

Production, Manufacturing and Logistics

An index heuristic for transshipment decisions in multi-location inventory systems based on a pairwise decomposition

T.W. Archibald ^{a,*}, D. Black ^a, K.D. Glazebrook ^b

^a *University of Edinburgh Management School, 50 George Square, Edinburgh EH8 9JY, UK*

^b *Lancaster University Management School, Lancaster, UK*

Received 26 May 2006; accepted 1 September 2007

Available online 19 September 2007

Abstract

In multi-location inventory systems, transshipments are often used to improve customer service and reduce cost. Determining optimal transshipment policies for such systems involves a complex optimisation problem that is only tractable for systems with few locations. Consequently simple heuristic transshipment policies are often applied in practice. This paper develops an approximate solution method which applies decomposition to reduce a Markov decision process model of a multi-location inventory system into a number of models involving only two locations. The value functions from the subproblems are used to estimate the fair charge for the inventory provided in a transshipment. This estimate of the fair charge is used as the decision criterion in a heuristic transshipment policy for the multi-location system. A numerical study shows that the proposed heuristic can deliver considerable cost savings compared to the simple heuristics often used in practice.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Dynamic programming; Inventory; Stochastic models; Transshipment; Decomposition

1. Introduction

The transshipment of inventory between locations in the same echelon of a multi-location inventory system is a strategy that is often used to improve customer service while attempting to control cost. Although many models of transshipments have been proposed (see for example the reviews by [Burton and Banerjee, 2005](#); [Kukreja et al., 2001](#)), past research does not completely satisfy the modern requirements of the retail industry. Often this is because the models proposed limit the number of locations (e.g. [Archibald et al., 1997](#); [Rudi et al., 2001](#)), impose restrictions on the timing and frequency of transshipments (e.g. [Jönsson and Silver, 1987](#); [Tagaras and Cohen, 1992](#); [Burton and Banerjee, 2005](#)), or require continuous review (e.g. [Axsäter, 2003](#); [Minner et al., 2003](#)), and possibly one-for-one (e.g. [Lee, 1987](#); [Kukreja et al., 2001](#); [Wong et al., 2005](#)), replenishment policies.

The models analysed in this paper have been developed in conjunction with a tyre retailer that has a network of 50 service depots covering Scotland and Northern England. The company uses periodic review replenishment policies and considers the use of transshipments any time a stockout (i.e. demand at a location that requires more items than are held there at the time) arises between review epochs. We believe that this scenario is quite common in the retail industry, and

* Corresponding author. Tel.: +44 131 650 4604; fax: +44 131 668 3053.

E-mail address: tarchibald@ed.ac.uk (T.W. Archibald).

hence that the models developed in this paper provide practical insight on the use of transshipments within many multi-location inventory systems.

In the absence of suitable models, companies have typically employed simple heuristics (Tagaras and Cohen, 1992) such as:

- no pooling – where transshipments are never used,
- complete pooling – where transshipments are used to satisfy stockouts whenever feasible,
- partial pooling – where a location will only agree to transship while its inventory level is above a fixed threshold.

However as the following examples show, these heuristics do not capture the full complexity of the transshipment decision. If the inventory level at a location is high relative to the expected demand at the location in the remainder of a period, it is unlikely the location will need to call upon all its inventory to meet local demand before the next replenishment. In such cases agreeing to a transshipment request would have the advantage of reducing holding cost as well as improving customer service. On the other hand if the inventory level at a location is relatively low, it may be almost certain that simply meeting local demand will exhaust the location's inventory before the next replenishment. Transshipment of an item from such a location in response to a stockout would then effectively move the location of the stockout as well as the item, and any reduction in holding cost (as the item would be used earlier in the period) is likely to be dominated by the transshipment cost. These examples suggest that the most effective transshipment policies are likely to be those in which the transshipment decision depends on the distribution of inventory held in the system and the time until the next replenishment. The aim of this paper is to develop a practical methodology to determine effective transshipment policies with these characteristics, and so help to bridge the gap between the performance of the simple heuristics and optimal performance. In a recent paper, Archibald (2007) uses a decomposition method to develop a structured heuristic transshipment policy for a periodic review multi-location inventory system with transshipment in response to stockouts. While our approach uses a similar decomposition based on pairs of locations in the system, the two methods tackle the transshipment decision in completely different ways. Archibald (2007) creates three structured heuristic transshipment policies by exploiting the structure of the optimal transshipment policies for the two-location subproblems. These heuristics use the assumption that a location facing a stockout will request transshipments from other locations in the system in order of increasing transshipment cost. Our method uses the value functions of the two-location subproblems to estimate a fair charge for inventory provided in a transshipment and uses this to inform the transshipment decision. Hence, the resulting heuristic attempts to allow for the transshipment cost and the expected value of the inventory used. The different methods yield very different transshipment policies. An advantage of the approach of Archibald (2007) is that the heuristic transshipment policies have a simple structure. However, our computational results show that for systems with many locations the performance of these simple heuristics can be significantly poorer than the heuristic we develop in this paper.

