Proceedings of the 2009 Winter Simulation Conference
M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin and R. G. Ingalls, eds.

SAMPLE AVERAGE APPROXIMATION APPROACH TO MULTI-LOCATION TRANSSHIPMENT PROBLEM WITH CAPACITATED PRODUCTION

Deniz Özdemir

Graduate Program in
Systems Engineering
Universidad Autónoma de Nuevo León
San Nicolás de los Garza
N.L., 66450, MEXICO

Enver Yücesan

Technology and Operations
Management Area
INSEAD
Boulevard de Constance
77305 Fontainebleau Cedex, FRANCE

Yale T. Herer

Faculty of Industrial Engineering
And Management
Technion
Israel Institute of Technology
Haifa 32000, ISRAEL

ABSTRACT

We consider a supply chain, which consists of *N* stocking locations and one supplier. The locations may be coordinated through replenishment strategies and lateral transshipments, i.e., transfer of a product among locations at the same echelon level. The supplier has limited production capacity. Therefore, the total amount of product supplied to the *N* locations is limited in each time period. When total replenishment orders exceed total supply, not all locations will be able to attain their base stock values. Therefore, different allocation rules are considered to specify how the supplier rations its limited capacity among the locations. We team up the modeling flexibility of simulation with sample path optimization to address the multi-location transshipment problem. We solve the sample average approximation problem by random search and by gradient search. With this numerical approach, we can study problems with non-identical costs and correlated demand structures.

1 INTRODUCTION

Physical pooling of inventories has been widely used in practice to reduce cost and improve customer service. However, information pooling, which entails the sharing of inventory among stocking locations through lateral transshipments, has been less frequent. Transshipments, the monitored movement of material between locations at the same echelon, provide an effective mechanism for correcting discrepancies between the locations' observed demand and their available inventory. As a result, transshipments lead to cost reductions and improved service without necessarily increasing system-wide inventories. In our research, we focus on collaborative planning and replenishment policies via information pooling and, in particular, on transshipments as a way to improve both cost and service.

Our study is motivated by observations from various industries. For example, inventory-pooling strategies to hedge against the risk of supply disruption are quite common in retailing. Various retailers such as Foot Locker and distributors such as Ingram Micro pool inventory to increase the safety stock of their products. The recent turmoil in the financial markets triggered by the credit crunch provides another timely example of transshipments. In the banking system, there are two principal ways for banks to balance their cash reserves. Banks may borrow cash from the central bank (the supplier), which typically controls the money supply very tightly. Within the European Union, the European Central Bank (ECB) "replenishes" banks either through one-week lending on every Wednesday or lending for three months on the last Wednesday of each month. Alternatively, banks may borrow (transship) cash from one another on an overnight basis. To provide more liquidity for cash-starved markets, ECB has recently injected 40 billion euros into the system via three-month lending (increasing supplier capacity) to banks, which were still reluctant to (transship) lend to one another (Werdigier and Dougherty 2007).

Driven by significant advances in information and communication infrastructure, transshipments are becoming more popular. Mirroring practice, there is growing literature on transshipments. The literature, however, has generally addressed either problems with two retailers, e.g., Tagaras (1989), Tagaras and Cohen (1992), Robinson (1990), and Herer and Rashit (1999), or problems with multiple identical retailers, e.g., Krishnan and Rao (1965), Jönsson and Silver (1987), and Robinson (1990). Other recent work on transshipments includes Archibald, Sassen, and Thomas (1997), Bertrand and Bookbinder (1998), Tagaras (1999), Herer and Tzur (2001, 2003), Slikker, Fransoo, and Wouters (2004), and Bendoly (2004).

In contrast, we are concerned with multiple retailers, which may differ both in their cost structures and in their demand parameters. Herer, Tzur, and Yücesan (2006) have introduced a simulation-based optimization method to deal with such sys-

978-1-4244-5771-7/09/$26.00 ©2009 IEEE

tem complexity. On the one hand, simulation models capture arbitrary levels of operational detail; on the other hand, optimization via simulation also provides an efficient way to optimize system design. This set up has already been extended to cases with capacitated transportation, thereby limiting the flexibility of a transshipment system (Özdemir, Yücesan and Herer 2006). In this paper, we extend the original set up in another direction by considering a supplier with limited production capacity as Jönsson and Silver (1987) did with identical retailers.

We propose a simulation-based optimization approach for solving the multi-location transshipment problem with supplier capacity. To minimize the total system costs, the objective is to find the appropriate inventory policies, which are typically a base stock order-up-to policy. Given a modified order-up-to-S policy, we then determine a myopically optimal transshipment policy between any pair of stocking locations.

As in Herer, Tzur and Yücesan (2006), we use the sample average approximation (SAA) method (Kleywegt, Shapiro and Homem-de-Mello 2001). In this setting, the performance criterion is typically expressed as an expected value. A random sample is generated and the expected value is approximated by the corresponding sample average. The resulting sample average optimization problem is solved; the procedure is repeated until a stopping criterion is satisfied. The method exploits the fact that the function we wish to optimize is the limit, along almost every sample path, of a sequence of approximating functions; the basic idea is to generate a sufficient number of sample paths to have a good estimate of the performance criterion and optimize the resulting *deterministic* function. We then take the result as the estimate of an optimizer of the true function. Robinson (1996) refers to this approach as sample path optimization. Simulation-based optimization helps the search for an improved policy while allowing for complex features that are typically outside of the scope of analytical models. To solve the (deterministic) optimization problem, we use random search (as a benchmark) and gradient search with infinitesimal perturbation analysis (IPA), an efficient gradient estimation technique (Ho, Eyler, and Chien 1979). Unlike Herer, Tzur and Yücesan (2006), however, the presence of constraints on supply capacity renders these gradient estimators more complex.

