IIE Transactions (2006) **38**, 185–200

ISSN: 0740-817X print / 1545-8830 online
DOI: 10.1080/07408170500434539

The multilocation transshipment problem

YALE T. HERER[1]*, MICHAL TZUR[2] and ENVER YÜCESAN[3]

[1]*Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel*
E-mail: yale@technion.ac.il
[2]*Department of Industrial Engineering, Tel Aviv University, Tel Aviv 69978, Israel*
E-mail: tzur@eng.tau.ac.il
[3]*Technology and Operations Management Area, INSEAD, 77305 Fontainebleau, France*
E-mail: yucesan@insead.edu

Received April 1999 and accepted July 2005

We consider a supply chain, which consists of several retailers and one supplier. The retailers, who possibly differ in their cost and demand parameters, may be coordinated through replenishment strategies and transshipments, that is, movement of a product among the locations at the same echelon level. We prove that in order to minimize the expected long-run average cost for this system, an optimal replenishment policy is for each retailer to follow an order-up-to S policy. Furthermore, we demonstrate how the values of the order-up-to quantities can be calculated using a sample-path-based optimization procedure. Given an order-up-to S policy, we show how to determine an optimal transshipment policy, using a linear programming/network flow framework. Such a combined numerical approach allows us to study complex and large systems.

1. Introduction

Physical pooling of inventories (Eppen, 1979) has been widely used in practice to reduce costs and improve customer service. For example, Xerox has consolidated all of its country-based warehouses in Europe into a single European Logistics Center in the Netherlands. Alternately, the practice of *transshipment*, the monitored movement of material between locations at the same echelon (e.g., among retailers), may entail the sharing of stock through enhanced visibility, but without the need to put the stock physically in the same location. To emphasize the requirement for supply chain transparency at the same echelon, we will refer to this practice as *information pooling*. Such information pooling through transshipments has been less frequent. Transshipments provide an effective mechanism for correcting discrepancies between the locations' observed demand and their available inventory. As a result, transshipments may lead to cost reductions and improved service without increasing system-wide inventories. In this paper, we study transshipments as an effective materials management policy.

Consider the following examples. Suppose that you go shopping at Foot Locker in Hamburg, Germany. You find a pair of Avanti Leather shoes, but, to your disappointment, they do not have your size. Knowing that it would take at least a few weeks to get the shoes that you desire from the Italian manufacturer, you prepare to leave the store disappointed. However, a sales representative quickly determines, through a simple check on the store's computer, that the Foot Locker in Antwerp, Belgium, has the shoes in your size. As she arranges to have the shoes sent overnight, she suggests that you come back the next day to try them on.

FNAC is the leading retailer of cultural and leisure products in France. The company has recently opened an on-line channel, fnac.com, in addition to its vast network of retail shops. Upon the receipt of an order from the Internet, there are several options for order fulfillment: fnac.com's own stock, stock kept at a central distribution center, and stock from nearby FNAC stores. The last option represents one-way transshipments since the physical inventory held at a store is used to satisfy the demand at fnac.com instead of ordering the item from the central distribution center or from its supplier. Although fulfilling customer demand through transshipments has a higher short-term operational cost, the supply chain manager of the company asserts that exercising the transshipment option expands the portfolio of items they can offer through the Internet threefold without having to carry the associated stock (Yücesan, 2003).

In the above examples, transshipments are sometimes used in a *reactive* mode (in response to an actual stock-out). Alternatively, companies may realize that increased benefits can be achieved by *proactively* incorporating the transshipment option into the planning phase. Planned and

*This research was performed, in part, when the author was at the Department of Industrial Engineering, Tel Aviv University.

systematic transshipments represent a relatively novel idea. They replace physical consolidation with virtual integration through information sharing. In this paper, we propose a model that allows the exploitation of the advantages of deploying transshipments in a proactive fashion.

There are two key reasons why information pooling has not yet been widely adopted in practice: the inadequacy of the IT infrastructure and the lack of realistic models to exploit the benefits of this policy. Whereas the past decade has seen significant investment in IT infrastructure (e.g., implementation of Enterprise Resource Planning (ERP) systems and other web-based technologies) enabling transparency within supply chains, new business models of transshipments have not been developed as rapidly. The literature on transshipments has generally addressed either problems with two retailers, e.g., Tagaras (1989), Robinson (1990), Tagaras and Cohen (1992), and Herer and Tzur (2001) or problems with multiple, mainly identical retailers, e.g., Krishnan and Rao (1965) and Robinson (1990). Herer and Tzur (2003) considered nonidentical multiretailers in a deterministic setting. In contrast, we consider multiple retailers, who are allowed to differ from one another both in their cost structure and in their demand parameters, in a stationary infinite-horizon setting. In addition we allow demand to be dependent across retailers within any particular period. Other recent work on transshipments includes Archibald *et al.* (1997), Herer and Rashit (1999), Tagaras (1999), Rudi *et al.* (2001) and Dong and Rudi (2004).

The paper most closely related to ours is that of Robinson (1990). For the model considered, it provides an analytical solution when there are two nonidentical locations and when there are multiple identical (in cost) locations. Additionally, it contains a heuristic for the multilocation nonidentical case, which contains a stochastic integration based on Monte Carlo sampling. Robinson also contains a mathematical program that is close to the mathematical program presented here in Section 3.2. Despite these similarities there are important differences both in our models and approaches. The cost parameters that we allow are more general. Also, Robinson considers minimizing expected discounted cost and we consider expected long-run average cost, where cost includes holding, shortage, transshipment, and replenishment costs. Most importantly, our approach is guaranteed to converge to the optimal values, whereas Robinson's heuristic, although it performs very well, provides no such guarantees.

Some of the recent papers have incorporated high levels of complexity into transshipment models. Unfortunately, such models become intractable rather quickly, leaving simulation as the only tool to investigate interesting policies. Crude simulation, however, can be very time consuming. We therefore propose to combine the modeling flexibility of simulation with stochastic optimization approaches. Simulation-based optimization techniques help the search for an improved policy while allowing for complex features that are typically outside of the scope of analytical models.

In particular, we show that an optimal policy for the system we consider is for each retailer to follow an order-up-to policy. The optimality of the order-up-to policy takes into consideration the use of transshipments among retailers, to be performed once demand is observed. While we also show how to find optimal transshipment quantities, an order-up-to policy remains optimal under any (even nonoptimal) stationary transshipment policy. This result is useful when considering what-if scenarios, for example, when transshipments are performed only within clusters of locations. We also demonstrate how the values of the order-up-to quantities can be calculated using a procedure that is based on Infinitesimal Perturbation Analysis, IPA (Ho *et al.*, 1979). Whereas an optimal order-up-to quantity has to be found only once for the entire system, an optimal transshipment strategy has to be determined on a period-by-period basis, given the period's demand realization. We also show how these transshipment quantities can be found using a linear programming (LP)/network flow framework.

The contribution of this paper is twofold. First is the development of an integrated IPA/LP algorithm for a system that allows transshipments. The system we consider differs from many previously studied systems with transshipments in that we consider multiple retailers, who differ both in their cost structure and in their demand parameters. Moreover, we show that we can find an optimal inventory replenishment policy for any stationary transshipment policy that may arise from practical considerations. This enables the comparison of several such alternatives, as well as a comparison of each alternative with the optimal solution. Second is a methodological contribution obtained by formulating and validating IPA derivative estimators for the transshipment problem. The estimators are based solely on data from the operation of a system at a single set of parameter values. Therefore, they are easily computed from the sample path generated by a simulation run. Formulating these estimates means introducing appropriate algorithms; validating them calls for showing that they converge to the correct values, where convergence is over the number of independent simulation replications (obtained, in our case, over regenerative cycles) used to estimate the derivative information.

The paper is organized as follows. In Section 2 we describe the multilocation transshipment problem and introduce notation. In Section 3 we present the form of a combined optimal policy for the replenishment and transshipment strategies, together with our solution technique. In Section 4 we discuss the numerical study, which illustrates the solution technique. Section 5 concludes the paper.

2. Problem description

2.1. *The model*

In the system being investigated, there is one supplier and also N nonidentical retailers, associated with N distinct stocking locations, facing customer demand. The demand distribution at each retailer in a period is assumed to be known and stationary over time. The system inventory is reviewed periodically, and replenishment orders are placed with the supplier. In any period, transshipments provide a means to reconcile demand-supply mismatches.

Within each period, events occur in the following order: first, replenishment orders placed with the supplier in the previous period arrive. These orders are used to satisfy any outstanding backlog and to increase the inventory level. Next in the period is the occurrence of demand. Since demand realization represents the only uncertain event of the period, once it is observed all the decisions of the period, i.e., transshipment and replenishment quantities, are made and paid for. The transshipments are then made immediately, and subsequently the demand is satisfied. Unsatisfied demand is backlogged. At this point, backlogs and inventories are observed, and penalty and holding costs, respectively, are incurred. The inventory is carried, as usual, to the next period.

The goal is to find the transshipment and replenishment quantities that minimize the expected long-run average cost over an infinite horizon. The cost is the sum of the replenishment, transshipment, holding, and penalty costs. Note that items, which are supplied through transshipments, satisfy demand immediately whereas backlogged items have to wait until the beginning of the next period. Thus, the advantage of using transshipments is in gaining a source of supply whose reaction time is shorter than that of the regular supply.