Section 2 presents a detailed description of the problem we consider, as well as a discussion of the modelling assumptions and a Markov decision process (MDP) model of the problem. While it is straightforward to model the problem as a MDP, the model is not a practical tool to support management decisions in multi-location inventory systems in general. This is due to the dimension of the state space (and some action spaces) which is equal to the number of locations in the system, making it impractical to solve the model for systems with more than a few locations. The challenge, therefore, is to develop an approximate solution method for this model which is practical for large systems and yields policies that are close to optimal.

In Section 3 we develop an approximate solution method for the MDP based on a “pairwise” decomposition. More precisely, we solve one two-location problem for every possible pair of locations in the system and use the resulting value functions to construct an index to guide the transshipment decision. The index is related to the notion of a “fair charge” for the inventory provided by the transshipment. For an m -location system, this approach involves the solution of $m(m - 1)$ simple two-location problems and so is feasible for systems with many locations.

Section 4 presents the results of a simulation study comparing the performance of the proposed approach with that of a number of simple policies, including no pooling and complete pooling, and, in a limited number of cases, the optimal policy. These results show that our approach performs remarkably well in a wide range of problem instances. Finally in Section 5, we present our conclusions.

2. A model of a multi-location inventory system

Consider a periodic review multi-location inventory system with m locations. Periodic review systems are common in the retail industry where the time between review epochs may be fixed by the contract with the supplier. Let T be the time between successive review epochs and assume that the replenishment policy restores the inventory level at location k to S_k items at each review epoch. (Actually following the approach of Archibald et al. (1997), it is easy to show that this is an optimal replenishment policy for the model considered in this paper.) Fixed replenishment times allow the supplier

to plan for the replenishment of the system and, hence, offer a low unit cost for replenishment orders. We assume that replenishments are completed during intervals in which the locations are closed to customers (e.g. overnight or during weekends). Hence, the effective lead time for replenishment is zero. As the maximum number of items that can be stored, and hence ordered, at each location must be finite, there are a finite number of possible replenishment policies. Let M_k be the maximum number of items that can be stored at location k .

The occurrence of demand is modelled as a compound Poisson process with Bernoulli splitting. Let λ be the rate at which demand occurs in the system and ϕ_k be the probability that an instance of demand occurs at location k . The number of items required to satisfy an instance of demand is a random variable. Let $F(d)$, for $1 \leq d \leq D$, be the distribution function of the number of items required to satisfy a demand. For the purposes of calculation, we divide the period between successive review epochs into $N \geq \lambda T$ intervals of length $\delta = T/N$ and assume that there can be at most one demand in the system in any interval of length δ . Define $p_k = \lambda \delta \phi_k$ to be the probability that a demand occurs at location k during an interval of length δ and $p_0 = 1 - \lambda \delta$ to be the probability that there is no demand in the system in an interval of length δ .

When demand arises, a decision must be taken about how it is to be satisfied. We assume that demand will be met from local stock (i.e. inventory held at the location at which the demand occurs) whenever possible, as there can be no practical justification for any other action. When this is not possible, that part of the demand which cannot be satisfied from local stock is sourced either by a single transshipment from one other location or by an emergency order placed with the supplier. Hence there are at most m ways of satisfying a demand. Note that transshipments and emergency orders are only ever used to satisfy a shortage when a stockout has arisen. While in some cases transshipments and emergency orders in anticipation of stockouts may be worthwhile (e.g. when the cost of transshipment involves a high fixed cost), we believe that many companies only use these actions as a last resort and, in such cases, our assumption is appropriate. We also believe that most of the companies that use transshipments actually use a policy of complete pooling. One of our objectives is to investigate how much such a company might gain by implementing a more sophisticated transshipment policy.

When demand cannot be satisfied entirely from local stock, it would be possible to source the additional items required by a number of transshipments from different locations or by a combination of transshipment and emergency order. We do not consider this possibility for two reasons. From a practical point of view, such an approach is hard to coordinate and increases the chance of errors (e.g. the transshipment of the wrong number or type of items) leading to delays in satisfying the demand. Also regarding the complexity of the model, this approach greatly increases the size of the action spaces, as the number of dimensions increases from one to the number of locations with stock plus one in many cases.

Customers are assumed always to accept supply by a transshipment or emergency order. This assumption is only appropriate if the delay in satisfying the demand is insignificant and so we assume that lead times for transshipments and emergency orders are zero. As the supplier is unable to plan for emergency orders, the unit cost of such orders is much higher than the unit cost for replenishment orders. It is therefore impractical to use emergency orders to replenish locations between review epochs. The supplier is assumed to always be able to fully satisfy any practical replenishment or emergency order.

