



Production, Manufacturing and Logistics

Enhanced lateral transshipments in a multi-location inventory system

Colin Paterson^{a,*}, Ruud Teunter^{b,1}, Kevin Glazebrook^{a,2}^a Lancaster University Management School, Lancaster, LA1 4YX, UK^b University of Groningen, PO Box 800, 9700 AV, Groningen, The Netherlands

ARTICLE INFO

Article history:

Received 25 November 2010

Accepted 4 March 2012

Available online 12 March 2012

Keywords:

Inventory control

Lateral transshipments

Dynamic programming

ABSTRACT

In managing an inventory network, two approaches to the pooling of stock have been proposed. Reactive transshipments respond to shortages at a location by moving inventory from elsewhere within the network, while proactive stock redistribution seeks to minimize the chance of future stockouts. This paper is the first to propose an enhanced reactive approach in which individual transshipments are viewed as an opportunity for proactive stock redistribution. We adopt a quasi-myopic approach to the development of a strongly performing enhanced reactive transshipment policy. In comparison to a purely reactive approach to transshipment, service levels are improved while a reduction in safety stock levels is achieved. The aggregate costs incurred in managing the system are significantly reduced, especially so for large networks. Moreover, an optimal policy is determined for small networks and it is shown that the enhanced reactive policy substantially closes the gap to optimality.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Lateral transshipments (LTs) are stock movements between locations in the same echelon of an inventory system. They provide a valuable tool to supply chain managers who are looking to reduce the penalties associated with a lack of stock at one or more inventory points. By strategically reallocating excess stock it can be possible to improve the system wide service levels and/or lower the cost of operating the system. These goals have traditionally been sought within spare part networks, where there is a high penalty attached to a shortage. However the benefits of LTs have also been realized in sectors ranging from retail to energy generation. The challenge that LTs bring is in managing when and where it is beneficial to instigate a stock movement. An LT may reduce the short term shortage risk at the receiving location but it inevitably increases the longer term risk at the sending location. A transshipment policy must therefore balance these contrasting risks and decide when the cost of transshipment is outweighed by the benefit it is expected to deliver.

The suitability of a given LT policy will often depend on the attributes of the inventory system in which it is employed. However, a key distinction within the literature on LTs is that between reactive and proactive policies. Reactive LTs are performed when a shortage or potential shortage occurs, by shipping either the whole

customer demand or the number of units short from a different location. Proactive transshipments are performed periodically to rebalance the whole system's stock levels. This paper's principle motivation is in considering an enhanced reactive policy which falls between these two distinct sets so as to maximize the benefit each transshipment can deliver. Rather than merely looking to meet the excess demand, the proposed policy views each triggered transshipment as an opportunity to proactively rebalance the two interacting locations' inventory.

Often when a transshipment occurs the cost associated with the stock movement will primarily be a fixed cost, independent of the size of the transshipment. The reason for this is that regardless of whether one item is transported or several, the costs such as using a vehicle and the associated fuel cost of instigating the journey will be highest for the first item. The marginal cost for subsequent items will typically be much lower. When such cost structures exist it is important to know how best to carry out transshipments. Economies of scale are considered throughout inventory management, from ordering in batches to centralizing warehouses. It is therefore natural to want to know how best to operate a transshipment policy when the opportunity to extract similar benefits exists.

Within the existing literature Reactive LTs have been studied under both a periodic and continuous inventory review setting. For periodic review models, [Krishnan and Rao \(1965\)](#) develop optimal transshipments in a single period for a system with two locations. This is expanded to a multi-location, multi-period setting by [Robinson \(1990\)](#), although here the optimal reactive solution can only be determined for either two locations or multiple identical locations and when the transshipment cost structure does not

* Corresponding author. Tel.: +44 1524 594673.

E-mail addresses: c.paterson@lancaster.ac.uk (C. Paterson), r.h.teunter@rug.nl (R. Teunter), k.glazebrook@lancaster.ac.uk (K. Glazebrook).¹ Tel.: +31 50 3638617.² Tel.: +44 1524 592697.

include a fixed element. This highlights the complexity of determining optimal transshipment policies. These papers perform LTs once all demand is known but before it has to be satisfied. In contrast, Archibald (2007) and Archibald et al. (2009, 2010) develop approximately optimal policies which can respond to continuous demand within each period. The former proposes heuristics to deal with the transshipment decision process, while the latter papers look to improve upon this and relax some of the restrictions using dynamic programming policy improvement techniques. The results obtained from these policies show them to be reasonably close to optimal when used in small networks. This method of validation is one which this paper looks to emulate. The above models focus on single echelon centralized models. However, additional research in the periodic setting considers the benefit of LTs within two echelon models (e.g. Dong and Rudi, 2004), decentralized models (e.g. Rudi et al., 2001) and production based models (e.g. Zhao et al., 2008). The latter of these is closely related to this paper as it also considers an enhanced transshipment policy. However, it considers a production model where stock can be reallocated upon production whilst also allowing reactive transshipments. It therefore falls between a production allocation model and a reactive transshipment model.

Much of the literature on reactive LTs in a continuous order review setting is motivated by applications in the spare parts industry. Here, practical settings include electronic component manufacturing and electricity generation companies. Building on the METRIC repairs model of Sherbrooke (1968), Lee (1987) proposes a model which uses complete pooling within preset groups of identical locations. This shows the benefit of LTs within the area and the model is expanded by Axsäter (1990) to allow non-identical locations. Several papers have been written which further expand these ideas by relaxing or tightening some constraints such as making repair capacity finite (Jung et al., 2003), using lost sales rather than backordering (Dada, 1992) or considering a model where backorders have to be minimized rather than costs (Sherbrooke, 1986). In addition to this, inventory systems that supply more than one type of item are investigated by papers such as Wong et al. (2006b) and Kranenburg and van Houtum (2009). The latter examines the benefits of partial pooling, where only certain transshipments are performed. All of these papers assume an order-up-to replenishment policy for each location. Wijk et al. (2009) consider a single item system where parts are repaired at each location and use dynamic programming to determine an optimal transshipment policy. Kukreja and Schmidt (2005) consider a more general (s, S) policy, but have to resort to a simulation based approach to determine the optimal order policy.

Away from spare parts, Archibald et al. (1997) shows that in a periodic review model without fixed order costs, an order-up-to policy is optimal. However, positive order costs or minimum order quantities often suggest that an (R, Q) policy is more appropriate in practice. Several papers take this approach. Evers (2001) and Minner et al. (2003) develop heuristics that can be used to determine when and how much to transship for systems with lost sales. Axsäter (2003) does the same, but for a model with backorders. He proposes a decision rule which is constructed to make optimal decisions under an assumption that no further transshipments will be made. This assumption enables the exact myopic benefit of transshipping to be calculated and optimized. A related model by Axsäter et al. (2010) considers an (R, Q) inventory system in practice. They determine approximately optimal replenishment policy parameters when transshipments are sourced from a support warehouse.

Research on proactive LTs explores their use to rebalance an entire system's stock on hand. This rebalancing is done at a set point during a review period and before all demand has been realized. Allen (1958) and Agrawal et al. (2004) consider this problem

independently of replenishment decisions. Allen (1958) looks to perform the transshipments at the start of the demand period, whilst Agrawal et al. (2004) devise a method to calculate the best time to redistribute stock during the period.

Other authors study proactive transshipment and replenishment decisions together. However, due to the periodic nature of proactive redistribution all known studies only consider their use alongside a periodic review replenishment policy. Gross (1963) provides optimality results for a two-location system, where both ordering and redistribution decisions take place at the beginning of the review period. This idea is further developed by Das (1975), who allows the redistribution point to occur at an arbitrary time during the review period. Gross and Das both assume negligible transshipment times. Jönsson and Silver (1987) and Bertrand and Bookbinder (1998) allow positive transshipment times. The main difference between these two studies is that Jönsson and Silver (1987) consider how best to meet service levels whilst Bertrand and Bookbinder (1998) examine the goal of cost reduction.