To describe the operation of the system, we use the following notation.

N = number of retailers;
D_i = random variable associated with the periodic demand at retailer i with $E[D_i] < \infty$;
$f(\mathbf{D})$ = joint probability density function for the demand vector $\mathbf{D}$;
d_i = actual demand at retailer i in an arbitrary period;
h_i = holding cost incurred at retailer i per unit held per period;
p_i = penalty cost incurred at retailer i per unit backlogged per period;
c_i = replenishment cost per unit at retailer i;
$\hat{c}_{ij}$ = direct transshipment cost per unit transshipped from retailer i to retailer j;
c_{ij} = effective transshipment cost, or simply the transshipment cost, per unit transshipped from retailer i to retailer j, $c_{ij} = \hat{c}_{ij} + c_i - c_j$.

We will represent the vector of quantities described above, as well as the ones that we will introduce later in the paper, by dropping the subscripts, thus, $\mathbf{d} = (d_1, \ldots, d_N)$.

Note that c_{ij} is considered as the effective transshipment cost because when a unit is transshipped from retailer i to retailer j we pay, in addition to the direct transshipment cost, a cost of c_i instead of c_j to replenish the unit. We restrict our attention to situations where $c_{ij} \geq 0$. If this condition were violated it would mean that the replenishment costs at the two locations would differ by more than the direct transshipment cost; in fact, we would expect that in most situations $c_i = c_j$ is satisfied, that is, $c_{ij} = \hat{c}_{ij}$. In this case the differences, if any, between various h_i values result solely from the retailers' physical and geographical characteristics. For example, the size of the warehouse and its material-handling efficiency, or whether the retailer is in an expensive business area or in a rural suburb, may affect the cost structure. We consider base stock policies, where S_i represents the order-up-to level at retailer i and $\mathbf{S} = (S_1, S_2, \ldots, S_N)$. Given d_i, the actual demand at retailer i in a given period, the dynamic behavior of the system is captured through the following auxiliary variable:

I_i = inventory level at retailer i immediately after transshipments and demand satisfaction
$= S_i - \sum_{j=1}^{N} F_{B_iM_j} + \sum_{j=1}^{N} F_{B_jM_i} - d_i$, where $F_{B_iM_j}$ represents the transshipment quantity from retailer i to retailer j.

The motivation for this notation will become apparent below and a concise definition will be given in Table 1 later in this paper. Note that I_i may be either positive or negative, and we denote:

$$I_i^+ = \max\{I_i, 0\}, \quad I_i^- = \max\{-I_i, 0\} .$$

Thus, the realized cost of the system *in a given period* is equal to:

$$TC = \sum_{i=1}^{N}\sum_{j=1}^{N} c_{ij}F_{B_iM_j} + \sum_{i=1}^{N} h_i I_i^+ + \sum_{i=1}^{N} p_i I_i^- + \sum_{i=1}^{N} c_i d_i . \tag{1}$$

We show, in Section 3.1, that base stock policies minimize the expected long-run average cost. Since the optimal policy is to order up to S_i units at each retailer i, the beginning of each period, after orders arrive and backorders are satisfied, is a *regeneration point*. That is, the system returns to the same state (S_i units at each retailer). Thus, we can view the multiperiod problem as a series of single-period problems. In particular, minimizing the expected cost in an arbitrary period will also minimize our objective function, the expected long-run average cost. Furthermore, this regenerative structure enables the construction of an efficient algorithm to compute the optimal order-up-to values. The algorithm is introduced in Section 3.3.

In Equation (1), the term $\sum_{i=1}^{N} c_i d_i$ is needed to fully account for the replenishment costs. Since we are using an "order-up-to S" replenishment policy at each retailer, the total amount replenished system-wide will be exactly equal to the system-wide demand. Since this term is independent of our decision variables, it is omitted below. Recall that the replenishment cost differentials were included in the definition of c_{ij}.

2.2. *Modeling assumptions*

We will make mild assumptions, one regarding the replenishment policy and two regarding the transshipment policy, but first we need three definitions.

Definition 1. A replenishment policy is shortage *inducing* if and only if the beginning inventory, after orders arrive and backorders are satisfied, at some retailer can be strictly negative. Moreover, a replenishment policy, which is not shortage inducing, is termed *nonshortage inducing*.

Definition 2. A transshipment policy is *stationary* if and only if the transshipment quantities decision is independent of the period in which it is made. That is, it depends only on the pre-transshipment inventory and the observed demand. Similarly, a replenishment policy is stationary if and only if the replenishment decision is independent of the period in which it is made.

Definition 3. A transshipment policy is a *no-buildup* transshipment policy if and only if transshipments are never made to buildup inventory at the receiving location, that is, transshipments are only made to satisfy actual current demand.

We consider only replenishment policies that are nonshortage inducing and transshipment policies that are both stationary and no-buildup. The nonshortage inducing assumption is needed to eliminate some pathological situations where the order-up-to quantity is negative; moreover this assumption is easily justified from a service-level standpoint. A customer may accept a shortage from time to time, but not ordering enough to satisfy an existing shortage (as a shortage inducing policy may do) would not be a sustainable business decision. The stationary assumption is made without loss of generality since our planning horizon is infinite and both demand and the cost parameters are stationary, implying that we need only consider replenishment and transshipment policies that are stationary. The no-buildup property is guaranteed (see Corollary 1 to Theorem 1 below) if we assume (as was assumed in Tagaras (1989), Robinson (1990), and Herer and Rashit (1999) as well as others) the following relationship regarding the problem parameters:

$$h_i \leq c_{ij} + h_j \quad \text{for all } i \text{ and } j \; . \tag{2}$$

Intuitively, this inequality means that it is not economic to transfer a unit from retailer i to retailer j, so that it would be held in inventory at retailer j rather than at retailer i. Several other assumptions that are often made in the literature on transshipments and/or appear to be natural are not required here; see Section 3.4.

3. Optimal policies

Two decisions need to be made each and every period: replenishment and transshipment quantities. Those are discussed, respectively, in Section 3.1, where an order-up-to policy is proven to be optimal for the replenishment decision, and in Section 3.2, where an LP/network flow formulation is developed for the transshipment decision. In Section 3.3 we discuss how the optimal values of the order-up-to policy may be found. Finally, in Section 3.4 we discuss some relaxations of restrictions on the parameters.

3.1. *Optimality of an order-up-to policy*

The optimal form of the replenishment policy is based on the following definition.

Definition 4. A replenishment policy is an order-up-to $\mathbf{S} = (S_1, S_2, \ldots, S_N)$ replenishment policy if at retailer i the beginning inventory, after orders arrive and backorders are satisfied, is S_i in every period.

Note that due to the no-buildup assumption of the transshipment policy, an order-up-to $\mathbf{S}$ replenishment policy is regenerative whenever the replenishment policy is nonshortage inducing. On the other hand, if for some i, $S_i < 0$, then, at the end of the period, another retailer may make a transshipment to retailer i causing the prereplenishment inventory level at retailer i to be strictly greater than S_i. Since reducing inventory levels during the replenishment stage in our model is not allowed, and in fact, reducing inventory levels below zero has no obvious physical meaning, we cannot guarantee that a shortage inducing order-up-to $\mathbf{S}$ replenishment policy is regenerative.

Theorem 1. *There exists an order-up-to* $\mathbf{S} = (S_1, S_2, \ldots, S_N)$ *replenishment policy which is optimal within the class of nonshortage inducing replenishment policies for any stationary no-buildup transshipment policy.*

Proof. We begin the proof by defining and then analyzing a system, which is virtually identical to the system described above. In fact, the new system differs in only two aspects:

1. At the end of the period, after holding and shortage costs are incurred, a retailer can either purchase or sell stock back to the supplier for the same price the stock can purchased at the beginning of the period.
2. The stock level at each retailer at the end of the period is constrained to be zero, i.e., no inventory and no backorders are allowed.

In all other aspects the two systems are identical in every way.

Claim 1. Every replenishment policy in the original system has a corresponding replenishment policy in the new system with identical cost.

If, in the original system, the end-of-period inventory level at retailer i is I_i and the replenishment quantity is r_i (thus incurring a replenishment cost of $c_i r_i$), then in the new system retailer i would, at the end of the previous period, sell back to the supplier I_i units (or, if $I_i < 0$, purchase $-I_i$ units) and order $r_i + I_i$ units during the replenishment stage of the current period, thus incurring a cost of $c_i(-I_i + r_i + I_i) = c_i r_i$. The other aspects of the two systems are identical in every way.

Note that the converse of the claim is not true. In particular, using the supplier to reduce inventories in the new system is possible, whereas using the supplier to reduce inventories in the original system is impossible.

Now let us examine the replenishment policy in the new system. In this newly defined system, since demands are stationary and independent across time periods and because the transshipment policy is stationary, the end of each period is a regeneration point. This means that, even though the planning horizon is infinite, the optimal replenishment decision in each and every time period is the same. In particular, we let S_i be the optimal order quantity at retailer i in the new system and we also note that this replenishment policy is an order-up-to $\mathbf{S} = (S_1, S_2, \ldots, S_N)$ replenishment policy.