The contribution of the current paper is thus two-fold: for the transshipment literature, we extend existing transshipment models in a non-trivial fashion to include finite supply capacity. For the simulation optimization literature, we construct an SAA algorithm for a non-trivial application. The remainder of the paper is organized as follows: In the following section, we introduce the capacitated transshipment problem and the notation used in the paper. Section 3 is devoted to determining the replenishment quantities incorporating various allocation rules under limited supplier capacity. The policy for replenishments and transshipments together with the formulation is explained in Section 4. Section 5 presents the technical details of SAA. We illustrate the solution technique with a numerical study and discuss the findings in Section 6. We conclude with final remarks in Section 7.

2 MODEL

We consider a supplier serving N retailers, or stocking locations, which face random customer demand. The demand distribution of each stocking location in a period is assumed to be known and stationary over time. The stocking locations review their inventory periodically and place replenishment orders with the supplier that has a finite total production capacity, C. In any period, transshipments provide a means to reconcile demand-supply mismatches.

Within each period, events occur in the following order: the first event in each period is the arrival of replenishment orders placed in the previous period. These orders are used to satisfy any outstanding backlog and to increase inventory. Next in the period is the occurrence of demand. Since the realization of demand represents the only uncertain event of the period, once it is observed, all the decisions of the period, namely, the determination of the transshipment and replenishment quantities, are taken. The transshipment transfers are then made immediately, and demand is subsequently satisfied. Different from Özdemir, Yücesan and Herer (2006), we assume that unsatisfied demand is backlogged. At this point, backlogs and inventories are observed, and resulting penalty and holding costs, are incurred. The remaining inventory, if any, is carried to the next period. Note that items in stock elsewhere in the system are supplied immediately through transshipments while backlogged items have to wait until the beginning of the next period. Thus, transshipments provide an additional source of supply whose reaction time is shorter than that of regular supply.

We consider *modified* base stock policies for replenishment. The policy is "modified" in the following sense. In a base stock policy, when the supplier does not have a capacity constraint, the inventory positions at all stocking locations are raised up to S_i units at the beginning of each period. Given the finite supplier capacity, however, the locations may not receive the full replenishment quantity ordered in the previous period. Therefore, order-up-to levels may not be attained at the beginning of each period. Furthermore, when replenishment through the supplier is capacitated, different allocation rules must be specified to reflect how the supplier rations its limited capacity among the locations.

2.1 Notation

In developing our model, we use the following parameters:

c_i = unit procurement cost at stocking location i ;

$\hat{t}_{ij}$ = direct transshipment cost per unit transshipped from stocking location i to stocking location *j;* this is the additional administrative and logistics costs (packaging, re-labeling, transferring, etc.) per unit due to transshipment.

t_{ij} = effective transshipment cost, or simply transshipment cost, per unit transshipped from stocking location i to stocking location *j*, $t_{ij} = \hat{t}_{ij} + c_i - c_j$;

h_i = holding cost incurred at stocking location i per unit held per period;

p_i = penalty cost incurred at stocking location i per unit backlogged per period.

We assume (as was assumed in Tagaras (1989), Robinson (1990), and Herer and Rashit (1999) as well as others) the following relationships regarding the problem parameters:

$$h_i < h_j + t_{ij} \qquad i, j = 1,\ldots,N$$

$$p_j < p_i + t_{ij} \qquad i, j = 1,\ldots,N$$

$$t_{ij} < h_i + p_j \qquad i, j = 1,\ldots,N$$

The first relationship reflects the fact that shipping items to a stocking location is not allowed, if there is already a surplus item there. Similarly, stocking locations don't transship goods to other locations if they already have a shortage, as reflected by the second inequality. Finally, we assume that if there is a shortage at one of the stocking locations and surplus at another, lateral transshipment is (myopically), cost advantageous. These inequalities ensure that transshipments from location *i* to location *j* are economically justifiable only if location *i* has excess inventory and location *j* has a shortage. These inequalities imply that the *complete pooling* transshipment policy (Tagaras 1989) is optimal with no constraints on supplier capacity (Herer, Tzur, and Yücesan 2006); hence, we will continue to use complete pooling here.

In addition, we have

D_i = random variable associated with demand at location i in each period with $E[D_i] = \mu_i < \infty$;

d_i^n = actual realization of demand at stocking location i in period *n*;

$F_i(d_i^n)$ = cumulative distribution function of demand at location i , i.e. the probability that $D_i < d_i^n$;

C = total production capacity per period ($C > E[\sum_i D_i]$);

I_{i0}^n = net inventory level at stocking location i at the end of period n;

I_i^n = net inventory level at stocking location i at the beginning of period n after replenishment. I_i^n is the net inventory level in period *n* after the arrival of replenishment orders from the previous period (and before demand is observed). When we consider quantities in an arbitrary period, time superscripts are dropped.

Two decisions need to be made for each stocking location every period: Transshipment quantities between any pair of stocking locations and replenishment quantities. The associated decision variables are the following:

S_i = target inventory level (or order-up-to level) at stocking location i at the beginning of each period;

X_{ij}^n = number of items transshipped from stocking location i to stocking location *j* in period *n*;

R_i^n = number of items received from the supplier by stocking location i in period *n*+1 that were ordered from the supplier in period *n*. Note that, when production is capacitated, the number of items received is not necessarily equal to the number of items ordered.