For a detailed overview of the literature we refer to Paterson et al. (2011). However, the highlighted literature shows that both reactive and proactive LTs provide cost benefits, but the cost benefits of proactive LTs have only been exploited in a periodic review setting. In this study, we analyze the first 'hybrid' transshipment policy which tries to secure the benefits of both under a continuous review replenishment policy by enhancing a traditional reactive approach. Our policy can quickly react to shortages by allowing transshipments at any time when they occur, as for previously proposed reactive LT policies. However, the policy also seeks to proactively redistribute stock between the sending and receiving locations whenever such an LT is triggered. This will allow maximum benefit to be extracted from each transshipment instance and will be especially beneficial in systems where there is a significant fixed cost involved in carrying out a transshipment.

The specific setting that we consider is as in Axsäter (2003), with backordering and an arbitrary number of stocking locations which all apply (R, Q) ordering policies. Axsäter (2003) derives an algorithm that determines near-optimal reactive transshipment decisions. These are shown in a simulation study on small networks (with two and three locations) to provide a significant cost benefit compared to not transshipping at all and to applying a simpler transshipment policy. In this paper, we generalize this algorithm with the goal of determining an approximately optimal enhanced reactive transshipment policy that allows additional stock redistribution when reacting to a stock out. The results of a comparative numerical study show that, for small networks, the enhanced policy significantly outperforms the original Axsäter reactive proposal, achieving an average 1.6% cost saving over 600 experiments. Such a recurrent saving is of major practical importance, considering that inventory costs typically account for a substantial proportion of a business's total turnover. To analyze the closeness to optimality of our enhanced policies, we also develop a dynamic programming (DP) approach to finding an ϵ -optimal transshipment policy which also allows for a proactive element in each transshipment. More significantly we show through numerical results that the optimality gap is closed by over 95% on average compared to a policy of not transshipping and by 88% compared to the original reactive policy. This is strong evidence that our development of an enhanced reactive approach makes an important contribution to the application of transshipments and enables benefits which are close to optimal.

In a further numerical study, we compare the traditional and enhanced algorithms for larger networks with 5–20 locations. The exact DP algorithm is too numerically intensive to be applied in these experiments. The results of a comparison of the policies show that the improvement of the enhanced reactive policy over the traditional reactive policy is even larger than for small

networks, with a cost reduction of over 6.4% on average. It also provides an average saving of 14.4% over not transshipping at all. A sensitivity study provides further insights into when the cost reduction is most significant. The study also highlights the additional benefits that an improved transshipment policy can deliver. The average service level within the large network study is improved by 1.5% points by the enhanced policy over the purely reactive policy and the amount of safety stock required is reduced by over half when compared to a policy of no transshipments.

The remainder of the paper is organized as follows. In the next section, we describe the model and the algorithm for computing our approximately optimal enhanced reactive transshipment policy. The cost benefit of the enhanced reactive policy is tested numerically in Section 3 for networks with two locations. Section 4 describes the exact DP algorithm, and the optimality gap for small networks is investigated in Section 5. Larger networks are explored in Section 6. We end with conclusions and directions for further research in Section 7.

2. An enhanced transshipment policy

An approximately optimal policy for determining purely reactive transshipment decisions in an inventory system, in which all locations follow a continuous review (R, Q) ordering policy is derived by Axsäter (2003). The algorithm determining the policy is constructed by making use of an assumption that the considered transshipment will be the last one ever made. Whenever a location experiences a shortage, the algorithm calculates the most cost efficient amount and location to transship from under this assumption. However, the derived policy only looks to react to a shortage, not to be proactive in future shortage prevention. The policy proposed in this paper allows more stock to be transshipped than is needed to meet the immediate shortage. This permits the two locations which are parties to a transshipment to redistribute their stock and balance future risk. For the sake of clarity in the remainder of the paper we refer to the proposed policy as the *enhanced* policy, with the existing policy referred to as the *reactive* policy. Table 1 provides a list of notation needed to define the inventory system.

Table 1
List of Notation.

N	Number of demand locations: location $i \in \{1, \dots, N\}$
R_i	Reorder point at i
Q_i	Batch size of orders placed at i
$t_{s,i}$	Time until s th unit of stock becomes available at i
IL_i	Inventory level at i
IP_i	Inventory position at i
L_i	Lead time of orders placed at i
$D_i(L_i)$	Stochastic lead time demand at i
A_i	Order cost at i (per order)
h_i	Holding cost at i (per item per unit time)
b_i	Backorder cost at i (per item per unit time)
$c_{i,j}(y)$	Cost of transshipping y units from i to j
$c^f_{i,j}$	Fixed cost of transshipping from i to j
$c^u_{i,j}$	Cost per unit of transshipping from i to j
λ_i	Arrival rate of demand instances at i
μ_i	Average size of each demand at i
$f_{i,j}$	Probability that a demand at i will be of size j
$f^n_{i,j}$	Probability that j units are demanded by n customers at i
$P^n_{i,j}$	Probability that n th customer demands the j th unit at i
C_i	Steady state cost of location i
$C_i(k)$	Expected cost rate at time L_i given current $IP_i = k$
X_i	State variable at location i
$X'_i(d)$	State variable at location i less d units of stock
$V_i(X_i, t)$	Expected total cost at i during the interval $[0, t]$
$\alpha_i(X_i)$	Expected lead time bias at i given starting state X_i
$\beta_i(IP_i)$	Expected bias after lead time at i given current IP_i
$\gamma_i(X_i)$	Total expected bias at i given starting state X_i

The system has N stocking locations which all experience independent compound Poisson demand with location i having arrival intensity λ_i . Modeling demand as a compound Poisson process enables the policy to be applied in a wide range of inventory systems that allow customers to demand more than one item at a time. The probability that a demand at location i is of size j is denoted by $f_{i,j}$ and we do not prescribe any distributional form. It should also be noted that $P^n_{i,j}$, the probability that the n th customer demands at location i the j th unit of stock, can be calculated using the recursion $P^n_{i,j} = P^n_{i,j-1} - f^n_{i,j-1} + f^n_{i,j-1}$.

To replenish its stock location $i \in \{1, \dots, N\}$ places an order of size Q_i at the central supplier whenever its inventory position drops to or below reorder level R_i . This supplier is assumed to have sufficient inventory capacity so that this order can always be met and takes a fixed lead time L_i to arrive. If a location does not have sufficient stock on hand to satisfy a demand, then items may be transshipped to that location from a different location with negligible lead time. This setup is valid as stocking points are often geographically closer than that of the supplier and do not have the external logistical delay of placing an order with a supplier. Any demand that cannot be met immediately (after transshipping) is backordered. Costs are incurred for ordering (A_i), holding stock on hand (h_i per item and per time unit), backordering (b_i per item and per time unit), and for transshipping. The cost for transshipping y units from location i to j is given by $c_{i,j}(y) = c^f_{i,j} + yc^u_{i,j}$, where $c^f_{i,j}$ is the fixed cost per transshipment and $c^u_{i,j}$ is the cost per unit transshipped.

Under the assumption that no transshipments take place, Axsäter (2003) derives an expression for the *bias* associated with each system state. This measures the transient effect on costs of starting the system in that state. Performing a transshipment will instantly move the system to a new state, so the benefit of a given transshipment can be identified by comparing the bias of the current state with the aggregate of the bias if the transshipment is enacted and the cost to enact it. By maximizing the difference between these quantities over all possible locations and transshipment quantities, the best myopic decision can be identified. The limitation of the formulation given by Axsäter (2003) is that it does not allow the size of the transshipment to be larger than the shortage. This restriction is mathematically convenient in that it ensures that the inventory position of any location never exceeds $R_i + Q_i$ and thus provides a closed range on which to consider the future bias. However, when there is a significant fixed cost per transshipment it seems intuitive that allowing larger transshipments may well deliver greater cost benefits. Not every such enhanced transshipment will take the inventory position above $R_i + Q_i$, but to provide a completely general account and to calculate the benefit associated with every possible decision it is necessary to extend the methodology to account for such scenarios. Please note that, while this extension is certainly needed in general, we shall encounter in Section 4 two location set-ups where almost all of the (very considerable) benefits from enhanced transshipments can actually be achieved while keeping the inventory below $R_i + Q_i$.

