The Multi-Location Transshipment Problem with Positive Replenishment Lead Times

Yeming Gong and Enver Yucesan

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2006-048-LIS
Publication	August 2006
Number of pages	39
Persistent paper URL	
Email address corresponding author	ygong@rsm.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS

Abstract	Transshipments, monitored movements of material at the same echelon of a supply chain, represent an effective pooling mechanism. With a single exception, research on transshipments overlooks replenishment lead times. The only approach for two-location inventory systems with non-negligible lead times could not be generalized to a multi-location setting, and the proposed heuristic method cannot guarantee to provide optimal solutions. This paper uses simulation optimization by combining an LP/network flow formulation with infinitesimal perturbation analysis to examine the multi-location transshipment problem with positive replenishment lead times, and demonstrates the computation of the optimal base stock quantities through sample path optimization. From a methodological perspective, this paper deploys an elegant duality-based gradient computation method to improve computational efficiency. In test problems, our algorithm was also able to achieve better objective values than an existing algorithm.
Free Keywords	Transshipment, Simulation Optimization, Infinitesimal Perturbation Analysis (IPA)
Availability	The ERIM Report Series is distributed through the following platforms: Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal Social Science Research Network (SSRN), SSRN ERIM Series Webpage Research Papers in Economics (REPEC), REPEC ERIM Series Webpage
Classifications	The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems: Library of Congress Classification, (LCC) LCC Webpage Journal of Economic Literature, (JEL), JEL Webpage ACM Computing Classification System CCS Webpage Inspec Classification scheme (ICS), ICS Webpage

The Multi-Location Transshipment Problem with Positive Replenishment Lead Times

YEMING GONG* and ENVER YUCESAN**

*RSM Erasmus University, P.O. box 1738, 3000 DR Rotterdam, The Netherlands

**INSEAD, Technology and Operations Area, Boulevard de Constance, 77305 Fontainebleau Cedex, France

Abstract

Transshipments, monitored movements of material at the same echelon of a supply chain, represent an effective pooling mechanism. With a single exception, research on transshipments overlooks replenishment lead times. The only approach for two-location inventory systems with non-negligible lead times could not be generalized to a multi-location setting, and the proposed heuristic method cannot guarantee to provide optimal solutions. This paper uses simulation optimization by combining an LP/network flow formulation with infinitesimal perturbation analysis to examine the multi-location transshipment problem with positive replenishment lead times, and demonstrates the computation of the optimal base stock quantities through sample path optimization. From a methodological perspective, this paper deploys an elegant duality-based gradient computation method to improve computational efficiency. In test problems, our algorithm was also able to achieve better objective values than an existing algorithm.

Key words: Transshipment, Simulation Optimization, Infinitesimal Perturbation Analysis (IPA)

1. Introduction

Physical pooling of inventories (Eppen 1979) has been widely used in practice to reduce cost and improve customer service. For example, *CIBA Vision* has consolidated all of its country-based warehouses in Europe into a single European Logistics Center near Frankfurt, Germany. On the other hand, the practice of *transshipment,* the monitored movement of material between pairs of 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, this practice is typically referred to as *information pooling*. 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.

Although they are often overlooked in the literature, replenishment lead times constitute one of the critical factors in a transshipment system. Consider, for example, the Normandy landing where we can view the military logistics system as a two-echelon supply chain with the main base as a "supplier" in England and five bases on Normandy beaches in France. When the Allied Forces landed on Utah Beach, they met much less Nazi resistance than those landing on Omaha Beach, which enabled them to move troops and material from Utah Beach to Omaha Beach. This flow can simply be viewed as transshipment. In this case, ignoring replenishment lead times, i.e., the time to move new troops and material across the English Channel, would have disastrous consequences.

Similarly, ASML, a Dutch manufacturer of photolithography equipment, reports that its customers in Japan, which manufacture electronic components, regularly tranship spare parts among themselves in order to avoid downtime –hence, lost throughput– due to replenishment lead times from Holland.

Transshipments have the advantage of improved flexibility and responsiveness without increasing total inventories. Replenishment lead times, however, will weaken the responsiveness and the flexibility of a supply chain by reducing the attractiveness of transshipments. To the best of our knowledge, with the exception of Tagaras and Cohen (1992), replenishment lead times have not been incorporated in transshipment models. Hence, in terms of positive replenishment lead times, this paper extends Herer et al. (2005), who studied the multi-location transshipment without replenishment lead times. In terms of a multi-location setting, this paper generalizes Tagaras and Cohen (1992), who considered non-negligible replenishment lead times in two-location inventory systems. However, their method has not proved to be generalizable to a multi-location setting. Furthermore, their heuristic algorithm cannot guarantee optimal solutions.

In order to compute the optimal values for multi-location system with positive replenishment lead times, one of the most efficient methods is simulation optimization, which can help the search for an improved policy while allowing for complex features that are typically outside of the scope of analytical models. Sample path optimization (SPO), also called the stochastic counterpart method, is a simulation optimization method that has the significant advantages of high efficiency and convenience. However, SPO requires a technique to estimate the gradient.

There exist a large number of gradient estimation techniques such as Infinitesimal Perturbation Analysis (IPA), Likelihood Ratios (LR), Finite Differences (FD), Symmetric Difference (SD), and Simultaneous Perturbation (SP) (Fu 2002). IPA is an efficient gradient estimation technique (Ho et al. 1979). 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).

To study the multi-location transshipment problem with positive replenishment lead times, this paper deploys an LP/network flow model, uses sample path optimization and infinitesimal perturbation analysis techniques, and demonstrates the computation of the optimal base stock quantities. In contrast with the existing literature, this paper uses an elegant duality-based gradient computation method to improve algorithm efficiency.

The remainder of the paper is organized as follows: In the following section, we introduce the multi-location transshipment model with the positive replenishment lead times and its network flow and LP representations. Section 3 is devoted to the details of the algorithm, its implementation, and its verification and validation. In Section 4, we present the results of our extensive numerical experimentation. We conclude with final comments in Section 5.

2. Model

2.1 The Model Description

We consider a system with one supplier and N retailers, associated with N distinct stocking locations that face customer demand. The retailers may differ in their cost and demand parameters. 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. The replenishment order will arrive after a positive replenishment lead time L. In the presence of a positive replenishment lead time, the system needs a bigger safety inventory, with a significant effect on transshipment.

In each period, the replenishment and transshipment quantities must be determined in order to minimize the expected average total cost. The total cost is the sum of the replenishment, transshipment, holding, backlog penalty, and lost sales costs. Herer et al. (2005) prove that, in the absence of fixed costs, if transshipments are made to compensate for an actual shortage (instead of building up inventory at another stocking location), there exists an optimal base stock $S = (S_1, S_2, \ldots, S_N)$ policy for all possible stationary transshipment policies. In our case, since the transshipment policy is stationary, we will continue to adhere to the base stock replenishment policy.

In period t, events occur in the following order, as illustrated in Figure 1: First, retailers observe demands. Demand realizations represent the only uncertain event of the period. Once demand is observed, decisions about transshipment quantities are made. The transshipment transfers are then made immediately; subsequently, demand is satisfied. Any unsatisfied demand will be backlogged or lost. At this point, backlogs and inventories are observed, and penalty and holding costs, respectively, are incurred. Second, replenishment orders placed at the supplier in period t-L arrive. These orders are used to satisfy the backlog in period t-L and, if possible, to increase the inventory level in period t. The decision on the replenishment quantity is then made. Any remaining inventory is carried to the next period, t+1.

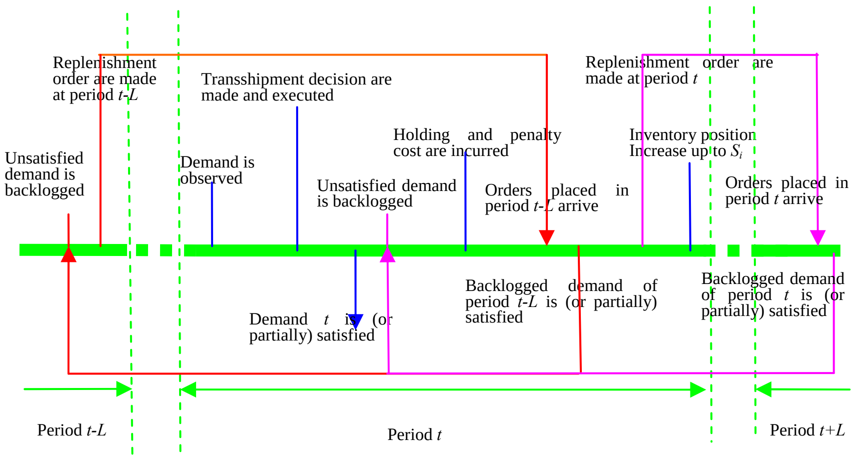

FIGURE 1: Sequence of events in a period

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

L = positive replenishment **l**ead time;

T = the **t**ime horizon;

N = **n**umber of retailers;

$D_i(t)$ = random variable associated with demand at retailer i, $i=1,2,\ldots,N$, $t=1,2,\ldots,T$;

$d_i(t)$ = actual **d**emand at retailer i and an arbitrary period t, $i=1,2,\ldots,N$, $t=1,2,\ldots,T$;

S_i = base stock quantity at the location i, $i=1,2,\ldots,N$;

h_i = **h**olding cost incurred at retailer i per unit held per period, $i=1,2,\ldots,N$;

p_i = **p**enalty cost incurred at retailer i per unit backlogged per period in the first T-L periods, $i=1,2,\ldots,N$;

l_i = penalty for **l**ost sales at retailer i per unit of unmet demand per period in the last L periods, $i=1,2,\ldots,N$. During the last L periods, it is impossible for replenishment orders to arrive on time. The unmet demand cannot therefore be backordered but is lost;

c_i = replenishment cost per unit at retailer i, $i=1,2,\ldots,N$;

c_{ij} = effective transshipment cost, or simply the transshipment cost, per unit transshipped from retailer i to retailer j, $i,j=1,2,\ldots,N$;

We consider base stock policies, where $IOH_i(t)$ is **i**nventory **o**n **h**and at location i and the beginning of period t. When t=1, $IOH_i(t)$ is $S_i(t)$, the base stock at retailer i. Given $d_i(t)$, the actual demand at retailer i in period t, the dynamic behavior of the system is captured through the following auxiliary variable:

$I_i(t)$: inventory level at retailer i immediately after transshipments and demand satisfaction

$$= S_i(t) - \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + \sum_{j=1, j\neq i}^{N} B_j(t)M_i(t) - d_i(t), \text{ for } t=1$$

$$= IOH_i(t) - \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + \sum_{j=1, j\neq i}^{N} B_j(t)M_i(t) - d_i(t), \text{ for } t=2,\ldots T,$$

where $B_i(t)M_j(t)$ represents the transshipment quantity from retailer *i* to retailer *j*. We denote: $I_i^+(t) = \max\{I_i(t),0\}$, $I_i^- = \max\{-I_i(t),0\}$. Thus, the realized average cost per period of the system over a horizon *T* is equal to:

$$AC = \frac{1}{T}[\sum_{t=1}^{T} (\sum_{i=1}^{N} \sum_{j=1, j\neq i}^{N} c_{ij} B_i(t) M_j(t) + \sum_{i=1}^{N} h_i I_i^+(t) + \sum_{i=1}^{N} c_i d_i(t)) + \sum_{t=1}^{T-L} \sum_{i=1}^{N} p_i I_i^-(t) + \sum_{t=T-L+1}^{T} \sum_{i=1}^{N} l_i I_i^-(t))]$$

Term $\sum_{i=1}^{N} c_i d_i$ fully accounts for replenishment costs. Since this term is independent of our decision variables, it is omitted below.

