

SUPPLY CONTRACTS WITH SERVICE LEVEL REQUIREMENTS ¹

Carmen Del Vecchio ^{*,2}
Ioannis Ch. Paschalidis ^{**,3}

^{*} *Università degli Studi del Sannio*
^{**} *Boston University*

Abstract: We consider the problem of optimizing the terms of a contract between a supplier and a buyer to minimize overall system costs. The supplier seeks to minimize inventory costs in a supply chain which includes assembly operations and features stochastic models of available capacity at its various stages. The buyer needs to satisfy demand which is modeled as a rather general stochastic process and requires from the supplier the enforcement of certain Quality of Service constraints. We build on our earlier work for solving the supplier's inventory control problem and develop an algorithm that enables supplier and buyer to jointly optimize contract parameters. This can be done in a distributed fashion without revealing to each other their corresponding cost structures. *Copyright*© 2005 IFAC.

Keywords: Manufacturing systems, inventory control, large deviations, optimization problems.

1. INTRODUCTION

We study a contract model between one supplier and one buyer. The supplier provides a product to the buyer *under explicit Quality of Service (QoS) constraints* that provide probabilistic guarantees on satisfying demand. The buyer provides to the supplier a stochastic model of demand based on which the supplier organizes the supply chain. Our objective is to find a mutually beneficial demand model, namely, a model that minimizes overall system costs incurred by both the supplier and the buyer. There is a large literature on supply chain contracts; contributions can be roughly split into two classes; see (Corbett and Tang (1999)) for a

review. The first assumes that the contract terms are fixed and explores agents' (i.e., supplier's and buyer's) optimal actions under these terms. The impact of contract parameters on agents' costs and profits are mostly explored. The second class assumes agents' optimal policies under a contract as given and explores whether the terms of the contract can be modified to improve the supply chain. Our model lies within this second category.

Under the contract model we consider the supplier inventory problem is solved minimizing inventory costs while providing QoS guarantees to the buyer. In two previous works an approach for solving the supplier's inventory problem has been developed: in (Paschalidis and Liu (2003)) for a serial supply chain and in (DelVecchio and Paschalidis (2003)) for the more general case of supply chains involving assembly operations. In this paper, we analyze the economic incentives the contract can provide. Our analysis allows demand and production to be modeled by *autocorrelated stochastic processes*, in order to capture strong *temporal dependencies* that are often observed in these processes (e.g., due to seasonal effects,

¹ Research partially supported by the NSF under a CAREER award ANI-9983221 and grants DMI-0330171, DMI-0300359, CNS-0435312, and by the ARO under the MURI 2001 Grant DAAD19-01-1-0465 to the Center for Networked Communicating Control Systems.

² Piazza Roma, 24 - 82100 Benevento, Italy, e-mail: c.delvecchio@unisannio.it

³ Corresponding author. Center for Information and Systems Engineering, Boston University, Boston, MA 02215, e-mail: yannisp@bu.edu, url: <http://ionia.bu.edu/>.

sales events, weather patterns, machine failures, etc.). Moreover, we capture QoS guarantees by requiring stockout probabilities to stay bounded below a given desirable level. We will refer to such constraints as *QoS* or *service-level* constraints. We believe that they provide a more direct way of taking QoS considerations into account rather than penalizing backorders or lost sales (as it is done in most of the literature). Under our modeling assumptions analyzing supplier stockout probabilities exactly is intractable. In our previous work we relied on *large deviations techniques* which yield asymptotically tight approximations of the stockout probabilities as they approach zero. These approximations also gave rise to an analytical approximation of expected inventory costs. Then, the overall inventory problem can be transformed to a deterministic (nonlinear) optimization problem whose objective function (inventory cost) and constraints (service-level constraints) are given by closed-form analytical expressions. Solving this problem using standard optimization techniques yields the supplier stock levels we seek.

