h_j}\right\}\right) - \hat{q}_j, \tag{30}$$

$$Z_k(\lambda) \equiv \sum_{j=1}^{k}\left[F_j^{-1}\left(\min\left\{1, \frac{-\lambda}{h_j}\right\}\right) - \hat{q}_j\right] \equiv \sum_{j=1}^{k} \zeta_j(\lambda). \tag{31}$$

Note that the iterative procedure shown in **Figure 2** is initiated at $i = \eta$. This gives an advanced start relative to the algorithms employed in [6] and [7], which are always initialized at $i = 1$.

The proof that the procedure presented in Figure 2 terminates and then generates the global optimum solution to the problem (8)–(10) after at most $N - \eta + 1$ iterations is verified in the Appendix. This allocation policy must now be coupled with an order-up-to level. The determination of a nearly optimal level by a heuristic search procedure is developed in the next section.

• *Order-up-to by simulation*

The order-up-to level Ω that satisfies a given aggregate service level β is determined through Monte Carlo simulation. A brief description of the experimental framework (**Figure 3**) is first presented to facilitate following the steps in the framework. The value of Ω that achieves the specified aggregate service level is determined by a bisection procedure. This is the outer **WHILE** loop in the experiment. For each iteration of this bisection procedure, the aggregate service achieved is estimated through simulation of the procurement and allocation procedure over a horizon of 2500 periods. The bisection procedure is initialized with an estimate, which can be the sum of mean demands of all the products, for example. In each period in the simulation, allocation of the Ω components is done to the N products considering the current state of the system according to the allocation policy described in the previous section. The current state of the system is described by the backlog or excess product inventory of all of the products. After release of the allocated components, the procurement decision for the component is made to bring up to Ω the level of the components available for the next period. At the end of the period, the releases become available and are added to the product inventory. Then the demand for the products is realized. The excess product inventory or backlog for each product is updated. For each of the products, a record is kept of whether or not demand was met. The simulation then steps into the next time period. At the beginning of the next period, the component procurement arrives. This set of events is repeated over the entire horizon. At the end of the simulation, the aggregate service achieved is computed as the fraction of all the time periods in which there was no backlog for any product. The bisection procedure continues until the service achieved equals the target aggregate service within a specified tolerance.

When commonality is not considered, the process of procurement and allocation is modeled separately for each product. Components procured for a product are allocated only to that product. In this case, if the demands for the products are independent of one another, the event of satisfaction of demand for each product is also independent of that of the others. The same is not true when commonality is considered; the allocation procedure introduces a correlation among the demand satisfaction events. Therefore, in order to meet an aggregate service level of β, the service β_j for product j on each of the individual products should be such that $\beta_1 \beta_2 \cdots \beta_N \geq \beta$. If the product prices and demand distributions are identical for all products, then $\beta_1 = \beta_2 = \cdots = \beta_N = \beta^{1/N}$. Consequently, the procurement and allocation problem for N products can be decomposed into N identical problems. It is therefore sufficient to solve one of them and obtain the inventory costs for the N products by simply multiplying the results by N. Each individual problem is then solved using the experimental framework presented earlier in this section by setting N to 1. In this case, the allocation procedure reduces to a simple policy of releasing components up to the difference between maximum demand for the product and the excess inventory/backlog for the product.

- *Computational study*

A computational study was conducted of the performance of the method that considers component commonality and the one that does not. Several experiments were performed in the study by varying the number of products, target aggregate service level, and type of demand distributions. The number of products was varied from 2 to 44 for a given target service level. The set of experiments was then repeated for service levels ranging from 0.5 to 0.9 for a given number of products. Symmetric triangular distribution characterized by low product demand of 0 units with probability of 0, medium product demand of 500 with the maximum probability, and high product demand of 1000 with probability 0 was used. Uniform distribution of the same range, i.e., 1 to 1000, was then used to repeat the same set of experiments. Each experiment was then run for a horizon of 2500 periods.