We now describe how the above ideas can be developed to yield near-optimal decisions in our enhanced approach. In steady state under an (R_i, Q_i) replenishment policy without transshipments it is known that location i 's inventory position is uniformly distributed over the range $[R_i + 1, \dots, R_i + Q_i]$. With current (time zero) inventory position k and stochastic lead time demand $D_i(L_i)$ then the mean inventory cost rate at L_i is given by

$$\begin{aligned}
 C_i(k) &= h_i \mathbb{E}(k - D_i(L_i))^+ + b_i \mathbb{E}(k - D_i(L_i))^- \\
 &= (h_i + b_i) \mathbb{E}(k - D_i(L_i))^+ + b_i \mathbb{E}(D_i(L_i) - k) \\
 &= (h_i + b_i) e^{-\lambda_i L_i} \sum_{j=0}^{k-1} (k-j) \sum_{n=0}^j \frac{(\lambda_i L_i)^n}{n!} f^n_{i,j} + b_i (\lambda_i \mu_i L_i - k). \quad (2.1)
 \end{aligned}$$

We infer that the steady state cost rate for location i is given by

$$C_i = \frac{1}{Q_i} \sum_{k=R_i+1}^{R_i+Q_i} C_i(k). \tag{2.2}$$

The current state X_i of location i incorporates information on the current inventory position (IP_i) together with the times at which each inventory item becomes available. We write $V_i(X_i, t)$ for the total expected cost incurred under an (R_i, Q_i) ordering policy with no transshipments during the time interval $[0, t]$ when the system is in state X_i at time 0. The bias associated with X_i is defined as

$$\gamma_i(X_i) \triangleq \lim_{t \rightarrow \infty} \{V_i(X_i, t) - tC_i\}. \tag{2.3}$$

Suppose that $t > 0$ and the random state of location i at time L_i is denoted by $X_i^{L_i}$; we then have

$$V_i(X_i, L_i + t) = V_i(X_i, L_i) + \mathbb{E}V_i(X_i^{L_i}, t). \tag{2.4}$$

We now easily see that

$$\gamma_i(X_i) = \lim_{t \rightarrow \infty} \{V_i(X_i, L_i + t) - (L_i + t)C_i\}, \tag{2.5}$$

has a decomposition given by

$$\gamma_i(X_i) = \alpha_i(X_i) + \beta_i(IP_i), \tag{2.6}$$

where

$$\alpha_i(X_i) \triangleq V_i(X_i, L_i) - L_i C_i, \tag{2.7}$$

and

$$\beta_i(IP_i) \triangleq \lim_{t \rightarrow \infty} \{\mathbb{E}V_i(X_i^{L_i}, t) - tC_i\}. \tag{2.8}$$

The calculation of $\alpha_i(X_i)$ is given in the Appendix and is unchanged by the fact that here we may have $IP_i > R_i + Q_i$. However the earlier method for calculating β_i is dependent on the inventory position at time L_i being in the range $[R_i + 1, \dots, R_i + Q_i]$. A particularly large proactive transshipment may invalidate this assumption. Therefore in our formulation we must take account of the additional bias that will be arise if the inventory position is taken above $R_i + Q_i$. We now describe the calculation of β_i .

We first observe that, since $X_i^{L_i}$ is stochastically independent of all information in X_i save only IP_i , then β_i depends upon X_i only through IP_i . More specifically, $IL_i^{L_i}$, the inventory level at time L_i , can be calculated through the relation

$$IL_i^{L_i} = IP_i - D_i(L_i). \tag{2.9}$$

We also note from Axsäter (2003) the following notation to show the relationship between IL_i and IP_i at any point in time. For a given $IL_i = k$ the corresponding inventory position can be given by

$$\langle k \rangle = k + nQ_i, \tag{2.10}$$

where n is the smallest integer such that $\langle k \rangle \geq R_i + 1$.

As location i evolves under an (R_i, Q_i) ordering policy with no transshipments, the associated inventory position process regenerates upon every entry into state $R_i + 1$ (or indeed any other state – $R_i + 1$ is chosen for convenience). In (2.8) we now deem $X_i^{L_i}$ to be the location i state at time zero and write T for the first subsequent time at which the corresponding inventory position, given by rearranging (2.9), is in the regeneration state $R_i + 1$. Standard theory, for example in Tijms (1986), partitions the infinite horizon considered in (2.8) as $[0, \infty) = [0, T) \cup [T, \infty)$ and demonstrates that the contribution to $\beta_i(IP_i)$ from the latter interval is zero. This reasoning yields both

$$\beta_i(IP_i) = \mathbb{E}V_i(X_i^{L_i}, T) - E(T)C_i, \tag{2.11}$$

and

$$\beta_i(R_i + 1) = 0. \tag{2.12}$$

We now fix $IP_i = k > R_i + 1$ and develop an expression for $\beta_i(k)$ by conditioning upon the size of the first demand to occur after time zero. The expected time at which this first demand occurs is λ_i^{-1} and from (2.11) the cost rate contribution to $\beta_i(k)$ prior to then is $C_i(k) - C_i$. Following demand d the inventory position at i will be changed to $\langle k - d \rangle$ and hence we obtain

$$\beta_i(k) = \frac{\{C_i(k) - C_i\}}{\lambda_i} + \sum_{d=1}^{\infty} f_{i,d} \beta_i(\langle k - d \rangle). \tag{2.13}$$

We now use (2.13) to recover β_i as the limit of an iterative scheme which is akin to DP value iteration. It proceeds as follows:

$$\beta_i^0(k) = 0, \quad k \in [R_i + 1, \dots, R_i + Q_i, \dots, S], \tag{2.14}$$

$$\beta_i^n(R_i + 1) = 0, \quad n \geq 0, \tag{2.15}$$

$$\beta_i^{n+1}(k) = \frac{\{C_i(k) - C_i\}}{\lambda_i} + \sum_{d=1}^{\infty} f_{i,d} \beta_i^n(\langle k - d \rangle), \tag{2.16}$$

$$k \in [R_i + 2, \dots, R_i + Q_i, \dots, S],$$

where S is a large inventory position state such that no higher state is reached. The scheme must converge geometrically fast, with β_i the limit.

We can now use the complete bias functions γ_i for each location i to identify the transshipment that delivers the most cost benefit. We use the notational shorthand $X_i'(d)$ for the resulting state of a location i once d units of inventory have been withdrawn, either through customer demand or for transshipment. Suppose that some demand d occurs at location i when in state X_i and that this causes a shortage. The long term cost benefit of transshipping y units from location j (in current state X_j) to location i in comparison with performing no transshipment and allowing all demand to be absorbed at i is identified as the quantity

$$\Delta(j, y) = \lim_{t \rightarrow \infty} \{V_i(X_i'(d), t) + V_j(X_j, t) - V_i(X_i'(d - y), t) - V_j(X_j'(y), t)\} - c_{j,i}(y) - y(A_j/Q_j - A_i/Q_i). \tag{2.17}$$

From (2.3) we have

$$\Delta(j, y) = \gamma_i(X_i'(d)) + \gamma_j(X_j) - \gamma_i(X_i'(d - y)) - \gamma_j(X_j'(y)) - c_{j,i}(y) - y(A_j/Q_j - A_i/Q_i). \tag{2.18}$$

We note that, prior to (2.17), we have taken no account of the order costs in our calculations. A transshipment of y units from j to i has the effect of adjusting the long-term cost burden from orders by $y(A_j/Q_j - A_i/Q_i)$. We write

$$\Delta^* = \max_{j \neq i} \max_{1 \leq y \leq IL_j} \Delta(j, y). \tag{2.19}$$

Our enhanced policy is as follows: If $\Delta^* \leq 0$, do not transship. If $\Delta^* > 0$ transship y^* units from j^* to i , where y^* and j^* are the maximizers in (2.19). This decision rule is invoked every time a demand occurs at a location which would result in a shortage.

This policy is quasi-myopic in that it is cost minimizing if no further transshipment is permitted after the current shortage is dealt with. As the reactive policy of Axsäter (2003) is a more constrained case of our general enhanced policy, under the myopic assumption any enhanced decision will be at least as good as that of the purely reactive policy.

3. Two location simulation study: reactive vs. enhanced policy

To analyze the performance of the enhanced transshipment policy, an initial simulation study is conducted. This study is restricted to networks with two identical locations. In a second study, of

which the details and results will be presented in later sections, larger and more varied networks will be considered. The two reasons for the more restricted initial exploration are as follows. First, there are many model parameters and varying all of them for larger networks is time consuming. This initial exploration allows us to observe how the parameters impact on costs so that the larger network can focus on key issues. Secondly, the restriction will allow us to determine the optimal transshipment policy using DP in Section 5 for the same set of experiments, and therefore to study the relative reduction of the optimality gap from the original reactive policy to the new enhanced transshipment policy.