Building on this work, we estimate the derivative of the supplier inventory cost with respect to demand model parameters and use it to evaluate contract incentives in order to achieve a demand configuration that can lead to lower supply chain costs. The algorithm we propose to that end can be followed in a distributed fashion with neither the supplier nor the buyer revealing to each other their corresponding cost structures.

The remainder of this paper is organized as follows. In Sec. 2 we introduce the manufacturer inventory cost problem and present the large deviations approach. In Sec. 3 we present the contract model. Numerical results are in Sec. 4. We conclude with some brief remarks in Sec. 5.

2. SUPPLY CHAIN CONTROL

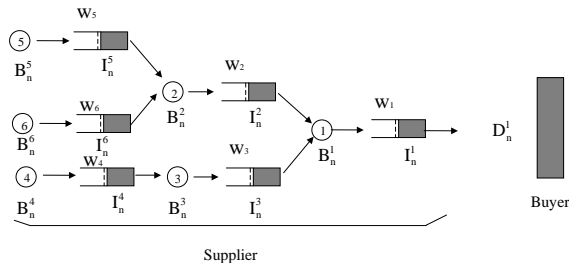


Fig. 1. A supply chain with assembly operations. The system forms a tree rooted at stage 1.

The supplier's production process consists of an assembly network that produces a single product and is formed by N production facilities or stages (see Fig. 1). Buyer demand is met from the finished goods inventory maintained in front of stage 1 and is backordered if inventory is not

available. Each stage has exactly one successor stage so that the assembly system forms a tree rooted at stage 1. Except from the most upstream stages (the leaves of the tree) which are fed with an infinite supply of raw material, every stage requires one unit of intermediate product (i.e., a component) from each immediate predecessor stage in order to assemble one product unit (e.g., in the system of Fig. 1 one unit of finished product requires one intermediate product unit from stage 2 and one from stage 3). That is, we assume (without loss of generality) that exactly one unit of each component is required for producing one unit of finished product.

We use the following notation. $S(i)$ denotes the unique immediate successor of stage i , $i = 1, \dots, N$, where $S(1) = 0$. $A(i)$ denotes the set of all successors (immediate and more downstream) of stage i , $i = 1, \dots, N$, where $A(1) = \{0\}$. We let $P(i)$ denote the set of immediate predecessors of stage i , $i = 1, \dots, N$, where $P(i) = \emptyset$ for all stages i that are leaves of the tree. $B(i)$ denotes the set of all predecessors (immediate and more upstream) of stage i , where again $B(i) = \emptyset$ for all leaf stages i . \mathcal{F} denotes the set of all leaf stages, that is stages i for which $P(i) = \emptyset$. We assume that \mathcal{F} contains at least stage N . Finally, $\mathcal{N} = \{1, \dots, N\}$ denotes the set of all stages of the assembly network.

We assume a periodic review policy where time is divided into time slots of equal duration. Let B_n^i , $i = 1, \dots, N$, denote the amount that the facility of stage i can produce during time slot n . D_n^1 denotes the amount of external orders arriving at stage 1 during time slot n . Each stage has an inventory buffer for its final product and from which the successor stage draws material for its production. We let I_n^i , $i = 1, \dots, N$, denote the inventory in front of stage i at the beginning of time slot n . In intermediate stages, $i = 2, \dots, N$, the inventory is constrained to be nonnegative. Inventory at stage 1 can take negative values to denote backordered demand, in particular, when I_n^1 is negative $-I_n^1$ is equal to the amount of backordered demand. We denote by X_n^i the sum of the inventory of stage i and of all its successors at time n , i.e., $X_n^i = I_n^i + \sum_{k \in A(i)} I_n^k$. We will refer to this quantity as the stage- i *echelon inventory*. We assume that all stages adopt an echelon base-stock policy: stage i produces when X_n^i falls below a certain threshold w^i and idles otherwise. It is implied that $w^i \leq w^k$ for all $i \in \mathcal{N}$ and $k \in P(i)$.