Replenishment costs C_k per item at location k , emergency orders cost $R_{E,k}$ per item at location k , transshipment of d items from location j to location k costs $R_{j,k}(d)$ and the holding cost rate per item per time unit at location k is h_k . The replenishment cost and emergency order cost rates depend on location to allow for the situation where the cost of delivery from the supplier depends on location. It is assumed that $R_{E,k} > C_k$, as otherwise it might be optimal to satisfy all demand at location k by emergency orders. Transshipments are often delivered by staff from the source location and the dominant cost is often the opportunity cost due to the staff being unavailable for the time required. In such cases the cost of transshipment will depend on the distance between the locations involved and will be largely independent of the number of items in the transshipment. In other cases a courier may be used and then the cost might be directly proportional to the number of items in the transshipment. Hence $R_{j,k}(1) \leq R_{j,k}(d) \leq dR_{j,k}(1)$. In all cases one would expect the transshipment cost to be symmetric, i.e. $R_{j,k}(d) = R_{k,j}(d)$. It is assumed that $dC_k < dC_j + R_{j,k}(d)$, as otherwise it might be optimal to order items for location k at location j and then transship them to location k . The main element of the holding cost is generally the cost of capital, so the holding cost can be expressed as a rate per item per time unit. Since the unit cost of replenishment depends on location, the holding cost rate will depend on location also.

Throughout the paper we use bold font to denote vectors and adopt the convention that u_k denotes the value in position k of a vector \mathbf{u} . We also use $\mathbf{0}$ to denote the zero vector and \mathbf{e}_k to denote the vector with 1 in position k and 0 elsewhere. Define $v^n(\mathbf{i})$ to be the minimum expected cost of satisfying demand arising in the inventory system when the time until the next replenishment is $n\delta$ and the inventory level at location k is i_k . Consider the costs that may be incurred in the next interval of length δ . We assume a holding cost of δh_k is incurred for each item currently held at location k . If a demand for $d \leq i_k$ items arrives at location k , this can be met from local stock at no extra cost. However a demand for $d > i_k$ items at location k will result in a shortage of $d - i_k$ items at that location. This shortage can be satisfied by an emergency order, at a cost of $(d - i_k)R_{E,k}$, or by a transshipment from any other location j satisfying $i_j \geq d - i_k$, at a cost of $R_{j,k}(d - i_k)$. Using a standard dynamic programming approach (Puterman, 1994), it can be shown that $v^n(\mathbf{i})$, where $1 \leq n \leq N$ and $\mathbf{0} \leq \mathbf{i} \leq \mathbf{M}$, is the value function of a finite state, finite action, finite horizon MDP with the following optimality equation:

$$v^n(\mathbf{i}) = \sum_{k=1}^m i_k \delta h_k + p_0 v^{n-1}(\mathbf{i}) + \sum_{k=1}^m p_k \left(\sum_{d=1}^{i_k} F(d) v^{n-1}(\mathbf{i} - d\mathbf{e}_k) + \sum_{d=i_k+1}^D F(d) \min \left\{ (d - i_k) R_{E,k} + v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k) \min_{j \neq k: i_j \geq d - i_k} (R_{j,k}(d - i_k) + v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k - (d - i_k) \mathbf{e}_j)) \right\} \right). \tag{1}$$

As the replenishment policy restores the inventory at location k to the same level at each review epoch, each item of inventory remaining at location k at a review epoch will reduce the cost of a replenishment order by C_k . Hence $v^0(\mathbf{i}) = -\sum_{k=1}^m i_k C_k$. An optimal replenishment policy is given by any \mathbf{S} , satisfying $\mathbf{0} \leq \mathbf{S} \leq \mathbf{M}$, that minimises:

$$\sum_{k=1}^m S_k C_k + v^N(\mathbf{S}). \tag{2}$$

3. An index heuristic for transshipment based on a pairwise decomposition of the system

We develop a method to approximate a fair charge for the inventory provided in a transshipment between two locations in the system. There are two elements to this charge: the immediate cost of moving the inventory and the change in the expected future cost as a result of having less inventory at the location at the source of the transshipment. The method of approximation only considers demand arising at the two locations in question, and assumes that all such demand must be met from local stock, by an emergency order or by a transshipment between the two locations.

Consider the model of an m -location inventory system developed in Section 2 and suppose a demand for d items has arisen at location k when the time until the next replenishment is nd and the inventory levels at the locations are given by \mathbf{i} . Suppose further that $d > i_k$. Note that under the assumptions of the model, the demand that has arisen will exhaust the inventory at location k regardless of the decision taken. Assume for some location $j \neq k$, $i_j \geq d - i_k$, so that it is possible to satisfy the shortage at location k by transshipment from location j . The immediate cost of using transshipment from location j to satisfy the shortage at location k is $R_{j,k}(d - i_k)$. However this ignores the fact that, compared to using an emergency order to satisfy the shortage at location k , there will be $d - i_k$ fewer items available at location j to meet demand arising in the system before the next replenishment. Hence the true cost to the system of using transshipment from location j to meet the shortage at location k is

$$R_{j,k}(d - i_k) + v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k - (d - i_k) \mathbf{e}_j) - v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k). \tag{3}$$

Dividing (3) by $d - i_k$ gives

$$\frac{R_{j,k}(d - i_k) + v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k - (d - i_k) \mathbf{e}_j) - v^{n-1}(\mathbf{i} - i_k \mathbf{e}_k)}{d - i_k}, \tag{4}$$

which can be interpreted as a per item fair charge for meeting the shortage at location k by transshipment from location j . The choice of optimal action in (1) can be determined from (4) as follows. The best location to source the transshipment from is the one with the smallest per item fair charge. However if the smallest per item fair charge is greater than $R_{E,k}$, it is better to satisfy the shortage at location k by an emergency order. When the number of locations, m , is large, it is not practical to use (4) to determine the optimal transshipment decision, because it involves the value function of a MDP with an m -dimensional state space. However we can still use this approach if we can find a good approximation to the fair charge for transshipment using low-dimensional value functions that are easy to calculate. Our proposal has echoes of an approach by Whittle (1988) to the development of index-based solutions to an intractable class of MDPs called restless bandits. Whittle’s indices also had an interpretation as fair charges for the decision options under consideration.