Since the two locations are assumed identical, we will drop the location identifying subscript for cost parameters in the remainder of this section. We will do the same for the policy parameters, and assume that both locations use the same replenishment and transshipment policy.

Although the transshipment policy might have some effect on the optimal ordering quantity, in practice there are often fixed or minimal order sizes. For this reason and in line with previous studies, we fix the order cost (\$100) and use the EOQ formula $(\sqrt{(2A\lambda\mu)/(h)})$ to determine the order size. The validity of this decision is discussed below. For the sizes of successive demands within the study we use a geometric distribution such that $f_j = p(1 - p)^{j-1}$, $j \geq 1$. The holding cost rate (\$1) is used as the unit cost and all other model parameters are varied in a full factorial study. The full range of parameters examined is shown in Table 2. There are 600 parameter combinations and hence 600 experiments in total.

For both the original reactive and enhanced transshipment policies, the transshipment decisions are determined by the algorithms developed in Section 2. To obtain further insights we also consider the no transshipment policy as a benchmark. For all three (no, reactive and enhanced) transshipment policies, the optimal value of R is found post hoc by conducting simulation studies on the full range of possible values.

Over the 600 experiments an average saving of 1.59% is observed for the enhanced policy when compared to the reactive policy. This is broken down by each parameter in Table 3. For example, the 200 experiments which have an arrival rate of 2.4 customers per unit time yield an average saving of 1.77%.

The results in Table 3 confirm intuition in several ways. Once a transshipment is instigated, a policy which looks to transship more will see a greater benefit if the marginal cost of adding an extra unit is low. This is supported by the c^u results. Further, if a policy is performing each transshipment more efficiently then it is natural that the saving will increase when there are more transshipment opportunities. By increasing parameters λ or L , or decreasing parameter p we raise the lead time demand variability and thus the chance of a shortage. Unsurprisingly, increasing the shortage risk through varying these parameters in the manner described leads to an observed gain in savings.

Considering the backorder costs, we see that the savings are greater as the penalty for not immediately meeting a demand increases. However, as the fixed transshipment cost increases the dynamics are less straightforward. As the cost per transshipment

Table 3
Enhanced rule vs reactive rule.

λ	p	L	b	c^f	c^u
0.8	0.85%	0.6	1.72%	2	1.23%
2.4	1.77%	0.8	1.46%	3	1.95%
4.0	2.14%				
				10	0.84%
				20	1.47%
				30	1.76%
				40	1.91%
				50	1.98%
				10	1.48%
				20	1.75%
				30	1.72%
				40	1.59%
				50	1.41%

increases then naturally so does the reward for carrying out more effective transshipments but at the same time the benefit associated with carrying out any transshipment decreases. This can be observed in Table 3 where the benefit conferred by the enhanced approach increases up to a point before decreasing. This is supported by the average number of transshipments per 100 units of time, which falls from 11.2 when $c^f = 10$ to only 3.3 when $c^f = 50$.

More detailed results are provided in the Appendix for a sample of parameter values (27 experiments). All statistical comparisons have used paired t -tests at a 95% confidence level, with common random numbers used for all policies. Overall, the average saving of the new enhanced policy compared to one which uses no transshipments is 4.21%, reinforcing the view that transshipping is worthwhile. The standard errors of the costs for the reactive and the new enhanced policies are also given in Table A.1 (in brackets). There are a few low demand cases where there is no statistically significant difference between the reactive and enhanced policies but no cases were found where the enhanced policy was outperformed by the existing policy. We are not able to prove that the enhanced policy always outperforms the reactive policy but we have studied a very large number of problem instances and have as yet found no counterexamples.

A comparison of the safety stock levels R provides insight into how the cost reduction is achieved. The average reorder level is 9.8 units without transshipments, 8.9 units for the reactive transshipment policy and 8.2 units for the new enhanced transshipment policy. Better transshipment decisions reduce the negative cost effects of shortages, thereby allowing the system to function with lower safety stocks. Further exploration also showed that in general the enhanced transshipment policy is more cost efficient, with transshipments happening less frequently but with larger quantities (4.7 units for the enhanced policy compared to 2.1 units for the reactive policy on average).

To examine the effect of the batch size on the relative performance of the policy, the reorder point R , which was optimised using a value of Q set by the EOQ, was fixed and used to determine the corresponding optimal batch size in a similar fashion. It should be noted that this pair may still not be the global optimum. The results from this showed that the difference between the local optimal value of Q and the EOQ value is relatively consistent between the policies, with a difference of 2 units for the reactive policy and 2.2 for the enhanced policy. In fact the relative performance of the two policies remained constant, with the performance gain remaining at 1.6%. This indicates that the observed gains are robust when the batch size is varied and that using the EOQ value gives a suitable overview of the policies' performance.

The results in this section have clearly shown that the new enhanced transshipment policy significantly outperforms the original reactive transshipment policy. However, it remains of real interest to discover how close (in cost terms) it is to an optimal transshipment policy. To explore this, we will develop a dynamic programming (DP) formulation for finding the optimal transshipment policy in the next section, and then compare its cost to that of the new enhanced transshipment policy in Section 5, using the same set of experiments that were investigated in this section.

Table 2
Parameter values.

Arrival rate (λ)	0.8, 2.4, 4.0
Geometric distribution parameter for demand size (p)	0.6, 0.8
Lead time (L)	2, 3
Backorder cost (b)	10, 20, 30, 40, 50(\$)
Transshipment cost [per item] (c_i)	1, 2(\$)
Transshipment cost [per transshipment] (c_f)	10, 20, 30, 40, 50(\$)

Table 4
Additional notation.

δ	Time quantum between each state transition
\tilde{f}_{ij}	Truncated probability of demand j at location i such that $R_i - Q_i < IL_i$ is always true
Y_i	Number of quanta until the outstanding order arrives at location i
Z	Location of most recent demand where $Z = 0$ indicates no demand occurred in the preceding period
I_H	Indicator function where H is a logical statement: if true $I_H = 1$ else $I_H = 0$

4. Dynamic programming formulation

In this section, we provide a DP formulation that can be used to find the optimal transshipment policy using value iteration. In our system, the resulting policy determines how much to transship, given the locations' inventory levels and remaining lead times of outstanding orders, in order to minimize the overall cost rate incurred. We remark that our DP formulation utilizes a discrete time approximation of the actual continuous review system. However, for a sufficiently small time quantum the approximation is very good. This will be verified in the numerical investigation of Section 5. For presentational ease, we will refer to the transshipment policy that minimizes the cost under the DP formulation as the optimal transshipment policy.

Our formulation is for a two location system setup. While it is possible to generalize the approach to larger systems, value iteration becomes computationally intractable very quickly due to the rapid growth in the number of states. In order to limit the state space, only one order is allowed to be outstanding at each location at any time. This assumption is supported by real world practice and by the simulation results obtained in Section 3. These showed that conditional upon a location having orders outstanding the probability of multiple orders outstanding was less than 1%. Our simulations also showed that limiting the enhanced policy in a two location network so that a location could only ever have a maximum stock of $R_i + Q_i$ resulted in no statistically significant cost difference. The reader is referred to the related discussion in Section 2. This result is due to the fact that it is unlikely that one location would have enough excess stock to be able to completely refill the other's safety stock, and thus make $IP_i > R_i + Q_i$. Therefore it is feasible to limit the inventory state space to a range of $R_i \pm Q_i$. It should be noted that in a multi-location network, and in particular with non-identical locations, this assumption will not hold. Recall from Section 3 that we will apply the DP algorithm for the same set of experiments as described there. To develop the DP model we introduce additional notation and formulae in Table 4. For sufficiently small time quantum δ , we may assume that in a single time

slot the system experiences an instance of demand at either one location or neither location. The probability of a demand of size j at location i is $\delta\lambda_i\tilde{f}_{ij}$ during each period while the probability of no demand in the system is $1 - \delta(\lambda_1 + \lambda_2)$. In the limit $\delta \rightarrow 0$ these probabilities converge to the exact Poisson probabilities.