Recall that any order-up-to $\mathbf{S}$ replenishment policy is also feasible in the original system. Moreover, since the new system has strictly more feasible solutions this replenishment policy is also optimal in the original system, which completes the proof of the theorem. ∎

Corollary 1. *If Equation* (2) *holds for all retailers i and j, then the optimal transshipment policy has the no-buildup property.*

Building up inventory in the new system when Equation (2) holds is clearly suboptimal. Since every no-buildup transshipment policy is feasible in the original system, we know that the optimal transshipment policy has the no-buildup property.

Note that the transshipment policy need not be optimal (or even reasonable) for Theorem 1 to hold. In the next section, we show how to compute the optimal transshipment policy. However, if for some reason another transshipment policy is desired, e.g., grouping retailers into (possibly overlapping) pooling groups such that retailers only transship to other retailers in the same group, then Theorem 1 still holds.

3.2. *Determining the optimal transshipment quantities*

Given an order-up-to $\mathbf{S}$ policy for the replenishment quantities, the optimal transshipment quantities need to be determined each period between every pair of retailers. To this end, we develop a linear-cost network flow model of an arbitrary single period. The network flow model we develop is not the only one possible; indeed there exists a network flow representation with N fewer nodes and N fewer arcs than the one we present[1]. We choose to present this particular representation because it clearly reflects the events and actions in a period, implicitly showing the flow of time.

Let us recall the events in this arbitrary period; in particular, let us examine the movement of material. At the beginning of the period, after orders arrive and backorders are satisfied, there are S_i units in stock at each retailer i. These units can be used in one of three different ways: (i) satisfy demand at retailer i; (ii) satisfy demand at retailer j (i.e., a transshipment from retailer i to j), and (iii) hold in inventory at retailer i. Whereas it is true that it is physically possible to move stock from retailer i to another retailer, e.g., j, for storage, this is precluded by the no-buildup assumption.

At the end of the period units are on order from the supplier. These units will be used in two different ways: (i) to satisfy a backorder at a retailer; or (ii) to buildup inventory at a retailer so that the retailer will start the next period, after the order arrives and backorders are satisfied, with S_i units in stock. The stock at the beginning of the period, after the order from the previous period arrives and backorders are satisfied, and the replenishment made during the current period are the only two sources of material.

Let us now examine the material flow from the demand side (i.e., the sinks). The demand at retailer i, d_i, can be satisfied in one of three different ways: (i) from the inventory at retailer i; (ii) from the inventory at another retailer j (i.e., through a transshipment from retailer j to retailer i); or (iii) from replenishment during the current period (that arrives at the start of the next period). Another sink for material is the requirement that each retailer i begins the next period, after orders arrive and backorders are satisfied, with S_i units. These units can come from one of two sources: (i) the inventory at retailer i; or (ii) replenishment during the period. As discussed above, inventory from another retailer will not be used to buildup inventory levels at retailer i.

Using the observations above, we model the movement of stock during a period as a network flow problem. In particular, we have a source node, B_i, to represent the beginning, i.e., initial inventory at retailer i, after orders arrive and backorders are satisfied, and a source node, R, to represent the replenishment that occurs in the period that arrives at the start of the next period. The sink node associated with the demand at retailer i will be denoted M_i. Similarly, we will denote by E_i the ending inventory at retailer i, including units on order from the supplier. Note that this is equal to the inventory at the beginning of the next period, after orders arrive and backorders are satisfied. The arcs in the network flow problem are exactly those activities described above and are summarized (with their associated cost per unit flow) in Table 1. We use the letter "F" to denote the flow in the network and subscripts to indicate the starting

[1]We would like to thank the anonymous referee for pointing out the existence of the alternative network flow representation.

Table 1. The definition of the arcs in the network flow problem

Arc	*Variable*	*Cost per unit flow*	*Meaning*
(B_i, E_i)	$F_{B_iE_i}$	h_i	Inventory is held at retailer i
(B_i, M_i)	$F_{B_iM_i}$	0	Stock at retailer i is used to satisfy demand at retailer i
(B_i, M_j)	$F_{B_iM_j}$	c_{ij} ($c_{ii} = 0$)	Stock at retailer i is used to satisfy demand at retailer j, i.e., Transshipment from retailer i to retailer j
(R, M_i)	F_{RM_i}	p_i	Shortage at retailer i is satisfied through replenishment
(R, E_i)	F_{RE_i}	0	Inventory at retailer i is increased through replenishment

and ending nodes of the flow, thus $F_{B_iM_j}$ is the flow in the network from node B_i to M_j.

The complete network flow representation of the problem can be found in Fig. 1 for four retailers. Note that the graph is bipartite, although our representation of the graph, which was chosen to show the connection to the underlying inventory problem, does not emphasize this characteristic. The LP formulation associated with this network flow problem is as follows:

Problem (P)

$$Z(\mathbf{S}, \mathbf{d}) = \min \sum_{i=1}^{N} h_i F_{B_iE_i} + \sum_{i=1}^{N} \sum_{j=1}^{N} c_{ij} F_{B_iM_j} + \sum_{i=1}^{N} p_i F_{RM_i},$$

subject to

$$S_i = F_{B_iM_i} + \sum_{\substack{j=1 \\ j \neq i}}^{N} F_{B_iM_j} + F_{B_iE_i} \quad i = 1, \ldots, N, \tag{3}$$

$$F_{B_iM_i} + \sum_{\substack{j=1 \\ j \neq i}}^{N} F_{B_jM_i} + F_{RM_i} = d_i \quad i = 1, \ldots, N, \tag{4}$$

$$\sum_{i=1}^{N} d_i = \sum_{i=1}^{N} F_{RM_i} + \sum_{i=1}^{N} F_{RE_i}, \tag{5}$$

$$F_{B_iE_i} + F_{RE_i} = S_i \quad i = 1, \ldots, N, \tag{6}$$

$$F_{B_iE_i}, F_{B_iM_j}, F_{RM_i}, F_{RE_i} \geq 0 \quad i = 1, \ldots, N, \quad j = 1, \ldots, N. \tag{7}$$

Equations (3), (4), (5), and (6), respectively, represent the inventory balance constraint at the B_i, M_i, R, and E_i nodes.

3.3. *Finding the optimal order-up-to levels*

In the most general setting, exact computation of optimal order-up-to levels by analytical methods is difficult. We therefore use an approach based on Monte Carlo simulation: a set of demand realizations are sampled at random. For each demand realization, optimal transshipment quantities are computed along with gradient values. These values are then averaged across all samples and the sample average of the gradient value is used in the optimization. In particular, we deploy a stochastic approximation technique to compute the optimal order-up-to levels.

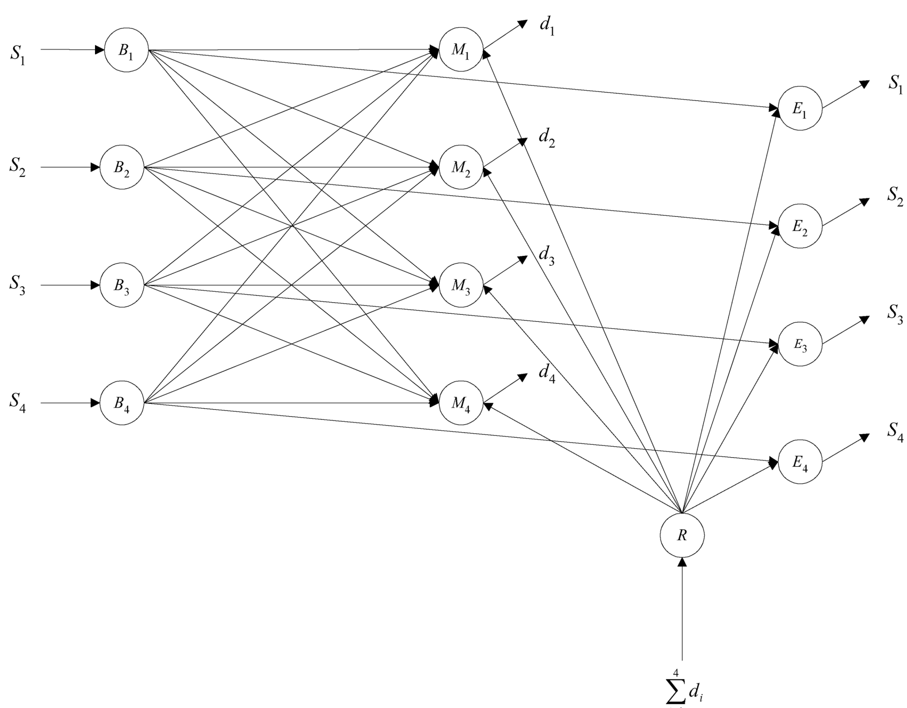

Fig. 1. Network flow representation of a single period.

Glasserman (1991) established the general conditions for the unbiasedness of the IPA estimator. 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).

The idea is to use the expected value of the sample path derivative obtained via simulation instead of using the derivative of the expected cost in a gradient search method. In other words, the gradient of interest is $\mathrm{d}E[TC]/\mathrm{d}S$ whereas our numerical procedure computes $E[\mathrm{d}TC/dS]$. To validate this approach, that is, to justify the interchange of the derivative and the integral, we need to show that the objective function is jointly convex and "smooth" in the S variables.