2.2 Modeling Assumptions

We will make the following assumptions, which are necessary to avoid pathological cases.

Assumption 1 (Lead time): Replenishment lead times are both positive and deterministic.

Assumption 2 (Lateral transshipment): Lateral transshipment lead times are negligible between any pair of stocking locations.

Assumption 3 (Demand): Customer demand at each retailer is generated by a stochastic process. Demand is backlogged when a retailer is out of stock in t=1,2,…*T-L*, but is lost in the last *L* periods since the replenishment orders cannot arrive on time within the finite horizon. Demand has a continuous CDF, but is not necessarily independent across retailers.

Assumption 4 (Replenishment policy): The base stock quantity is nonnegative, which also implies a non-shortage inducing replenishment policy (Herer et al. 2005). A replenishment quantity ordered at period *t-L* arrives at period *t* and satisfies the backorder at period *t-L*; any remaining units go to the next period, *t*+1.

Assumptions 5 (Transshipment policy): The transshipment policy is stationary, that is, the transshipment quantities are independent of the period in which they are made; they depend only on the pre-transshipment inventory and the observed demand. As stated earlier, we assume that transshipments are never made to build up inventory at the receiving location, and only made to satisfy current actual shortage.

2.3 Model formulation

We present the network flow formulation first, and then give stochastic programming and its determinant counterpart, i.e., the LP formulation based on the network flow formulation.

2.3.1 Network Flow Representation

Given a base stock policy for the replenishment quantities, the optimal transshipment quantities need to be determined each period between every pair of retailers. We develop a linear cost network flow model as follows. In the presence of positive lead times, the inventory position will not always be equal to inventory on hand since there exists inventory on order in the pipeline. Proposition 1 establishes that, in the presence of positive lead times, $\{IOH(t)\}$ is not a regenerative process; hence, this transshipment system cannot be reduced to a one-period problem. We therefore formulate the system as a finite horizon system. In Proposition 1, $\Re^+$ denotes the set of non-negative real numbers.

Proposition 1: *Let* $\Xi = \{1,2,...,T\}$ *be time index set from 1 to finite horizon T,* $\{IOH(t), t \in \Xi\}$ *be the stochastic inventory-on-hand process with* $IOH(t) \in \Re^+$. *Then in the presence of positive replenishment lead times,* $\{IOH(t), t \in \Xi\}$ *is not a regenerative process, and the regenerative epoch* t_1 *with* $t_1 \in \Xi$ *do not always exist in this transshipment system.*

Proof The proof is presented in Appendix A.

Let us recall the events in a period *t*; in particular, let us examine the material flows. At the beginning of the period, the excess inventory from the previous period is available.

This stock can be used in one of three different ways: satisfy demand at retailer *i*, satisfy demand at retailer *j* (i.e., transshipment from retailer *i* to *j*), and hold in inventory at retailer *i*. At the end of the period, the material will be used in two ways: to satisfy backorder or to build up inventory at a retailer. Note that the stock at the beginning of the horizon, and the replenishment made during the first *T-L* periods 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* at period *t*, $d_i(t)$, can be satisfied in one of two different ways: from the inventory at retailer *i*, or from the inventory at another retailer *j* (i.e., through a transshipment from retailer *j* to retailer *i*). Another sink for material is the requirement that each retailer *i* begin the next period with inventory position equal to S_i. These units can come from one of two sources: the inventory at retailer *i* or replenishment arrival during the period.

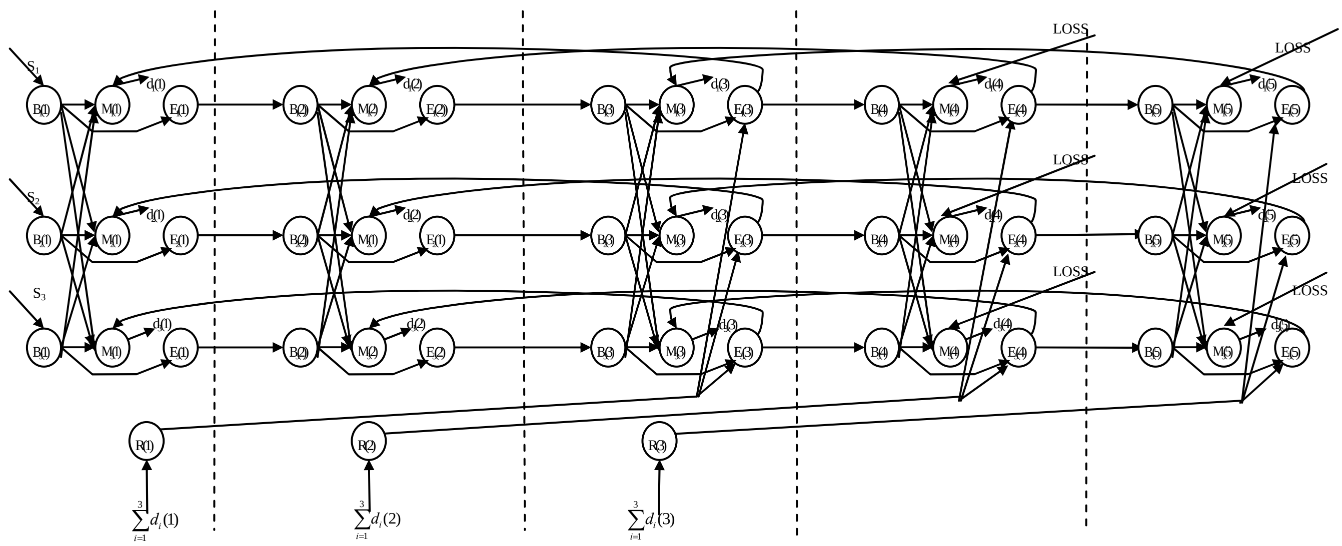

Figure 2: Network flow representation for a 5-period horizon and 2-period lead time system

Using the observations above, we model the movement of stock during the planning horizon as a network flow problem. Figure 2 presents a network with a 5-period horizon and 2-period replenishment lead time. In each period *t* we have a source node, $B_i(t)$, to represent the **b**eginning, i.e., initial inventory at retailer *i* and period *t*. The middle sink node associated with the de**m**and at retailer *i* in period *t* will be denoted by $M_i(t)$. Similarly, we will denote by $E_i(t)$ the **e**nding inventory at retailer *i* in period *t*. Note that this is equal to the inventory at the beginning of the next period. Finally, we have a node $R(t)$ to represent the **r**eplenishment requested in period *t* to be delivered in period *t*+*L*.

The arcs in the network flow problem are exactly those activities described above and are summarized (with the associated cost per unit flow) in Table 1. We use such variables as $B_i(t)M_j(t)$ to denote the flow in the network, indicating the starting and ending nodes. For example, $B_i(t)M_j(t)$ is the flow in the network from node $B_i(t)$ to $M_j(t)$.

Table 1: Definition of the arcs in the network flow problem

Arc	Variable	Unit cost	Meaning
$(B_i(t), E_i(t))$	$B_i(t)E_i(t)$	h_i	Inventory held at retailer *i* at period *t*.
$(B_i(t), M_i(t))$	$B_i(t)M_i(t)$	0	Stock at retailer *i* used to satisfy demand at retailer *i* in period *t*.
$(B_i(t), M_j(t))$	$B_i(t)M_j(t)$	$c_{ij}(c_{ii}=0)$	Transshipment from retailer *i* to retailer *j* in period *t*.
$(E_i(t+L), M_i(t))$	$E_i(t+L)M_i(t)$	p_i	In the first *H-L* periods, shortages backlogged at retailer *i*.
		l_i	In the last *L* periods, lost demand at retailer *i*.
$(E_i(t), B_i(t+1))$	$E_i(t)B_i(t+1)$	0	Inventory on hand at the end of period *t* carried to the next period *t*+1.
$(R(t), E_i(t+L))$	$R(t)E_i(t+L)$	0	Inventory at retailer *i* increased through replenishment at period *t*+*L*.

Replenishment order quantities can be computed as indicated in Lemma 1. Based on Lemma 1, Proposition 2 reformulates the flow balance equations at nodes $R(t)$, which significantly simplifies our network flow representation.

Lemma 1: *For a base stock policy, replenishment orders at location i in period t can be computed by the formula below:*

$$R_i(t)E_i(t+L) = [S_i - (E_i(t)B_i(t+1) + \sum_{m=t-L}^{t-1} R_i(m)E_i(m+L) - E_i(t+L)M_i(t))]^+, when \quad t > L$$

$$= [S_i - (E_i(t)B_i(t+1) + \sum_{m=1}^{t-1} R_i(m)E_i(m+L) - E_i(t+L)M_i(t))]^+, when \quad 1 \le t \le L$$

Proof The proof is presented in Appendix B.

Proposition 2: *In a base stock policy following the above assumptions, the sum of the replenishment orders at all the locations is equal to the sum of the demands in the period. This relationship can be expressed as the formula below:*

$$\sum_{i=1}^{N} R_i(t)E_i(t+L) = \sum_{i=1}^{N} d_i(t)$$

Proof The proof is presented in Appendix B.

We can observe that the system states in the first *L* periods, the last *L* periods, and the middle [*L*+1*,T*-*L*] periods are different. We present the characteristics of four different stages below.

i) The first period: t=1. There are no replenishment order arrivals. For the period 1, inventory on hand at the beginning of period $IOH_i(1)$ is just equal to the base stocks S_i.

ii) The first L-1 periods: t=2,...,L. There are no replenishment order arrivals. Inventory on hand at the beginning of period $IOH_i(t)$ is the inventory from the previous period $E_i(t-1)B_i(t)$.

iii) The middle periods: t=L+1,...,T-L. This is the typical period; inventory on hand at the beginning of period $IOH_i(t)$ is just the inventory from the last period $E_i(t-1)B_i(t)$. Unmet demand is backlogged, and replenishment ordered at the period t will arrive at the period t+L.

iv) The last L periods: t=T-L+1,...,T. Different from the middle period, the unmet demand is lost because the replenishment orders $R(t)E_i(t+L)$ in the last L periods cannot arrive in time within the finite horizon.