If it is optimal to use transshipment from location j to satisfy the current shortage at location k , one might expect that transshipment from location j will also be a good way of satisfying any further demand arising at location k before the next replenishment (which will inevitably result in further shortages). For example it could be that the cost of transshipment between the locations is relatively low or that location j has a relatively high inventory level. Hence one of the main consequences of using inventory from location j now will be having less inventory available to meet subsequent demand arising at locations j and k before the next replenishment. One way to approximate the knock on effect of transshipment is to assume that locations j and k operate independently of the other locations in the system and use an additive decomposition to approximate the value function as follows:

$$v^n(\mathbf{i}) = v^n(i_j, i_k) + v^n(\{i_\ell : 1 \leq \ell \leq m, \ell \neq j \text{ or } k\}). \tag{5}$$

With such an approximation,

$$v^n(\mathbf{i} - i_k \mathbf{e}_k - (d - i_k) \mathbf{e}_j) - v^n(\mathbf{i} - i_k \mathbf{e}_k) = v^n(i_j - (d - i_k), 0) - v^n(i_j, 0) \tag{6}$$

and the knock on effect of transshipment is approximated by the value function of a two-location inventory system in which one of the locations has no inventory. Hence this method of approximation only requires the solution of a one-dimensional dynamic programming problem for each pair of locations in the system.

Define $\omega_{j,k}^n(i)$ to be the minimum expected cost of satisfying the demand arising at locations j and k , without any transshipments between these locations and other locations in the system, when the time until the next replenishment is $n\delta$ and the inventory levels at locations j and k are i and 0 , respectively. This cost is not affected by demand that occurs at other locations, so we can define $p_0 = 1 - p_j - p_k$ and calculate $\omega_{j,k}^n(i)$ from (1) as follows:

$$\begin{aligned} \omega_{j,k}^n(i) = & i\delta h_j + (1 - p_j - p_k)\omega_{j,k}^{n-1}(i) + p_j \left(\sum_{d=1}^i F(d)\omega_{j,k}^{n-1}(i-d) + \sum_{d=i+1}^D F(d) \left((d-i)R_{E,k} + \omega_{j,k}^{n-1}(0) \right) \right) \\ & + p_k \left(\sum_{d=1}^i F(d) \min\{dR_{E,k} + \omega_{j,k}^{n-1}(i), R_{j,k}(d) + \omega_{j,k}^{n-1}(i-d)\} + \sum_{d=i+1}^D F(d) \left(dR_{E,k} + \omega_{j,k}^{n-1}(i) \right) \right) \end{aligned} \quad (7)$$

for $1 \leq n \leq N$ and $0 \leq i \leq M_j$, where $\omega_{j,k}^0(i) = -iC_j$. The per item fair charge for meeting the shortage at location k by transshipment from location j can then be approximated by

$$\frac{R_{j,k}(d - i_k) + \omega_{j,k}^{n-1}(i_j + i_k - d) - \omega_{j,k}^{n-1}(i_j)}{d - i_k} \quad (8)$$

This suggests the following heuristic transshipment policy which we refer to as the pairwise heuristic. Consider the m -location system and suppose a shortage of x items has arisen at location k when the time until the next replenishment is $n\delta$ and the inventory levels at the locations are given by \mathbf{i} . If

$$R_{E,k} < \min_{j \neq k: i_j \geq x} \left\{ \frac{R_{j,k}(x) + \omega_{j,k}^{n-1}(i_j - x) - \omega_{j,k}^{n-1}(i_j)}{x} \right\} \quad (9)$$

then the shortage at location k should be met by an emergency order, otherwise it should be met by the transshipment of x items from the location that minimises the RHS of (9). The pairwise heuristic requires the solution of $m(m - 1)$ MDPs each with a one-dimensional state space. As the computational complexity of the pairwise heuristic only increases as the square of the number of locations in the system, it is practical for systems with many locations.

Compared to the simple heuristic transshipment policies considered in this paper, the pairwise heuristic is dynamic (in the sense that decisions depend upon the time until the next replenishment) and is more sensitive to the distribution of inventory held in the system. For example often in implementations of complete or partial pooling, ties are broken by choosing the nearest (measured in terms of transshipment cost) location with sufficient inventory. Hence provided the nearest location to the stockout has enough inventory to be a candidate, it will always be chosen as the source of the transshipment, regardless of the inventory levels at other locations. However the location that minimises the RHS of (9), and hence the transshipment decision for the pairwise heuristic, could change if the inventory level at any location changes. The approximation to a fair charge for a transshipment from a location, Eq. (8), is non-increasing in the inventory level at that location. Hence with the pairwise heuristic, a location is generally more likely to be called upon to transship as the inventory level at the location increases.