To show that the expected cost is jointly convex in the decision variables, we first show that for a given demand, $\mathbf{d}$, $Z(\mathbf{S},\mathbf{d})$ is jointly convex in $\mathbf{S}$. This is done by rewriting problem (P) such that all the S_i variables appear on the "right-hand side". We then apply the result that the objective functions of linear programs are convex piecewise linear functions of their right-hand sides (see, e.g., Bradley *et al.* (1977, p. 697)). Since the convolution of a convex function is itself convex we know that the expected cost in a single period, is itself jointly convex in $\mathbf{S}$.

It remains to show that the objective function is "smooth", i.e., the derivatives are both continuous and bounded to validate our IPA estimators (which we formulate below). As illustrated in Lemma 3.2 of Glasserman and Tayur (1995), continuity and boundedness can be verified by establishing that inventories are, with a probability of one, Lipschitz functions of the order-up-to levels, which is clearly the case here. Since the Lipschitz property is preserved by min/max and addition operators, the derivatives of the total cost are also both bounded and continuous functions of the order-up-to levels. To summarize, since we established the smoothness of the objective function, our IPA estimators are guaranteed to be unbiased.

3.3.1. *Description of the IPA procedure*

The procedure starts with an arbitrary value for the order-up-to levels, $\mathbf{S}$. An instance of the demand is generated at each retailer. Note that any covariance structure is allowed in $f(\mathbf{D})$. Once the demand is observed, problem (P) is solved in a deterministic fashion to compute the minimum-cost solution. The gradient of the total cost (derivatives with respect to the order-up-to levels) is estimated and accumulated over regenerative cycles; the average gradient value is then used to update the values of $\mathbf{S}$. A thorough review of simulation-based stochastic optimization techniques can be found in Shapiro (2001).

The procedure is summarized in a pseudo-code format, where K denotes the number of steps taken in a path search, U represents the number of regenerative cycles, a_k represents the step size at iteration k, and S_i^k represents the order-up-to level for retailer i at the kth iteration:

Algorithm 1

```
Initialize K
Initialize U
Set k ← 1
For each retailer, set initial order-up-to levels, S_i^0,
   possibly based on demand distribution
Repeat
Set dTC ← 0
Set u ←0
          Repeat
          i. Generate an instance of the demand at
             each retailer, d, from f(D)
          ii. Solve problem (P) to determine optimal
             transshipment quantities
          iii. Accumulate the desired gradients
             (derivatives) of the total cost, dTC
          iv. u ← u + 1
          Until u = U
     v.        Calculate the desired gradient(s),
                dTC/U
     vi.       Update the order-up-to levels, S_i:
                S_i^k ← S_i^(k-1) − a_k(dTC_i/U)
     vii.      k ← k + 1
Until k = K
```

In Step (iii) of the algorithm, we use IPA to compute the gradient. To illustrate the sample-path derivative idea, suppose that we end a period with inventory at retailer i. In this case, raising S_i by one unit would result in increasing the total cost by h_i. In the computer implementation, for each retailer i, we could partially code Step (iii) as:

$\mathrm{d}TC_i = \mathrm{d}TC_i + h_i$, if inventory at retailer i is positive, at the end of Step (ii).

Starting with $\mathrm{d}TC_i = 0$ for all i at the beginning of the simulation and dividing $\mathrm{d}TC_i$ by U in Step (v) yields the derivative estimates.

Our network flow formulation greatly simplifies computations. Increasing S_i corresponds to increasing the supply at source node B_i and the demand at sink node E_i. From a network flow perspective, $\mathrm{d}TC/\mathrm{d}S_i = h_i$, if the arc (B_i, E_i) is basic or, equivalently, the flow $F_{B_iE_i}$ is positive. If the arc is nonbasic, then since any basic solution corresponds to a tree in the network, there exists a unique augmenting path from B_i to E_i whose total cost yields the gradient value. For example, the augmenting path may go from B_i to M_j to R to E_i, with an associated cost of $c_{ij} - p_j$. Such a path represents a transshipment from retailer i to retailer j (with a cost of c_{ij}), a reduction in backorders at retailer j (with a savings of p_j) and a purchase of another unit at retailer i (cost of zero).

Furthermore, our implementation of the derivative computation in Step (iii) is very efficient. Since the value of the gradient is equal to the total cost along the unique path from B_i to E_i for each retailer i, this quantity can be calculated directly as the difference between the holding cost at retailer i and the reduced cost of the arc (B_i, E_i), which is readily available from the LP solution in Step (ii).

In Step (vi) of the algorithm, one typically imposes conditions on the step size a_k such that:

$$\sum_{k=1}^{\infty} a_k = \infty \quad \text{and} \quad \sum_{k=1}^{\infty} a_k^2 < \infty .$$

For instance, $a_k = 1/k$ satisfies these requirements. The first condition facilitates convergence by ensuring that the steps do not become too small too fast. However, if the algorithm is to converge, the step sizes must eventually become small, as ensured by the second condition. Note that when the gradient estimator is unbiased (as is the case here), Step (vi) represents a Robbins-Monro algorithm (Robbins and Monro, 1951) for stochastic search.

3.3.2. *Sensitivity analysis*

In a similar fashion, we can compute the derivative of the total cost with respect to other model parameters such as holding cost, penalty cost, transshipment cost, and replenishment cost. Furthermore, we can conduct performance analysis for service levels, expressed in terms of fill rate both at a single retailer and system-wide. Some of these gradient estimators are illustrated in Table 2.

The derivative estimators are quite intuitive. For example, suppose that in the optimal solution to the network flow problem retailer i holds inventory at the end of a period ($F_{B_iE_i} > 0$). An increase in the holding cost would therefore increase the total cost by the amount of excess stock being held. Similarly, if no excess inventory is held at retailer i at the end of a period, an increase in holding cost would have no impact on the total cost.

Finally, we should point out that as long as the transshipment policy preserves the smoothness of the cost function with respect to the order-up-to levels, Algorithm 1 (with an appropriately defined method of obtaining the per-period gradient information) can be used without modification. That is, the transshipment policy need not be optimal (as was also the case with the correctness of Theorem 1) if for some reason another transshipment policy is desired.

Table 2. Other gradient estimators

Derivative	*Estimator*
$\mathrm{d}TC/\mathrm{d}h_i$	$F_{B_iE_i}$
$\mathrm{d}TC/\mathrm{d}p_i$	F_{RM_i}
$\mathrm{d}TC/\mathrm{d}\hat{c}_{ij}$	$F_{B_iM_j}$
$\mathrm{d}TC/\mathrm{d}c_i$	$F_{B_iM_i} + \sum_{j\neq i} F_{B_iM_j} - \sum_{j\neq i} F_{B_jM_i}$

3.4. *Relaxing the restrictions on the parameters*

Several assumptions that are often made in the literature on transshipments and/or appear to be natural are *not* required for our model and analysis. These assumptions, some of which are typically referred to as triangle inequalities, are:

1. $c_{ij} \leq h_i + p_j$: Not requiring this inequality, i.e., allowing $c_{ij} > h_i + p_j$, means that when one retailer has an inventory surplus and the other has backlog before transshipments, it is not necessarily economic to transfer a unit from the former to the latter. With two-location models, as well as with identical-location models, this inequality is needed to ensure that transshipments are economic (otherwise, no transshipments will ever occur). However, since we have a multilocation model with possibly nonidentical costs this restriction is no longer natural. (Clearly, if this inequality is not satisfied for *all* pairs of i and j, no transshipment will occur.)
2. $p_j \leq c_{ij} + p_i$: Not requiring this inequality, i.e., allowing $p_j > c_{ij} + p_i$, means that it may be economic to transship a unit from retailer i to retailer j even when retailer i herself has a shortage. Such a cost structure may occur when different retailers have different priorities, and therefore a retailer with a higher priority might have a (possibly significant) higher unit shortage cost. We would expect this inequality to hold in most practical situations.
3. $c_{ik} \leq c_{ij} + c_{jk}$: Not requiring this inequality, i.e., allowing $c_{ik} > c_{ij} + c_{jk}$, means that it may be economic to use retailer j as an intermediary point between retailer i and retailer k, rather than to transship it directly from retailer i to retailer k. We envision such a situation when transshipments have to be accomplished within a limited time. Then, retailers i and j may be close enough to allow transshipments, and similarly retailers j and k. However, the time to transfer goods between retailers i and k may be so large that c_{ik} is in essence infinite.

   When retailer j is used as an intermediary point the amount transshipped through it is limited to S_j. Thus, it is incorrect to set $c_{ik} = c_{ij} + c_{jk}$. This point is illustrated in our computational study where this is the only difference between systems 3 and 4 (which are defined in Section 4.2 below).
4. Demands at different retailers in the same period are independent of one another. Not requiring this assumption means that in our model the demands among retailers in a given period may be correlated. Some of the existing transshipment literature could easily be extended to incorporate correlated demand, but the subject, in general, is not considered.

As mentioned, our model and analysis can handle all the above relaxations and generalizations without any modification.