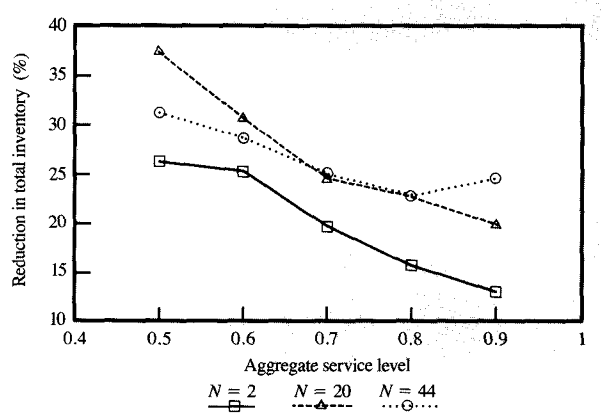

Figure 4

Service vs. inventory reduction for triangular distribution.

During each of the experiments, the component inventory after releases, the leftover product inventory or backlog, and the total inventory in the system (the sum of component and product inventories) were tracked in order to compute averages over the horizon. For each period, it was determined whether demand was satisfied for each of the products in order to compute the aggregate service and individual product service achieved over the horizon.

Total inventory as a function of target aggregate service level is graphed in **Figure 4** for 2, 20, and 44 products with triangular distributions for product demands. **Tables 1**, **2**, and **3** contain the data used in these graphs and a breakdown of the total inventory into component inventory and product inventory.

Several empirical observations are made from the trends indicated by the data. The reduction in total inventory is significant, ranging from 12 to 37%. This reduction decreases in most situations when the target used for aggregate service is increased, since the individual service rapidly approaches 1 as the aggregate service increases. For example, when aggregate service is 0.5 and $N = 20$, the individual service level measured is 0.82; when the aggregate service is 0.8 and $N = 20$, the individual service measured is 0.97. Therefore, the releases into the final

Table 1 Reduction in total inventory (I), finished-goods inventory (I × II), and component inventory (I × III) when $N = 2$ and demand follows triangular distribution.

Aggregate service	Total inventory (I)	Finished goods (II) (%)	Component inventory (III) (%)	I × II	I × III
0.5	26.29	82.01	17.99	21.56	4.73
0.6	25.29	78.12	21.88	19.76	5.53
0.7	19.73	69.93	30.07	13.8	5.93
0.8	15.72	59.30	40.7	9.32	6.4
0.9	12.99	43.09	56.9	5.6	7.39

Table 2 Reduction in total inventory (I), finished-goods inventory (I × II), and component inventory (I × III) when $N = 20$ and demand follows triangular distribution.

Aggregate service	Total inventory (I)	Finished goods (II) (%)	Component inventory (III) (%)	I × II	I × III
0.5	37.38	64.29	35.71	24.03	13.35
0.6	30.61	54.92	45.08	16.81	13.8
0.7	24.6	41.24	58.76	10.15	14.45
0.8	22.8	31.1	68.9	7.09	15.71
0.9	19.86	15.65	84.35	3.11	16.75

Table 3 Reduction in total inventory (I), finished-goods inventory (I × II), and component inventory (I × III) when $N = 44$ and demand follows triangular distribution.

Aggregate service	Total inventory (I)	Finished goods (II) (%)	Component inventory (III) (%)	I × II	I × III
0.5	31.24	50.67	49.33	15.83	15.41
0.6	28.73	42.88	57.12	12.32	16.41
0.7	25.07	31.53	68.47	7.9	17.17
0.8	22.77	21.56	78.44	4.91	17.86
0.9	24.57	13.57	86.43	3.33	21.24

assembly process in both approaches are made such that the probability of meeting the maximum demand for each of the products is very close to 1. In this instance, the reduction in excess finished-goods inventory decreases in significance. Therefore, the difference between the two approaches arises mainly from a decrease in the component inventory.

It is interesting to note that even at high levels of service, the reduction in total inventory between the two approaches remains significant. The reduction is more significant when the number of products is high. For example, when $N = 44$ and aggregate service is equal to 0.9, the reduction in total inventory is 24.6% (refer to column 1 in Table 3). This is explained as follows. At high levels of service, as argued in the previous paragraph, both approaches release to meet the maximum demand with a probability very close to 1, resulting in a significant amount of leftover product inventory. However, in the approach that considers commonality, the sum of product inventories has a coefficient of variation that is much less than that of the other approach, which considers each product separately. For the problem under study, the coefficient of variation for the approach that considers commonality is roughly of the order of $1/\sqrt{N}$ of that of the other approach. Consequently, the variability in the total release for N products, where commonality is considered, is insignificant when N is large. However, in the single-product case the variability remains significant. Higher component inventory must therefore be maintained to respond to this uncertainty. Thus, even in an ATF system there is considerable potential for component inventory reduction from risk pooling.