The assumption of zero lead time for replenishment orders is important for the development of the pairwise heuristic, as review epochs are then renewal points of the process. However, the heuristic can still be applied to guide the transshipment decision even when the lead time for replenishment orders is non-zero. Assume that the lead time for replenishment orders is known and that, at each review epoch, location k orders the number of items required to restore its inventory level to \tilde{S}_k . Let L , satisfying $0 < L < T$, be the lead time for replenishment orders. Such a system can be approximated by a system in which the lead time for replenishment orders is zero and, at each review epoch, the inventory level at location k is restored to $S_k = \tilde{S}_k - \lambda L \phi_k$. Now the pairwise heuristic can be applied to determine how to meet a shortage of items at a location in the system as in (9). Although this approximation will work best when the lead time is short and the variance of demand is low, the heuristic might be expected to perform well relative to simple transshipment heuristics in many other circumstances as well.

4. Numerical investigation

We have conducted an extensive numerical investigation to compare the performance of the pairwise heuristic developed in Section 3 with that of complete pooling, no pooling, the α -heuristic of (Archibald 2007) and, for small problems, the optimal transshipment policy. This section presents a summary of our results, concentrating on the comparison between complete pooling and the pairwise heuristic. We have chosen this focus because we believe that complete pooling is the most commonly applied heuristic and therefore is a suitable benchmark for the new heuristic we have developed. We have

used the α -heuristic in our investigation, because Archibald (2007) concludes this is the best of the three heuristics developed in that paper.

In the tables below, NP, CP, α , PW and OPT denote the no pooling, complete pooling, α -heuristic, pairwise heuristic and optimal transshipment policies, respectively. We use G_H to denote the expected cost rate under transshipment policy H and define $\Delta_{H'}^H = 100(G_H - G_{H'})/G_{H'}$, the percentage difference in the expected cost rates under transshipment policies H and H' .

Due to the computational complexity of the model, optimal transshipment policies have not been determined for systems with more than three locations. For the same reason, simulation has been used to estimate the expected cost rate for all heuristic transshipment policies for systems with more than three locations. Our implementation of complete pooling chooses the source of a transshipment to be the “nearest” location with sufficient inventory. More precisely, if i describes the inventory levels at the m locations in the system and location k has a shortage of x items, transshipment from the location j that minimises $R_{j,k}(x)$ subject to $i_j \geq x$ is used to meet the shortage at location k . We have found that the pairwise heuristic outperforms the α -heuristic in all cases. However, for three-location inventory systems, the difference in the expected cost rates for these heuristics can be small. For this reason, results for the α -heuristic are only included for larger problems.

Table 1 compares the performance of the heuristic policies and the optimal policy for an example of a three-location system with a range of emergency order costs. As one would expect, relative to the optimal policy, the performance of complete pooling improves and the performance of no pooling deteriorates as the cost of emergency orders increases. However it is worth noting that even when an emergency order costs less than any transshipment, no pooling is not optimal and even when an emergency order costs more than any transshipment, complete pooling is not optimal. In these extreme cases, more efficient use of transshipments reduces the cost rate by up to 3% for this example.

In Table 1, the expected cost rate under the pairwise heuristic is always within 0.5% of optimal and the pairwise heuristic is clearly the best of the heuristics considered. As one might expect, the savings from more efficient use of transshipments are greatest when the emergency order cost lies in the range of the transshipment costs. When $R_{E,k} = 25$, the expected cost rates for complete pooling and no pooling are, respectively, 5.6% and 7.7% higher than the expected cost rate for the pairwise heuristic.

It is just possible to see from Table 1 that the performance of the pairwise heuristic slowly deteriorates as the cost of emergency orders increases. This appears to be a general trend and, for extremely high emergency order costs ($R_{E,k} > 200$ for this example), the pairwise heuristic is actually slightly inferior to complete pooling. The pairwise heuristic does not consider the inventory held in the system as a whole and so tends to over-estimate the value of inventory held at a location. As a result the pairwise heuristic tends to use emergency orders more frequently than optimal. As the cost of emergency orders increases, the effect of this becomes more significant. However for reasonable emergency order costs, the pairwise heuristic yields a robust transshipment policy that delivers near optimal performance. Further evidence of this is provided in Fig. 1 which compares the performance of the heuristics on 25 three-location inventory systems with different emergency order costs and transshipment costs. In all 25 problems, $N = 1000$, $D = 1$, $\lambda\phi_k = 20$, $M_k = 24$, $C_k = 0$ and $h_k = 1$. The emergency order costs $R_{E,k}$ are independent of k and chosen, independently for each problem, from the uniform $U(15,40)$ distribution. The transshipment costs $R_{j,k}(1)$ are symmetric and chosen, independently for each pair of locations and problem instance, from the uniform $U(10,30)$ distribution.