4.1. State definition

A five dimensional system state incorporates the inventory level and the time until the outstanding order arrives (if there is one) at both locations. The fifth dimension indicates the location where any current demand has occurred. We write state s as follows

$$s = \langle IL_1, IL_2, Y_1, Y_2, Z \rangle \tag{4.1}$$

Recall that we do not allow the inventory position to go above $R_i + Q_i$. Further, the decision to allow a maximum of one outstanding order bounds the inventory level below. Hence we have that $R_i - Q_i < IL_i \leq R_i + Q_i$. Moreover, we clearly have $0 \leq Y_j < \frac{L_j}{\delta}$.

4.2. Action space definition

The action a ($-a$) is the amount to transship from locations $1(2)-2(1)$. For a given state s , the set of actions $Act(s)$ is bounded by zero, and by the minimum of the amount available to transship and the amount that can be stored at the receiving location (i.e. that does not take the inventory position above its maximum). We summarize these constraints by.

If $IL_2 > 0, IL_1 < 0$ and $Z = 1$:

$$a \in \{-\min(IL_2, R_1 + I_{(Y_1=0)} * Q_1 - IL_1), \dots, 0\}$$

If $IL_1 > 0, IL_2 < 0$ and $Z = 2$:

$$a \in \{0, \dots, \min(IL_1, R_2 + I_{(Y_2=0)} * Q_2 - IL_2)\}$$

Else: $a \in \{0\}$.

Hence, for example, if $IL_1 > 0, IL_2 < 0, Z = 2$ and $Y_2 > 0$ then a is an integer in the range from 0 to $\min(IL_1, R_2 - IL_2)$.

4.3. Cost function

The cost ($\zeta_s(a)$) associated with being in state s and choosing action a is obtained by aggregating the cost of holding (or backordering) the current level of inventory after the transshipping action has taken place with the cost of the action itself. There is also the additional cost of placing any replenishment order

Table 5
Optimality gap analysis.

λ				p				L			
Parameter value	Reactive vs. Opt. (%)	Enhanced vs. Opt. (%)	Improvement (%)	Parameter value	Reactive vs. Opt. (%)	Enhanced vs. Opt. (%)	Improvement (%)	Parameter value	Reactive vs. Opt. (%)	Enhanced vs. Opt. (%)	Improvement (%)
0.8	1.00	0.14	86	0.6	2.01	0.29	86	2	1.35	0.12	91
2.4	1.99	0.23	90	0.8	1.61	0.15	92	3	2.28	0.33	86
4.0	2.45	0.31	90								
b				c^f				c^H			
10	0.94	0.11	90	10	1.96	0.49	76	1	1.99	0.26	88
20	1.64	0.18	90	20	2.00	0.26	88	2	1.64	0.19	89
30	1.99	0.24	89	30	1.88	0.16	92				
40	2.19	0.28	87	40	1.71	0.12	93				
50	2.30	0.32	86	50	1.52	0.10	94				

which is included in the period directly after the order has been instigated. Hence we have

$$\zeta_s(a) = h_1(IL_1 - a)^+ + h_2(IL_2 + a)^+ + b_1(IL_1 - a)^- + b_2(IL_2 + a)^- + A_1 I_{(Y_1 = \frac{L_1}{\delta} - 1)} + A_2 I_{(Y_2 = \frac{L_2}{\delta} - 1)} + c^u |a| + c^f I_{|a| > 0} \quad (4.2)$$

4.4. State transitions

If the current state is $s = \langle IL_1, IL_2, Y_1, Y_2, Z \rangle$, action a is undertaken and demand d occurs at location 1 then the new state will be $s' = \langle IL'_1, IL'_2, Y'_1, Y'_2 \rangle$, where

$$\begin{aligned} IL'_1 &= IL_1 - a - d + Q_1 I_{(Y_1=1)}, \\ IL'_2 &= IL_2 + a + Q_2 I_{(Y_2=1)}, \\ Y'_1 &= (Y_1 - 1) I_{(1 \leq Y_1 \leq \frac{L_1}{\delta} - 1)} + \left(\frac{L_1}{\delta} - 1\right) I_{(IL_1 - a - d \leq R_1 \& Y_1 = 0)}, \\ Y'_2 &= (Y_2 - 1) I_{(1 \leq Y_2 \leq \frac{L_2}{\delta} - 1)} + \left(\frac{L_2}{\delta} - 1\right) I_{(IL_2 + a \leq R_2 \& Y_2 = 0)}. \end{aligned} \quad (4.3)$$

Similar transitions can be identified for the cases when a demand occurs at location 2 and when no demand occurs anywhere in the system.

4.5. Value iteration

The above are deployed in a value iteration in which the value function $(V_n(s))$ is the minimal cost incurred over an n -period horizon from initial state s . If we write $\psi(s, s')$ for the probability of moving from state s to s' then the optimality principle gives

$$V_n(s) = \min_{a \in Act(s)} \left(\zeta_s(a) + \sum_{s' \in S} \psi(s, s') V_{n-1}(s') \right). \quad (4.4)$$

We develop the $(V_n)_{n \geq 1}$ using backwards induction and utilize the stopping criterion recommended by Tijms (1986). The minimizing actions in the final iteration yield the ϵ -optimal policy.

5. Optimality gap

For each of the 600 experiments, the dynamic programming model was used to determine the optimal transshipment policy. This was then used within the simulation model of Section 3 along with the two approximate policies and the policy of not transshipping. This ensured that like for like comparisons could be made using the same set of randomly generated events. It was necessary to gather results for a range of values of the reorder

point R so that the optimal value could be used for each parameter set.

5.1. Results

The benefit of using the enhanced policy can now be seen in the large step it takes towards the performance of the ϵ -optimal policy. Table 5 shows that over a range of parameter values there is a consistent level of improvement in cost performance achieved by including a proactive stock rebalancing element to a reactive policy. Compared to a policy of no transshipment the enhanced policy closes the suboptimality gap by over 95% over the entire data set. Compared to the reactive policy it is closed by 89%; it is this relative improvement over the optimality gap which is broken down in Table 5.

The enhanced policy appears to perform strongly in comparison to the optimal policy. In fact in 90 of the 600 experiments there is no statistical difference between the two policies. It also displays very similar characteristics to the dynamic programming based optimal policy. Both its average reorder point (8.2 vs. 8.1) and average transshipment size (both 4.7) are similar. The one difference is the number of transshipment events. For every 100 time units the ϵ -optimal policy carries out one extra transshipment on average (6.4 vs. 7.5). This implies that the lack of future information results in a slightly more conservative approach in the enhanced policy. A subsection of the full results can be found in Table A.2.

Fig. 1 clearly demonstrates the superiority in performance of the enhanced policy over the reactive policy, more so for larger values of the fixed transshipment cost. It also gives an indication of the greater consistency the enhanced policy achieves in performance.

5.2. Accuracy of Discrete Time Assumption & Computation Time

The above results arising from the DP implementation all use discrete time quantum $\delta = \frac{1}{8}$. We confirm the acceptability of this choice by resolving with $\delta = \frac{1}{16}$ and observing that the resulting changes in cost rates are minimal and nowhere statistically significant.

While using the optimal policy may reduce inventory costs its development is computationally expensive when compared to the heuristic policies, which can be obtained very rapidly in real time. The time taken to compute each experiments optimal policy was recorded.

Table 6 gives a breakdown of the time needed to develop an optimal policy by arrival rate. It displays how halving the size of δ from $\delta = \frac{1}{8}$ increases the computational time by a factor of 5 or

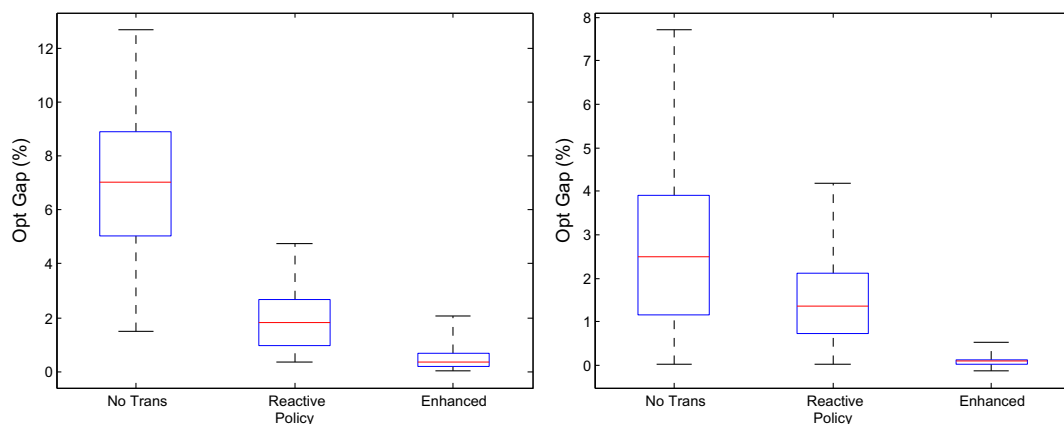


Fig. 1. Sensitivity Analysis for $c^f = 10$ (left) and $c^f = 50$ (right).