3 DETERMINING THE REPLENISHMENT QUANTITIES

In any period *n*, the net inventory level at stocking location i at the end of the period is the sum of the inventory level in period *n*, immediately after demand is observed, and the difference between the total quantity received (via transshipments from other locations) and sent (via transshipments to other locations) during period *n*. Furthermore, in period *n*+1, the net inventory level at stocking location i immediately before demand is observed, is equal to the sum of the inventory level at the end of period *n* and items received from the supplier in period *n*. In each period, the replenishment quantity R_i^n is the minimum of remaining production capacity and the difference between the order-up-to value (S_i) and the inventory level at the end of the period at location i. Therefore, the sample path of the system in any period n can be described as follows:

$$I_{i0}^n = I_i^n + \sum_{j:j\neq i} X_{ji}^n - \sum_{j:j\neq i} X_{ij}^n - d_i^n, \qquad i = 1,\ldots,N$$

$$I_i^{n+1} = I_{i0}^n + R_i^n, \qquad i = 1,\ldots,N \tag{1}$$

$$R_i^n = \min\left(\left(C^{prod} - \sum_{j\neq i} R_j^n\right), \left(S_i - I_{i0}^n\right)\right) \qquad i = 1,\ldots,N \tag{2}$$

and, under complete pooling, X_{ij}^n, the transshipment quantity from stocking location i to j is equal to

$$X_{ij}^n = \min\left\{\left(I_i^n - d_i^n - \sum_{k\neq j;k\neq i} X_{ik}^n\right)^+, \left(d_j^n - I_j^n - \sum_{k\neq j;k\neq i} X_{kj}^n\right)^+\right\} \quad i\neq j. \tag{3}$$

The total cost of the system in period n is given by:

$$TC_n = \sum_{i=1}^{N} \sum_{j=1}^{N} t_{ij} X_{ij}^n + \sum_{i=1}^{N} h_i I_{i0}^+(n) + \sum_{i=1}^{N} p_i I_{i0}^-(n) + \sum_{i=1}^{N} c_i d_i^n,$$

where $I_{i0}^+(n) = \max\{0, I_{i0}^n\}$ and $I_{i0}^-(n) = \max\{0, -I_{i0}^n\}$. Therefore, $I_{i0}^n = I_{i0}^+(n) - I_{i0}^-(n)$. The unit purchase cost at location i is multiplied by the demand at location i and not by the replenishment quantity at location i since the procurement cost differentials are included in the transshipment costs.

When total replenishment orders exceed total supply capacity, not all locations will be able to attain their base stock levels. We will refer to this difference between the order-up-to level (S_i) and the inventory level at location i at the beginning of period $n+1$ as the *shortfall* at location i at the end of period n. We will use the shortfall values later in the analysis. Moreover, for the allocation of the available supply among stocking locations, we propose to implement and test five allocation rules. These allocation rules are described in the next section.

3.1 Replenishment Allocation Rules

3.1.1 Shortfall Balancing Rule

In a multi-location setting, with identical cost structures, the location of the shortfall does not affect the service level in each location nor the total cost of the system in that particular period. However, to provide balanced service at each location in the subsequent period, it is desirable to distribute the shortfall at the end of the period evenly across the locations. We therefore distribute the available stock in such a way so as to balance the shortfall or, if this is not possible, to minimize the maximum shortfall. If all locations are identical, i.e., all locations have the same cost structure and their demand distributions are identical, then all locations will have the same base stock levels. Thus, balancing the beginning inventory at each location (I_i^{n+1}) is equivalent to balancing shortfall ($S_i - I_i^{n+1}$) at each location. On the other hand, in the case of non-identical cost and demand structures, since base stock levels will be different for each stocking location, reflecting not only the mean and standard deviation of the demand but the safety stock levels based on underage and overage costs, the allocation among stocking locations will also be different.

The allocation scheme that minimizes total expected backorders, proposed by Jönsson and Silver (1987), uses expected backordered units to calculate a unique safety stock factor for all locations. Although in the shortfall balancing rule original safety stock factors are not necessarily identical, for demand structures with identical standard deviations, both allocation rules suggest similar allocations.

3.1.2 Equal Allocation Rule

We also implement a simple allocation rule whereby an equal amount is allocated to each location, unless the replenishment quantity needed to reach the base stock level is less than the allocated amount. This unallocated quantity is then shared equally among those locations, which did not yet reach their base stock levels. This rule was tested mainly for benchmarking purposes and not as a proposal for implementation.

3.1.3 Pretransshipment Service Balancing Rule

The above rules while naively intuitive, do not take into account all the important available information. In particular they do not consider the distribution of the demand in general and in the probability of a stockout in particular. The pretransshipment service balancing rule minimizes the maximum probability of having a pretransshipment shortage at a location in the next period, i.e. it minimizes the maximum value of $1-F_i(I_i^n)$. In other words we allocate inventory to locations with the greatest chance of a pretransshipment stockout next period. Of course we constrain ourselves not to allocate a location more than its order-up-to quantity S_i.

3.1.4 Prioritized Pretransshipment Service Balancing Rule

Whereas the Pretransshipment Service Balancing Rule takes into account more information (the demand distribution) it does not consider the cost information. The Prioritized Pretransshipment Service Balancing Rule adds this information to the priority rule by weighting the pretransshipment stockout probability by the cost of a stockout at the location. That is we minimize the maximum $p_i(1-F_i(I_i^n))$. In other words we allocate inventory to the location with the greatest expected marginal pretransshipment stockout cost.

3.1.5 Pretransshipment Cost Balancing Rule

Our final rule is similar to the Prioritized Pretransshipment Service Balancing Rule except that it recognizes that when a location is allocated inventory, in addition to reducing the risk of a stockout, it increases the risk of having excess inventory. The Pretransshipment Cost Balancing Rule thus reduces the marginal expected benefit of decreasing the pretransshipment stockout cost by the marginal expected loss of the inventory being unused before transshipments, i.e. by $h_i F_i(I_i^n)$. In other words we allocate inventory such as to minimize the maximum $p_i(1-F_i(I_i^n))-h_i F_i(I_i^n)$. Note that this is the very marginal expected benefit that is examined by the single location periodic review stochastic inventory problem.