4. Computational study

In this section, we report on our numerical study. We first report in Section 4.1 on a study conducted to validate our results and to fine-tune our algorithm. In Section 4.2, we describe the experimental design, which serves as the base case for all our experiments. In Section 4.3, we describe and analyze the results obtained for this basic experiment. In Sections 4.4 and 4.5, we describe two other experiments, for correlated demand and nonidentical costs, respectively, and describe and analyze their results.

4.1. *Validation and fine-tuning*

Recall that, in Step (v) of the algorithm, we incorporate our derivative estimates in a stochastic version of a gradient search technique. More specifically, for each retailer i we compute $S_i^k \leftarrow S_i^{k-1} - a_k\,(\mathrm{d}TC_i/U)$, where S_i^k is the order-up-to level for retailer i at the kth iteration, a_k is the step size, and $(\mathrm{d}TC_i/U)$ is the estimate of the gradient of the average cost when S_i^{k-1} is the order-up-to level at retailer i.

Finding effective values for the algorithm parameters, that is, starting values for the order-up-to levels, S_i^0, the step sizes, a_k, and the termination criteria, is generally a difficult problem. We conducted a thorough search, experimenting with different strategies using the illustrative examples from Krishnan and Rao (1965) and Tagaras (1989), where optimal solutions are available.

Based on this experimentation, we set the total number of steps for the path search $K = 10\,000$, the number of independent replications at each step $U = 50\,000$, and the step size $a_k = 1000/k$ for the validation examples. As a stopping criterion, we compared the order-up-to levels over 1000 iterations and required that these values do not differ by more than one. In all of our experiments, the convergence criterion was satisfied long before 10 000 iterations. Each experiment has also been replicated. The reported results reflect the averages across ten independent replications.

During the execution of the algorithm, the path search may push the order-up-to levels, S_i^k, below zero. This is due to the step size, a_k, being too large. Since a negative order-up-to level is not allowed by our assumption that the replenishment policy is nonshortage inducing, our algorithm simply resets their value to zero. We now illustrate our algorithm through example 1 in Krishnan and Rao (1965) with seven retailers. The characteristics of the retailers are summarized in Table 3 along with the optimal order-up-to levels calculated in Krishnan and Rao (1965). Recall that all retailers have identical cost structures with a holding cost of \$1 per unit, shortage cost of \$4 per unit, and a transshipment cost of \$0.10 per unit.

The last two rows of Table 3 depict the order-up-to levels computed by our algorithm. The half-width of a 95% confidence interval based on ten independent replications is also reported to show the low variability of the IPA estimators. The initial values for the order-up-to levels were $S_i^0 = 100$ for all retailers. The experiments were conducted on a personal computer with a 3-GHz Pentium IV microprocessor. Figure 2(a) shows the convergence of the algorithm for the seven-retailer network. Figure 2(b) illustrates the convergence of the order-up-to level for retailer 7, to depict the convergence rate more clearly. Figure 3 shows the run times (expressed in terms of the wall clock time) for networks ranging from two to seven retailers.

Note, from Fig. 2(a), that convergence to the correct order-up-to levels is very rapid. Quick convergence was also observed in all network configurations with two to seven retailers. Also, note that the results computed by the algorithm never deviate by more than 0.5% from the values reported in Krishnan and Rao (1965). Similar convergence behavior was observed with the test problem taken from Tagaras (1989). We should point out that the computational time, between 2 to 7 minutes for different numbers of retailers, is quite reasonable for a planning problem. Moreover, to obtain a rough estimate of the results even faster, e.g., for the purpose of a "what-if" type of analysis, a limited number of iterations may be conducted (see Fig. 2(b)).

4.2. *Experimental design*

To show the flexibility afforded by our modeling and analysis framework, we have experimented with large networks, with retailers whose demand is correlated, and with an arbitrary cost structure. We consider systems with $N + 1$ retailers, where $N \in \{7, 9, \ldots, 21\}$. An illustrative example of the system with four retailers is shown in Fig. 4. Let us call retailer 0 the *central retailer* and all the other N retailers the *remote retailers.* We begin by considering the case

Table 3. Example 1 from Krishnan and Rao (1965)

Retailer	*1*	*2*	*3*	*4*	*5*	*6*	*7*
Normal demand (μ, σ)	100,20	200,50	150,30	170,50	180,40	170,30	170,50
S_i^*	106.7	216.7	160.0	186.7	193.4	180.0	186.7
Computed avg S_i^*	106.72	216.21	160.14	186.81	193.42	180.11	186.82
Half width of a 95% confidence interval	0.065	0.058	0.073	0.064	0.066	0.064	0.046

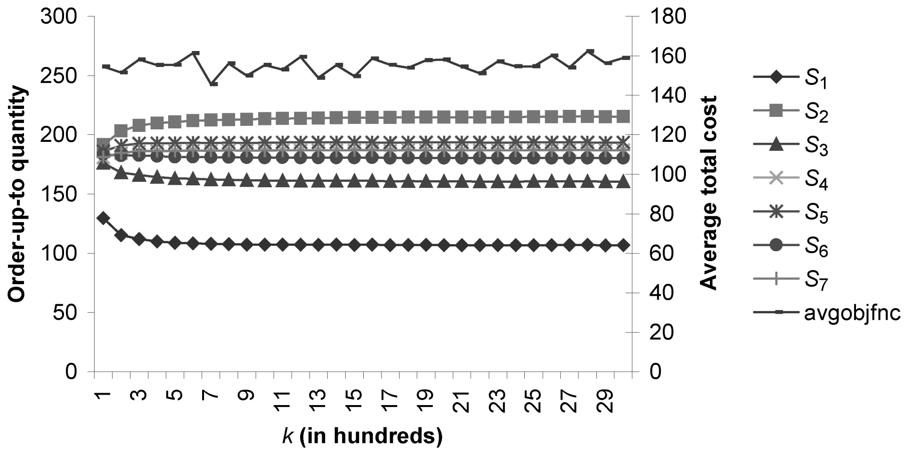

(a)

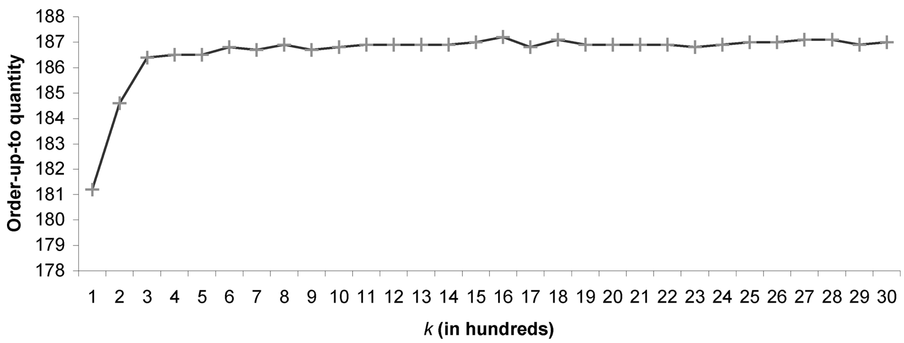

(b)

Fig. 2. (a) Convergence of the algorithm for a seven-retailer configuration; and (b) computation of the order-up-to level for retailer 7.

of identical retailers, the cost parameters are as follows: $h_i \equiv \mathbf{h} = \$1$ per unit, $p_i \equiv \mathbf{p} = \$4$ per unit, and the basic direct transshipment cost, $c_t = \$0.5$ per unit, when transshipments are allowed. Each retailer faces an independent demand stream distributed uniformly over (0, 200).

Note that $c_{0i}, i = 1, 2, \ldots, N$, represents the transshipment cost from the central retailer to remote retailers, $c_{i0}, i = 1, 2, \ldots, N$, represents the transshipment cost from the remote retailers to the central retailer, and $c_{ij}, 1, j = 1, 2, \ldots, N$, denotes the transshipment cost from remote retailer i to remote retailer j. As summarized in Table 4, we consider five alternative system configurations and we denote by S_i^s the order-up-to level for retailer i under system $s, s = 1, \ldots, 5$. Note that $c_{ij} = \infty$ implies that transshipments are not allowed between retailers i and j.

System 1, where no material movement is allowed among retailers, represents $N + 1$ independent newsvendor problems. It thus serves as a benchmark. In system 2, transshipments are allowed only from the central retailer to the remote retailers. System 3 extends the scenario in system 2 by also allowing transshipments from the remote retailers to the central retailer as well. In system 4, all material movement is possible. However, transshipments between

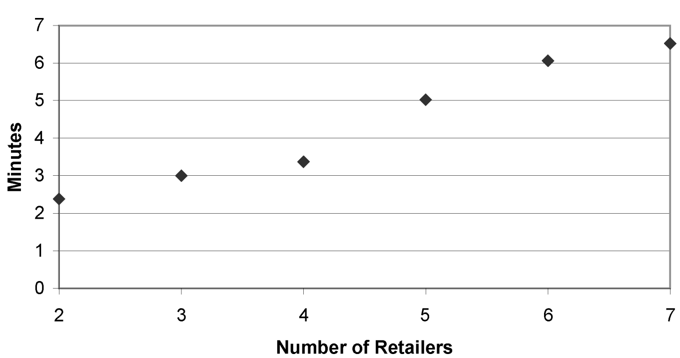

Fig. 3. Run time of the algorithm.