To summarize, the following observations were made. There is significant reduction in total inventory when commonality is considered. There are two parts to the total inventory, the finished-goods inventory and the component inventory. The reduction in finished-goods inventory decreases as aggregate service level increases (column 5 of Tables 1–3). The reduction in component inventory increases as aggregate service level or number of products increases (column 6 of Tables 1–3). These two effects are counteracting.

The above observations help explain two anomalies. In the graph corresponding to N equal to 44 in Figure 4, the percentage reduction in total inventory decreases for aggregate service levels up through 0.8 and then increases at a service level of 0.9. The reduction in total inventory increases as N increases to 20 and then decreases when N is increased to 44 for some service levels (column 1 of Tables 1–3).

For the case of uniform distributions for product demands, the reduction in total inventory as a function of aggregate service level is graphed in **Figure 5**. The reduction in total inventory decreases as a function of aggregate service level in this case also. However, the reduction in total inventory goes up as N increases from 2 to 20 and also to 44, unlike the case of triangular distributions.

Conclusions

The risk-pooling and allocation policy effects in the context of procuring and assembling to forecast were investigated in the special case where one component is used in N different products under the requirement of an aggregate service level. The empirical results obtained in this study indicate significant potential for savings in total inventory levels. Contributions to the savings by finished-goods inventory and component inventory were analyzed. While the individual contributions of these two quantities seem to vary in opposite directions as a function of aggregate service, the total savings are still significant at high levels of aggregate service. The actual problem to be solved in the industry would typically involve hundreds of components being used in tens of products. The savings in procurement in such cases may not be directly estimated from the results obtained here. However, the significant savings reported here strongly suggest that component commonality should be explicitly incorporated in the formulation of the procurement planning problem. Also, the results reported here may be applicable to products that share a single expensive component. This would be the case, for instance, in a personal computer where the microprocessor is very expensive compared to the rest of the components. Careful planning for the procurement of the microprocessor alone may achieve significant benefits.

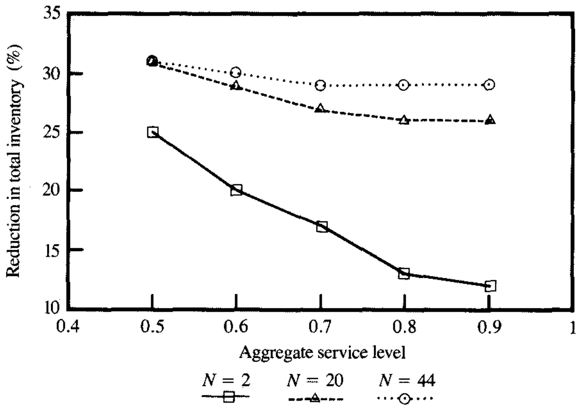

Figure 5

Service vs. inventory reduction for uniform distribution.

Appendix

An explanation is provided here to support the assertion made in arriving at the objective function (8). The expression

$$E\left(\sum_{j=1}^{N} h_j \cdot [q_j + \hat{q}_j - d_j]^+\right)$$

can be rewritten as

$$\sum_{j=1}^{N} h_j \cdot E([q_j + \hat{q}_j - d_j]^+),$$

since in general the expectation of the sum of a set of random variables is equal to the sum of expectations of each of the random variables. This can be expanded as

$$\sum_{j=1}^{N} h_j$$

$$\cdot \int_0^{u_1} \cdots \int_0^{u_N} [\hat{q}_j + q_j - d_j] \cdot f(d_1, \cdots, d_N) d\, d_1 \cdots d\, d_N,$$

which can also be rewritten using Bayes' theorem as

$$\sum_{j=1}^{N} h_j \cdot \int_0^{u_1} \cdots \int_0^{u_N} [\hat{q}_j + q_j - d_j]$$