Table 6
Computational time (mins).

Arrival rate	$\delta = \frac{1}{8}$	$\delta = \frac{1}{16}$	Multiple
0.8	3.2	24.2	7.6
2.4	33.6	222.7	6.6
4.0	154.9	759.7	4.9
Overall	63.9	335.5	5.3

more on average. Moreover, the computation time increases rapidly with the arrival rate.

6. Large network study

Having shown that the enhanced policy improves the original reactive policy and makes a large step towards closing the optimality gap while reducing some of the variability in performance, the next step is to consider its performance in larger networks. In a simulation study, designed in a similar way to the above small network study, inventory systems with 5, 10 and 20 location are considered. Rather than a full factorial study the focus is now put on how the size of the network and the arrival rate (and hence the lead time demand variability) influences policy performance. Eight different arrival rates are considered in networks with identical locations. Additionally, networks which had two different levels of arrivals are considered. These networks have 40% of the locations at a higher demand rate than the other 60%. In the latter case, the overall system arrival rate is set equal to a corresponding identical location configuration so that comparisons could be fairly drawn (see Table 7).

A full set of results is given in Tables A.3 and A.4 within the Appendix. The average percentage savings, broken down by assigned values of the arrival rate, are given in Table 8. The overall average results show a saving over the reactive policy of 6.38% for identical locations and 6.42% for the networks with a two tier arrival rate structure. As for smaller networks, the greatest savings occur when the arrival rate is large.

The difference in results between the identical location setup and the two tier setup is small. Ideally a system with many different arrival rates could be considered but this is challenging to implement due to the necessity of determining suitable values of R for each location via a post hoc optimization.

The enhanced policy offers a consistent level of cost improvement. Average costs are reduced by between 11% and 17% in comparison with no transshipment. In the case of the reactive policy, the cost saving can be as little as 6%. The greater consistency in performance of the enhanced policy against increasing arrival rate is illustrated by Fig. 2.

While costs are an important part of inventory systems it is not the only performance measure of interest. Service levels within a system are also a key consideration and the fill rates (the percentage of demand filled immediately from stock on hand or via transshipment) were also recorded for the large net-

Table 7
Parameter values.

Number of locations	5, 10, 20
Arrival rate [Identical Locations](λ)	1.4, 2.2, 3.0, 3.8, 4.6
Low arrival rate [Different Locations](λ_l)	1.0, 1.8, 2.6, 3.4, 4.2
High arrival rate [Different Locations](λ_h)	2.0, 2.8, 3.6, 4.4, 5.2
Distribution of order size (p)	0.8
Lead time (L)	3
Backorder cost (b)	30(\$)
Transshipment cost [per item] (c^u)	1(\$)
Transshipment cost [per transshipment] (c^t)	10(\$)

Table 8
Percentage cost savings for the enhanced policy over the reactive policy: large networks.

Average λ	Identical (%)	Two tier (%)
1.4	4.57	4.64
2.2	5.93	5.82
3.0	6.69	6.60
3.8	7.17	7.33
4.6	7.55	7.69




Fig. 2. Savings achieved by the enhanced policy over no transshipment.

work study. For the reactive policy a service level of 96.8% was achieved, but the enhanced policy increased this to 98.3%.

The large network study also reinforces other findings from the smaller network results. For identical locations, safety stock is reduced from 12.6 units on average with no transshipments to 9.3 units under the reactive policy and to 5.6 units under the new enhanced policy. Similar results are obtained for the two tier networks. The average size of transshipment again increases, from 1.3 units to 6.0 units. These results illustrate the greater efficiency possible from anticipating shortages rather than merely responding to them.

7. Conclusions

We have shown that the benefits of reactive transshipments can be enhanced by the development of a new type of policy, which incorporate a proactive element. System costs can be reduced and the efficiency of the transshipment process improved. This has been observed through an extensive study of both small and large (R, Q) replenishment policy inventory networks, with the benefits growing with the number of stock holding locations. Moreover the improvements that this enhanced transshipment policy can bring have been shown to significantly reduce the optimality gap.

The comparison to optimality has been achieved through a dynamic programming model that enables the calculation of an ϵ -optimal transshipment policy and the resulting costs. Whilst this formulation is restricted to small systems it is an important step in understanding the transshipment process and in evaluating the performance of the more easily developed enhanced policy. More importantly, the numerical results show that the enhanced policy performs close to optimal over the examined parameters, with some scenarios showing no significant difference between the enhanced policy and the ϵ -optimal policy.

One possible way to further improve the enhanced policy is to relax the myopic assumption that underpins it. Another avenue would be to develop the redistribution element by considering transshipments at times other than at those when shortages occur.

One clear limitation within batch ordering systems with transshipments is the challenge to find appropriate replenishment policy parameter values, even more so in systems with non-identical locations. Our results have shown that the enhanced transshipment policy can significantly alter the optimal reorder point when compared to no-transshipments. Indeed part of the savings achieved is a consequence of being able to lower the amount of safety stock required throughout the system. Future work could develop analytical approaches to the determination of reorder

points. This would enable the full benefits of the improvements in the transshipment policy to be realized in more complex inventory systems with a larger number of stocking locations and non-identical demand rates.

Appendix A. Detailed results

Tables A.1–A.4.

Table A.1

Two location results: $p = 0.8, L = 3, c^d = 1$.

λ	b	c^f	Q	R	No tran.	R	Reactive (Err)	Saving (%)	R	Enhanced (Err)	Saving (%)	Improve (%)
0.8	10	10	15	1	29.96 (0.02)	1	29.08 (0.02)	2.96	1	28.86 (0.02)	3.66	0.73
0.8	10	30	15	1	29.96 (0.02)	1	29.79 (0.02)	0.56	1	29.61 (0.02)	1.17	0.61
0.8	10	50	15	1	29.96 (0.02)	1	29.91 (0.02)	0.18	1	29.84 (0.02)	0.41	0.23
0.8	30	10	15	3	33.41 (0.03)	2	31.04 (0.02)	7.10	2	30.75 (0.02)	7.95	0.92
0.8	30	30	15	3	33.41 (0.03)	3	32.29 (0.02)	3.34	3	31.91 (0.02)	4.47	1.17
0.8	30	50	15	3	33.41 (0.03)	3	32.84 (0.02)	1.70	3	32.38 (0.02)	3.06	1.38
0.8	50	10	15	4	34.83 (0.03)	3	31.80 (0.02)	8.70	3	31.66 (0.02)	9.11	0.46
0.8	50	30	15	4	34.83 (0.03)	3	33.11 (0.03)	4.95	3	32.56 (0.02)	6.52	1.65
0.8	50	50	15	4	34.83 (0.03)	4	33.81 (0.02)	2.93	3	33.26 (0.02)	4.51	1.63
2.4	10	10	25	7	51.76 (0.02)	6	50.93 (0.02)	1.59	5	49.78 (0.02)	3.83	2.27
2.4	10	30	25	7	51.76 (0.02)	7	51.66 (0.02)	0.19	6	51.16 (0.02)	1.15	0.96
2.4	10	50	25	7	51.76 (0.02)	7	51.73 (0.02)	0.05	7	51.55 (0.02)	0.41	0.36
2.4	30	10	25	10	57.14 (0.03)	9	54.32 (0.02)	4.94	8	53.08 (0.02)	7.11	2.28
2.4	30	30	25	10	57.14 (0.03)	10	56.31 (0.02)	1.45	9	54.81 (0.02)	4.09	2.68
2.4	30	50	25	10	57.14 (0.03)	10	56.81 (0.02)	0.59	9	55.67 (0.02)	2.57	2.00
2.4	50	10	25	12	59.51 (0.03)	9	55.53 (0.03)	6.70	9	54.39 (0.02)	8.61	2.05
2.4	50	30	25	12	59.51 (0.03)	11	57.82 (0.02)	2.85	10	56.11 (0.02)	5.72	2.96
2.4	50	50	25	12	59.51 (0.03)	11	58.60 (0.03)	1.54	10	57.02 (0.02)	4.18	2.68
4	10	10	32	13	66.82 (0.02)	12	66.02 (0.02)	1.20	10	64.09 (0.02)	4.08	2.92
4	10	30	32	13	66.82 (0.02)	13	66.76 (0.02)	0.10	12	65.93 (0.02)	1.34	1.24
4	10	50	32	13	66.82 (0.02)	13	66.82 (0.02)	0.01	12	66.49 (0.02)	0.50	0.49
4	30	10	32	17	73.55 (0.03)	15	70.52 (0.02)	4.13	14	68.44 (0.02)	6.95	2.94
4	30	30	32	17	73.55 (0.03)	16	72.81 (0.03)	1.00	15	70.60 (0.02)	4.01	3.04
4	30	50	32	17	73.55 (0.03)	17	73.33 (0.03)	0.30	16	71.74 (0.02)	2.47	2.17
4	50	10	32	18	76.38 (0.04)	16	72.03 (0.03)	5.70	15	70.15 (0.03)	8.16	2.60
4	50	30	32	18	76.38 (0.04)	17	74.84 (0.03)	2.02	16	72.23 (0.03)	5.43	3.48
4	50	50	32	18	76.38 (0.04)	18	75.68 (0.03)	0.91	17	73.48 (0.02)	3.79	2.91