2.3.2 SP and LP Representations

We are now ready to introduce a stochastic programming model. When demand is generated, we give its determinant counterpart, i.e., a linear programming model. The reason for building two models here is that we will use a stochastic counterpart algorithm to compute the optimal base stock value. In this algorithm, we need to know the determinant counterpart of stochastic programming model.

Since demand is stochastic, our problem is built as a stochastic programming model. We formulate this stochastic programming model in problem (S). The objective is to minimize the expected average cost per period in the system.

Problem (S)

$$\min \quad E[AC(S,D)] = \frac{1}{T} E[\sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{j=1, j\neq i}^{N} c_{ij} B_i(t) M_j(t) + \sum_{t=1}^{T} \sum_{i=1}^{N} h_i I_i^+(t) + \sum_{t=1}^{T-L} \sum_{i=1}^{N} p_i I_i^-(t) + \sum_{t=T-L+1}^{T} \sum_{i=1}^{N} l_i I_i^-(t))]$$

Subject to

$$I_i(t) = S_i(t) - \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + \sum_{j=1, j\neq i}^{N} B_j(t)M_i(t) - D_i(t)$$, when t=1

$$I_i(t) = IOH_i(t) - \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + \sum_{j=1, j\neq i}^{N} B_j(t)M_i(t) - D_i(t)$$, when t=2, ... T

$$B_i(t)M_j(t) \geq 0 \quad i, j = 1,..., N, \quad t = 1,..., T$$

$$I_i^+(t) \geq 0, I_i^-(t) \geq 0, i = 1,..., N, \quad t = 1,..., T$$

Based on the stochastic programming problem (S), and the network flow model presented in Figure 3, we construct an LP formulation (D). When demand is realized, problem (D) is a determinant counterpart of problem (S). Through this LP formulation, we compute the derivative using duality –hence, avoiding cumbersome derivative recursions and decision tree methods found in current literature, simplifying the computation, and improving algorithm efficiency. In addition, highly efficient LP packages exist to solve large-scale LP problems to support our sample-path-based algorithm.

Recall that the system states in the first L periods, the last L periods, and in the middle [L+1, T-L] periods are different. This is reflected in the formulation below:

Problem (D)

$$\min AC(S,d) = \frac{1}{T}\Big(\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1, j\neq i}^{N} c_{ij}B_i(t)M_j(t) + \sum_{t=1}^{T}\sum_{i=1}^{N} h_i B_i(t)E_i(t) + \sum_{t=1}^{T-L}\sum_{i=1}^{N} p_i E_i(t+L)M_i(t) + \sum_{t=T-L+1}^{T}\sum_{i=1}^{N} l_i E_i(t+L)M_i(t)\Big)$$

Subject to

$$B_i(t)M_i(t) + \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + B_i(t)E_i(t) = S_i \quad i = 1,..., N, \quad t = 1$$

$$B_i(t)M_i(t) + \sum_{j=1, j\neq i}^{N} B_i(t)M_j(t) + B_i(t)E_i(t) - E_i(t-1)B_i(t) = 0 \quad i = 1,..., N, \quad t = 2,..., T \qquad \text{......(A-1)}$$

$$B_i(t)M_i(t) + \sum_{j=1, j\neq i}^{N} B_j(t)M_i(t) + E_i(t+L)M_i(t) = d_i(t) \quad i = 1,..., N, t = 1,..., T \qquad \text{......(A-2)}$$

$$\sum_{i=1}^{N} R(t)E_i(t+L) = \sum_{i=1}^{N} d_i(t) \quad i = 1,..., N, \quad t = 1,..., T-L \qquad \text{......(A-3)}$$

$$B_i(t)E_i(t) - E_i(t)B_i(t+1) = 0 \quad i = 1,..., N, t = 1,..., L$$

$$B_i(t)E_i(t) + R(t-L)E_i(t) - E_i(t)M_i(t-L) - E_i(t)B_i(t+1) = 0 \quad i = 1,..., N, \quad t = L+1,..., T \qquad \text{......(A-4)}$$

$$B_i(t)E_i(t), B_i(t)M_i(t), B_i(t)M_j(t), E_i(t)M_i(t), R(t)E_i(t) \geq 0 \quad i, j = 1,..., N, \quad t = 1,..., T$$

Equations (A-1), (A-2), (A-3), and (A-4) represent the physical inventory balance constraints at the B_i (t), M_i (t), R (t), and E_i (t) nodes, respectively. There are $(N^2+4N)T$ decision variables, $(N^2+4N)T$ components in cost vector **c**, $(3N+1)T$ components in right hand column **b**, and the parameter matrix is a $(N^2+4N)T\times(3N+1)T$ matrix.

This LP formulation will be at the heart of our algorithm, so its feasibility is a necessary condition for successful implementation. If all cost parameters, demand, and base stock levels are finite, then problem (D) is feasible and has a finite optimum. This is established by Proposition 3.

Proposition 3: *Let the location index be* $\mathcal{I} =\{1,2,...,N\}$. *If demand* $D_i\ \forall i\in \mathcal{I}$ *has a density on* $(0,\infty)$ *and* $E[D_i]<\infty\ \ \forall i\in \mathcal{I}$, *unit cost* $h_i,c_{ij},p_i,l_i\in\Re^+$ *and* $h_i,c_{ij},p_i,l_i<\infty$ $\forall i,j\in \mathcal{I}$, *base stock* $S_i\in\Re^+$ *and* $S_i<\infty\ \forall i\in \mathcal{I}$, *then problem (D) is feasible, and has a finite optimum with probability 1.*

Proof The proof is presented in Appendix C.

It should be pointed out that our formulation can be easily generalized to solve variants of our current problem, including most models such as two-location transshipment, two-location transshipment with positive lead times, multiple location transshipment with negligible lead times, and no-transshipment problem. Furthermore, our formulation can also be generalized to solve problems with different system configurations and pooling policies.

3 Algorithms and Implementation

3.1 Algorithms

To compute the optimal base stock values, we adopt a sample-path-based optimization algorithm, where we use IPA to compute the gradient value. In particular, we start with an arbitrary base stock level, S_i, for each stocking location. After randomly generating an instance of the demand for each location, we construct and solve problem (D) in a deterministic fashion. Then, we can compute the gradient values by invoking duality. In other words, the LP is used not only to compute the optimal transshipment quantities, but also to help accumulate IPA gradients ($\partial AC / \partial S_i$). The latter are used in the path search algorithm to determine the optimal base stock levels.

The procedure is summarized in a pseudo-code format in Figure 3, where K denotes the total number of steps taken in a path search, U represents the total number of inner cycles, a_k represents the step size at the each iteration k, and S_i^k represents the base stock level for retailer i at the k^{th} step.

3.2 Explanation and Justification of the algorithm

I) Initialization

The algorithm starts with an arbitrary value for the base stock levels, S_i^0. K and U should also be specified by the experimenter and can be determined, for instance, by a pilot study to mitigate the following trade-off: with a small K, the experiment cannot provide sufficient data, and output will have a big variance. A K that is too big is inefficient in improving the optimal value.

II) Outer loop

The outer loop includes the inner loop computations, the desired gradient calculation, and the updating of order-up-to-levels.

(II.1)*A*. The demand is generated at each retailer. Note that any covariance structure is allowed in $f(D)$.

(II.1)*B*. Once the demand is observed, problem (D) is solved in a deterministic fashion to compute the optimal transshipment quantities and the minimum-cost flow.

(II.1)*C*. The gradient of the average cost (derivatives with respect to the base stock levels) is computed. Our LP formulation greatly simplifies these computations. The implementation of the derivative computation in this step is very efficient, as established in Theorem 1.

(I)Initialization

(I.1) Initialize K

(I.2) Initialize U

(I.3) For each retailer i, set initial base stock levels, S_i^0

(II) Repeat

Set k←1

(II.1) *Repeat*

Set u←1

(II.1)A. Generate the demand at each location from f(D)

(II.1)B. Solve problem (D) to determine optimal transshipment quantities

(II.1)C. Compute/Accumulate the desired gradients of the average cost, dAC_u

u←u+1, until u=U

(II.2) Calculate the desired gradient(s), $\frac{1}{U}\sum_{u=1}^{U} dAC_u$

(II.3) Update the order-up-to-levels, $S_i^k \leftarrow S_i^{k-1} - \alpha_k \frac{1}{U}\sum_{u=1}^{U} dAC_u$

k←k+1, until k=K

(III)Return the S_i *and objective function value.*

FIGURE 3: Description of the sample-path-based optimization procedure

Theorem 1: *Based on the special LP structure in our problem and infinitesimal perturbation assumption of base stock, the gradient of average cost with respect to base stock* $\partial AC/\partial S_i$ *is just the corresponding dual optimal solution* $p^*_{w,}$ *, where w is determined by the position of* S_i *in the LP formulation. For an N-location problem, w=N+i.*

Proof The proof is presented in Appendix D.

The fact that, for a linear program, the dual value of a constraint is the derivative of the objective function with respect to the right-hand side of that constraint was first used by Swaminathan and Tayur (1999). From Proposition 3 and Theorem 1, we have the Corollary 1. This corollary will subsequently support Proposition 4.

Corollary 1 *If demand* $D_i \ \forall i \in I$ *has a density on* $(0,\infty)$ *and* $E[D_i] < \infty \ \forall i \in I$*, unit cost* $h_i, c_{ij}, p_i, l_i \in \Re^+$ *and* $h_i, c_{ij}, p_i, l_i < \infty \ \forall i \in I$*, base stock* $S_i \in \Re^+$ *and* $S_i < \infty \ \forall i, j \in I$*, then the gradient of average cost with respect to base stock* $\partial AC/\partial S_i$ *exists and is bounded.*

Proof The proof is presented in Appendix D.

(II.2) We estimate the desired gradient(s) by the formula $\frac{1}{U}\sum_{u=1}^{U} dAC_u$, which is just the IPA technique. With IPA, we need to establish the unbiasedness of the gradient estimator. Recall that the implicit assumption of IPA is that the average of the changes represents the change in expectations, which yields an unbiased estimator. This assumption is true only under a commuting condition (Glasserman 1991). For our setting, in order to prove Proposition 4 below, we need to first prove several basic properties of the average cost function. Lemmas 2 and 3 provide one of the basic conditions directly required by Proposition 4.