$$\cdot f_j(d_j) \cdot f(d_1, \cdots, d_N | d_j) d\, d_1 \cdots d\, d_N,$$

where $u_j = \hat{q}_j + q_j$. This is equal to

$$\sum_{j=1}^{N} h_j \cdot \int_0^{u_j} [\hat{q}_j + q_j - d_j] \cdot f_j(d_j)$$

$$\cdot \int_0^{u_1} \cdots \int_0^{u_N} f(d_1, \cdots, d_N | d_j) d\, d_1 \cdots d\, d_N\, d\, d_j.$$

The conditional probability term has the value one; therefore, the above expression is equal to

$$\sum_{j=1}^{N} h_j \cdot \int_0^{u_j} [\hat{q}_j + q_j - d_j] \cdot f_j(d_j) d\, d_j.$$

The fact that the Hessian of the objective function (8) is a positive semidefinite matrix is now established. In this regard, recall the formula for differentiating an integral of the following form:

$$\text{if } h(y) \equiv \int_{A(y)}^{B(y)} f(x, y)\, dx, \tag{A1}$$

then

$$\frac{dh(y)}{dy} = f(B(y), y) \frac{dB}{dy} - f(A(y), y) \frac{dA}{dy}$$

$$+ \int_{A(y)}^{B(y)} \frac{\partial f(x, y)}{\partial y} dx.$$

For the objective function,

$$I(\cdot) \equiv \sum_{j=1}^{N} h_j \cdot E([\hat{q}_j + q_j - d_j]^+)$$

$$= \sum_{j=1}^{N} h_j \cdot \int_0^{\hat{q}_j+q_j} (\hat{q}_j + q_j - d_j) f_j(d_j) d\, d_j . \tag{A2}$$

This implies

$$\frac{\partial I}{\partial q_j} = h_j \cdot \int_0^{\hat{q}_j+q_j} f_j(d_j) d\, d_j = h_j \cdot F_j(\hat{q}_j + q_j),$$

$$\frac{\partial^2 I}{\partial q_i \partial q_j} = 0 = \frac{\partial^2 I}{\partial q_j \partial q_i},$$

$$\frac{\partial^2 I}{\partial^2 q_j}(q_j) = h_j \cdot f_j(\hat{q}_j + q_j) \geq 0.$$

This concludes the verification.

The proof that the procedure presented in the section on the allocation problem terminates and then generates the global optimum solution to problem (8)–(10) after at most $N - \eta + 1$ iterations is now pursued. First recall that if the optimality procedure is required, the current state of the system has failed to satisfy the conditional (12) and thus is initialized with the assurance that

$$Z_\eta(0) \equiv \sum_{j=1}^{\eta} [F_j^{-1}(0) - \hat{q}_j] = \sum_{j=1}^{\eta} (-\hat{q}_j) = \sum_{j=1}^{N} [-\hat{q}_j]^+ < \Omega.$$

Thus, if

$$Z_\eta(-\gamma_{\eta+1}) \geq \Omega,$$

then, by the continuity of $Z_\eta(\cdot)$,

$$\exists\ \lambda^* \in [-\gamma_{\eta+1}, -\gamma_\eta) \ni Z_\eta(\lambda^*) = \Omega,$$

which may be obtained via binary search restricted to this interval. Otherwise, we continue with the knowledge that

$$Z_{\eta+1}(-\gamma_{\eta+1}) = Z_\eta(-\gamma_{\eta+1}) < \Omega,$$

since

$$Z_{\eta+1}(-\gamma_{\eta+1})$$

$$= Z_\eta(-\gamma_{\eta+1}) + \left[F_{\eta+1}^{-1}\left(\min\left\{1, \frac{\gamma_{\eta+1}}{h_{\eta+1}}\right\}\right) - \hat{q}_{\eta+1}\right]$$

$$= Z_\eta(-\gamma_{\eta+1}) + 0,$$

since

$$\frac{\gamma_{\eta+1}}{h_{\eta+1}} = F_{\eta+1}(\hat{q}_{\eta+1}) \text{ implies } F_{\eta+1}^{-1}\left(\min\left\{1, \frac{\gamma_{\eta+1}}{h_{\eta+1}}\right\}\right) - \hat{q}_{\eta+1} = 0.$$