The demand process $\{D_n^1; n \in \mathbb{Z}\}$ and the production capacity processes $\{B_n^i; n \in \mathbb{Z}\}$, $i = 1, \dots, N$, are mutually independent, possibly autocorrelated, arbitrary stationary stochastic processes satisfying certain mild technical conditions (a large deviations principle) that hold for renewal processes, Markov-modulated processes, and gen-

eral stationary processes with mild mixing conditions. For stability purposes assume

$$\mathbf{E}[D_n^1] < \min_{i=1,\dots,N} \mathbf{E}[B_n^i], \quad (1)$$

which by stationarity carries over to all time slot n . Stability can be shown under the base-stock policy we consider by using techniques from (Baccelli and Liu (1992)). On a notational remark, in the sequel we will denote by $\Lambda_{D^1}^{*+}(a)$ (resp. $\Lambda_{D^1}^{*-}(a)$) the large deviations rate function associated with the tail probability $\mathbf{P}[\sum_{i=1}^n D_i^1 \geq na]$ (resp. $\mathbf{P}[\sum_{i=1}^n D_i^1 \leq na]$). We adopt similar notation for all production processes as well.

In what follows we will show how to find the safety stocks w^i for each stage that minimize the manufacturer's expected inventory costs and guarantee that for all stages i the steady-state stockout probabilities $\mathbf{P}[X_n^i \leq 0]$ do not exceed desirable small values ϵ_i , respectively.

Echelon inventory levels in the assembly network evolve as follows:

$$\begin{aligned} X_{n+1}^i &= \min \{X_n^i - D_n^1 + B_n^i, w^i, \min_{k \in P(i)} \{X_n^k \\ &\quad - D_n^1 + I_n^k\}\}, \quad i \in \mathcal{N} \setminus \mathcal{F}, \\ X_{n+1}^i &= \min \{X_n^i - D_n^1 + B_n^i, w^i\}, \quad i \in \mathcal{F} \end{aligned}$$

Define the shortfall of the echelon inventory as:

$$Y_n^i = w^i - X_n^i, \quad i \in \mathcal{N}.$$

In terms of the shortfalls the dynamics of the system can be written as:

$$\begin{aligned} Y_{n+1}^i &= \max \{0, Y_n^i + D_n^1 - B_n^i, \max_{k \in P(i)} \{Y_n^k \\ &\quad + D_n^1 - (w^k - w^i)\}\}, \quad i \in \mathcal{N} \setminus \mathcal{F}, \quad (2) \\ Y_{n+1}^i &= \max \{0, Y_n^i + D_n^1 - B_n^i\}, \quad i \in \mathcal{F}. \quad (3) \end{aligned}$$

2.1 Large Deviation Analysis of the Stockout Probability

In this section we briefly present the large deviations result for the steady state probability $\mathbf{P}[Y^1 \geq w^1]$, which is equal to the steady-state stockout probability $\mathbf{P}[X^1 \leq 0]$. In the interest of space we only state the main theorem. A proof can be found in (DelVecchio and Paschalidis (2003)).

In what follows we will be using

$$\mathcal{O}^i = \left\{ (\xi_i, \boldsymbol{\xi}_{A(i)}) \mid \xi_i \in [0, 1], \boldsymbol{\xi}_{A(i)} \in [0, 1]^{|A(i)|}, \right. \\ \left. \xi_i + \boldsymbol{\xi}'_{A(i)} \mathbf{e} = 1 \right\}, \quad \forall i \in \mathcal{N},$$

where $|A(i)|$ denotes the cardinality of the set $A(i)$, \mathbf{e} is the vector of all ones, and prime denotes transpose. Moreover for any stage i let $A(i) = \{1, \dots, |A(i)|\}$ and we will be writing $\boldsymbol{\xi}_{A(i)}$ for the vector $(\xi_1, \dots, \xi_{|A(i)|})$.