From Fig. 1 it is clear that the pairwise heuristic is far better than the other heuristics for these problems. The expected cost rate under the pairwise heuristic is always within 0.3% of the optimal expected cost rate, while the expected cost rates under complete pooling and no pooling can be as much as 22% and 26%, respectively, higher than the optimal expected cost rate. Although, for these problems, no pooling is inferior to complete pooling in terms of worst case performance, the average performance of no pooling is slightly better than that of complete pooling. This is due to the range of emergency order cost used in this experiment. When R_E is near the low end of the range, no pooling is close to optimal (the best case of

Table 1
Performance of the heuristics as the emergency order cost varies

$R_{E,k}$	G_{OPT}	Δ_{OPT}^{PW} (%)	Δ_{OPT}^{CP} (%)	Δ_{OPT}^{NP} (%)	Δ_{PW}^{CP} (%)
15	63.5	0.0	18.4	0.9	18.4
20	70.4	0.0	8.2	1.4	8.2
25	73.1	0.0	5.6	7.7	5.6
30	75.2	0.1	4.0	14.4	3.9
35	76.5	0.2	3.5	22.0	3.3
40	77.7	0.4	3.2	29.6	2.9

Problem parameters: $m = 3$, $N = 1000$, $D = 1$, $\lambda\phi_k = 20$, $C_k = 0$, $h_k = 1$, $M_k = 24$, $R_{1,2}(1) = 19.5$, $R_{1,3}(1) = 30.0$, $R_{2,3}(1) = 20.5$ and $R_{i,j}(1) = R_{j,i}(1)$.

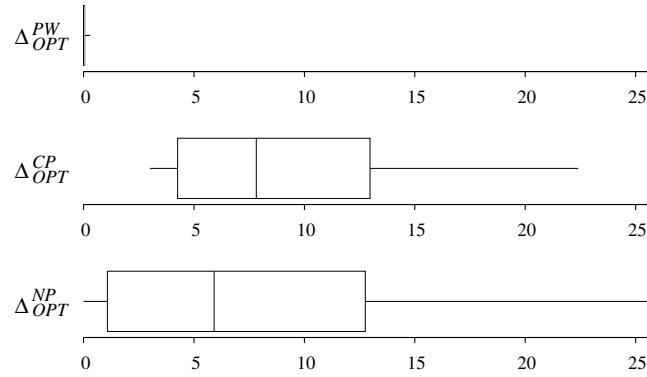


Fig. 1. Box and whisker plot showing the maximum, minimum, medium and interquartile range of Δ_{OPT}^{PW} , Δ_{OPT}^{CP} and Δ_{OPT}^{NP} for 25 three-location inventory systems with $D = 1$.

no pooling is within 0.01% of optimal). However when R_E is near the high end of the range, complete pooling is still some way from optimal (complete pooling is never within 3% of optimal).

To investigate the effect of batch demand on the performance of the heuristics, we used 25 three-location inventory systems with the same characteristics as those above except for the demand rates, the number of items required to satisfy a demand and the transshipment costs. In all cases the maximum number of items required to satisfy a demand, D , is 4 and the distribution of the number of items required to satisfy a demand, $F(d)$, is 0.4, 0.3, 0.2 and 0.1 for d equal to 1, 2, 3 and 4, respectively. Hence the average number of items required to satisfy a demand is 2. We then chose $\lambda\phi_k = 10$, so that the average demand at each location between successive review epochs is 20 as before. We consider two transshipment cost structures for each problem. Firstly we use $R_{j,k}(d) = Xd$ and then we use $R_{j,k}(d) = X$ where in both cases X is chosen, independently for each pair of locations and problem instance, from the uniform $U(10,30)$ distribution. In the first case, transshipment cost is proportional to the size of the transshipment and in the other case it is independent of the size of the transshipment. Figs. 2 and 3 compare the performance of the heuristics on these problems. Again it is clear that the pairwise heuristic outperforms the other heuristics. The expected cost rate under the pairwise heuristic is never more than 2% higher than the optimal expected cost rate indicating that the heuristic is robust to changes in the structure of the transshipment costs. As one might expect, complete pooling outperforms no pooling when the transshipment cost is independent of the transshipment size and the opposite is true when the transshipment cost is proportional to transshipment size. Hence, not only is the performance of these simple heuristics far inferior to that of the pairwise heuristic, it is also highly sensitive to the structure of the problem.

Having demonstrated that the pairwise heuristic can give close to optimal performance for small inventory systems, we turn our attention to larger systems for which it is not practical to compute an optimal transshipment policy due to the high dimensionality of the MDP formulation. For such large systems, we estimate the expected cost rates for the heuristics by simulating 1000 instances of the interval between successive review epochs. Tables 2 and 3 compare the performance of the heuristics on a set of 20-location inventory systems with two different transshipment cost structures. In Table 2 transshipment cost is proportional to transshipment size and in Table 3 it is independent of transshipment size. As in all previous problems, $N = 1000$, $C_k = 0$ and $h_k = 1$. Each problem allows batch demand, with $D = 4$ and $F(d) = 0.4, 0.3, 0.2$ and 0.1 for $d = 1, 2, 3$ and 4, respectively. Hence the average number of items required to satisfy a demand, \bar{D} , is 2. The emergency




Fig. 2. Box and whisker plot showing the maximum, minimum, medium and interquartile range of Δ_{OPT}^{PW} , Δ_{OPT}^{CP} and Δ_{OPT}^{NP} for 25 three-location inventory systems with $D = 4$ and transshipment cost proportional to transshipment size.