Table 4. System configurations

System	c_{0i}	c_{i0}	c_{ij}
1	∞	∞	∞
2	c_t	∞	∞
3	c_t	c_t	∞
4	c_t	c_t	$2c_t$
5	c_t	c_t	c_t

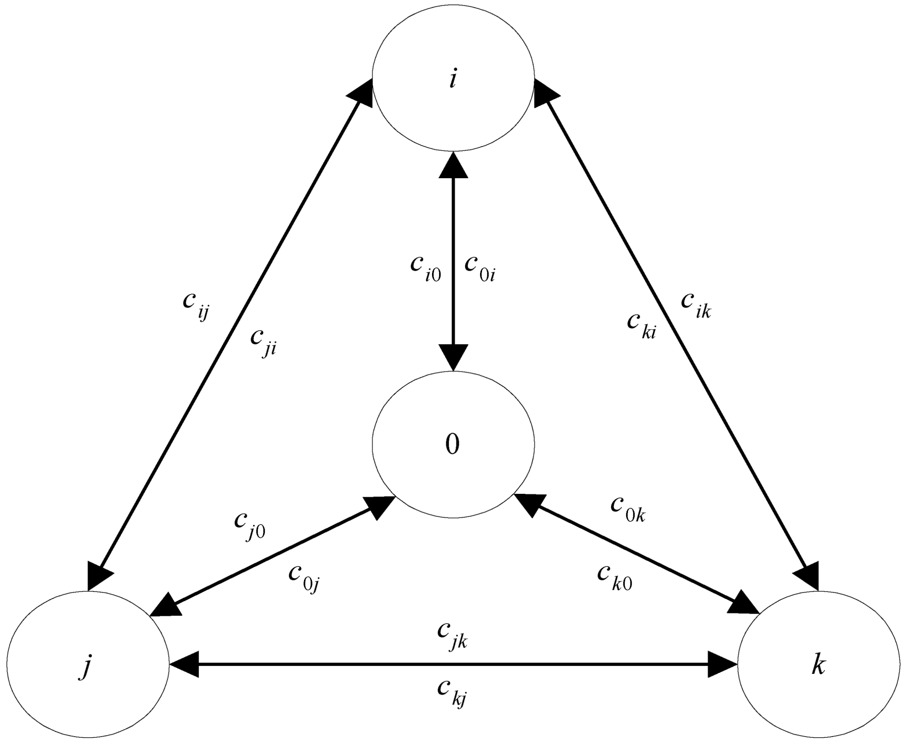

Fig. 4. Configuration with four retailers.

any two remote retailers are twice as expensive as the transshipments from/to the central retailer. Finally, all transshipment costs are identical in system 5.

4.3. *Results and analysis for the base case*

The first set of experiments consists of configurations where the N remote retailers have identical cost parameters, and all retailers have independent and identically distributed demand. The order-up-to levels computed by our algorithm for the ten-retailer configuration are listed in Table 5 and are depicted in Fig. 5. As locations with identical characteristics are converging to the same number, we present the base stock level of the central retailer and the average base stock level of the remote retailers. The average total cost for the optimal configuration is also shown in Table 5.

Table 5. Optimal order-up-to levels in a ten-retailer configuration

System	S_0	*Average* $S_1 - S_9$	*Total Inv.*	*Cost*
1	159.9	160.0	1600	800.64
2	481.7	87.1	1266	486.95
3	319.2	100.3	1222	410.67
4	172.8	113.3	1193	383.81
5	117.1	117.1	1171	334.66

The results of this set of experiments confirm the intuition about the behavior of the systems, as follows: in system 2, the central retailer carries considerably more inventory than the other retailers, since this stock can be transshipped to other retailers to meet the demand they face. Given the possibility of transshipments to/from the central retailer in system 3, we observe a reduction in inventory in the central retailer together with an increase in inventory at the other retailers. In system 4, this phenomenon is further accentuated. System 5, where transshipments are allowed among all retailers, distributes the inventory evenly throughout the system as in system 1, but at a lower cost than the newsvendor benchmark of system 1. Comparing system 1, where we have ten independent newsvendors, with system 5, where transshipments are allowed among all retailers at the basic cost, we note that system-wide inventory is significantly reduced. For the ten-retailer configuration, this reduction in inventory leads to a 58% reduction in total costs, as shown in Fig. 6. Note, however, that a large part of this benefit, a 39% reduction in total cost, is obtained when moving from system 1 to system 2, thus demonstrating that a little bit of flexibility goes a long way.

Jordan and Graves (1995) showed results which are qualitatively similar to ours in the context of process flexibility, defined as the ability to build different types of products in the same plant at the same time. In particular, they showed that limited flexibility, configured as a chain that connects

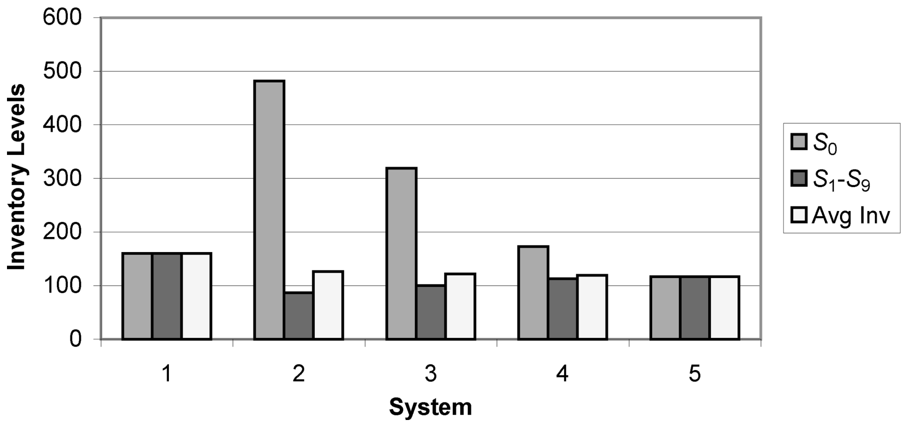

Fig. 5. Optimal order-up-to levels in a ten-retailer configuration.

products and plants, yields most of the benefits of total flexibility.

Figure 7 depicts the value of transshipments: for varying number of retailers considered, we observe that the expected total system cost decreases significantly as transshipments become more flexible and less expensive.

4.4. *Correlated demand*

To study the impact of correlated demand, we consider a ten-retailer configuration, with the same cost structure as described in the previous section. We experiment with scenarios of high (± 0.9), medium (± 0.5), and low (± 0.2) values of the demand correlation coefficient. A case with zero correlation is also added for reference. Unlike the previous section, the demand faced by the retailers is modeled as a multivariate normal random variable with a mean of 100 and a standard deviation of 20. The (i,j)th entry of the variance-covariance matrix is given by $\sigma_i \sigma_j \rho_{ij}$, where ρ_{ij} denotes the demand correlation coefficient being investigated when $i \neq j$ and one when $i = j$. Thus, for example, when we investigate medium negative correlation the diagonal elements of the variance-covariance matrix are all 400 and the off-diagonal elements are all –200.

Correlated demand can be found in many real situations. For example, positive correlation can be caused by some event common to all locations, e.g., rain causes demand for umbrellas to increase at all locations. Negative correlation, on the other hand, can be due to the fact that sometimes there is only one winner. Thus, the demand for alcohol in the hometown of two competing football teams is negatively correlated, as only one of the team's supporters will have something to celebrate.[2]

For system 1, where no transshipments take place, positively or negatively correlated demand has no impact on base stock levels or total cost as each retailer solves his own newsvendor problem, the solution of which is to order-up-to 116.8 units. When transshipments are allowed, however, correlated demand does have a sizeable impact. In general, positive correlation reduces the effectiveness of transshipments whereas negative correlation enhances it. In particular, with a high positive correlation, the difference among the five systems under consideration is relatively small. In particular, every system behaves similarly to system 1, in

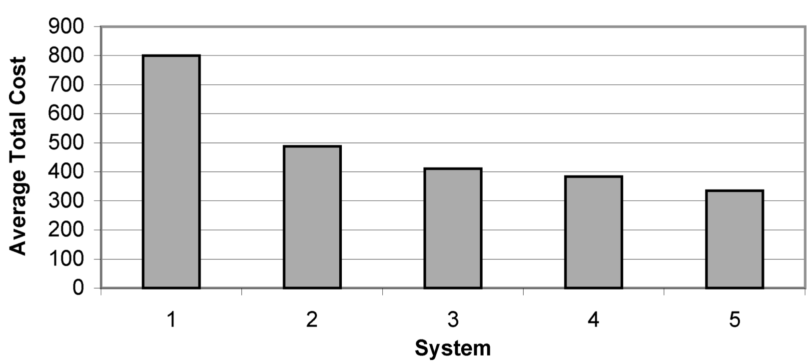

Fig. 6. Average total cost under different systems for the ten-retailer configuration.

[2]We thank an anonymous referee for suggesting this example.