Theorem 1. Assume the hedging points w^1, w^2, \dots, w^N in the assembly system satisfy

$$w^i = \beta_i w^1, \quad i = 2, \dots, N,$$

where β_i are constants and $1 \leq \beta_h \leq \beta_i$ for all $h \in A(i)$ ($\beta_1 \triangleq 1$). The steady state shortfall Y^1 of the assembly system satisfies

$$\lim_{w^1 \rightarrow \infty} \frac{1}{w^1} \log \mathbf{P}[Y^1 \geq w^1] = -\theta_{G,1}^*, \quad (4)$$

where $\theta_{G,1}^*$ is given by

$$\begin{aligned} \theta_{G,1}^* &= \min \left[\inf_{a>0} \frac{1}{a} \inf_{x_0 - x_1 = a} (\Lambda_{D^1}^{*+}(x_0) + \Lambda_{B^1}^{*-}(x_1)), \right. \\ &\quad \min_{2 \leq i \leq N} \left\{ \inf_{a>0} \frac{1}{a} \inf_{x_0 - \sum_{h \in A(i)} \xi_h x_h - \xi_i x_i = a \beta_i} \right. \\ &\quad \left. \left. (\xi_i, \boldsymbol{\xi}_{A(i)}) \in \mathcal{O}^i \right. \right. \\ &\quad \left. \left. \left(\Lambda_{D^1}^{*+}(x_0) + \sum_{h \in A(i)} \xi_h \Lambda_{B^h}^{*-}(x_h) \right. \right. \right. \\ &\quad \left. \left. \left. + \xi_i \Lambda_{B^i}^{*-}(x_i) \right) \right) \right]. \quad (5) \end{aligned}$$

Thm. 1 provides the asymptotic decay rate for the overflow probability of the shortfall at stage 1 as its base-stock level grows large. More intuitively, the stockout probability can be approximated as:

$$\mathbf{P}[X^1 \leq 0] = \mathbf{P}[Y^1 \geq w^1] \sim e^{-\theta_{G,1}^* w^1}. \quad (6)$$

The result of Thm. 1 can be easily generalized to other stages as follows.

Corollary 2. Assume the safety stock levels w^k , for $k \in B(i)$, in the assembly network satisfy $w^k = \beta_{i_k} w^i$, $k \in B(i)$, where β_{i_k} are constants and $1 \leq \beta_{i_l} \leq \beta_{i_k}$ for all $l \in A(k)$ and $k \in B(i)$. The steady-state shortfall Y^i satisfies:

$$\lim_{w_i \rightarrow \infty} \frac{1}{w_i} \log \mathbf{P}[Y^i \geq w^i] = -\theta_{G,i}^* \quad (7)$$

where $\theta_{G,i}^*$ is determined by

$$\begin{aligned} \theta_{G,i}^* &= \min \left[\inf_{a>0} \frac{1}{a} \inf_{x_0 - x_i = a} (\Lambda_{D^1}^{*+}(x_0) + \Lambda_{B^i}^{*-}(x_i)), \right. \\ &\quad \min_{k \in B(i)} \left\{ \inf_{a>0} \frac{1}{a} \inf_{x_0 - \sum_{h \in A(k) \setminus A(i)} \xi_h x_h - \xi_k x_k = a \beta_{i_k}} \right. \\ &\quad \left. \left. (\xi_k, \boldsymbol{\xi}_{A(k) \setminus A(i)}) \in \mathcal{O}^k \right. \right. \\ &\quad \left. \left. \left(\Lambda_{D^1}^{*+}(x_0) + \sum_{h \in A(k) \setminus A(i)} \xi_h \Lambda_{B^h}^{*-}(x_h) \right. \right. \right. \\ &\quad \left. \left. \left. + \xi_k \Lambda_{B^k}^{*-}(x_k) \right) \right) \right]. \quad (8) \end{aligned}$$