4 DETERMINING THE TRANSSHIPMENT QUANTITIES

In each period, the replenishment and transshipment quantities must be determined. Herer, Tzur, and Yücesan (2006), who focused on the uncapacitated version of our problem, proved that, if transshipments are only made to compensate for an actual shortage (instead of building up inventory at another stocking location), there exists an optimal order-up-to $S = (S_1, S_2, \ldots, S_N)$ policy for all possible stationary transshipment policies**.** Order-up-to S policies are widespread in practice because they are not only simple to implement but are also shown to be very effective. For the capacitated case, the determination of the optimal replenishment policy is an open problem. Nevertheless, since the transshipment policy is stationary and the fixed ordering cost is negligible, we will continue to adhere to an order-up-to S replenishment policy.

Once demand is observed, for a given base stock level, it is possible to solve the transshipment decision problem via a network flow formulation. After solving for the myopically optimal transshipment decision, instead of using the R_i values obtained through the LP (which represent a lexicographic priority rule imposed by the LP algorithm), we will make the replenishment decisions in accordance with one of the allocation rules introduced in Section 3.1.

We adapt the complete network flow problem and the associated LP formulation approach of Herer, Tzur, and Yücesan (2006) and Özdemir, Yücesan, and Herer (2006) to model the multi-location capacitated transshipment problem. We will use the following decision variables in our LP formulation:

X_{ij}^n : transshipment quantity from stocking location i to stocking location j in period n;

$I_{i0}^+(n)$: inventory held at stocking location i in period n;

$I_{i0}^-(n)$: shortage at stocking location i satisfied through replenishment in period n. Therefore $I_{i0}(n) = I_{i0}^+(n) - I_{i0}^-(n)$.

R_i^n : total replenishment for stocking location i in period n;

$I_i^+(n)$: on-hand inventory at stocking location i at the beginning of period n, after arrival of orders and flushing of backlogs;

$I_i^-(n)$: backordered inventory at stocking location i at the beginning of period n, after arrival of orders and flushing of backlogs.

We formulate the LP for the transshipment decision in period n as follows:

$$(P_n)\ \min \quad TC_n = \sum_{i=1}^{N}\sum_{j=1}^{N} t_{ij} X_{ij}^n + \sum_{i=1}^{N} h_i I_{i0}^+(n) + \sum_{i=1}^{N} p_i I_{i0}^-(n)$$

subject to

$$\sum_{j=1}^{N} X_{ij}^n + I_{i0}^+(n) = I_i^+(n), \qquad i = 1, \ldots, N \tag{4}$$

$$\sum_{j=1}^{N} X_{ji}^n + I_{i0}^-(n) - I_i^-(n) = d_i^n, \qquad i = 1, \ldots, N \tag{5}$$

$$I_{i0}^+(n) + R_i^n - I_{i0}^-(n) + I_i^-(n+1) = I_i^+(n+1), \qquad i = 1, \ldots, N \tag{6}$$

$$I_i^+(n+1) \le S_i \qquad i = 1, \ldots, N \tag{7}$$

$$\sum_{i=1}^{N} R_i^n \le C, \tag{8}$$

$$I_{i0}^+(n), I_{i0}^-(n), X_{ij}^n, R_i^n \ge 0, \qquad i,j = 1, \ldots, N$$

$$I_i^+(n+1), I_i^-(n+1) \ge 0 \qquad i = 1, \ldots, N \tag{9}$$

While on-hand inventory or backordered inventory levels at stocking location i at the beginning of period $n+1$ ($I_i^+(n+1)$ and $I_i^-(n+1)$, respectively) are decision variables, the values of on-hand or backordered inventory at the beginning of period n ($I_i^+(n)$ and $I_i^-(n)$, respectively) are calculated from (P_{n-1}) and are thus known at period n. Therefore, $I_i^+(n)$ and $I_i^-(n)$ are parameters in (P_n).

The constraint sets (4) and (6) ensure the balance of the inventory position of each stocking location at the beginning and at the end of each period, respectively. Constraint set (5) guarantees that the observed demand at location i (d_i^n) and the shortfall from previous period ($I_i^-(n)$) will be satisfied either from the location's own inventory (X_{ii}^n), transshipped from another location (X_{ji}^n) or backlogged and satisfied through the replenishment from the supplier ($I_{i0}^-(n)$). Moreover, due to the supplier capacity constraint, the inventory position may not attain the order-up-to levels, S_i, which is captured by the constraint set (7). Finally, constraint (8) guarantees that total replenishment to all stocking locations will be at most C units, reflecting supplier capacity. Non-negativity constraints (9) are also included. The objective is to minimize the cost of demand-supply mismatch (inventory holding and shortage penalty costs) and transshipment costs.

5 THE SOLUTION ALGORITHM

For the capacitated transshipment problem, determining the exact order-up-to levels is analytically difficult. To determine the optimal order-up-to-S values, we therefore use the *sample average approximation* method to minimize the average total cost per period. More specifically, the stochastic optimization problem we are addressing has the following form:

$$\min_{\theta\in\Theta} f(\theta) = E\left[f(\theta,\xi)\right]$$

for some random variable ξ (representing, in our case, stochastic demand faced by the retailers) and parameter $\theta\in\Theta$ is the set of possible values for the parameter θ (representing, in our case, the retailer order-up-to levels). For a fixed θ, we generate the independent random sample $\xi_1, \xi_2, \ldots, \xi_U$ (of demand realizations) to define the sample mean over $(f(\theta,\xi_i) : 1 \le i \le U)$ as $\overline{f}_U(\theta) = \frac{1}{U}\sum_{i=1}^{U} f(\theta,\xi_i)$.