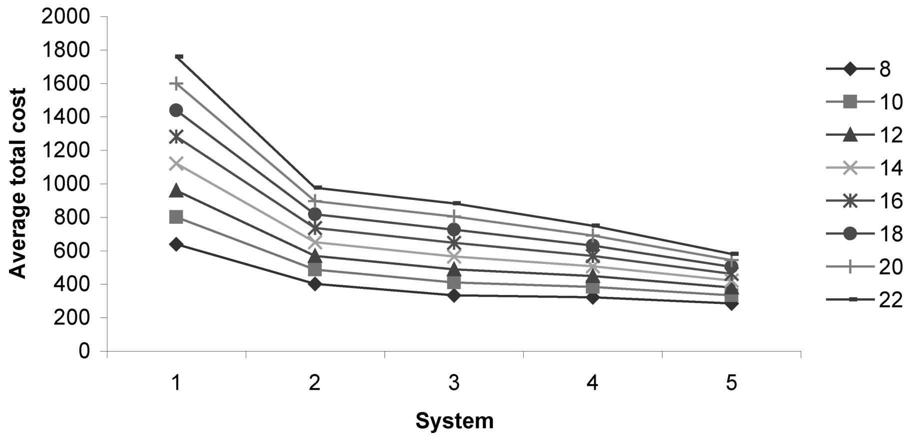

Fig. 7. Average total cost under different systems for varying numbers of retailers.

which transshipments are not allowed, and the objective function values of all systems are practically indistinguishable. The base stock and the average inventory levels in systems 2 through 5 are shown in Fig. 8(a–d).

In system 2, positive correlation limits the role of the central retailer as a clearinghouse for the remote retailers. Any level of negative correlation, on the other hand, reinforces the central retailer's clearinghouse role.

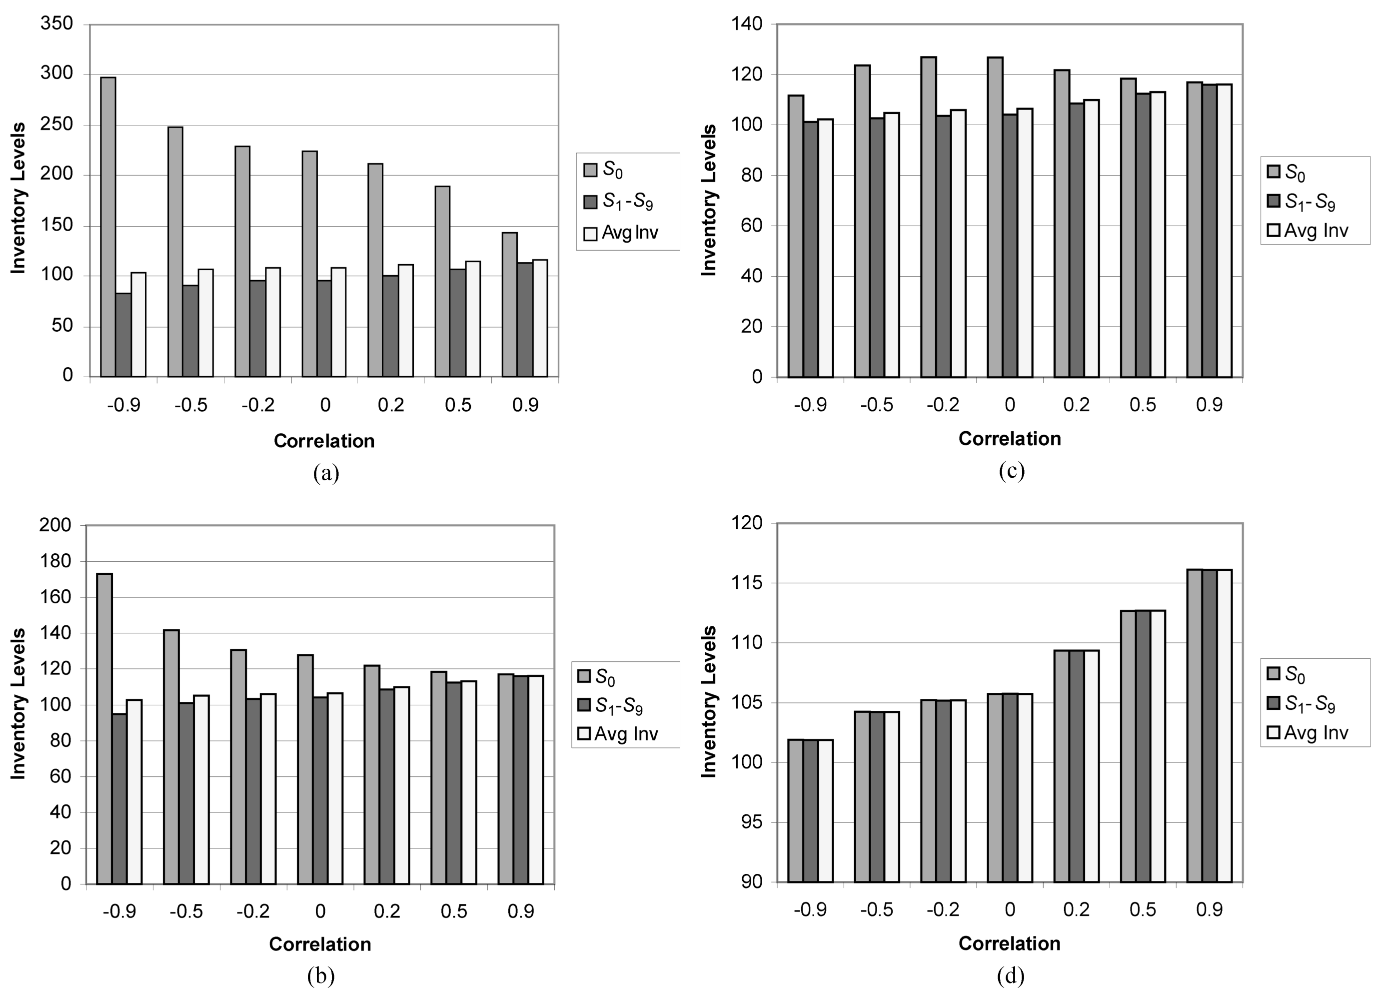

Fig. 8. Optimal order-up-to and average inventory levels for ten retailers with correlated demand: (a) system 2; (b) system 3; (c) system 4; and (d) system (5).

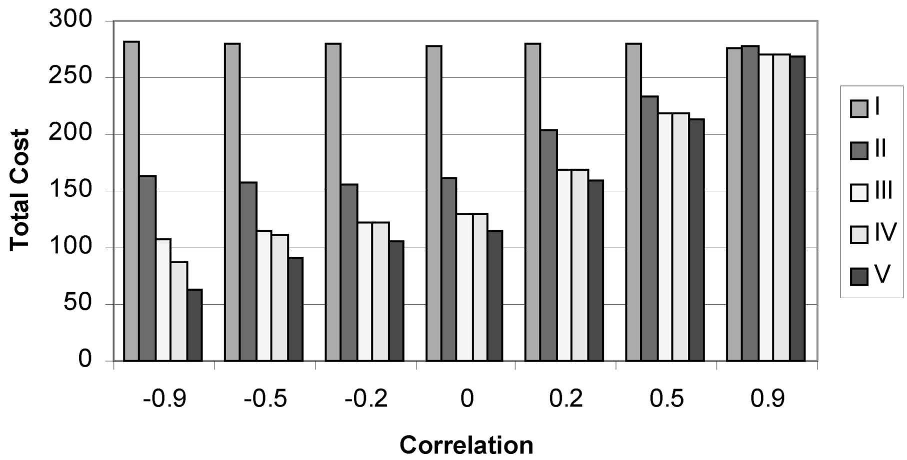

Fig. 9. Average total cost for ten retailers with correlated demand.

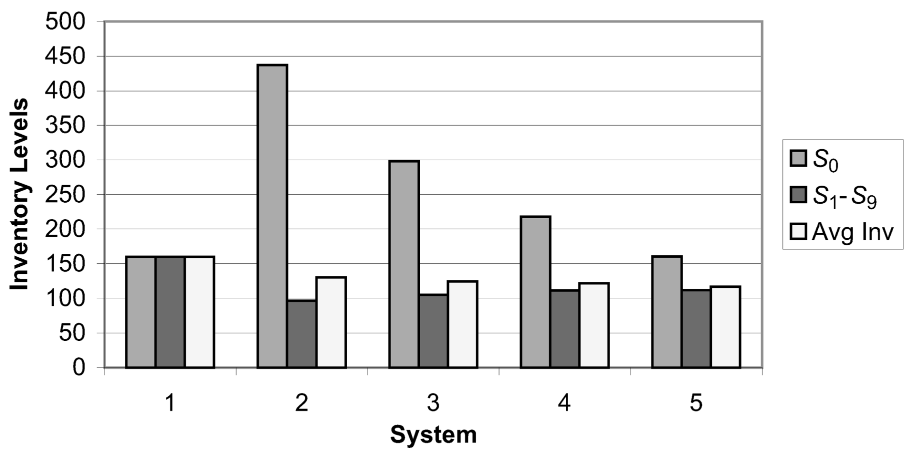

Fig. 10. Optimal order-up-to levels for ten nonidentical retailers.