To make this approximation more accurate, especially for relatively large stockout probabilities (i.e., small safety stock w^1), we will introduce a prefactor in front of the exponential. We will use the following refined approximation:

$$\mathbf{P}[Y^1 \geq w^1] \approx f_1(w^1, \boldsymbol{\beta}) e^{-\theta_{G,1}^* w^1}, \quad (9)$$

where the prefactor $f_1(w^1, \beta)$ is a function of w^1 and $\beta = (\beta_2, \dots, \beta_N) = (\frac{w^2}{w^1}, \dots, \frac{w^N}{w^1})$. Notice that both the prefactor and the exponent are functions of the hedging point vector $\mathbf{w} = (w^1, \dots, w^N)$.⁴ We elect to use a function $f_1(w^1, \beta)$ which is piecewise linear in w^1 and β because it is easier to obtain and yields very satisfactory results. Some computational remarks and details on the prefactor evaluation can be found in (DelVecchio and Paschalidis (2003)) and in (Paschalidis and Liu (2003)). Similar arguments also lead to an appropriate prefactor in approximating stockout probabilities at stages $2, \dots, N$ using Corollary 2.

2.2 Approximating the expected inventory cost

We will now turn our attention to approximating expected inventory costs. We assume a linear cost structure. Let h_i be the holding cost for echelon- i inventory for all $i = 1, \dots, N$. Noting that expected echelon- i inventory is given by $h_i(E[I^i] + \sum_{h \in A(i) \setminus \{1\}} E[I^h] + E[(I^1)^+])$ where $(I^1)^+ = \max(I^1, 0)$, the total expected inventory cost is given by:

$$\left(\sum_{i=1}^N h_i \right) \mathbf{E}[(I^1)^+] + \sum_{i=2}^N h_i \left(\mathbf{E}[I^i] + \sum_{j \in A(i) \setminus \{1\}} \mathbf{E}[I^j] \right). \quad (10)$$

We have

$$\begin{aligned} \mathbf{E}[(I^1)^+] &= \mathbf{E}[(w^1 - Y^1)^+] \\ &= w^1 - \mathbf{E}[Y^1] + \mathbf{E}[\max(0, Y^1 - w^1)]. \end{aligned} \quad (11)$$

Using the tail distribution of Y^1 given in Eq. (9)

$$\begin{aligned} \mathbf{E}[\max(0, Y^1 - w^1)] &= \int_0^\infty \mathbf{P}[Y^1 - w^1 > y] dy \\ &\approx f_1(w^1, \beta) \frac{e^{-\theta_{G,1}^* w^1}}{\theta_{G,1}^*}. \end{aligned} \quad (12)$$

For all $i \geq 2$ we have $I^i = (w^i - Y^i) - (w^{S(i)} - Y^{S(i)})$, which implies:

$$\mathbf{E}[I^i] = (w^i - \mathbf{E}[Y^i]) - (w^{S(i)} - \mathbf{E}[Y^{S(i)}]). \quad (13)$$

Thus, combining (10), (11), (12), and (13) we obtain the following expression for expected inventory costs:

$$\sum_{i=1}^N h_i (w^i - \mathbf{E}[Y^i]) + \left(\sum_{i=1}^N h_i \right) f_1(w^1, \beta) \frac{e^{-\theta_{G,1}^* w^1}}{\theta_{G,1}^*}. \quad (14)$$

To obtain an analytical approximation for the inventory cost we are left with computing $\mathbf{E}[Y^i]$. We will establish some structural properties for $\mathbf{E}[Y^i]$. To this end for each stage i ($i \in \mathcal{N}$) we define $\Delta_i \triangleq (w^i - w^{S(i)})$; this is a well defined quantity as each node has only one successor (we set $\Delta_1 \triangleq 0$). In the following proposition we state that for each stage i , $\mathbf{E}[Y^i]$ is a convex, monotonically nonincreasing, function of all Δ_k with $k \in B(i)$. Again, the proof is omitted due to space limitations.