The sample average approximation problem is then defined as one of minimizing the sample average, i.e.,

$$\min_{\theta\in\Theta} \overline{f}_U(\theta), \tag{10}$$

since, by the strong law of large numbers, $\overline{f}_U(\theta)$ converges to $f(\theta)$ w.p. 1 as $U \to \infty$. Large deviations results (e.g., Dai, Chen, and Birge 2000) further show that one may not need a large sample in order to solve the original problem exactly with a high probability by solving the SAA problem. The required sample size is, however, problem dependent, which may be difficult to estimate in an *a priori* fashion. We solve the minimization problem in (10) by random search (as a benchmark) and by gradient search.

One of the attractive features of random search is its simplicity of construction as no potentially costly computation of derivative or Hessian matrix is involved. Properties of simulation-based optimization via random search have been studied extensively. A comprehensive review along with new convergence results is given by Chia and Glynn (2007). We have deployed two versions of random search. The naïve version randomly generates a vector of order-up-to levels and estimates the associated average total cost. The vector that yields the lowest average cost is reported as the minimizer of (10) provided that it has been evaluated a minimum number of times (Andradóttir 1999). In the second version of random search, we first randomly generate a vector of order-up-to levels and estimate the associated average total cost; we then explore the neighborhood of that vector by randomly picking one of its components, say j, and randomly perturbing its order-up-to level as $S_j \pm 1$. The associated average total cost is then estimated with the resulting vector. This approach strives to maintain a balance between global search, or exploration, and local search, or exploitation (Andradóttir and Prudius 2009). As before, the vector that yields the lowest average cost is reported as the minimizer of (10) provided that it has been evaluated a minimum number of times. Figure 3 summarizes our algorithm.

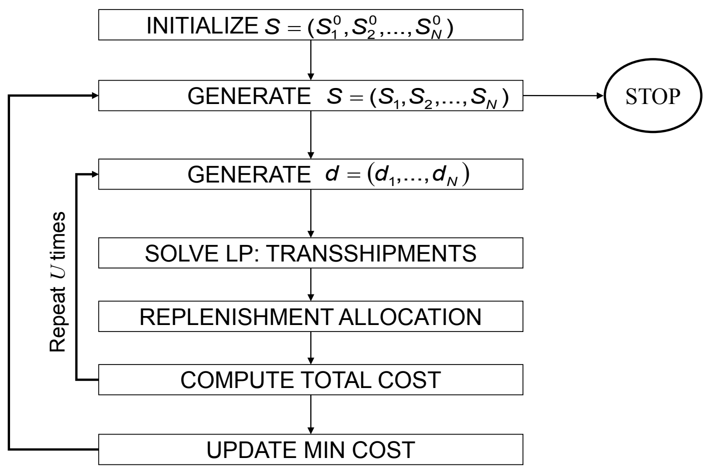

Figure 1: Random search

The fact that the dual of a constraint in an LP formulation is the derivative of the objective function with respect to the right-hand side of that constraint motivated us to solve (10) using gradient search guided by IPA. Under the assumption that $\Theta \subset \mathfrak{R}^N$, we have a convex, possibly non-smooth, stochastic performance function. With IPA, instead of using finite differences in a gradient search method, we use the mean value of the sample path derivative, which is obtained through a single simulation (Glasserman 1991). Applications of perturbation analysis have been reported in simulations of Markov chains (Glasserman 1992), inventory models (Fu 1994), manufacturing systems (Glasserman 1994), finance (Fu and Hu 1997), and control charts for statistical process control (Fu and Hu 1999). IPA-based methods have also been introduced to analyze supply chain problems (Glasserman and Tayur 1995). In a set-up similar to ours, Plambeck Dai, L., Chen, C.-H., and Birge. (1996) use the sample path optimization method with IPA gradient estimates for the optimization of stochastic PERT problems with respect to parameters of the activity length distributions.

IPA gradients will be unbiased provided that the objective function of (P_n) is convex and smooth with respect to the order-up-to levels. The objective functions of linear programs are convex (piecewise linear) functions of their right-hand sides (Theorem 5.1 in Bertsimas and Tsitsiklis 1997). Note that, for a given demand realization, all S_i variables appear on the right-hand side of (P_n), establishing convexity. On the other hand, the objective function is piecewise linear and differentiable almost everywhere except on a set of Lebesgue measure zero assuming that demand is a continuous random variable. As a result, the generalized gradient is well defined everywhere (Clarke 1990). Complete technical details are also given by Shapiro Dai, L., Chen, C.-H., and Birge (2002).

In addition, we need to calculate the IPA gradient over one complete regenerative cycle, which will most likely last for more than a single period. In our algorithm, the period in which all stocking locations simultaneously reach their order-up-to levels, S_i, is a regeneration point. Therefore, we need to propagate the gradients through the periods in the cycle. That is, we use the LP to compute the transshipment quantities, but not the replenishment quantities; we further use the LP output to accumulate the IPA gradients ($\partial TC / \partial S_i$), which are used in sample path optimization to determine the optimal order-up-to levels This is indeed the specialization of stochastic approximation to the capacitated transshipment problem. The stochastic optimization algorithm we used exploits this property; the technical details are presented in Özdemir (2004).