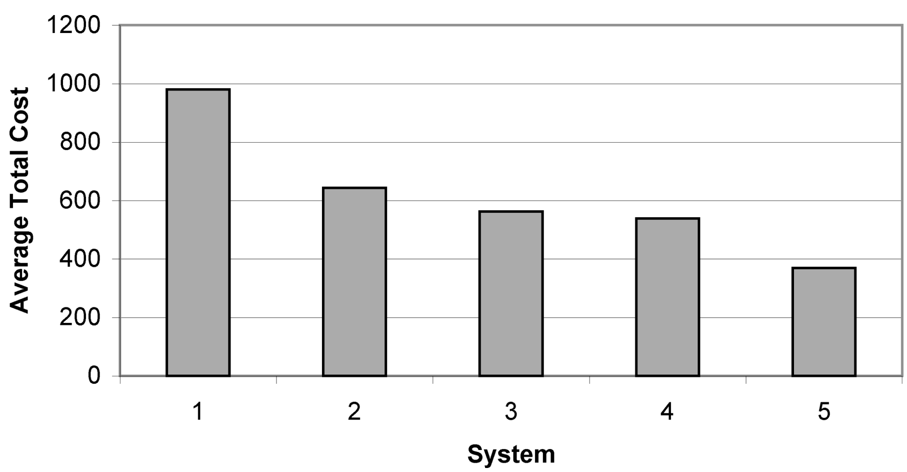

Fig. 11. Average total cost for ten retailers with nonidentical costs.

This type of behavior is also observed in systems 3 and 4. The graph for system 5 is the most drastic illustration of how positive the correlation reduces effectiveness of transshipments. When demand has a high positive correlation then the average inventory is 116.1 units, which is very close to the system 1 level of 116.8 units. As the demand becomes negatively correlated, however, the ability to match demand with supply through transshipments is further enhanced.

Figure 9 illustrates the impact of demand correlation on the average total cost for a ten-retailer configuration. High levels of positive correlation eliminate the value of transshipments making all five systems quite costly to operate. As demand correlation gets smaller (or negative), the effectiveness of transshipments in matching demand and supply is enhanced, which is reflected by the significantly lower average total cost of system 5.

4.5. *Nonidentical costs*

In all configurations considered thus far, all remote retailers have had identical cost parameters. These cost parameters differ from the central retailer's cost parameters with respect to the transshipment cost, as our solution technique can handle nonidentical cost parameters. To further emphasize this ability, we consider a ten-retailer configuration, where we modify the cost parameters without violating Equation (1), $h_i \leq c_{ij} + h_j$ for all *i,j*. In particular, we set $h_0 = \$1$ as before, and $h_i = h_{i-1} + 0.05$, $i = 1, \ldots, 9$. Similarly, $p_0 = \$4$ as before, and $p_j = p_{j-1} + 0.20$, $j = 1, \ldots, 9$. For system 1, where no transshipments are allowed, $c_{ij} = +\infty$. For system 2, $c_{01} = \$0.5$ and $c_{0j} = c_{0,j-1} + 0.1$, for $j = 2, \ldots, 9$, and $c_{ij} = +\infty$ otherwise. For system 3, $c_{j0} = c_{0j}$ of system 2, and $c_{ij} = +\infty$ otherwise. For system 4, c_{j0} and c_{0j} as in system 3 and $c_{12} = c_{21} = \$1.2$, $c_{ij} = c_{i-1,j-1} + 0.2$, $i, j = 2, \ldots, 9$, $i \neq j$. Finally, for system 5, $c_{ij} = \$0.5$ for all *i,j*.

We observe that transshipments maintain their positive impact on the overall performance, as systems $1, \ldots, 5$ lead gradually to lower stock levels (Fig. 10) and lower total cost (Fig. 11). The magnitude of this improvement, however, is heavily dependent on the relative cost parameters.

5. Summary

In this paper, we considered the multilocation dynamic transshipment problem. First, an arbitrary number of nonidentical retailers was considered with possibly dependent stochastic demand. Second, we modeled the dynamic behavior of the system in an arbitrary period as a network flow problem. Finally, we employed a simulation-based method using IPA for optimization. Our simulation-based optimization approach therefore provides a flexible platform to analyze transshipment problems of arbitrary complexity. An interesting generalization to the problem addressed in this paper is the case of positive replenishment lead times. In this case, it is not immediately clear how to find the optimal transshipment policy, since it may be beneficial for a retailer to hold back some of her own inventory rather than transship it. As a result, it is also not clear whether an order-up-to policy remains optimal. These will be interesting issues for future research.

Acknowledgements

This work has been supported by a grant of the Israeli Ministry of Science and AFIRST, Association Franco-Israélienne pour la Recherche Scientifique et Technologique. The authors would also like to express their sincere appreciation for the programming support of Ann Larsson.

References

Archibald, T.W., Sassen, A.A.E. and Thomas, L.C. (1997) An optimal policy for a two-depot inventory problem with stock transfer. *Management Science*, **43**, 173–183.

Bradley, S.P., Hax, A.C. and Magnanti, T.L. (1977) *Applied Mathematical Programming*, Addison-Wesley, Reading, MA.

Dong, L. and Rudi, N. (2004) Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Science*, **50**, 645–657.

Eppen, G. (1979) Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, **25**, 498–501.

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, Norvell, MA.

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, Norvell, MA.

Glasserman, P. (1992) Derivative estimates from simulation of continuous-time Markov chains. *Operations Research*, **40**, 292–308.

Glasserman, P. (1994) Perturbation analysis of production networks, in *Stochastic Modeling and Analysis of Manufacturing Systems*, Yao (Ed.), Springer-Verlag, New York.

Glasserman, P. and Tayur, S. (1995) Sensitivity analysis for base stock levels in multi-echelon production-inventory systems. *Management Science*, **41**, 263–281.

Herer, Y.T. and Rashit, A. (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. and Tzur, M. (2001) The dynamic transshipment problem. *Naval Research Logistics*, **48**, 386–408.

Herer, Y.T. and Tzur, M. (2003) Optimal and heuristic algorithms for the multi-location dynamic transshipment problem with fixed transshipment costs. *IIE Transactions*, **5**(5), 419–432.

Ho, Y.C., Eyler, M.A. and Chien, T.T. (1979) A gradient technique for general buffer storage design in a serial production line. *International Journal of Production Research*, **17**, 557–580.

Jordan, W.C. and Graves, S.C. (1995) Principles on the benefits of manufacturing process flexibility. *Management Science*, **41**, 577–594.

Krishnan, K.S. and Rao, V. R. K. (1965) Inventory control in *N* warehouses. *Journal of Industrial Engineering*, **16**, 212–215.

Robbins, H. and Monro, S. (1951) A stochastic approximation method. *Annals of Mathematical Statistics*, **22**, 400–407.

Robinson, L.W. (1990) Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Operations Research*, **38**, 278–295.

Rudi, N., Kapur, S. and Pyke, D. (2001) A two-location inventory model with transshipment and local decision making. *Management Science*, **47**, 1668–1680.

Shapiro, A. (2001) Stochastic programming by Monte Carlo simulation methods. Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

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*, **27**, 39–59.

Tagaras, G. and Cohen, M. (1992) Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Science*, **38**, 1067–1083.

Yücesan, E. (2003) Private communication with FNAC supply chain manager.

Biographies

Yale T. Herer , B.S. (1986), M.S. (1990), Ph. D. (1990), Cornell University, Department of Operations Research and Industrial Engineering. Yale joined the Faculty of Industrial Technology Engineering and Management at the Tehnion—Israel Institute of Technology in 1990 immediately after the completion of his graduate studies. In 1997 Yale joined the Department of Industrial Engineering at Tel- Aviv University and in 2001 he returned to the Technion. During the 2004–2005 academic year he spent a sabbatical at Northwestern University in the Department of Industrial Engineering and Management Sciences. Yale also briefly visited Cornell University in the Department of Operations Research and Industrial Engineering. Between 2003–2004, Yale is the Head of the Industrial Engineering Area at the Technion. He has worked for several industrial concerns, both as a consultant and as an advisor to project groups. In 1996 Yale received the IIE Transactions best paper award. Yale is a member of the Institute for Operations Research and Management Sciences (INFORMS), Institute of Industrial Engineers (IIE), and the Operations Research Society of Israel (ORSIS). He serves on the Editorial Board of IIE Transactions. His research interests include supply chain management, especially when integrated with transshipments. He is also interested in production control and production system design.

Michal Tzur is a senior lecturer in the department of Industrial Engineering at Tel Aviv University, Israel, and is currently the department chair. Recently she visited the IEMS department at Northwestern University. Michal joined Tel Aviv University in 1994 after spending 3 years at the Operations and Information Management Department at Wharton School at the University of Pennsylvania. She received her B.A. from Tel Aviv University and her M. Phil and Ph.D. from Columbia University. Her research interests are in the areas of supply chain management, multi-echelon inventory management, production planning and operations scheduling.

Enver Yücesan is a professor of Operations Management at INSEAD in Fontainebleau, France. He received his undergraduate degree in Industrial Engineering from Purdue University and his PhD in Operations Research from Cornell University. His research interests include discrete event simulation modeling, simulation optimization, ranking and selection methods, and supply chain management.

Contributed by the Supply Chains/Production-Inventory Systems Department