Proposition 3. Consider the assembly system and let w^1, \dots, w^N be the corresponding hedging points satisfying $0 \leq w^k \leq w^i$ with $k \in A(i)$. For all stages i which do not have any predecessor (i.e., $i \in \mathcal{F}$), $\mathbf{E}[Y^i]$ is a constant function of $(\Delta_2, \dots, \Delta_N)$. Furthermore, for all stages which have at least one predecessor (i.e., $i \in \mathcal{N} \setminus \mathcal{F}$), $\mathbf{E}[Y^i]$ is a function of Δ_k for all $k \in B(i)$, and it is convex and monotonically nonincreasing in every coordinate.

Motivated by these properties of $\mathbf{E}[Y^i]$ we will approximate it using a piecewise linear convex function g_i which is a function of Δ_k for all $k \in B(i)$. More specifically we will use the following approximation

$$\mathbf{E}[Y^i] \approx g_i(\Delta_{B(i)(1)}, \dots, \Delta_{B(i)(|B(i)|)}), \quad (15)$$

where $B(i)^{(j)}$ denotes the j th element of $B(i)$. As before, we refer to (DelVecchio and Paschalidis (2003)) for details on how to obtain $g_i(\cdot)$.

2.3 The inventory optimization problem

We can now pose the problem of optimizing expected inventory costs subject to maintaining service level constraints. Using the approximating expression for the expected inventory cost in (14) we have the following optimization problem.

$$\begin{aligned} \min \quad & \sum_{i=1}^N h_i (w^i - g_i(\Delta_{B(i)(1)}, \dots, \Delta_{B(i)(|B(i)|)})) \\ & + \left(\sum_{i=1}^N h_i \right) f_1(w^1, \dots, \frac{w^N}{w^1}) \frac{e^{-\theta_{G,1}^* w^1}}{\theta_{G,1}^*} \\ \text{s.t.} \quad & \mathbf{P}[Y^i \geq w^i] = e^{-\theta_{G,i}^* w^i} \\ & f_i(w^i, \frac{w^{B(i)(1)}}{w^i}, \dots, \frac{w^{B(i)(|B(i)|)}}{w^i}) \leq \epsilon_i, \quad \forall i \\ & w^i \leq w^k, \quad \forall i \text{ and } \forall k \in P(i). \end{aligned} \quad (16)$$

The decision variables are the hedging points w^i and the constraints correspond to service level constraints imposed at all nodes of the assembly network. This problem can be solved analytically using standard nonlinear programming techniques.

⁴ $\theta_{G,1}^*$ (cf. (5)) is a function of \mathbf{w} also.

3. THE CONTRACT MODEL

Next we put forth our contract model. The buyer and the supplier interact as follows. The buyer determines (a model of) the demand D_n^1 according to market conditions and associates a cost (not necessarily monetary) to any change in demand parameters. The supplier production process is an assembly network following a base stock production policy. Safety stocks are set to minimize inventory costs – the cost function in problem (16) – using as input the buyer’s demand model, the service level agreement – the constrain in problem (16) – and (a model of) the production capacities B_n^i , $i = 1, \dots, N$. Supplier and buyer share information on their demand parameters and service level constraints, as well as on their cost functions. We will explore how demand parameters can be adjusted according to both supplier and buyer cost structures so that any generated savings can be shared among the two.

We assume that the buyer’s demand is an m -state Markov-modulated process, with transition probability matrix \mathbf{P}_D and demand levels at each state given by the vector \mathbf{r}_D . The manufacturer’s production capacity at each node i of the assembly network is modeled by an m^i -state Markov-modulated process with transition probability matrix \mathbf{P}_{B^i} and capacities at each state given by the vector \mathbf{r}_{B^i} , $i = 1, \dots, N$. We will also assume that any mutually agreed adjustments in \mathbf{r}_D keep the mean demand $\mathbf{E}[D]$ constant since otherwise the retailer would be unable to satisfy demand in the long term. We denote by $\overline{\mathbf{E}[D]}$ this constant value.