6 COMPUTATIONAL STUDY

We analyze the impact of different factors on transshipment relations with limited production capacity under 5 configurations of 10 stocking locations. In all configurations,

a) We consider stocking locations with identical cost parameters. In particular, we set the holding cost to h_i = \$ 1 and penalty cost to p_i = \$ 4 for all ten locations.

b)We consider stocking locations with nonidentical cost parameters. Particularly, split the stocking locations into two groups, each having 5 retailers. We define two values for the holding and shortage cost: we set high holding cost to h_i^H = \$ 2 and low holding cost to h_i^L = \$ 1. Similarly, we set high penalty cost to p_i^H = \$ 8 and low penalty cost to p_i^L = \$ 4. We generate 6 scenarios with following cost combinations:

Table 1: Nonidentical Cost Parameter Setting

Locations 1 to 5 (h_i, p_i)	Locations 6 to 10 (h_i, p_i)
(h_i^L, p_i^L)	(h_i^H, p_i^H)
(h_i^H, p_i^L)	(h_i^L, p_i^H)
(h_i^H, p_i^L)	(h_i^H, p_i^H)
(h_i^L, p_i^L)	(h_i^L, p_i^H)
(h_i^L, p_i^H)	(h_i^H, p_i^H)
(h_i^L, p_i^L)	(h_i^H, p_i^L)

Each location faces an independent demand with Gaussian distribution with mean 100 and standard deviation 20. As summarized in Table 2, we consider five alternative system configurations with different unit transshipment costs, t_{ij}, for units transshipped from stocking location i to stocking location j. Note that $t_{ij} = \infty$ implies that transshipments are not allowed between locations i and j.

Table 2: System Configurations

System	t_{1i}	t_{i1}	t_{ij}
1	∞	∞	∞
2	0.5	∞	∞
3	0.5	0.5	∞
4	0.5	0.5	1.0
5	0.5	0.5	0.5

As a base case, in system #1, no material movement is allowed among stocking locations, turning the system into 10 independent newsvendors. In system #2, only the first stocking location can transship to the other stocking locations. In system #3, transshipments from all stocking locations to the first stocking location are also allowed. In systems #4 and #5, all material movement is allowed. In system #4, however, transshipments between any two stocking locations (which do not include location #1) are twice as expensive. Note that, in systems #2, 3, and 4, we can view the first stocking location as a distributor for the entire network, while in systems #1 and 5, all locations are identical. For each system, we generate three scenarios with different supplier capacity. The capacity values used are: C= {1100, 1250, 1400}

In addition to sample average approximation algorithm with IPA gradients, we implement 2 variants of random local search algorithms: Random search with restart and Random search with local search. Moreover, for all cost and capacity scenarios five replenishment allocation rules discussed above are implemented.

In general, the selection of effective values for algorithm parameters is a difficult problem. After conducting an extensive search with different strategies, we set the total number of steps for the path search to K= 10000, the number of regenerative cycles at each step to U= 2000, and the step size to $\alpha_k = 1000/k$ for IPA gradient search. As a stopping criterion, we compare the computed order-up-to levels over 1000 iterations and require that these values do not differ by more than 0.1. The total average cost of 2000 periods is calculated and reported for each scenario.

For random search with restart, we restart the search K = 100 000 times, and to evaluate quality of the random basestock configuration, we calculated the total average cost of 1000 periods. For the best 5 basestock configuration, total average cost of 2000 periods is calculated and reported.

For random neighborhood search, we restart the search K = 1000 times, and generate a new "neighbor configuration" 500 times. To evaluate quality of the random basestock configuration, we calculated the total average cost of 50 periods. For the best 10 basestock configuration, total average cost of 2000 periods is calculated and reported.

In all systems, independent of the replenishment allocation rule implemented, we observe an increase in the total cost and in the total inventory level of the configuration when, we observe an increase in the total cost and in the total inventory level of the configuration when the supplier capacity constraint is incorporated. As the number of items that can be replenished per period increases, the total cost and the total inventory levels decrease. We observe that even a little extra capacity leads to great cost savings by allowing the order-up-to levels to decrease.

Comparing the simulation-based optimization algorithms, as it can be seen in Figure 2, we observe that gradient based search algorithm always converges to a better objective function, given approximately identical run times. The average total cost estimation calculated by naïve random search algorithm with restart improves if the number of simulation increases, e.g. for 1 million repetition estimation is very close to the IPA estimator, however this ends up with very long (more than 30 hours) runs compared to IPA estimations. The random search with neighborhood exploration on the other hand results in a reasonable cost estimate with small number of simulations however, we observe that after a good estimate the further improvement is very slow.