Table A.2

Two location results: optimality comparison $p = 0.8, L = 3, c^d = 1$.

λ	b	c^f	Q	Opt.	No tran.	Gap (%)	Reactive	Gap (%)	Enhanced	Gap (%)
0.8	10	10	15	28.84	29.96	3.75	29.08	0.81	28.86	0.09
0.8	10	30	15	29.58	29.96	1.29	29.79	0.73	29.61	0.12
0.8	10	50	15	29.82	29.96	0.49	29.91	0.31	29.84	0.08
0.8	30	10	15	30.68	33.41	8.17	31.04	1.16	30.75	0.24
0.8	30	30	15	31.89	33.41	4.55	32.29	1.25	31.91	0.08
0.8	30	50	15	32.39	33.41	3.05	32.84	1.37	32.38	0.00
0.8	50	10	15	31.55	34.83	9.43	31.80	0.80	31.66	0.34
0.8	50	30	15	32.53	34.83	6.60	33.11	1.73	32.56	0.08
0.8	50	50	15	33.22	34.83	4.61	33.81	1.73	33.26	0.10
2.4	10	10	25	49.62	51.76	4.12	50.93	2.57	49.78	0.31
2.4	10	30	25	51.12	51.76	1.24	51.66	1.05	51.16	0.09
2.4	10	50	25	51.53	51.76	0.44	51.73	0.39	51.55	0.04
2.4	30	10	25	52.70	57.14	7.77	54.32	2.97	53.08	0.71
2.4	30	30	25	54.72	57.14	4.23	56.31	2.82	54.81	0.15
2.4	30	50	25	55.62	57.14	2.67	56.81	2.09	55.67	0.09
2.4	50	10	25	53.88	59.51	9.47	55.53	2.97	54.39	0.93
2.4	50	30	25	55.93	59.51	6.03	57.82	3.27	56.11	0.32
2.4	50	50	25	56.93	59.51	4.34	58.60	2.84	57.02	0.17
4.0	10	10	32	63.68	66.82	4.71	66.02	3.55	64.09	0.65
4.0	10	30	32	65.86	66.82	1.44	66.76	1.34	65.93	0.10
4.0	10	50	32	66.47	66.82	0.53	66.82	0.52	66.49	0.03
4.0	30	10	32	67.44	73.55	8.31	70.52	4.36	68.44	1.46
4.0	30	30	32	70.40	73.55	4.28	72.81	3.31	70.60	0.28
4.0	30	50	32	71.63	73.55	2.61	73.33	2.31	71.74	0.15
4.0	50	10	32	69.07	76.38	9.57	72.03	4.11	70.15	1.54
4.0	50	30	32	71.91	76.38	5.85	74.84	3.91	72.23	0.44
4.0	50	50	32	73.29	76.38	4.04	75.68	3.16	73.48	0.26

Table A.3
Large network results: identical locations.

No.	λ	Q	R	No tran.	R	Reactive	Saving (%)	R	Enhanced	Saving (%)	Improve (%)
5	1.4	19	6	109.70	3	98.66	10.06	2	95.70	12.76	3.00
5	2.2	24	9	136.79	6	125.13	8.53	4	119.55	12.61	4.46
5	3.0	28	13	159.31	10	147.22	7.59	6	139.42	12.49	5.29
5	3.8	31	16	178.90	13	167.06	6.62	9	157.43	12.00	5.76
5	4.6	34	19	196.75	16	184.77	6.09	12	173.50	11.81	6.09
10	1.4	19	6	219.54	3	195.23	11.08	1	186.30	15.14	4.57
10	2.2	24	9	274.11	6	247.89	9.56	3	232.83	15.06	6.08
10	3.0	28	13	319.13	9	292.21	8.44	5	272.45	14.63	6.76
10	3.8	31	16	358.32	12	332.20	7.29	7	307.81	14.10	7.34
10	4.6	34	19	394.16	16	367.49	6.77	10	339.12	13.96	7.72
20	1.4	19	6	438.92	3	389.24	11.32	1	365.30	16.77	6.15
20	2.2	24	9	548.17	6	493.83	9.91	3	458.09	16.43	7.24
20	3.0	28	13	638.43	9	582.55	8.75	5	535.93	16.06	8.00
20	3.8	31	16	717.40	12	661.97	7.73	7	606.33	15.48	8.41
20	4.6	34	19	788.43	16	733.04	7.03	9	668.19	15.25	8.85

Table A.4
Large network results: two tier locations.

No.	λ	Q	R	No tran.	R	Reactive	Saving (%)	R	Enhanced	Saving (%)	Improve (%)
5	1.0–2.0	16–23	4–9	107.90	1–5	97.09	10.02	1–4	94.09	12.80	3.09
5	1.8–2.8	22–27	8–12	135.84	4–6	124.09	8.65	3–6	118.75	12.58	4.30
5	2.6–3.6	26–30	11–15	159.09	8–10	146.93	7.64	5–8	139.26	12.46	5.22
5	3.4–4.4	30–34	14–18	178.51	11–13	166.35	6.81	8–12	156.69	12.22	5.80
5	4.2–5.2	33–37	17–21	196.36	14–16	184.15	6.22	10–14	172.65	12.07	6.24
10	1.0–2.0	16–23	4–9	216.45	1–3	192.37	11.13	0–3	183.66	15.15	4.53
10	1.8–2.8	22–27	8–12	272.49	4–7	246.42	9.57	2–5	231.84	14.92	5.91
10	2.6–3.6	26–30	11–15	318.51	7–10	291.95	8.34	4–7	271.97	14.61	6.84
10	3.4–4.4	30–34	14–18	357.89	11–13	330.70	7.60	6–9	305.60	14.61	7.59
10	4.2–5.2	33–37	17–21	393.22	14–16	366.33	6.84	8–11	337.39	14.20	7.90
20	1.0–2.0	16–23	4–9	433.49	2–4	383.29	11.58	0–2	359.18	17.14	6.29
20	1.8–2.8	22–27	8–12	544.63	5–6	490.86	9.87	1–3	455.33	16.40	7.24
20	2.6–3.6	26–30	11–15	637.43	8–10	581.32	8.80	4–7	536.26	15.87	7.75
20	3.4–4.4	30–34	14–18	715.42	11–13	659.18	7.86	5–8	602.59	15.77	8.58
20	4.2–5.2	33–37	17–21	786.75	14–16	730.31	7.17	8–11	665.03	15.47	8.94