Fig. 3. Box and whisker plot showing the maximum, minimum, medium and interquartile range of Δ_{OPT}^{PW} , Δ_{OPT}^{CP} and Δ_{OPT}^{NP} for 25 three-location inventory systems with $D = 4$ and transshipment cost independent of transshipment size.

Table 2

Performance of transshipment policies as the cost of emergency orders varies for 100 20-location inventory systems with $D = 4$ and transshipment cost proportional to transshipment size

$R_{E,k}$	G_{PW}		G_{CP}		G_{NP}		G_z		Δ_{PW}^{CP} (%)
	AVE	SE	AVE	SE	AVE	SE	AVE	SE	
45	950	1.4	1005	1.7	1410	2.0	972	1.5	5.8
50	957	1.5	1006	1.7	1535	2.2	983	1.5	5.1
55	962	1.5	1006	1.7	1660	2.4	992	1.6	4.6
60	966	1.5	1007	1.7	1786	2.7	999	1.6	4.2
65	969	1.6	1007	1.7	1911	2.9	1005	1.7	4.0
70	971	1.6	1008	1.7	2037	3.1	1010	1.7	3.7
75	974	1.6	1008	1.7	2162	3.3	1015	1.7	3.5
80	976	1.6	1009	1.8	2288	3.5	1020	1.7	3.4
85	978	1.6	1010	1.8	2413	3.8	1024	1.8	3.2
90	981	1.7	1010	1.8	2539	4.0	1028	1.8	3.0

Notation: AVE denotes average cost and SE denotes standard error.

Table 3

Performance of transshipment policies as the cost of emergency orders varies for 100 20-location inventory systems with $D = 4$ and transshipment cost independent of transshipment size

$R_{E,k}$	G_{PW}		G_{CP}		G_{NP}		G_z		Δ_{PW}^{CP} (%)
	AVE	SE	AVE	SE	AVE	SE	AVE	SE	
45	628	0.8	649	0.8	1410	2.0	659	0.8	3.3
50	631	0.8	649	0.8	1535	2.2	664	0.9	3.0
55	633	0.8	650	0.9	1660	2.4	668	0.9	2.7
60	635	0.8	650	0.9	1786	2.7	673	0.9	2.5
65	637	0.8	651	0.9	1911	2.9	677	1.0	2.3
70	638	0.9	652	0.9	2037	3.1	681	1.0	2.0
75	640	0.9	652	0.9	2162	3.3	686	1.0	1.8
80	642	0.9	653	0.9	2288	3.5	690	1.0	1.6
85	644	0.9	653	0.9	2413	3.8	694	1.1	1.4
90	646	0.9	654	1.0	2539	4.0	699	1.1	1.2

Notation: AVE denotes average cost and SE denotes standard error.

order cost $R_{E,k}$ is assumed to be independent of k and, to examine the effect of increasing $R_{E,k}$, we consider 10 values for $R_{E,k}$ in the range 45–90. For each emergency order cost, we consider 100 cases by randomly generating 10 sets of demand rates and 10 sets of coordinates for the locations. The demand rates for the locations, $\lambda\phi_k$, are chosen independently from the uniform $U(5,15)$ distribution and each location is positioned in the unit square independently according to a uniform distribution. We assume that the replenishment policy restores the inventory level at location k to $\lambda\phi_k\bar{D} + \sqrt{\lambda\phi_k\bar{D}}$ at each review epoch. Hence each location has enough inventory to meet the expected demand at that location between successive review epochs plus a safety stock of $\sqrt{\lambda\phi_k\bar{D}}$ items. The transshipment costs are symmetric and defined as a function of $e_{j,k}$,

Table 4

Performance of pairwise heuristic and complete pooling as the lead time for replenishment orders varies for 100 20-location inventory systems with $D = 4$ and transshipment cost proportional to transshipment size

Lead time L	G_{PW}		G_{CP}		Δ_{PW}^{CP} (%)
	AVE	SE	AVE	SE	
0.05 T	1198	0.3	1229	0.3	2.6
0.10 T	1336	0.3	1382	0.4	3.4
0.15 T	1498	0.4	1565	0.4	4.4
0.20 T	1681	0.4	1773	0.5	5.5
0.25 T	2002	0.5	2147	0.6	7.3

Notation: AVE denotes average cost and SE denotes standard error.

the Euclidean distance between locations j and k . For Table 2, $R_{j,k}(d) = (10 + 70e_{j,k})d$ and, for Table 3, $R_{j,k}(d) = 10 + 70e_{j,k}$.

It is apparent from both tables that the expected cost rate under complete pooling is largely unaffected by the increase in the emergency order cost. This indicates that there is almost enough inventory in the system as a whole to satisfy all demand arising between successive review epochs and so, with complete pooling, emergency orders are rarely required. However it is also apparent from the tables that the planned use of emergency orders can result in significant savings in expected cost rate. When transshipment cost is proportional to transshipment size (Table 2), the expected cost rate for complete pooling is between 6%, when $R_{E,k} = 45$, and 3%, when $R_{E,k} = 90$, higher than the expected cost rate for the pairwise heuristic. As before, the advantage of the pairwise heuristic over complete pooling decreases as the emergency order cost increases. However, the average per item transshipment cost for these problems is approximately 45, so the pairwise heuristic still offers considerable savings when the emergency order cost is significantly higher than the average transshipment cost.