We denote by $C_R(\mathbf{r}_D)$ the retailer cost associated with any change of \mathbf{r}_D . In particular,

$$C_R(\mathbf{r}_D) = \sum_{i=1}^m \zeta_i (r_{D,i} - \bar{r}_{D,i})^2, \quad (17)$$

where ζ_i expresses the retailer’s cost for changing the demand level in the i^{th} state of the demand process. $\bar{\mathbf{r}}_D$ is the vector of demand levels initially determined and to which the retailer associates zero cost. The manufacturer’s inventory cost function, $C_M(\mathbf{r}_D)$, is the optimal value function of problem (16). Our objective is to find the vector \mathbf{r}_D that solves the following problem:

$$\begin{aligned} \min f(\mathbf{r}_D) &= C_R(\mathbf{r}_D) + C_M(\mathbf{r}_D) & (18) \\ \text{s.t. } \mathbf{E}[D] &= \overline{\mathbf{E}[D]} \\ r_{D,i} &\geq 0, \quad i = 1, \dots, m. \end{aligned}$$

We indicate by Ξ the polyhedron containing all feasible \mathbf{r}_D for problem (18), i.e.,

$$\Xi = \{\mathbf{r}_D \mid \mathbf{E}[D] = \overline{\mathbf{E}[D]}, r_{D,i} \geq 0 \forall i\}.$$

Problem (18) is a nonlinear optimization problem over a convex set and can be solved using the conditional gradient method in (Bertsekas (1999)).

The method consists of generating sequences of feasible points \mathbf{r}_D^k by searching along descent directions. Hence, the vector of demand levels \mathbf{r}_D changes according to the following algorithm:

$$\mathbf{r}_D^{k+1} = \mathbf{r}_D^k + \alpha^k (\bar{\mathbf{r}}_D^k - \mathbf{r}_D^k), \quad k = 0, 1, \dots, \quad (19)$$

where $\bar{\mathbf{r}}_D^k - \mathbf{r}_D^k$ is a feasible direction of descent. More specifically, it satisfies the descent condition $\nabla f(\mathbf{r}_D^k)'(\bar{\mathbf{r}}_D^k - \mathbf{r}_D^k) < 0$ and, assuming that \mathbf{r}_D^k is feasible, the vector \mathbf{r}_D^{k+1} remains feasible for small enough stepsizes α^k . The most straightforward way to generate feasible descent directions is to solve the optimization problem:

$$\begin{aligned} \min \nabla f(\mathbf{r}_D^k)'(\mathbf{r}_D - \mathbf{r}_D^k) & & (20) \\ \text{s.t. } \mathbf{r}_D \in \Xi, & \end{aligned}$$

and obtain $\bar{\mathbf{r}}_D^k$ as the solution, that is,

$$\bar{\mathbf{r}}_D^k = \arg \min_{\mathbf{r}_D \in \Xi} \nabla f(\mathbf{r}_D^k)'(\mathbf{r}_D - \mathbf{r}_D^k).$$

Since Ξ is a polyhedron, problem (20) is a linear programming problem which can be solved very efficiently (in polynomial time). Finally, the stepsize α^k in expression (19) is determined using Armijo’s rule; see (Bertsekas (1999)) in order to guarantee convergence to a local minimum.

To apply the conditional gradient method to problem (18) we need the cost function gradient. The gradient of the buyer’s cost can be easily derived from (17). We evaluated the gradient of the manufacturer inventory costs using finite differences.

It is interesting to observe that the above algorithm can be used in a *distributed* fashion. In particular, the buyer can take charge of performing iteration (19) and solving subproblem (20). At each iteration, the buyer presents the manufacturer with the new demand levels \mathbf{r}_D^k and the manufacturer responds with the gradient $\nabla C_M(\mathbf{r}_D^k)$ of its cost function. This gradient information can be provided in the form of appropriate incentives to the buyer. The buyer uses this gradient information to solve subproblem (20) and compute \mathbf{r}_D^{k+1} . Note that neither the manufacturer nor the buyer need to know each other’s cost structures.