Appendix B. Calculating $\alpha_i(X_i)$

If t_s is the time when the s th unit becomes available at location i then at each point in time a location's state can be described by a variable X_i , where

$$X_i = (IP_i, t_{IP}, t_{IP-1}, \dots, t_1 - (IL_i)^-). \tag{B.1}$$

For location i the pdf and cdf of the distribution of the time when the n th customer arrival instant occurs can be respectively given as:

$$g_i^n(t) = \frac{\lambda_i^n t^{n-1} e^{-\lambda_i t}}{(n-1)!}, \tag{B.2}$$

$$G_i^n(t) = 1 - \sum_{k=1}^{n-1} \frac{(\lambda_i t)^k}{k!} e^{-\lambda_i t}. \tag{B.3}$$

Using these distributions it is possible to obtain the pdf and cdf of the time when the j th unit is demanded at location i :

$$r_i^j(t) = \sum_{n=1}^j P_{ij}^n g_i^n(t), \tag{B.4}$$

$$R_i^j(t) = \sum_{n=1}^j P_{ij}^n G_i^n(t), \tag{B.5}$$

It is now possible to calculate $\alpha_i(X_i)$. We let $x_i(s, t_s)$ be the expected holding and backorder costs, associated with the s th item of stock demanded during the lead time L_i . If $s \leq 0$ then the item has already been demanded, with $s=0$ the most recently demanded item. This gives four specific cases. The first two being:

For $s \leq 0, t_s \leq L_i$,

$$x_i(s, t_s) = b_i t_s. \tag{B.6}$$

For $s \leq 0, t_s > L_i$,

$$x_i(s, t_s) = b_i L_i. \tag{B.7}$$

We now define quantity $U_i^j(t)$ as

$$U_i^j(t) = \int_0^t r_i^j(u) u du = \sum_{n=1}^j P_{ij}^n G_i^{n+1}(t) \frac{n}{\lambda_i}. \tag{B.8}$$

This allows us to express the other two cases as:

For $s > 0, t_s \leq L_i$,

$$\begin{aligned} x_i(s, t_s) &= h_i \left(\int_{t_s}^{L_i} r_i^s(u) (u - t_s) du + (1 - R_i^s(L_i)) (L_i - t_s) \right) \\ &\quad + b_i \int_0^{t_s} r_i^s(u) (t_s - u) du, \\ &= h_i [U_i^s(L_i) - U_i^s(t_s) - t_s(1 - R_i^s(t_s))] \\ &\quad + L_i(1 - R_i^s(L_i)) + b_i [R_i^s(t_s)t_s - U_i^s(t_s)]. \end{aligned} \tag{B.9}$$

For $s > 0, t_s > L_i$,

$$x_i(s, t_s) = b_i [R_i^s(L_i)L_i - U_i^s(L_i)]. \quad (\text{B.10})$$

We calculate $\alpha_i(X_i)$ using

$$\alpha_i(X_i) = \sum_{s=1-(IL_i)^-}^{IP} x_i(s, t_s) + \sum_{s=IP+1}^{\infty} x_i(s, L_i) - L_i C_i, \quad (\text{B.11})$$

where $L_i C_i$ is the steady state cost incurred during the lead time period. The first summation calculates the expected cost over the period L_i associated with each unit of inventory either in stock, on order or already demanded. Similarly the second summation calculates the expected cost for any potential demand where the corresponding item of stock has not yet been ordered from the supplier. The final term subtracts the steady state cost for the given period of time so that overall $\alpha_i(X_i)$ gives the bias over period L_i .

References

- Agrawal, V., Chao, X., Seshadri, S., 2004. Dynamic balancing of inventory in supply chains. *European Journal of Operational Research* 159, 296–317.
- Allen, S.G., 1958. Redistribution of total stock over several user locations. *Naval Research Logistics Quarterly* 5 (4), 51–59.
- Archibald, T.W., Sassen, S.A.E., Thomas, L.C., 1997. An optimal policy for a two depot inventory problem with stock transfer. *Management Science* 43 (2), 173–183.
- Archibald, T.W., 2007. Modelling replenishment and transshipment decisions in periodic review multilocation inventory systems. *Journal of the Operational Research Society* 58 (7), 948–956.
- Archibald, T.W., Black, D., Glazebrook, K.D., 2009. An index heuristic for transshipment decisions in multi-location inventory systems based on a pairwise decomposition. *European Journal of Operational Research* 192, 69–78.
- Archibald, T.W., Black, D., Glazebrook, K.D., 2010. The use of simple calibrations of individual locations in making transshipment decisions in a multi-location inventory network. *Journal of the Operational Research Society* 61 (2), 294–305.
- Axsäter, S., 1990. Modelling emergency lateral transshipments in inventory systems. *Management Science* 36 (11), 1329–1338.
- Axsäter, S., 2003. A new decision rule for lateral transshipments in inventory systems. *Management Science* 49 (9), 1168–1179.
- Axsäter, S., Howard, C., Marklund, J., 2010. A Distribution Inventory Model with Transshipments from a Support Warehouse. Working Paper, Lund University.
- Bertrand, L.P., Bookbinder, J.H., 1998. Stock redistribution in two-echelon logistics system. *Journal of the Operational Research Society* 49, 966–975.
- Dada, M., 1992. A two-echelon inventory system with priority shipments. *Management Science* 38 (8), 1140–1153.
- Das, C., 1975. Supply and redistribution rules for two-location inventory systems: one period analysis. *Management Science* 21 (7), 765–776.
- Dong, L., Rudi, N., 2004. Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Science* 50 (5), 645–657.
- Evers, P.T., 2001. Heuristics for assessing emergency transshipments. *European Journal of Operational Research* 129 (2), 311–316.
- Gross, D., 1963. Centralized inventory control in multilocation supply systems. In: Scarf, H.E., Gilford, D.M., Shelly, M.W. (Eds.), *Multistage Inventory Models and Techniques*. Stanford University Press, Stanford, California, pp. 47–84.
- Jung, B., Sun, B., Kim, J., Ahn, S., 2003. Modeling lateral transshipments in multiechelon repairable-item inventory systems with finite repair channels. *Computers and Operations Research* 30, 1401–1417.
- Jönsson, H., Silver, E.A., 1987. Analysis of a two echelon inventory control system with complete redistribution. *Management Science* 33, 215–227.
- Kranenburg, A.A., van Houtum, G.J., 2009. A new partial pooling structure for spare parts networks. *European Journal of Operational Research* 199, 908–921.
- Krishnan, K., Rao, V., 1965. Inventory control in N warehouses. *Journal of Industrial Engineering* XVI (3), 212–215.
- Kukreja, A., Schmidt, C.P., 2005. A model for lumpy demand parts in a multi-location inventory system with transshipments. *Computers and Operations Research* 32 (8), 2059–2075.
- Lee, H.L., 1987. A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science* 33 (10), 1302–1316.
- Minner, S., Silver, E.A., Robb, D.J., 2003. An improved heuristic for deciding on emergency transshipments. *European Journal of Operational Research* 148, 384–400.
- Paterson, C., Kiesmüller, G., Teunter, R., Glazebrook, K., 2011. Inventory models with lateral transshipments: a review. *European Journal of Operational Research* 210 (2), 125–136.
- Robinson, L.W., 1990. Optimal an approximate policies in multiperiod multilocation inventory models with transshipments. *Operations Research* 38 (2), 278–295.
- Rudi, N., Kapur, S., Pyke, D.F., 2001. A two-location inventory model with transshipment and local decision making. *Management Science* 47 (12), 1668–1680.
- Tijms, H.C., 1986. *Stochastic Modelling and Analysis: A Computational Approach*. Wiley, Chichester, New York.
- Sherbrooke, C.C., 1968. METRIC: a multi-echelon technique for recoverable item control. *Operations Research* 16, 122–141.
- Sherbrooke, C.C., 1986. VARI-METRIC: improved approximations for multi-indenture, multi-echelon availability models. *Operations Research* 34 (2), 311–319.
- Zhao, H., Ryan, J.K., Deshpande, V., 2008. Optimal dynamic production and inventory transshipment policies for a two-location make-to-stock system. *Operations Research* 56 (2), 400–410.
- Wijk, A.C.C., Adan, I.J.B.F., van Houtum, G.J., 2009. Optimal Lateral Transshipment Policy for a Two Location Inventory Problem. Working Paper, EURANDOM.
- Wong, H., van Houtum, G.J., Cattrysse, D., van Oudheusden, D., 2006b. Multi-item spare parts systems with lateral transshipments and waiting time constraints. *European Journal of Operational Research* 171, 1071–1093.