As expected, the relative performance of complete pooling improves when the transshipment cost is independent of the transshipment size (Table 3). However, even in this case, savings in the expected cost rate of between 3%, when $R_{E,k} = 45$, and 1%, when $R_{E,k} = 90$, are achieved due to the selective use of emergency orders advocated by the pairwise heuristic.

The expected cost rate for the α -heuristic is always at least 2% higher than that for the pairwise heuristic for the range of emergency order costs shown in Tables 2 and 3, and in the worst case is 8% higher. Relative to the pairwise heuristic, the performance of the α -heuristic appears to deteriorate as the emergency order cost increases and as the average transshipment cost decreases. Our results demonstrate that the pairwise heuristic can deliver substantial savings in recurrent costs compared to the α -heuristic across a wide range of problem characteristics.

We note that no pooling is never competitive with the other heuristics for any of these problems. Even when $R_{E,k} = 45$ and transshipment cost is proportional to transshipment size, the expected cost rate for no pooling is more than 48% higher than for the pairwise heuristic and more than 40% higher than for complete pooling. These examples clearly illustrate the importance of transshipments to the management of inventory in a retail network.

Table 4 compares the performance of the pairwise heuristic and complete pooling on a set of 20-location inventory systems when the lead time for replenishment orders is non-zero. Lead times of between 5% and 25% of the time between successive review epochs are considered. In the test problems, $R_{E,k} = 70$ and all other problem parameters are the same as in the problems used for Table 2. To determine a heuristic transshipment policy using the pairwise heuristic, each problem is approximated by a problem with zero replenishment lead time as explained at the end of Section 3. Simulation is used to compare the performance of this policy with that of complete pooling for the problem with non-zero replenishment lead time. As replenishment lead time increases, the relative advantage of the pairwise heuristic appears to increase. This suggests that the planned use of emergency orders is even more important for non-zero replenishment lead time and that the pairwise heuristic is still a useful technique in such cases.

5. Conclusions

We have proposed a Markov decision process model of a multi-location inventory system that addresses the issues involved in the management of a large network of retail outlets operating in a highly competitive environment. It is not practical to use the model to determine the optimal transshipment policy due to the high dimension of the state space. Instead we have developed a decomposition approach that constructs two subproblems for each pair of locations in the original problem. These two subproblems correspond to the two possible directions of transshipments between the locations. The state space in each subproblem has only one dimension and so can be solved easily. The solution to a subproblem provides an approximation to a fair charge for the inventory supplied in a transshipment between the two locations in the original model. This approximate fair charge is then used to rank the possible ways of meeting a shortage at a

location in the original model in a heuristic transshipment policy. Our proposed heuristic is therefore simple to compute and based on an easily understood criterion. We have compared the performance of this heuristic with optimal transshipment policies for systems with three locations and simple heuristic policies for systems with 20 locations. We have considered a wide range of system characteristics including non-zero replenishment lead time. These results suggest that our proposed heuristic provides a robust solution to the problem of organising transshipments in large retail networks and can result in considerable cost savings compared to the frequently used policies of no pooling and complete pooling.

Acknowledgements

The authors gratefully acknowledge the support of EPSRC Grant GR/T08562/01 and thank the referees for helpful comments and suggestions.

References

- Archibald, T.W., 2007. Modelling replenishment and transshipment decisions in periodic review multilocation inventory systems. *Journal of the Operational Research Society* 58, 948–956.
- 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, 173–183.
- Axsäter, S., 2003. A new decision rule for lateral transshipments in inventory systems. *Management Science* 49, 1168–1179.
- Burton, J., Banerjee, A., 2005. Cost-parametric analysis of lateral transshipment policies in two-echelon supply chains. *International Journal of Production Economics* 93–94, 169–178.
- Jönsson, H., Silver, E.A., 1987. Analysis of a two echelon inventory control system with complete redistribution. *Management Science* 33, 215–227.
- Kukreja, A., Schmidt, C.P., Miller, D.M., 2001. Stocking decisions for low usage items in a multilocation inventory system. *Management Science* 47, 1371–1383.
- Lee, H.L., 1987. A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science* 33, 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.
- Puterman, M.L., 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley, New York.
- Rudi, N., Kapur, S., Pyke, D.F., 2001. A two-location inventory model with transshipment and local decision making. *Management Science* 47, 1668–1680.
- Tagaras, G., Cohen, M.A., 1992. Pooling in two location inventory systems with non-negligible lead times. *Management Science* 38, 1067–1083.
- Whittle, P., 1988. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability* A25, 87–298.
- Wong, H., Cattrysse, D., Van Oudheusden, D., 2005. Inventory pooling of repairable spare parts with non-zero lateral transshipment time and delayed lateral transshipments. *European Journal of Operational Research* 165, 207–218.