Proceeding in this manner, either one obtains

$$Z_c(-\gamma_{c+1}) \geq \Omega \tag{A3}$$

for some $c < N$, or after at most $N - \eta + 1$ iterations arrives at the valid conditional

$$Z_N(\nu) \equiv \sum_{j=1}^{N} \left[F_j^{-1}\left(\min\left\{1, \frac{\max\{h_j : j \in \mathcal{N}\}}{h_j}\right\}\right) - \hat{q}_j\right]$$

$$= \sum_{j=1}^{N} [F_j^{-1}(1) - \hat{q}_j] = \sum_{j=1}^{N} (\bar{d}_j - \hat{q}_j) \geq \Omega,$$

since the conditional (13) failed to be satisfied prior to initiating the optimization process.

Assuming that the conditional (34) is satisfied at the $i = c^{\text{th}}$ iteration, then again by the continuity of $Z_c(\cdot)$

$$\exists\ \lambda^* \in [-\gamma_{c+1}, -\gamma_c) \ni Z_c(\lambda^*) = \Omega,$$

which may be obtained by numerical inversion. Let us first consider the capacity constraint (20) required by the Karush–Kuhn–Tucker conditions,

$$\sum_{j=1}^{N} q_j^* = \sum_{j=1}^{c} \zeta_j(\lambda^*) = Z_c(\lambda^*) = \Omega$$

by construction.

We proceed to verify the additional conditions by considering the indices

$$1 \leq j \leq c \ni \lambda^* > -h_j.$$

The gradient relationships (19) of the Karush–Kuhn–Tucker conditions become

$$h_j \cdot F_j\left(F_j^{-1}\left(-\frac{\lambda^*}{h_j}\right) - \hat{q}_j + \hat{q}_j\right) + \lambda^* - 0 + 0 = 0.$$

For the lower bound constraint (21), we need to consider the two cases. First, for $j \in \Gamma^-$ we have

$$q_j^* \equiv F_j^{-1}\left(-\frac{\lambda^*}{h_j}\right) - \hat{q}_j \geq 0 - \hat{q}_j \equiv [-\hat{q}_j]^+,$$

and for $j \in \Gamma^+$ we have

$$q_j^* \equiv F_j^{-1}\left(-\frac{\lambda^*}{h_j}\right) - \hat{q}_j \geq 0 \equiv [-\hat{q}_j]^+,$$

since

$$\lambda^* < -\gamma_c \leq -\gamma_j = -h_j \cdot F_j(\hat{q}_j)\ \ \forall\ j \leq c \text{ and } j \in \Gamma^+.$$

The upper bound constraints (22) are easily seen to be satisfied, since

$$q_j^* \equiv F_j^{-1}\left(-\frac{\lambda^*}{h_j}\right) - \hat{q}_j \leq \bar{d}_j - \hat{q}_j .$$

The complementarity conditions (23) and (24), along with the nonnegativity requirements (26), are obviously satisfied, since $\lambda_j^* \equiv 0 \equiv \beta_j^*$.

We now consider the indices $1 \le j \le c \ni \lambda^* \le -h_j$. The gradient relationships (19) are now

$$h_j \cdot F_j(\bar{d}_j) + \lambda^* - 0 + (-\lambda^* - h_j)$$
$$= h_j \cdot 1 + \lambda^* - 0 + (-\lambda^* - h_j) = 0.$$

For the lower bound constraint (21), we need to consider the two cases. First, for $j \in \Gamma^-$ we have

$$q_j^* \equiv \bar{d}_j - \hat{q}_j \ge 0 - \hat{q}_j \equiv [-\hat{q}_j]^+,$$

and for $j \in \Gamma^+$ we have

$$q_j^* \equiv \bar{d}_j - \hat{q}_j \ge 0 \equiv [-\hat{q}_j]^+,$$

since by reformulation, if necessary, we were assured that the inequality (11) was not satisfied for any j. The upper bound constraints (22) are now satisfied with equality $\bar{d}_j - (q_j^* + \hat{q}_j) = 0$. This equality immediately implies that the complementarity conditions (24) are satisfied, and $\lambda_j^* \equiv 0$ validates the complementarity equations (23) along with the nonnegativity requirements (26) for the λ_j multipliers, while $\lambda^* \le -h_j$ immediately implies nonnegativity for the β_j multipliers $\beta_j^* \equiv -\lambda^* - h_j$, which completes the verification for this set of indices.