Lemma 2 $AC(S_i)$ *is a convex function.*

Proof The proof is presented in Appendix E.

Lemma 3 *If demand* $D_i \ \forall i \in I$ *has a density on* $(0,\infty)$ *and* $E[D_i] < \infty \ \forall i \in I$*, unit cost* $h_i, c_{ij}, p_i, l_i \in \Re^+$ *and* $h_i, c_{ij}, p_i, l_i < \infty \ \forall i \in I$*, base stock* $S_i \in \Re^+$ *and* $S_i < \infty \ \forall i, j \in I$*,* then $AC(S_i)$ *is a proper convex function.*

Proof The proof is presented in Appendix E.

In order to show that $AC(S_i)$ is smooth, we need to show that $AC(S_i)$ is continuously differentiable everywhere w.p.1. If $AC(S_i)$ is a proper convex function and CDF of demand F(D) is continuous, from Theorem 25.2 and Corollary 25.5.2 of Rockafellar (1970), we have Lemma 4.

Lemma 4 *If CDF of demand F(D) is continuous,* $AC(S_i)$ *is continuously differentiable every where w.p.1.*

Proof The proof is presented in Appendix E.

As shown by Glasserman (1991), provided that the objective function $AC(S_i)$ is convex and smooth with respect to the base stock levels, IPA estimators will be unbiased. We can now establish Proposition 4.

Proposition 4: *If demand* $D_i\ \forall i \in I$ *has a density on* $(0,\infty)$ *and* $E[D_i] < \infty\ \forall i \in I$*, then the gradient estimator* $\frac{1}{U}\sum_{u=1}^{U} dAC_u$ *is unbiased in the transshipment system with positive replenishment lead times. That is, we can interchange the integral and the derivative as the equation* $E[\nabla_S AC(S)] = \nabla_S E[AC(S)]$

Proof The proof is presented in Appendix E.

Here the term on the left-hand side $E[\nabla_S AC(S)]$ is what we obtain by averaging IID copies of the stochastic gradient and the term on the right-hand side $\nabla_S E[AC(S)]$ is what we want.

(II.3) The base stock level S_i is updated through $S_i^k \leftarrow S_i^{k-1} - \alpha_k \frac{1}{U}\sum_{u=1}^{U} dAC_u$. Also note that since the algorithm stops at $k=K$, we do not need an extra stopping rule. A key issue in this step is the selection of a suitable step size a_k, for which we have Condition 1 below:

Condition 1: *A criterion for choosing* a_k *is to let step size go to zero fast enough so that the algorithm actually converges to a value of S, but not so fast that it will induce a wrong value. One condition to meet that criterion is* $\sum_{k=1}^{\infty} a_k = \infty$ *and* $\sum_{k=1}^{\infty} a_k^2 < \infty$.

For instance, $a_k = a/k$ for some fixed $a>0$ satisfies Condition 1. The first part of this condition facilitates convergence by ensuring that the steps do not become too small too quickly. However, if the algorithm is to converge, the step sizes must eventually become small, as ensured by the second part of the condition. When the gradient estimator is unbiased (as is the case here), this step yields a Robbins-Monro algorithm (1951) for stochastic search. Although, theoretically, we can use any step size satisfying Condition 1, the practical implementation is more complicated. Section 3.3.2 addresses this problem to identify a suitable step size.

(III) Return the S_i and objective function value at each step. Then we can conduct the output analysis which will be investigated in section 3.3.

3.3 Implementation

We implement the algorithm in Matlab, as depicted in Figure 4, together with an LP solver. We have written the main program to implement the algorithm shown in Figure 3, two subroutines "tran_initial.m" and "trans_LP.m" to specify the LP problem, and input LP characteristic information to the LP solver. The LP solver returns the optimal base stock and optimal average cost at each step to the main program. Some subroutines like "drawbase.m", "drawcost.m", "hwmean.m" conduct the output analysis. In the implementation, we also need to note the three problems: setting the initial value for base stock levels, choosing the step size, and handling the initial transient.

3.3.1 Initial value

In order to explore the impact of initial values, we experimented with a 4-period lead time setting. With the same cost and simulation parameters, we tested different initial values.

The results were consistent in the sense that initial values do not affect the final result, which shows the robustness of our algorithm.

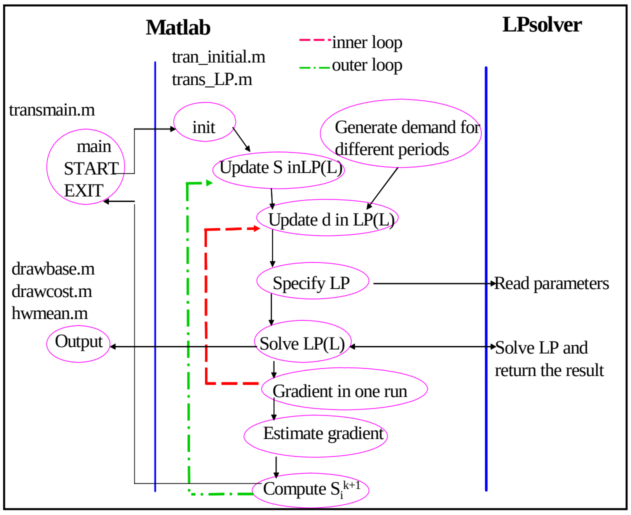

FIGURE 4: Implementation framework

3.3.2 *Step size*

Although all step sizes, which satisfy Condition 1, will theoretically lead to a correct result, different step sizes do affect the convergence rate of the implemented algorithm. In order to explore the impact of different step sizes, we considered a case with 12-period lead time. With the same unit cost and system parameters, we ran the simulation with K=2000 and an arbitrary initial value. Then we changed step size α_k from 0.003/k,0.03/k,0.3/k,1.5/k,3/k,6/k, 15/k, to 1000/k. The first 500 steps are presented in Figure 5. Based on this experiment, we have adopted α_k=3/k as our step size in subsequent analysis.

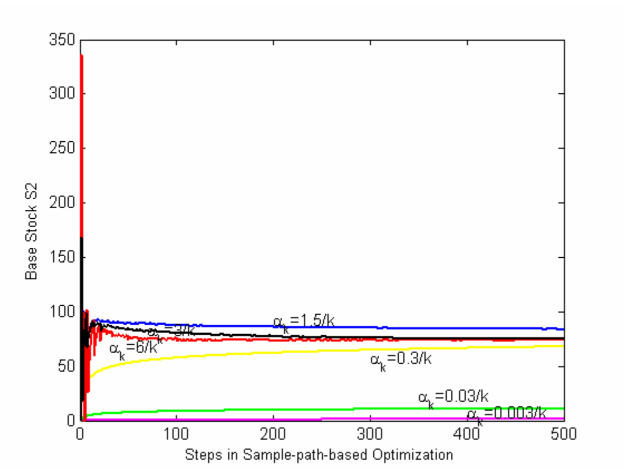

Figure 5 Search paths with different step sizes

3.3.3 Initial Transient Deletion

We use one case (Figure 6) to illustrate the output analysis. For a setting with 6 locations and a 2-period lead time along with the same simulation and unit cost parameters used in section 4.1.2, we have a main output: base stock in each step, from which we statistically compute the estimators for base stocks.

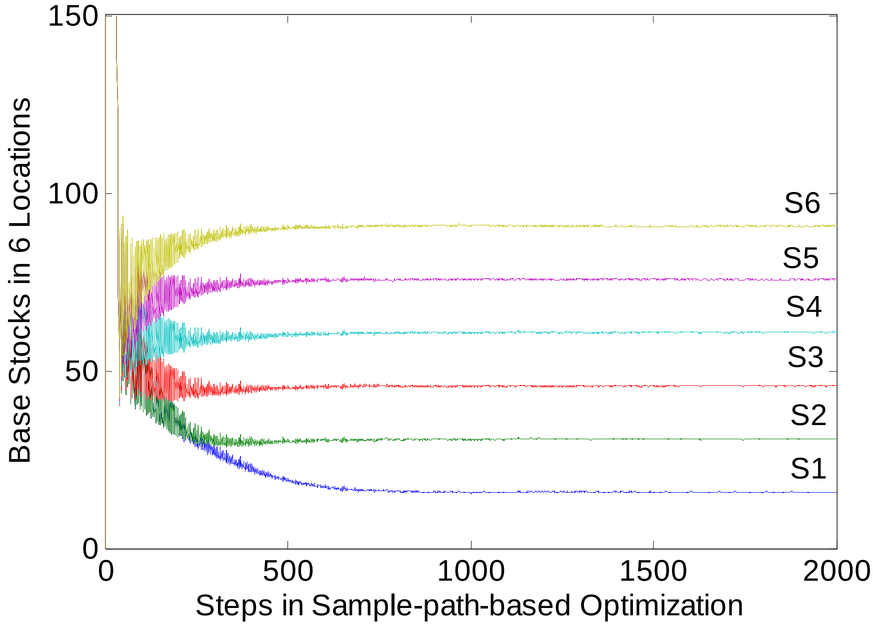

FIGURE 6 Initial transient in the output

Since the variance is too large during the transient stage, we use the transient deletion technique to eliminate the bias in the estimator. For all the experiments we conducted, we observed convergence within the first 1000 steps. We therefore delete the first 1000 data points and only use the remaining data in our algorithm.

4 Results

We present our solution and main experiments in section 4.1. In section 4.2, we conduct comparison studies. In order to validate our algorithm, we compare our results with the results from an existing algorithm, and the results from a system without transshipments. In section 4.3, we explore the impact of correlated demand on transshipment with positive replenishment lead times.

4.1. Experimental Setting

4.1.1 Varying the lead times

We consider a three-location problem with $D_1 \sim Norm(10,1)^+$, $D_2 \sim Norm(5,0.5)^+$, and $D_3 \sim Norm(5,0.5)^+$ over a horizon of T=50, and with unit holding cost h=2, unit penalty cost p=12, unit transshipment cost c_{ij}=5, and unit cost for lost demand in the last L periods l=20. We implement the simulation experiments with different lead times varying from 2 periods to 14 periods. After comprehensive pre-testing, we set the algorithm parameters to K=2000, α_k=3/k, and initial base stock values (0,0,0).

Experiments are conducted by a computer with Pentium 4 CPU 2.6GHz and 516MB of RAM. For each simulation run (the complete optimization process for each lead time case), the elapsed time ranges from 9 minutes to 12 minutes of wall clock time.