We notice that in the contract model we presented we evaluated changes in the demand parameters, while we can also consider changes in other contract parameters, such as the service levels. In that case the retailer cost function would be a function of the ϵ_i ’s and the manufacturer would evaluate inventory costs savings due to less strict service levels parameters. The contract model analysis for this case is similar and straightforward.

Table 1. The number of iterations is reported in the 1st column. Columns 2–4 report the demand levels. The 5th column reports the cost $f(\mathbf{r}_D)$ (cf. (18)). Columns 6–8 provide $\nabla f(\mathbf{r}_D)$. Finally, the last column lists the corresponding optimal hedging points obtained by solving (16).

Iteration	Demand Vector			Total costs	Cost derivative vector			\mathbf{w}^*
	r_1	r_2	r_3		$\nabla f_{r_{d1}}$	$\nabla f_{r_{d2}}$	$\nabla f_{r_{d3}}$	
1	2	8	0.5	230.94	126.00	141.73	161.13	(63.87, 123.56, 63.87)
5	2.37	5.89	1.31	161.13	-5.62	34.73	36.32	(37.66, 97.36, 37.67)
10	2.11	4.73	2.06	133.43	-29.88	-1.96	0.62	(26.76, 86.46, 26.77)
15	2.43	4.56	1.93	132.71	-5.57	3.94	6.57	(26.37, 86.05, 26.37)
20	2.19	4.06	2.35	126.64	0.19	-2.36	0.74	(23.96, 83.67, 23.88)
25	2.12	4.23	2.27	125.80	0.15	-2.29	0.72	(23.76, 83.46, 23.77)
30	2.11	4.32	2.26	125.71	0.13	-2.24	0.71	(23.76, 83.46, 23.77)




Fig. 2. Example 1.

4. NUMERICAL RESULTS

In this section we present numerical results we obtained solving the problem (18) by applying the algorithm (19). The supplier production process consists of three assembly phases. The initial retailer demand vector and the production capacity parameters of the supplier production stages are depicted in Fig. 2; all processes are Markov-modulated and a service level constraint with $\epsilon = 1e - 5$ has been imposed. The retailer's cost parameters are equal to $\zeta = (0.5, 0.3, 0.2)$. Numerical results after 30 iterations (where the algorithm reached a local minimum) are reported in Table 1. Clearly, our approach can lead to significant cost savings.

5. CONCLUSIONS

We presented an approach for a supplier and a buyer to interact in order to satisfy the buyer's demand and minimize the supply chain inventory costs. We assumed a rather general setting where (i) the supply chain includes assembly operations, (ii) demand and production capacity models are stochastic and accommodate strong temporal dependencies, and (iii) the supplier maintains probabilistic service level constraints. In this setting,

we prescribed how the supplier and the buyer can minimize their overall costs in order to satisfy demand and efficiently operate the supply chain. Our approach relies on asymptotically tight approximations of inventory stockout probabilities which were developed drawing from the theory of large deviations. These approximations transform an inherently stochastic problem into a deterministic one which enables the use of well-developed optimization techniques.

REFERENCES

- Baccelli, F. and Z. Liu (1992). On a class of stochastic recursive sequences arising in queueing theory. *Annals of Probability* **20**, 350–374.
- Bertsekas, D.P. (1999). *Nonlinear Programming*. 2nd ed.. Athena Scientific. Belmont, MA.
- Corbett, C.J. and C.S. Tang (1999). Designing supply contracts: contract type and information asymmetry. In: *Quantitative Models for Supply Chain Management* (S.R. Tayur, R. Ganeshan and M. Magazine, Eds.). pp. 269–297. Kluwer.
- DelVecchio, C. and I. Ch. Paschalidis (2003). Enforcing service-level constraints in supply chains with assembly operations. In: *Proceedings of the 42nd IEEE Conference on Decision and Control*. Maui, Hawaii.
- Paschalidis, I. Ch. and Y. Liu (2003). Large deviations-based asymptotics for inventory control in supply chains. *Operations Research* **51**(3), 437–460.