It remains to validate the Karush–Kuhn–Tucker conditions for the indices $j = c + 1, \cdots, N$. First observe that for this set of indices, we have that $j > \eta$, implying that $j \in \Gamma^+$ and that $\hat{q}_j > 0$. Also observe that if $c = N$, this **FOR** loop would never have been initiated and the procedure would have terminated. The gradient relationships (19) now become

$$h_j \cdot F_j(0 + \hat{q}_j) + \lambda^* - [\lambda^* + h_j \cdot F_j(\hat{q}_j)] + 0 = 0.$$

The lower bound constraint (21) is now satisfied with equality, since for $j \in \Gamma^+$ we have $q_j^* \equiv 0 \equiv [-\hat{q}_j]^+$, and thus the complementarity conditions (23) for the λ_j multipliers are also satisfied. Again, by reformulation if necessary, we were assured that the inequality (11) was not satisfied for any j. Therefore, the upper bound constraints (22) $\bar{d}_j - (0 + \hat{q}_j) \ge 0$ are satisfied. Since $\beta_j^* \equiv 0$, the complementarity equations (23), along with the nonnegativity requirements (26) for the β_j multipliers, are apparent. The expressions

$$-\gamma_N \le -\gamma_{N-1} \le \cdots \le -\gamma_{c+1} \text{ and } \lambda^* \in [-\gamma_{c+1}, -\gamma_c)$$

immediately imply nonnegativity for the λ_j multipliers

$$\lambda_j^* \equiv \lambda^* + h_j \cdot F_j(\hat{q}_j),$$

which completes the verification for this set of indices and consequently for the procedure as a whole. Thus, these values of q_j, λ, λ_j, and β_j are optimal for the convex programming problem.

Acknowledgments

The authors gratefully acknowledge the careful technical review of the paper by Dr. Brenda Dietrich and the directions suggested by Dr. Stanley Raider, which greatly enhanced the readability of the final version of the paper.

References

1. K. R. Baker, M. J. Magazine, and H. L. W. Nuttle, "The Effect of Commonality on Safety Stock in a Simple Inventory Model," *Manage. Sci.* **32**, No. 8, 982–988 (1986).
2. Y. Gerchak, M. J. Magazine, and A. B. Gamble, "Component Commonality with Service Level Requirements," *Manage. Sci.* **34**, No. 6, 753–760 (1988).
3. Y. Gerchak and M. Henig, "An Inventory Model with Component Commonality," *Oper. Res. Lett.* **5**, No. 3, 157–160 (1986).
4. Y. Gerchak and M. Henig, "Component Commonality in Assemble-to-Order Systems: Models and Properties," *Naval Res. Logist.* **36**, 61–68 (1989).
5. R. Srinivasan, R. Jayaraman, R. Roundy, and S. Tayur, "Procurement of Common Components in a Stochastic Environment," *Research Report RC-18580*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, December 11, 1992.
6. P. H. Zipkin, "Simple Ranking Methods for Allocation of One Resource," *Manage. Sci.* **26**, No. 1, 34–43 (1980).
7. H. Luss and K. Gupta, "Allocation of Effort Resources Among Competing Activities," *Oper. Res.* **23**, 360–366 (1975).