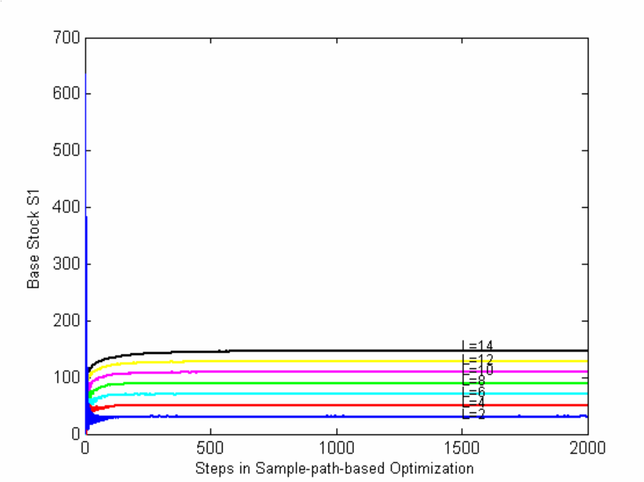

Figure 7.1 Search paths of base stock

We present the search paths for the base stock and for the average cost in Figure 7: Figure 7.1 illustrates the search for the optimal base stock levels. Figure 7.2 illustrates the paths for average cost. From the figures, we observe that, at the beginning, the step size is big, but all experiments rapidly converge before k=500. In order to get a more reliable estimator, we only use the last 1000 values to estimate the optimal base stock and the optimal average cost. We present our optimal base stock, optimal average cost in Table 2. Each estimator includes the mean and the half width for a 95% confidence interval.

Table 2 Optimal base stock and average cost for each lead time

	Location 1 Norm (10,1)$^+$	Location 2 Norm (5,0.5)$^+$	Location 3 Norm (5,0.5)$^+$	Average cost	Average IOH
$S^{L=2}$:	30.6526(±0.0041)	15.3704 (±0.0037)	15.3822(±0.0040)	8.2573(±0.0359)	22.1264(±0.0189)
$S^{L=4}$:	50.4545 (±0.0051)	25.2040 (±0.0049)	25.2452(±0.0048)	14.6677(±0.0446)	24.1594(±0.0298)
$S^{L=6}$:	70.0971 (±0.0049)	35.0081 (±0.0045)	35.0046(±0.0046)	23.8139(±0.0538)	27.6196(±0.0387)
$S^{L=8}$:	89.6009 (±0.0055)	44.7821 (±0.0055)	44.7901(±0.0054)	35.9557(±0.0598)	32.5638(±0.0432)
$S^{L=10}$:	108.9562 (±0.0060)	54.5304 (±0.0060)	54.5610(±0.0061)	50.9456(±0.0633)	39.0392(±0.0506)
$S^{L=12}$:	128.0369 (±0.0069)	64.2505 (±0.0066)	64.2461(±0.0068)	69.0652(±0.0703)	46.9169(±0.0497)
$S^{L=14}$:	146.8914 (±0.0141)	74.0360 (±0.0058)	74.0231(±0.0066)	90.1047(±0.0742)	56.2553(±0.0539)

HW: Half Width CI: Confidence Interval

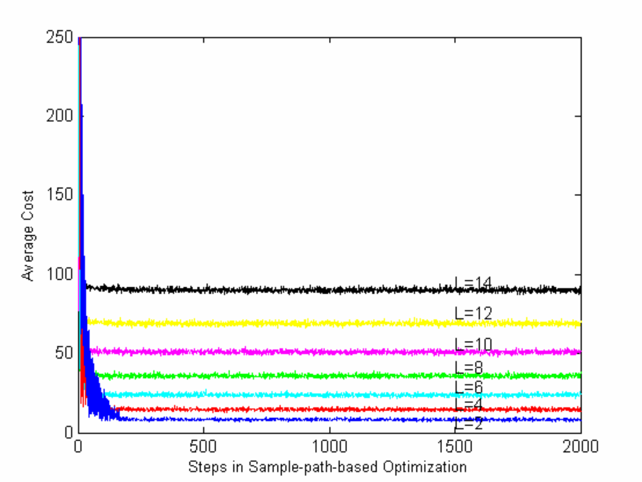

Figure 7.2 Paths of average cost

4.1.2 Varying the number of locations

For a setting with a 2-period lead time and 50-period horizon, and the same cost parameters as in Section 4.1.1, we experiment with the number of locations varying from

2 to 12. We present the search paths for the base stock with a 12-location scenario in Figure 8, which illustrates the search process for the optimal base stock levels in those 12 locations.

We present the optimal base stock and average cost estimates for scenarios with different number of locations in Table 3. Each estimator includes the mean and the half width for a 95% confidence interval. The first row also shows the demand distribution in those 12 locations.

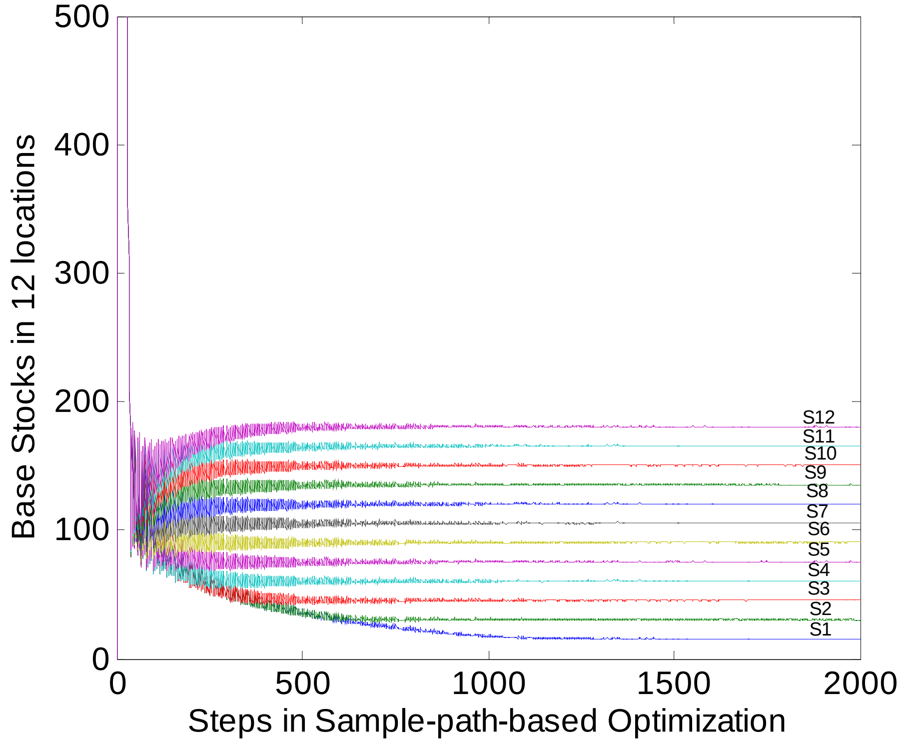

Figure 8 Search paths of base stock in a 12-location scenario

Table 3 The optimal base stock and average cost for the scenarios with different locations

Location →	No.1	No.2	No.3	No.4	No.5	No.6	No.7	No.8	No.9	No.10	No.11	No.12	Ave cost
Scenario ↓	N (5,1)	N (10,1)	N (15,1)	N (20,1)	N (25,1)	N (30,1)	N (35,1)	N (40,1)	N (45,1)	N (50,1)	N(55,1)	N (60,1)	
N=2	16.2371	31.1909											9.6973
	±0.0056	±0.0056											±0.0525
N=3	16.0321	30.9516	45.9934										13.3020
	±0.0054	±0.0053	±0.0054										±0.0610
N=4	15.8764	30.8455	45.8418	60.8674									17.2347
	±0.0078	±0.0079	±0.0084	±0.0088									±0.0723
N=5	15.8045	30.7582	45.7621	60.7869	75.7599								21.6681
	±0.0082	±0.0083	±0.0080	±0.0079	±0.0085								±0.0811
N=6	15.7853	30.7162	45.6681	60.7256	75.6719	90.6910							26.5204
	±0.0077	±0.0075	±0.0081	±0.0075	±0.0077	±0.0086							±0.0876
N=7	15.6355	30.6917	45.6773	60.6268	75.7098	90.6132	105.6723						31.9601
	±0.0117	±0.0109	±0.0117	±0.0114	±0.0118	±0.0116	±0.0114						±0.0999
N=8	15.6723	30.6057	45.5512	60.5967	75.6334	90.6355	105.5818	120.6668					37.8780
	±0.0127	±0.0120	±0.0113	±0.0119	±0.0116	±0.0115	±0.0125	±0.0121					±0.1096
N=9	15.5979	30.6593	45.5733	60.6261	75.5330	90.5500	105.5398	120.5954	135.5155				44.3246
	±0.0142	±0.0113	±0.0110	±0.0111	±0.0126	±0.0112	±0.0117	±0.0118	±0.0110				±0.1139
N=10	15.5503	30.5756	45.5659	60.5293	75.6409	90.6401	105.5278	120.5751	135.6223	150.5850			52.2959
	±0.0198	±0.0204	±0.0201	±0.0195	±0.0200	±0.0202	±0.0208	±0.0195	±0.0199	±0.0201			±0.2097
N=11	15.6799	30.5843	45.5969	60.5603	75.5368	90.5650	105.5320	120.5652	135.6044	150.4535	165.5541		60.1490
	±0.0244	±0.0211	±0.0212	±0.0220	±0.0215	±0.0220	±0.0219	±0.0208	±0.0212	±0.0220	±0.0214		±0.2527
N=12	15.6812	30.5948	45.5296	60.5509	75.4698	90.5736	105.5786	120.4727	135.5397	150.5619	165.5468	180.5179	68.5667
	±0.0383	±0.0213	±0.0213	±0.0217	±0.0218	±0.0220	±0.0217	±0.0213	±0.0224	±0.0213	±0.0217	±0.0213	±0.2723

Section 4.1 has shown that the algorithm works well with different lead times and different number of locations. Although not reported here, we have experimented with different demand distributions, different horizons, and different cost structure. In all these cases, the algorithm does converge fast, and provides a reliable estimate with a small variance.

4.2 Comparison studies

In section 4.2.1, we compare our optimal base stock value with that from the algorithm of Tagaras and Cohen (1992). In section 4.2.2, we compare the performance of two algorithms through their objective function value. To ensure "fairness" in comparison, we use common random numbers across the two algorithms, and finally show that the performance of our algorithm is better. In section 4.2.3, we also compare our result with that from a system without transshipment.

4.2.1 Comparison of optimal base stock values

We compare the optimal base stock values from our algorithm with those from Tagaras and Cohen (1992), which uses $S_I = \mu_i (L_i + 1) + k_i \sigma_i \sqrt{L_i + 1}$ as base stock values, where $\mu_i = E(D_i)$ is the mean and σ_i^2 is the variance of the single-period demand at location i, k_i is computed by $k_i = (S_i^0 - \mu_i)/\sigma_i$ and S_i^0 is the optimal value for the zero-lead-time problem. Since the Tagaras and Cohen algorithm is conceived for a two-location setting, we also apply our algorithm to the two-location setting.