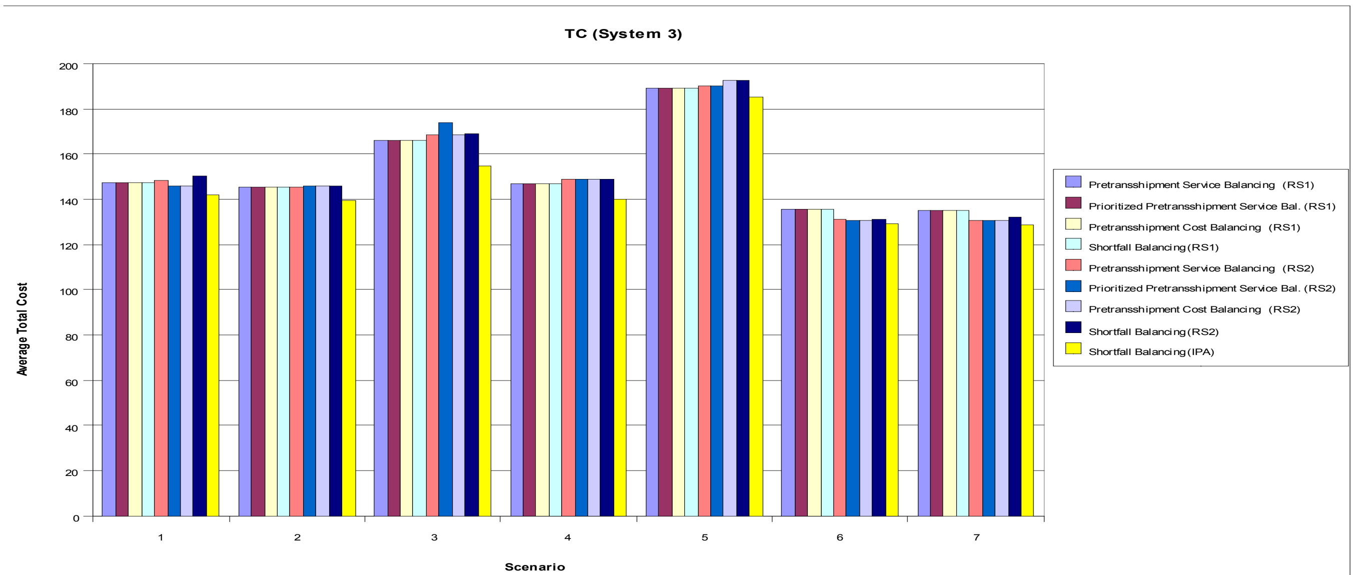

Figure 2: Total Average Cost for System 3

7 SUMMARY

We consider a supply chain, which consists of N stocking locations and one supplier. The locations may be coordinated through replenishment strategies and lateral transshipments. The supplier has limited production capacity. Therefore, the total amount of product supplied to the N locations is limited in each time period. When total replenishment orders exceed total supply, not all locations will be able to attain their base stock values. Therefore, different allocation rules are considered to specify how the supplier rations its limited capacity among the locations.

We team up the modeling flexibility of simulation with sample path optimization to address the multi-location transshipment problem. Under a modified base stock policy, we determine a myopically optimal transshipment policy using an LP/Network flow framework. Then using IPA estimators, we calculate the optimal values of base stock levels, employing simulation-based optimization. We analyze the outcomes on system behavior and performance measures of the stocking locations when the supplier may fail to satisfy all the replenishment order of locations. With this numerical approach, we can study problems with non-identical costs and correlated demand structures.

It is interesting to note that the sample path optimization framework deployed with IPA gradients established in Herer Dai, L., Chen, C.-H., and Birge. (2006) did not require any major technical modifications to analyze transshipment systems with transportation capacity (Özdemir Dai, L., Chen, C.-H., and Birge. 2006). Incorporation of limited supplier capacity, however, necessitated a significant overhaul of gradient calculations. In particular, instead of easily recovering the gradient information from the output of the LP in a single-period setting, we were obliged to propagate the gradient information throughout a regenerative cycle that might take several periods due to limited supplier capacity and various capacity allocation policies.

ACKNOWLEDGEMENTS

This research has been partially supported by a grant from *PriceWaterhouseCoopers* to INSEAD on high-performance organizations and a grant no PROMEP:103.5/08/4803.

REFERENCES

Andradóttir, S. 1999. Accelerating the Convergence of Random Search Methods for Discrete Stochastic Optimization. *ACM Transactions on Modeling and Computer Simulation*. 9: 349-380.

Andradóttir, S. and A.A. Prudius. 2009. Balanced Explorative and Exploitative Search with Estimation for Simulation Optimization. *INFORMS Journal on Computing*. Forthcoming.

Archibald, T.W., A.A.E. Sassen, and L.C. Thomas. 1997. An Optimal Policy for a Two-Depot Inventory problem with Stock Transfer. *Management Science* 43: 173-183.

Banerjee, A., J. Burton and S. Banerjee. 2003. A Simulation Study of Lateral Shipments in Single Supplier, Multiple Buyer Supply Chain Networks. *Int. Journal of Production Economics*. 81-82: 103-114.

Bendoly, E. 2004. Integrated Inventory Pooling for Firms Servicing Both On-line and Store Demand. *Computers & Operations Research*, 31: 1465-1480.

Bertrand, L.P. and J.H. Bookbinder. 1998. Stock Redistribution in Two-Echelon Logistics Systems. *Operational Research Society*, 49: 966-975.

Bertsimas, D. and J.N. Tsitsiklis. 1997. *Introduction to Linear Optimization*, Athena Scientific, Belmont.

Chia, Y.L. and Glynn, P.W. 2007. Optimal Convergence Rate for Random Search. In: *Proceedings of the 2007 INFORMS Simulation Society Workshop* (Chen, Henderson, Chick and Yücesan, eds.)

Clarke, F. 1990. *Optimization and Nonsmooth Analysis*. SIAM. Philadelphia.

Dai, L., Chen, C.-H., and Birge, J.R., 2000. Convergence properties of two-stage stochastic programming. *Journal of Optimization Theory and Applications* 106, 489-509.

Diks, E.B. and A.G. de Kok. 1996. Controlling a Divergent 2-Echelon Network with Transshipments Using the Consistent Appropriate Share Rationing Policy. *Int. Journal of Production Economics*, 45: 369-379.

Dong, L., and N. Rudi. 2004. Who Benefits from Transshipment? Exogenous vs. Endogenous Wholesale Prices. *Management Science* 50: 645-657.

Fu, M.C. 1994. Sample Path Derivatives for (s, S) Inventory Systems. *Operations Research* 42: 351-364.

Fu, M.C., and Hu, J.Q. 1997. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Kluwer Academic Publishers.

Fu, M.C., and Hu, J.Q. 1999. Efficient Design and Sensitivity Analysis of Control Charts using Monte Carlo Simulation. *Management Science* 45: 395-413.

Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers.

Glasserman, P. 1992. Derivative Estimates from Simulation of Continuous-Time Markov Chains. *Operations Research* 40: 292-308.

Glasserman, P. 1993. Regenerative Derivatives of Regenerative Sequences. *Advances in Applied Probability* 25:116-139.

Glasserman, P. 1994. Perturbation Analysis of Production Networks. In *Stochastic Modeling and Analysis of Manufacturing Systems* (Yao, Ed.). Springer-Verlag.