Received March 12, 1992; accepted for publication May 13, 1993

Stephen J. Grotzinger *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598.*

Dr. Grotzinger is a Research Staff Member in the Manufacturing Logistics Group of Manufacturing Research at the Thomas J. Watson Research Center. His research involves the development and analysis of mathematical models for automated manufacturing environments. His prior work has been in the areas of complementarity theory; characterizations of best approximates; isotonic regression, nonlinear programming; setup, sequencing, and Petri Net modeling of highly concurrent planar assembly and test systems; sampling strategies for contamination control; and assemble to forecast inventory models. Dr. Grotzinger received a B.S. in mathematics from the University of Delaware, and M.S. and Ph.D. degrees in mathematics from Rensselaer Polytechnic Institute. He was a member of the Department of Mathematical Sciences at George Mason University, a Postdoctoral Research Associate in the Center of Applied Mathematics of the National Bureau of Standards, and a Faculty Research Fellow at the Engineer Topographic Laboratories of the U.S. Army Corps of Engineers; he joined IBM in 1984. Dr. Grotzinger's publications have appeared in *Mathematical Programming, Applied Mathematics and Optimization, Operations Research, IEEE Transactions on Robotics and Automation, Environmental Science, Institute of Environmental Sciences, International Journal of Flexible Manufacturing Systems,* and *IIE Transactions.* He is a member of ORSA.

Ramesh Srinivasan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (SRAMESH at YKTVMV, sramesh@watson.ibm.com).* Dr. Srinivasan received a B.E. (Honours) in mechanical engineering from the University of Madras, India, in 1981. He received an M.S. in industrial and systems engineering from Ohio University in 1982 and a Ph.D in industrial engineering in 1986 from Purdue University. Dr. Srinivasan joined the IBM Thomas J. Watson Research Center in the Manufacturing Research Department in 1986 as a Research Staff Member. Since then he has worked in computer-aided process planning and modeling and optimization of manufacturing systems. His recent work has been in production planning under uncertainty and business process reengineering involving the development and deployment of new methodologies in IBM. He is a member of ORSA and IEEE.

Ramakrishna Akella *Graduate School of Industrial Administration, Pittsburgh, Pennsylvania 15213 (ra1c+@andrew.cmu.edu).* Professor Akella is an Associate Professor in the Graduate School of Industrial Administration and the Robotics Institute at Carnegie Mellon University in Pittsburgh. He is also associated with the Center for Integrated Manufacturing Decision Systems, and serves as Director of the Laboratory for Manufacturing, Design, and Enterprise Management. Professor Akella completed his B.S. in electronics at I.I.T. Madras in 1976, and received a Ph.D. in systems at I.I.Sc. Bangalore in 1982. He was a Postdoctoral Visitor at Harvard University during 1982 and a Postdoctoral Associate at M.I.T., Cambridge, from 1983 to 1985. He has visited with the Sloan/Engineering Leaders for Manufacturing Program at MIT, and is currently a Visiting Associate Professor at the University of California at Berkeley. Professor Akella's research interests are in computer-integrated manufacturing, and range from cell- and factory-level design and control to enterprise-wide coordination. He has received several awards, including the IBM Faculty Award, and has interacted extensively with industries, including corporations such as IBM, Digital, and TI, along with various Japanese and European companies. He has also lectured by invitation in Europe and Japan. Professor Akella serves as Associate Editor for *Operations Research* and *IEEE Transactions on Semiconductor Manufacturing.* He has also served as Guest Editor for the *IEEE Robotics and Automation* special issue on manufacturing systems. He has published in a variety of journals such as *Management Science, Operations Research, IEEE Transactions on Robotics and Automation, IEEE Transactions on Automatic Control, IEEE Transactions on Semiconductor Manufacturing, IEEE Transactions on Components, Hybrids and Manufacturing Technology, IBM Journal of Research and Development,* and *International Journal of Production Research.*

Srinivas Bollapragada *Graduate School of Industrial Administration, Pittsburgh, Pennsylvania 15213 (srin@isl1.ri.cmu.edu, sb4o@andrew.cmu.edu).* Mr. Bollapragada is a Ph.D. student in the interdisciplinary program between the Graduate School of Industrial Administration and the Robotics Institute, Carnegie Mellon University, in management of manufacturing and automation. He received a B.Tech degree in electrical engineering (electronics) from I.I.T. Madras in 1987, an M.E. in electrical communication engineering from I.I.Sc. Bangalore in 1989, and an M.S. in management of manufacturing and automation from Carnegie Mellon University in 1992. Mr. Bollapragada's research interests include inventory theory and simulation, stochastic models in manufacturing, computer vision, and distortion-invariant pattern recognition.