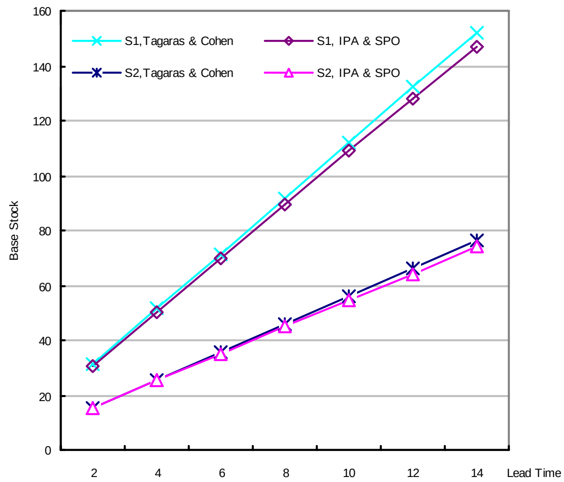

Figure 9 Comparison of base stock values

From Figure 9, we observe that: i) the Tagaras and Cohen (1992) heuristic tends to overestimate the optimal base stock values. All of their base stock quantities are consistently higher than ours. ii) With increasing lead times, the overestimation by the Tagaras and Cohen algorithm becomes even more significant.

4.2.2 Comparison of optimal objective function value: average cost

A lower base stock level does not always imply lower cost. We therefore need to compare the objective function values from both algorithms. To ensure fairness, we use common random numbers across the two algorithms and conduct the experiments as follows: i) For the Tagaras and Cohen algorithm, we input optimal base stock values given by the heuristic algorithm and run 1000 independent replications, and then estimate the average cost; ii) Similarly, for SPO and IPA algorithm, we also input the optimal base stock value given by SPO and IPA algorithm and run 1000 independent replications, and then compute the average cost.

We present the results in Table 4, and conclude that: i) for each lead time, the performance of our algorithm is better, consistently achieving lower average cost; ii) with increasing lead times, the relative performance of our algorithm becomes even more pronounced.

Table 4 Comparison of objective function values of the two algorithms

Lead time	Tagaras and Cohen algorithm	SPO and IPA algorithm
2	8.6007 (±0.0553)	8.2603(±0.0580)
4	15.8146(±0.0911)	14.5940(±0.0778)
6	25.9708 (±0.1179)	23.8311(±0.0896)
8	38.4691(±0.1535)	35.8688(±0.1021)
10	54.7811(±0.1846)	50.9376(±0.1164)
12	73.8404(± 0.2117)	68.9879(±0.1228)
14	96.2429(±0.2405)	90.0336(±0.1217)

The estimators include mean and HW for a 95% CI.

4.2.3 Comparing a system with transshipment and without transshipment

We also compare our results with the results from a no-transshipment system by a two-stage experiment. i) We first eliminate all transshipment flows in the network flow formulation (Figure 2), and compute the optimal base stock quantities for the system without transshipment. ii) For both systems, we then run 1000 independent replications under common random numbers and estimate the average cost.

From Figure 10, we conclude that: i) for each lead time, the performance of the system with transshipments is always better than that of the system with no transshipments. The average costs of system with no transshipments are always higher; ii) when the lead times are lower, the result from the Tagaras and Cohen algorithm is closer to that from our algorithm. But when lead times are longer, the result from the Tagaras and Cohen algorithm is closer to that from a system with no transshipments, failing to reflect the added value of information pooling.

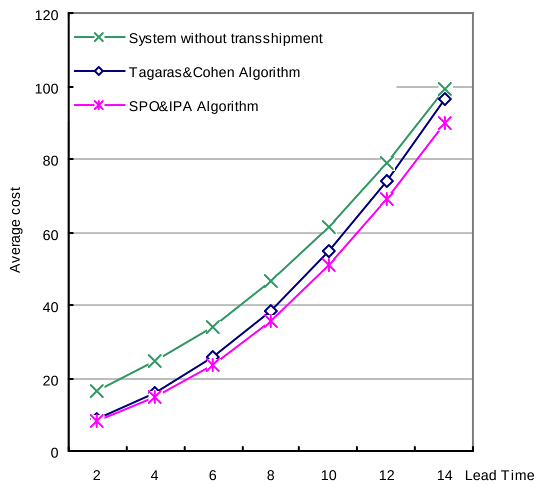

Figure 10 Comparing three settings

4.3 The Effect of Demand Correlation

To study the impact of correlated demand across retailers, we experiment with scenarios of different demand correlations (ρ=±0.5, ±0.25), and different lead times (L=2,5,8,11,14 periods). A case with zero correlation is also added for reference. Unlike the previous section, 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 level of demand correlation being investigated when $i \neq j$ and one when $i = j$.

We adopt the system configuration 5 (Herer et al. 2005), where transshipments are possible between any pair of locations, with unit holding cost h=1, unit penalty cost p=4, and unit transhipment cost c_{ij}=0.5; we set the unit loss cost l=10.

We examine the impact of demand correlation on the average cost per period. Figure 11 depicts the impact of demand correlation on the average total cost for a 3-retailer configuration for different lead times; Table 5 shows the estimated average costs. We observe that, for each lead time, when demand correlation gets smaller (or negative), the average cost always decreases, which implies that the effectiveness of transshipments in matching demand and supply is enhanced. In general, from the observation that smaller correlation significantly lowers average total cost, we can conclude that positive correlation reduces the effectiveness of transshipments while negative correlation enhances it.

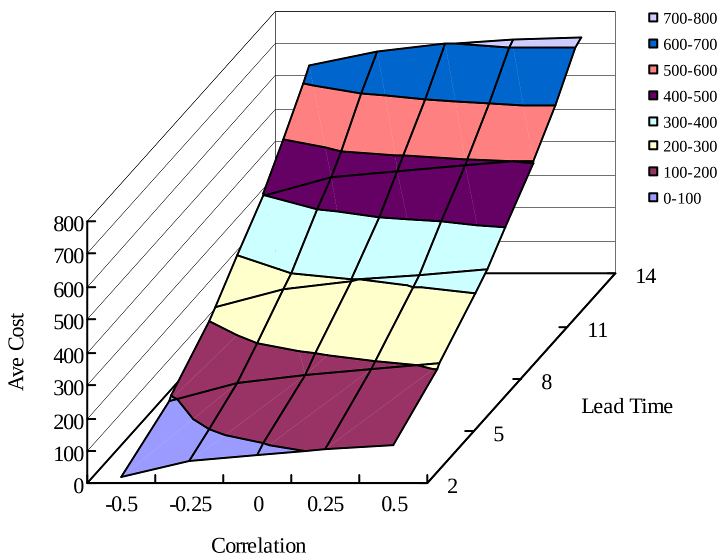

Figure 11 Average cost with different lead times L and correlation coefficients ρ

Table 5 Average cost with different lead times L and correlation coefficients ρ

L \ ρ	-0.5	-0.25	0	0.25	0.5
2	23.9439	68.3064	88.8866	104.3708	117.2655
5	93.7007	146.867	171.6306	188.9192	205.2476
8	217.5912	273.9231	298.1574	315.9017	331.671
11	397.9785	450.5808	472.6436	490.1562	503.9456
14	631.4524	677.1012	697.7432	712.8785	719.9255

We also explore the impact of demand correlation on the base stock levels. The optimal base stock levels S_i with different lead times L and correlation coefficients ρ are presented in Table 6. We conclude that: i) When lead time is low (e.g., L=2), base stock is lower when correlation gets smaller (or negative). This is similar to the 0-lead-time case reported by Herer et al (2005); ii) When lead time is high, base stocks will possibly increase when correlation gets smaller (or negative). Recall that our objective function is not inventory but average cost. A higher base stock may still possibly reduce the average cost. In settings with longer lead times, in order to meet the demand from other locations, when correlation gets smaller (or negative), base stocks may slightly increase to reduce the cost from backorders and lost demand.

Table 6 Base stock S_1 with different lead time L and correlation coefficient ρ

L \ ρ	-0.5	-0.25	0	0.25	0.5
2	300.773	305.8798	307.806	308.7214	310.0512
5	599.7815	596.0419	595.5098	594.3899	594.3217
8	899.7037	885.132	879.3808	873.9785	869.6306
11	1198.9	1171	1158.1	1147.4	1140.2
14	1498.3	1452.4	1433.8	1419.6	1405.6

5 Summary

In this paper, we consider the multi-location transshipment problem with positive replenishment lead times. The main contributions of this paper include the following: 1) using simulation optimization combined with an LP/network flow formulation and IPA, we provide a flexible and efficient algorithm to compute the optimal base stock quantities for the multi-location transshipment problem with positive replenishment lead times; 2) experimenting with scenarios of high and low levels of demand correlation along with different lead times, we show the negative relationship between the benefits of transshipments and demand correlation at different lead time settings; 3) our algorithm is also shown to be able to provide better objective function value than an existing algorithm; 4) we introduce an elegant duality-based gradient computation method to significantly improve computational efficiency.

References

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

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. 2002. Optimization for Simulation: Theory vs Practice. *INFORMS Journal on Computing*,14,3,192-215

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. 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.

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. and M. Tzur. 2003. Optimal and Heuristic Algorithms for the Multi-Location Dynamic Transshipment Problem with Fixed Transshipment Costs. *IIE Transactions,* 5(5), 419-432.

Herer, Y.T., Tzur, M., and Yücesan, E., 2005. The multi-location transshipment problem, (Forthcoming in *IIE Transactions*).

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.

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

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

Robbins H. and S. Monro. 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.

Rockafellar, R.T. 1970. *Convex Analysis.* Princeton University Press, Princeton, N.J.

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

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

Tijms H.C. 1994. Stochastic Models: An algorithmic approach. John Wiley & Sons Ltd, Chichester, England.