Glasserman, P. and S. Tayur. 1995. Sensitivity Analysis for Base Stock Levels in Multi-Echelon Production-Inventory Systems. *Management Science* 41: 263-281.

Glasserman P. and S. Tayur. 1994. The Stability of A Capacitated Multi-Echelon Production-Inventory System under A Base-Stock Policy. *Operations Research* 42: 913-925.

Herer, Y.T. and A. Rashit. 1999. Lateral Stock Transshipments in a Two-Location Inventory System with Fixed and Joint Replenishment Costs. *Naval Research Logistics* 46: 525-547.

Herer, Y.T., Tzur, M., and Yücesan, E., 2006. The multi-location transshipment problem, *IIE Transactions* 38: 185-200.

Herer, Y.T. and M. Tzur. 2001. The dynamic transshipment problem, *Naval Research Logistics* 48: 386-408.

Herer, Y.T. and M. Tzur. 2003. Optimal and heuristic algorithms for the multi-location dynamic transshipment problem with fixed transshipment costs, *IIE Transactions* 35: 419-432.

Ho, Y.C., M.A. Eyler, and T.T. Chien. 1979. A Gradient Technique for General Buffer Storage Design in a Serial Production Line. *International Journal of Production Research* 17: 557-580.

Jönsson, H. and E.A. Silver. 1987. Analysis of a Two-Echelon Inventory Control System With Complete Redistribution. *Management Science* 33: 215-227.

Kim, S. 2006. Gradient-Based Simulation Optimization. In *Proceedings of the 2006 Winter Simulation Conference*, eds. L. R. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 159-167. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc

Kleywegt, A.J., A. Shapiro, and T. Homem-de-Mello. 2001. The Sample Average Approximation Method for Stochastic Discrete Optimization. *SIAM Journal of Optimization* 12:479-502.

Krishnan, K.S. and V. R. K. Rao. 1965. Inventory Control in *N* Warehouses. *Journal of Industrial Engineering* 16**:** 212-215.

Loynes, R. M. 1962. The Stability of a Queue with Non-Independent Interarrival and Service Times. *Proceedings of Cambridge Philos. Society*. 58: 497-520.

Özdemir, D. 2004. *Collaborative Planning and Replenishment Policies*. PhD Dissertation. Technology (and Operations) Management Area, INSEAD.

Özdemir, D., E. Yücesan, and Y. T. Herer. 2006. Multi-Location Transshipment Problem with Capacitated Transportation. *European Journal of Operational Research* 175: 602-621.

Plambeck, E.L., B.R. Fu, S.M. Robinson, and R. Suri. 1996. Sample-Path Optimization of Convex Stochastic Performance Functions. *Mathematical Programming* 75:137-176.

Robinson, L.W. 1990. Optimal and Approximate Policies in Multiperiod, Multilocation Inventory Models with Transshipments. *Operations Research* 38: 278-295.

Robinson, S.M. 1996. Analysis of Sample Path Optimization. *Mathematics of Operations Research* 21: 513-528.

Rudi, N., S. Kapur, and D. Pyke. 2001. A Two-location Inventory Model with Transshipment and Local Decision Making. *Management Science* 47: 1668-1680.

Shapiro, A., Homem-de-Mello, T., and Kim, J. 2002. Conditioning on convex piecewise linear stochastic programs. *Mathematical Programming A* 94, 1-19.

Slikker M., J. Fransoo, and M. Wouters. 2004. Cooperation Between Multiple News-vendors with Transshipments*, European Journal of Operational Research*, In Press.

Swaminathan, J. and S. Tayur. 1999. Stochastic Programming Models for Managing Product Variety. In: *Quantitative Models for Supply Chain Management* (Tayur, Ganeshan, and Magazine, Eds.) Kluwer, 585-622.

Tagaras, G. 1989. Effects of Pooling on the Optimization and Service Levels of Two-Location Inventory Systems. *IIE Transactions* 21: 250-257.

Tagaras, G. 1999. Pooling in Multi-Location Periodic Inventory Distribution Systems. *Omega* 39-59.

Tagaras, G. and M. Cohen. 1992. Pooling in Two-Location Inventory Systems with Non-Negligible Replenishment Lead Times. *Management Science* 38: 1067-1083.

Werdigier, J. and C. Dougherty. ECB Acts to Dispel Clouds over European Credit Markets. *International Herald Tribune*, 23 August 2007.

Wolff, R.W. 1989 Stochastic Modelling and the Theory of Queueing. Prentice-Hall Int. Englewoodcliffs, NJ.

Wong W., Cattrysse D. and D. Van Oudheusden. 2005. Inventory Pooling of Repairable Spare Parts With Non-zero Lateral Transshipment Time and Delayed Lateral Transshipments, *European Journal of Operational Research*, 165: 207-218.

AUTHOR BIOGRAPHIES

DENIZ ÖZDEMIR is an Associate Research Professor in the Graduate Program in Systems Engineering at Universidad Autonoma de Nuevo Leon, Monterrey, Mexico. Her research interests are in the areas of supply chain management, reverse logistics, simulation methodology and applications. She has a Ph.D. in Production and Operations Management from INSEAD, Fontainebleau, France. Her e-mail is <deniz@yalma.fime.uanl.mx>

ENVER YÜCESAN is a Professor of Operations Management at the European Institute of Business Administration (INSEAD) at Fontainebleau, France. He is an Industrial Engineer with a PhD in Operations Research from Cornell University. He served as the Proceedings Co- Editor for WSC'02. His e-mail address is <enver.yucesan@insead.edu> and his web page is <www.insead.edu/faculty/yucesan>.

YALE T. HERER is an Associate Professor on the Faculty of Industrial Engineering and Management at Israel Institute of Technology, Technion. He holds a PhD in Operations Research from Cornell University. His e-mail address is <yale@technion.ac.il>.