Appendix A: Regenerative Processes

Proof of Proposition 1 For a regenerative process (Definition 1.3.1, Tijms 1994), there exists a regenerative epoch t_1 with $t_1 \in \Xi$ such that

(a) $\{IOH(t+t_1), t \in \Xi\}$ is independent of $\{IOH(t), 0 \le t < t_1\}$

(b) $\{IOH(t+t_1), t \in \Xi\}$ has the same distribution as $\{IOH(t), t \in \Xi\}$

Now we argue that condition (a) can not always be satisfied in the presence of positive lead times. Let $\sum_{m=t-L}^{t-1} R_i(m)E_i(m+L)$ be the inventory on order, $E_i(t+L)M_i(t)$ be the backorder, and $E_i(t)B_i(t+1)$ be the inventory on hand. Under a base stock policy, we can express the replenishment order quantity $R_i(t)E_i(t+L)$ as below. From our formulation, $E_i(t)B_i(t+1)$ is just $IOH_i(t)$, then we have

$$IOH_i(t) = S_i - \sum_{m=t-L}^{t-1} R_i(m)E_i(m+L) + E_i(t+L)M_i(t) - R_i(t)E_i(t+L), when \quad t > L$$

$$= S_i - E_i(t)B_i(t+1) - \sum_{m=1}^{t-1} R_i(m)E_i(m+L) + E_i(t+L)M_i(t) - R_i(t)E_i(t+L), when \quad 1 \le t \le L$$

From the above expression, we know that each $IOH_i(t)$ is dependent on $R_i(t-L)E_i(t)$ in the former L periods, which are dependent on $IOH_i(t-L)$ in the former L periods. The dependence of $IOH_i(t)$ here is transitive .Then for any epoch t_1 with $t_1 \in \Xi$, we have $\{IOH(t+t_1), t \in \Xi\}$ is dependent of $\{IOH(t+t_1-L), t \in \Xi\}$; $\{IOH(t+t_1-L), t \in \Xi\}$ is dependent of $\{IOH(t+t_1-2L), t \in \Xi\}$,..., is dependent of $\{IOH(t+t_1-kL), k \in \aleph, \ t \in \Xi\}$.

With the increasing of *k*, when $t-kL \le 0$, we have $t+t_1-kL \le t_1$ w.p.1, which means $\{IOH(t+t_1), t \in \Xi\}$ is dependent of $\{IOH(t), 0 \le t < t_1\}$ w.p.1. Therefore, the condition (a) always can not hold. In the presence of the positive replenishment lead times, $\{IOH(t), t \in \Xi\}$ is not a regenerative process.

Appendix B: Replenishment Order Quantity

Proof of Lemma 1 If the base stock is bigger than the inventory position, then this formula can be reduced to the form as below.

$R_i(t)E_i(t+L)=S_i-(E_i(t)B_i(t+1)+\sum_{m=t-L}^{t-1}R_i(m)E_i(m+L)-E_i(t+L)M_i(t))$ .It is easy to implementing this in linear programming, we can set the $R_i(t)E_i(t+L)\geq 0$.

From our assumption on the sequence of events (refer to fig 2), we make the replenishment order decision after we have observed the inventory on hand $E_i(t)B_i(t+1)$, the inventory on order $\sum_{m=t-L}^{t-1}R_i(m)E_i(m+L)$, and the backorder $E_i(t+L)M_i(t)$. So when $t>L$, we have

$$R_i(t)E_i(t+L)=S_i-IP_i(t)=S_i-(IOH_i(t)+IOR_i(t)-BO_i(t))$$
$$=S_i-(E_i(t)B_i(t+1)+\sum_{m=t-L}^{t-1}R_i(m)E_i(m+L)-E_i(t+L)M_i(t))$$

By the same logic we can get the proof when $1\leq t\leq L$, and the only difference is the expression of the inventory on order.

Proof of Proposition 2

(i) When *t*=1, Proposition 2 holds.

From Lemma 1, with a nonnegative $R_i(t)E_i(t+L)$ assumption, we have

$$R_i(t)E_i(t+L)=S_i-(E_i(t)B_i(t+1)-E_i(t+L)M_i(t)) \quad \text{for} \quad i=1,..,N \tag{B-1}$$

We have $E_i(t)B_i(t+1)=B_i(t)E_i(t)$ from the network flow balance at points $E_i(t)$. At points $M_i(t)$, we denote the backorder $E_i(t+L)M_i(t)=d_i(t)-\sum_{j=1}^{N}B_j(t)M_i(t)$.

Substitute the $E_i(t)B_i(t+1)$ and $E_i(t+L)M_i(t)$ into the formula (B-1), we have

$$R_i(t)E_i(t+L)=S_i-B_i(t)E_i(t)-\sum_{j=1}^{N}B_j(t)M_i(t)+d_i(t).$$

Summing the replenishment orders quantities in the different locations, we have

$$\sum_{i=1}^{N}R_i(t)E_i(t+L)=\sum_{i=1}^{N}[S_i-B_i(t)E_i(t)-\sum_{j=1}^{N}B_j(t)M_i(t)]+\sum_{i=1}^{N}d_i(t).$$

Note that $\sum_{i=1}^{N}\sum_{j=1}^{N}B_j(t)M_i(t)=\sum_{i=1}^{N}\sum_{j=1}^{N}B_i(t)M_j(t)$, then we have

$\sum_{i=1}^{N}[S_i-B_i(t)E_i(t)-\sum_{j=1}^{N}B_j(t)M_i(t)]=\sum_{i=1}^{N}[S_i-B_i(t)E_i(t)-\sum_{j=1}^{N}B_i(t)M_j(t)]=0$, which is from the network flow balance at points $B_i(t)$ .So we have $\sum_{i=1}^{N}R_i(t)E_i(t+L)=\sum_{i=1}^{N}d_i(t)$,when *t*=1.

(ii) If *t*=*m*, the Proposition 2 holds, it will also hold when *t*=*m*+1.

If $t=m$, the Proposition 2 holds, then the inventory position at the beginning of the period $t=m+1$ are the base stock levels, $S_i, \quad i=1,...,N$. Then we examine the decreasing of the inventory position after the demands $d_i, \quad i=1,...,N$ have been observed.

(A) Suppose that there are no backorders, the sum of decreasing physical inventory $\sum_{i=1}^{N} \Delta_i^{IOH}$ is equal to the sum of demands: $\sum_{i=1}^{N} \Delta_i^{IOH} = \sum_{i=1}^{N} d_i$.

(B) Suppose that there are backorders, the sum of decreasing physical inventory $\sum_{i=1}^{N} \Delta_i^{IOH}$ and backorder $\sum_{i=1}^{N} \Delta_i^{BO}$ is equal to the sum of demands: $\sum_{i=1}^{N} \Delta_i^{IOH} + \sum_{i=1}^{N} \Delta_i^{BO} = \sum_{i=1}^{N} d_i$.

In any case, the decrease of the inventory position is $\sum_{i=1}^{N} d_i$, and at the end of period, when the replenishment orders quantities are determined, the sum of inventory positions is $\sum_{i=1}^{N} S_i - \sum_{i=1}^{N} d_i$. In order to restore the inventory position to $S_i, \quad i=1,...,N$, one must order enough units so that $\sum_{i=1}^{N} R_i(t)E_i(t+L) = \sum_{i=1}^{N} d_i(t)$.

Appendix C: Problem (D) has a finite optimum.

Proof of Proposition 3

1) Problem (D) is feasible. We can always find a feasible solution. Let $R_i(t)E_i(t+L) = E_i(t+L)M_i(t) = d_i(t)$ for the first T-L periods and $E_i(t+L)M_i(t) = d_i(t)$ for the last L periods.

$S_i = B_i(1)E_i(1) = E_i(1)B_i(2) = B_i(2)E_i(2) =,..., B_i(T)E_i(T) = E_i(T)B_i(T+1), \forall i \in \mathcal{I}$.

All transshipment quantities $B_i(t)M_j(t)$ =0 with $\forall i, j \in \mathcal{I}$. This set of values can always satisfy all constraints in problem (D), so problem (D) is always feasible.

2) The optimum of the problem (D) is finite. Since cost vector $h_i, c_{ij}, p_i, l_i \in \Re^+$ and $h_i, c_{ij}, p_i, l_i < \infty$, and all decision variables are nonnegative, the objective value $c^T x \geq 0$. Since it is a minimum problem, the problem (D) has finite optimum.

Appendix D: Computation of Gradients

Proof of Theorem 1 From the Proposition 3, we have shown that problem (D) always has a finite optimum, and the optimal objective total cost value

$$AC^{*}=\boldsymbol{c}^{T}{}_{B}\boldsymbol{B}^{-1}\boldsymbol{b}^{*}\,. \qquad \text{(D-1)}$$

Here B is a basis matrix, $\boldsymbol{b}^*$ is the right-hand side column associated with the basis. Let $\boldsymbol{b}$ be the right-hand side column. $\boldsymbol{b}^{T}=(\boldsymbol{b1}^{T},\boldsymbol{b2}^{T},\ldots\boldsymbol{b4}^{T})$. From the structure of our problem there are $(3N+1)T$ components in right-hand side column $\boldsymbol{b}$. Then from the structure, for the parameter matrix, we have rank(A)= $(3N+1)T$, which is full rank. Since A is full rank, every component in right-hand side column is also in the $\boldsymbol{b*}$ associated with the basis $\boldsymbol{B}$, that is $\boldsymbol{b}^*=\boldsymbol{b}$. We have $AC^{*}=\boldsymbol{c}^{T}{}_{B}\boldsymbol{B}^{-1}\boldsymbol{b}$. (D-2)

So $\boldsymbol{b1}$ is also in the right hand column. $\boldsymbol{b1}$ is right-hand side column for the first period problem. Its $(N+1)^{th}$, $(N+2)^{th}$,…$(2N)^{th}$ components are respectively the base stock values $S_1, S_2, \ldots, S_N$. So $S_1, S_2, \ldots, S_N$ are in the optimal right hand column $\boldsymbol{b}$ associated with the basis $\boldsymbol{B}$**.** Besides, $S_1, S_2, \ldots, S_N$ only appear at these positions.

We also note that $\boldsymbol{c}^{T}{}_{B}\boldsymbol{B}^{-1}=\boldsymbol{p}$, and $\boldsymbol{p}$ is the dual optimal solution. So we obtain

$$AC^{*}=\boldsymbol{p}^{*T}\boldsymbol{b} = p^{*T}{}_{1}b_{1}+ p^{*T}{}_{2}b_{2}+ p^{*T}{}_{3}b_{3} +p^{*T}{}_{4}b_{4}+\ldots+p^{*T}{}_{(3N+1)T}b_{(3N+1)T} \qquad \text{(D-3)}$$

By checking LP formulation, we have $b_{N+1}=S_1$, $b_{N+2}=S_2,\ldots,b_{N+N}=S_N$.

Giving S_i an infinitesimal perturbation, $\boldsymbol{b}$ is perturbed to $\tilde{b}$, because the difference $b-\tilde{b}$ is sufficiently small, $\boldsymbol{B}^{-1}\boldsymbol{b}$ remains positive and we still have a basic feasible solution. The reduced costs $\bar{c}$ are not affected and remain nonnegative. Thus, the optimal basis $\boldsymbol{B}$ will not change, and the formula (D-2) still holds. Then we take the derivatives on both sides of formula (D-2) to obtain

$$\frac{\partial AC}{\partial S_1}=\frac{\partial AC}{\partial b_{N+1}}=p^{*}_{N+1},\frac{\partial AC}{\partial S_2}=\frac{\partial AC}{\partial b_{N+2}}=p^{*}_{N+2},\ldots,\frac{\partial AC}{\partial S_N}=\frac{\partial AC}{\partial b_{2N}}=p^{*}_{2N}\quad .$$

Therefore, the gradient $\partial AC/\partial S_i$ is just the corresponding dual optimal solution p^*_w, where w is determined by the position of S_i in the LP formulation. For N-location problem, $w=N+i$.

Proof of Corollary 1 From the Proposition 3 and Proposition 4, we have known that both problem (D) and its dual has finite optimum, gradient of average cost with respect to base stock $\partial AC/\partial S_i$ exists, and is just the corresponding dual optimal solution p^*_{w}. Note our demand is both continuous and bounded, $\partial AC/\partial S_i$ are differentiable up to period T with probability 1. From the Proposition 3, we also know problem (D) has finite optimum, and so its dual problem also has the finite optimum. So its dual solution p^*_w is bounded, and gradient of average cost with respect to base stock $\partial AC/\partial S_i$ is bounded.

Appendix E: Unbiasedness of IPA gradient estimation

Proof of Lemma 2: Convexity of $AC(S)$

Let $P(b)=\{x \mid Ax=b, x\geq 0\}$ be the feasible set, $S=\{b \mid P(b)\, is \quad nonempty\}$, and for any $b\in S$, we define $AC(b)=\min_{x\in P(b)} c^T x$, which is the optimal cost as a function ***b***. Bertsimas and Tsitsiklis (1997) have shown that the objective functions $AC(b)$ of linear programs are convex functions of their right-hand-sides ***b***. In our LP formulation, all S_i variables appear on the right-hand-side of the linear program. The convexity of the average cost in base stock levels S follows this property.

Proof of Lemma 3 In order to prove that $AC(S)$ is a proper convex function, we need to prove two points below: 1) $AC(S)>-\infty$ for all $S\in\Re^+$. If demand is finite and nonnegative, unit cost h_i, c_{ij}, p_i, l are finite and nonnegative, we have $AC(S)\geq 0$ for all $S_i\in\Re^+$. 2) $AC(S)<+\infty$ for some $S\in\Re^+$. (By contradiction) If it does not hold, problem (D) will have not finite optimum, which contradicts Corollary 1.

Proof of Lemma 4 We use Theorem 25.2 and Corollary 25.5.1 of Rockafellar (1970) to prove this lemma. We need two conditions: 1) $AC(S)$ is a proper convex function, which has already been proven by Lemma 3. 2) Partial derivatives $\partial AC(\cdot)/\partial S_i$ exist and are finite everywhere, which will be proven in the following. Recall that

$$AC(S)=\frac{1}{T}E[\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1,j\neq i}^{N}c_{ij}B_i(t)M_j(t)+\sum_{t=1}^{T}\sum_{i=1}^{N}h_iI_i^+(t,D_i(t))+\sum_{t=1}^{T-L}\sum_{i=1}^{N}p_iI_i^-(t,D_i(t))+\sum_{t=T-L+1}^{T}\sum_{i=1}^{N}l_iI_i^-(t,D_i(t))] \qquad \text{(E-1)}$$

With $I_i(t,D_i(t))=IOH_i(t)-\sum_{j=1,j\neq i}^{N}B_i(t)M_j(t)+\sum_{j=1,j\neq i}^{N}B_j(t)M_i(t)-D_i(t)$ and $IOH_i(1)=S_i$. We have

$$AC(S)=\int_{\Re_+^N}\frac{1}{T}[\sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1,j\neq i}^{N}c_{ij}B_i(t)M_j(t)+\sum_{t=1}^{T}\sum_{i=1}^{N}h_iI_i^+(t,D_i(t))+\sum_{t=1}^{T-L}\sum_{i=1}^{N}p_iI_i^-(t,D_i(t))+\sum_{t=T-L+1}^{T}\sum_{i=1}^{N}l_iI_i^-(t,D_i(t))]dF(D) \qquad \text{(E-2)}$$

Let $H(S,D)$ be term inside of the integral, we have $AC(S)=\int_{\Re_+^N}H(S,D)dF(D)$.

Then, we have its partial derivatives:

$$\frac{\partial AC(S_i)}{\partial S_i}=\lim_{\varepsilon\to 0}\int_{\Re_+^N}[\frac{H(S_i+\varepsilon e^i,D_i)-H(S_i,D_i)}{\varepsilon}]dF(D_i) \qquad \text{(E-3)}$$

Now we will show that the term inside the brackets has bounds in order to prove the integral in (E-3) is absolutely convergent. From Corollary 1, we know that the gradient of the average cost with respect to base stock exists and is equal to a finite p_w^* for any S_i.

We have

$$\left|\frac{H(S_i+\varepsilon e^i-D_i)-H(S_i-D_i)}{\varepsilon}\right|<\infty\,. \qquad \text{(E-4)}$$

Since the term inside the brackets has bounds, the integral in (E-3) is absolutely convergent. We can put the limit inside the integral in (E-3). Also note that $H(\cdot)$ is convex and its partials must exist everywhere except at a countable number of points (Theorem 25.3, Rockafellar 1970). Since $F(\cdot)$ is continuous, these points will have measure zero. Then we have $\frac{\partial AC(S_i)}{\partial S_i}=\int_{\Re_+^N-\Theta}\frac{\partial H(S_i,D_i)}{\partial S_i}dF(D_i)$ (E-5)

Where Θ is the countable points set at which the partials $H(\cdot)$ does not exist, and the measure of this set is zero. Then the partial derivatives exist everywhere outside of this zero-measure set Θ,i.e., partial derivatives exist everywhere w.p.1. For a proper convex function $AC(S)$, if the partial derivatives $\partial AC(\cdot)/\partial S_i$ exist and are finite everywhere w.p.1, then $AC(S)$ will be continuously differentiable everywhere w.p.1 from Theorem 25.2 and Corollary 25.5.1 of Rockafellar (1970).

Proof of Proposition 4 1) Lemma 2 has shown that $AC(S)$ is convex. 2) Lemma 4 has shown that the gradient of the average cost with respect to the base stock, $\partial AC(\cdot)/\partial S$, is continuous w.p. 1. 3) Corollary 1 has shown that $\partial AC(\cdot)/\partial S$ is finite. As shown by Glasserman (1991), IPA estimators will be unbiased provided that the objective function $AC(S)$ is convex and smooth (gradient is both continuous and finite) with respect to the base stock levels.

Publications in the Report Series Research* in Management

ERIM Research Program: “Business Processes, Logistics and Information Systems”

2006

Smart Business Networks Design and Business Genetics
L-F Pau
ERS-2006-002-LIS
http://hdl.handle.net/1765/7319

Designing and Evaluating Sustainable Logistics Networks
J. Quariguasi Frota Neto, J.M. Bloemhof-Ruwaard, J.A.E.E. van Nunen and H.W.G.M. van Heck
ERS-2006-003-LIS
http://hdl.handle.net/1765/7320

Design and Control of Warehouse Order Picking: a literature review
René de Koster, Tho Le-Duc and Kees Jan Roodbergen
ERS-2006-005-LIS
http://hdl.handle.net/1765/7322

A Theoretical Analysis of Cooperative Behavior in Multi-Agent Q-learning
Ludo Waltman and Uzay Kaymak
ERS-2006-006-LIS
http://hdl.handle.net/1765/7323

Supply-Chain Culture Clashes in Europe. Pitfalls in Japanese Service Operations
M.B.M. de Koster and M. Shinohara
ERS-2006-007-LIS
http://hdl.handle.net/1765/7330

From Discrete-Time Models to Continuous-Time, Asynchronous Models of Financial Markets
Katalin Boer, Uzay Kaymak and Jaap Spiering
ERS-2006-009-LIS
http://hdl.handle.net/1765/7546

Mobile Payments in the Netherlands: Adoption Bottlenecks and Opportunities, or… Throw Out Your Wallets
Farhat Shaista Waris, Fatma Maqsoom Mubarik and L-F Pau
ERS-2006-012-LIS
http://hdl.handle.net/1765/7593

Hybrid Meta-Heuristics for Robust Scheduling
M. Surico, U. Kaymak, D. Naso and R. Dekker
ERS-2006-018-LIS
http://hdl.handle.net/1765/7644

VOS: A New Method for Visualizing Similarities between Objects
Nees Jan van Eck and Ludo Waltman
ERS-2006-020-LIS
http://hdl.handle.net/1765/7654

On Noncooperative Games, Minimax Theorems and Equilibrium Problems
J.B.G. Frenk and G. Kassay
ERS-2006-022-LIS
http://hdl.handle.net/1765/7809

An Integrated Approach to Single-Leg Airline Revenue Management: The Role of Robust Optimization
S. Ilker Birbil, J.B.G. Frenk, Joaquim A.S. Gromicho and Shuzhong Zhang
ERS-2006-023-LIS
http://hdl.handle.net/1765/7808

Optimal Storage Rack Design for a 3D Compact AS/RS with Full Turnover-Based Storage
Yu Yugang and M.B.M. de Koster
ERS-2006-026-LIS
http://hdl.handle.net/1765/7831

Optimal Storage Rack Design for a 3-dimensional Compact AS/RS
Tho Le-Duc, M.B.M. de Koster and Yu Yugang
ERS-2006-027-LIS
http://hdl.handle.net/1765/7839

E-Fulfillment and Multi-Channel Distribution – A Review
Niels Agatz, Moritz Fleischmann and Jo van Nunen
ERS-2006-042-LIS
http://hdl.handle.net/1765/7901

Leveraging Offshore IT Outsoutcing by SMEs through Online Marketplaces
Uladzimir Radkevitch, Eric van Heck and Otto Koppius
ERS-2006-045-LIS
http://hdl.handle.net/1765/7902

Buyer Commitment and Opportunism in the Online Market for IT Services
Uladzimir Radkevitch, Eric van Heck and Otto Koppius
ERS-2006-046-LIS
http://hdl.handle.net/1765/7903

Managing Supplier Involvement in New Product Development: A Multiple-Case Study
Ferrie E.A. van Echtelt, Finn Wynstra, Arjan J. van Weele and Geert Duysters
ERS-2006-047-LIS

The Multi-Location Transshipment Problem with Positive Replenishment Lead Times
Yeming Gong and Enver Yucesan
ERS-2006-048-LIS

---

* A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship