THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

A semi-Markov model with holdout transshipment policy and phase-type exponential lead time

Citation for published version:

Archibald, T & Zhang, J 2011, 'A semi-Markov model with holdout transshipment policy and phase-type exponential lead time', *European Journal of Operational Research*, vol. 211, no. 3, pp. 650-657.
<https://doi.org/10.1016/j.ejor.2011.01.001>

Digital Object Identifier (DOI):

[10.1016/j.ejor.2011.01.001](https://doi.org/10.1016/j.ejor.2011.01.001)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

European Journal of Operational Research

Publisher Rights Statement:

This is an Author's Accepted Manuscript of the following article: Archibald, T. & Zhang, J. 16 Jun 2011, "A semi-Markov model with holdout transshipment policy and phase-type exponential lead time", in *European Journal of Operational Research*. 211, 3, p. 650-657. © Elsevier, the final publication is available at <http://dx.doi.org/10.1016/j.ejor.2011.01.001>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A semi-Markov model with holdout transshipment policy and phase-type exponential lead time

JIAQI ZHANG

Business School, The University of Edinburgh

Email: Jack.L.Zhang@gmail.com

THOMAS W ARCHIBALD

Business School, The University of Edinburgh

Email: T.Archibald@ed.ac.uk

Abstract: In this paper, a semi-Markov decision model of a two-location inventory system with holdout transshipment policy is reviewed under the assumption of phase-type exponential replenishment lead time rather than exponential lead time. The phase-type exponential lead time more closely approximates fixed lead time as the number of phases increases. Unlike past research in this area which has concentrated on the simple transshipment policies of complete pooling or no pooling, the research presented in this paper endeavors to develop an understanding of a more general class of transshipment policy. In addition, we propose an effective method to approximate the dynamic holdout transshipment policy.

Keywords: Inventory management, lateral transshipment policy, stochastic modeling and dynamic programming.

1. Introduction

Most studies of transshipment policy consider the system with the complete and no pooling policies only. Rather than that, we consider a system with a more general pooling policy which allows the system to make transshipments provided the inventory level at one specific location exceeds a given threshold. When the threshold value is equal to zero, it represents complete pooling. When the threshold is large enough, it represents no pooling. It represents a partial pooling policy when that threshold is between these two values. The threshold is referred as the **holdout** for a transshipment delivery.

In our previous study (Zhang, 2008, Section 6.3), we examined a semi-Markov decision process (**SMDP**) model of a two-location inventory system with exponential replenishment lead time.

Because of our assumption of exponential lead time, the optimal transshipment decision does not depend on the state of the replenishment order process due to the memory-less property of the exponential distribution. Intuitively one might expect that a location would become more willing to share inventory as its next replenishment approaches. One might also expect that, due to the form of the backorder cost (which increases with the time required to fill a backorder), the incentive for transshipment would decrease as the next replenishment at the location with inventory shortage approaches. These features cannot be captured by the **SMDP** model with exponential lead time. Hence, the results of the **SMDP** model with exponential lead time might not be applicable to, for example, the situation of fixed lead time. It is necessary to consider another **SMDP** model with phase-type lead time in which the

optimal transshipment decision is not only dependent on inventory level at the locations, but also dependent on the state of the replenishment order process.

Hence, we consider an advanced model with the assumption that the replenishment lead time is the sum of a fixed number of independent and identically distributed (**IID**) exponential random variables. The phase-type distribution is also convenient for a **SMDP** model because the time until the end of a phase does not depend on the time that has passed since the start of the phase. Hence, to model the time until the delivery of an order, it is only necessary to know how many phases have been completed since the order was placed.

Few published research works consider a **SMDP** approach assuming the phase-type distribution exponential lead time. Ours might be the first work to use this approach to investigate the transshipment decision in inventory systems. However, it is necessary to give a quick review of relevant studies of transshipment policy, particularly those using dynamic programming.

Much of the research on transshipment in multi-location inventory systems considers models which allow transshipment at a single point during a period. This point may fall after demand for the period is fully realised (see for example (Zou et al., 2010; Hu et al., 2008; Krishnan and Rao, 1965)) or before demand for the period is fully realised (see for example (Bertrand and Bookbinder, 1998; Karmarkar, 1987; Gross, 1963)). However, these models are not suitable when customers arrive sequentially and require immediate service. This situation often arises in a retail setting for example. One approach is to consider transshipment each time a stockout arises at a location. Many models of this type assume one-for-one replenishment and are intended for inventories of spare parts or other slow-moving items. Often such models assume complete pooling among locations in a neighbourhood, for example (Lee, 1987; Axsäter, 1990; Kutanoglu, 2008; Kutanoglu and Mahajan, 2009). Other approaches such as (Grahovac and Chakravarty, 2001) have considered partial pooling with a fixed holdout level. In contrast, our research considers a different class of replenishment policy and focuses on partial pooling policies that allow the holdout level to depend on remaining lead times of outstanding replenishment orders and the inventory levels of the locations.

Archibald et al. (1997) is one of the first papers to consider transshipment in response to stockouts for other types of replenishment policy. They develop a **SMDP** model of a periodic review inventory system with two-locations, a base-stock

replenishment policy and zero replenishment lead time. The optimal transshipment policy is shown to consist of a set of thresholds which depend on the inventory level at the location with stock on hand and the time until the next replenishment of the locations. The method adopted in our research follows a similar approach and proposes a heuristic transshipment policy with some similar features. However, our research assumes a continuous review system, non-zero replenishment lead time and a cost function that depends on the time taken to satisfy backorders. Consequently, the optimal transshipment policy for the model developed in this paper is a function of four variables rather than two.

Archibald (2007) proposes three heuristic methods for systems with more than two locations under the same assumptions as (Archibald et al., 1997). The heuristics exploit the structure of the optimal transshipment policy for the two-location system and differ in the willingness of the locations to meet transshipment requests.

Zhao et al. (2008) develop a **SMDP** model of a make-to-stock system with two locations in which transshipment can be triggered by a demand or a production completion at a location. In this way, transshipment allows both inventory and production capacity to be shared between the locations. Under the assumptions of Poisson demand and exponential production times, the structure of the optimal policy is derived. In our model, inventory is replenished by ordering from a supplier rather than production and the replenishment lead time has a phase-type distribution to approximate a fixed lead time.

Models with unidirectional transshipment have been proposed for situations with very different shortage costs at locations in the system (Axsäter, 2003), both direct and indirect sales channels (Seifert et al., 2006) or logistical barriers to bidirectional transshipment (Olsson, 2010). Seifert et al. (2006) assume that customers using the direct sales channel demand immediate service while customers using the indirect sales channel are willing to wait for service. Hence, transshipment can be used to satisfy customers using the indirect channel once the total demand for a period has been realised. As in our model, Olsson (2010) and Axsäter (2003) allow transshipment whenever a stockout arises. However, unlike our research, they assume complete pooling.

In this paper, we develop two **SMDP** models for a two-location inventory system: one is with a general transshipment policy and the other is with a holdout transshipment policy. In Section 2, we first define our assumptions and notation before

formulating **SMDP** models for the general and holdout transshipment policies. In Section 3, we present the results of a numerical analysis of the two **SMDP** models. Additionally, benefiting from **SMDP** technique, we propose a quick and reliable method to approximate the dynamic transshipment policy, which could be used as a powerful tool for other similar research.

2. MODEL FORMULATION

2.1 Model assumptions and notations

We consider a two-location inventory system with unidirectional transshipment from location 2 to location 1. At location k , $k=1, 2$, a Poisson demand process with rate λ_k is used to model the customer demand. Unmet demand at a location may be backordered. When a backorder is placed at location k , there is one-off stockout cost $\hat{b}_k$ and a further backorder cost of b_k per time unit until the backorder is satisfied. Alternatively unmet demands at location 1 may be met by transshipment from location 2. Transshipment is assumed to be instantaneous and involves a cost of ω per item. Further, at location k the fixed order cost is c_k and holding cost is h_k per item per time unit.

We assume that the inventory level at location k can not exceed the storage capacity M_k and the maximum number of backorders that can be placed is N_k . This assumption is necessary for a finite state model. If a demand occurs at a location after the maximum number of backorders has been reached and, in the case of location 1, transshipment is not possible, the demand results in a lost sale at a cost of B_k per item.

We assume a continuous review replenishment order policy with fixed order quantity Q_k at location k . We further assume that there can never be more than one outstanding order at each location. This is important for computational reasons (as it limits the dimension of the state space) and is reasonable when the order quantity is large relative to the lead time demand. We also assume that it is not possible to place a replenishment order at a location while the fixed order quantity for the location exceeds the available storage at the location. Otherwise it is possible (for example, if there is no demand during the replenishment lead time) that the storage capacity will be violated when the order is delivered. We note that it is straightforward to modify the model to allow for variable order quantities or (R, Q) replenishment order policies.

We use a phase-type distribution (see (Tijms, 1995, p162)). The lead time at location k is assumed to be the sum of W_k independent and identically distributed (**IID**) exponential random variables each with mean $\frac{1}{W_k\mu_k}$. We call this phase-type exponential lead time. Each random variable represents a phase of the lead time and W_k is referred to as the number of phases. As the number of phases increases, the model more closely approximates a fixed lead time.

The phase-type distribution is convenient for a **SMDP** model because the time until the end of a phase does not depend on the time that has passed since the start of the phase. Hence, to model the time until the delivery of an order, it is only necessary to know how many phases have been completed since the order was placed.

To help us to define relevant terms, we denote the indicator function $\delta(x) = \begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases}$ and its complement function $\hat{\delta}(x) = 1 - \delta(x) = \begin{cases} 1 & \text{if } x \leq 0 \\ 0 & \text{if } x > 0 \end{cases}$. Throughout this paper the terms decision and action are interchangeable.

2.2 General transshipment policy formulation

State space

The state of the process is defined to be the inventory levels and the status of replenishment orders at the locations. Let i_k denote the inventory level at location k , $k = 1, 2$, where negative values indicate outstanding backorders. Let $w_k = 0$ represent the situation where there is no outstanding replenishment order at location k . Let $w_k > 0$ represent the situation where there is one outstanding replenishment order for Q_k items at location k and that this order will arrive after w_k phases (i.e. after a time equal to the sum of w_k **IID** exponential random variables each with mean $\frac{1}{W_k\mu_k}$).

Definition 1. Under our assumptions, the state space is given by

$$I = \{(i_1, i_2, w_1, w_2): -N_k \leq i_k \leq M_k, \\ 0 \leq w_k \leq W_k \text{ for } k = 1, 2\}$$

Action space

Decisions have to be made concerning the replenishment of the locations and how to meet demand at location 1 when there is no local inventory. Let $X_k(i)$ denote the set of the possible replenishment order decisions at location k , $k = 1, 2$ when the system is in state $i = (i_1, i_2, w_1, w_2)$. When there is no outstanding order at location k

(i.e. $w_k=0$) and there is sufficient storage capacity for a delivery of Q_k items (i.e. $Q_k \leq M_k - i_k$), we can choose to place a replenishment order. The decision to place an order is modeled by setting the replenishment order decision equal to W_k and the decision not to place an order is modeled by setting the replenishment order decision equal to 0. Hence, if $w_k = 0$ and $i_k \leq M_k - Q_k$, $X_k(i) = \{0, W_k\}$. Under our assumption it is not possible to place a replenishment order at location k in any other states because either there already is an outstanding order at location k (i.e. $w_k > 0$) or the storage capacity is not sufficient for a delivery of Q_k items (i.e. $i_k > M_k - Q_k$). In such cases, it is convenient for the formulation to set the replenishment order decision equal to the number of phases remaining in the lead time (i.e. w_k). It follows that

$$\text{if } w_k > 0 \text{ or } i_k > M_k - Q_k, X_k(i) = \{w_k\}$$

Let w'_k represent the replenishment order decision at location $k, k=1, 2$. For any given state $i \in I$, w'_k can take any value in the set $X_k(i)$.

We next consider the decision of how to satisfy a demand at location 1, which we refer to as the transshipment decision. Let $D(i)$ denote the set of possible transshipment decisions when the system is in state i . Let 0 represent the decision to meet the demand at location 1 by transshipment and 1 represent the decision not to use transshipment. When the inventory level at location 1 is greater than zero, it is assumed that meeting the demand from inventory at location 1 is optimal. Further, when the inventory level at location 2 is less than or equal to zero, meeting the demand via transshipment from location 2 is not feasible. Hence, if $i_1 > 0$ or $i_2 \leq 0$ $D(i) = \{1\}$. In other situations, we can choose whether to use transshipment to meet the demand or not. Therefore, if $i_1 \leq 0$ and $i_2 > 0$, $D(i) = \{0, 1\}$.

Let d represent the decision with regard to transshipment from location 2 to location 1. For any given state $i \in I$, d can take any value in the set of $D(i)$.

Definition 2. Under our assumptions, the action space in state i is given by

$$A(i) = \{(w'_1, w'_2, d): w'_k \in X_k(i) \text{ for } k=1, 2, d \in D(i)\}$$

From the assumptions above, we conclude that the state space I and action space $A(i)$ are finite.

Decision epoch

Decisions need to be taken when a demand occurs at either location or when the system reaches the end of a phase of the lead time at either loca-

tion. We define an event to be any occurrence of a demand or the end of any phase of a replenishment order lead time in the system. With this definition, the time of the next event depends on the state of the system. When there is no outstanding order at location k , we introduce a fictitious event after a time that has an exponential distribution with mean $\frac{1}{W_k \mu_k}$. This can be thought of as the end of the final phase of the lead time of a zero replenishment order at location k and will leave the state unchanged. Hence, whatever the state and action, the next event can always be a demand at location 1, a demand at location 2, the end of a phase of the lead time of a (possibly fictitious) replenishment order at location 1 or location 2. The time until each of the possible events is exponential with scale parameter independent of the state and action. Hence, the expected time until next decision epoch no longer depends on state or decision. We can now define τ as the expected time until the next decision epoch where

$$\tau = \frac{1}{\lambda_1 + \lambda_2 + W_1 \mu_1 + W_2 \mu_2} \quad (1)$$

Transition probabilities

We define the transition probabilities for the model implicitly by considering which one of the four possible events occurs at the next decision epoch. Though the probability of the events do not always correspond to transition probabilities, the transition probabilities can be deduced from the probabilities of the events occurring.

Consider a decision epoch in which the system is in state $i = (i_1, i_2, w_1, w_2)$ and action $a = (w'_1, w'_2, d)$ is chosen. The probability that the next event is a demand at location 1 is $\lambda_1 \tau$. If $d=0$, this demand is met by transshipment from location 2 and the state of the process at the next decision epoch is $(i_1, i_2 - 1, w'_1, w'_2)$. Otherwise, if $d=1$, this demand is satisfied from location 1 (either from inventory or by backorder) and the state of the process at the next decision epoch is $(\max(-N_1, i_1 - 1), i_2, w'_1, w'_2)$. Similarly, the probability that the next event is a demand at location 2 is $\lambda_2 \tau$ and the state of the process at the next decision epoch is $(i_1, \max(-N_2, i_2 - 1), w'_1, w'_2)$.

The probability that the next event is the end of a phase of a replenishment order lead time at location 1 is $W_1 \mu_1 \tau$. If $w'_1 = 0$, this is a fictitious event and the state of the process at the next decision epoch is (i_1, i_2, w'_1, w'_2) . If $w'_1 = 1$, this is the delivery of a replenishment order at location 1 and the state of the process at the next decision epoch is $(i_1 + Q_1, i_2, 0, w'_2)$. Otherwise, if $w'_1 > 1$, this is

the end of a phase of the replenishment lead time and the state of the process at the next decision epoch is $(i_1, i_2, w'_1 - 1, w'_2)$. The probability that the next event is the end of a phase of a replenishment order lead time at location 2 is $W_2\mu_2\tau$ and, following a similar argument to the above, the state of the process at the next decision epoch is $(i_1, i_2 + Q_2, w'_1, 0)$ if $w'_2 = 1$ and $(i_1, i_2, w'_1, \min(0, w'_2 - 1))$ otherwise.

Immediate cost

We now define $c_i(a)$ the expected cost incurred until the next decision epoch when action a is chosen in state i at the current decision epoch. To reiterate, this cost consists of the fixed order cost, holding cost, backorder cost, stockout cost, lost sale penalty cost and transshipment cost.

Under our assumptions, a replenishment order is only placed at location k at a decision epoch when $w'_k > w_k$. Hence, the fixed order cost incurred until the next decision epoch is equal to $c_k\delta(w'_k - w_k)$ at location k , $k = 1, 2$. Inventory levels do not change between decision epochs, so the expected inventory cost at location k until the next decision epoch is $h_k i_k \tau$ if $i_k \geq 0$ and $-b_k i_k \tau$ if $i_k < 0$. When the next event is a demand at location 1, there is a cost of ω if it is met by transshipment ($d = 0$), $\hat{b}_1$ if it is met by backorder ($d = 1$ and $-N_1 < i_1 \leq 0$), and B_1 if it results in a lost sale ($d = 1$ and $i_1 = -N_1$). When the next event is a demand at location 2, there is a cost of $\hat{b}_2$ if it is met by backorder ($-N_2 < i_2 \leq 0$) and B_2 if it results in a lost sale ($i_2 = -N_2$). Hence, we conclude that the expected cost until the next decision epoch when action a is chosen in state i at the current decision epoch is as follows.

$$\begin{aligned} c_i(a) = & \left\{ \sum_{k=1}^2 \{c_k\delta(w'_k - w_k) + [h_k\delta(i_k) - b_k\hat{\delta}(i_k)]i_k\tau\} \right. \\ & + \lambda_1\tau d\{\hat{\delta}(i_1)\hat{b}_1 + \hat{\delta}(N_1 + i_1)(B_1 - \hat{b}_1)\} \\ & + \lambda_2\tau\{\hat{\delta}(i_2)\hat{b}_2 + \hat{\delta}(N_2 + i_2)(B_2 - \hat{b}_2)\} \\ & \left. + \lambda_1\tau(1 - d)\omega \right\} \end{aligned} \quad (2)$$

Value-iteration algorithm

Following standard techniques for **SMDP** models (see for example (Tijms, 2003)), the optimal long-run average cost per time unit can be calculated using the value iteration algorithm. This approach applies a data transformation to the **SMDP** model to create a discrete-time Markov decision process (**DMDP**) model. Under mild assumptions, which we will assume hold for the model being examined, the **DMDP** model has the same class of stationary policies and the same long-run average expected cost per time unit as the original **SMDP** model.

2.3 Holdout transshipment policy formulation

For the **SMDP** model with the holdout transshipment policy, most definitions including the state, decision epoch, transition probabilities, immediate cost and value-iteration algorithm are the same as those in the Section 2.2. Therefore, we just redefine the action space for this model as follows.

Action space

Let w'_k represent the decision taken regarding the replenishment order at location k . As before, for any given state $i \in I$, w'_k can take any value in the set $X_k(i)$.

We introduce a holdout threshold such that, when location 1 has no inventory, demand at location 1 will be met by transshipment from location 2 if and only if the inventory level at location 2 exceeds this threshold. Let I_2 denote the holdout threshold at location 2. Using the same notation as before, we define $D(i)$ as follows.

If $i_1 > 0$ or $i_2 \leq I_2$, $D(i) = \{1\}$

If $i_1 \leq 0$ and $i_2 > I_2$, $D(i) = \{0\}$

Let d represent the decision with regard to the transshipment decision from location 2 to location 1. For any given state $i \in I$, d takes the value in the set $D(i)$.

Definition 3. Under our assumptions, the action space in state i is given by

$$A(i) = \{(w'_1, w'_2, d); w'_k \in W_k(i) \text{ for } k=1, 2, d \in D(i)\}$$

From the assumptions, we conclude that the action space $A(i)$ is finite.

3. NUMERICAL ANALYSIS & DISCUSSION

3.1 Numerical experiments

For the **SMDP** model with phase-type exponential lead time, the computational overheads increase quadratically with an increase in the number of phases in the lead time distributions. We define $N_k = -80$, and $M_k = 70$ for $k = 1, 2$, as the limit on backorders and storage capacity in our numerical experiments. In Table 1, we compare the estimated computing time¹, the number of iterations and average total cost rate for different numbers of phases.

Unless otherwise stated we use the following parameters in our numerical experiments: $\lambda_k = 10$, $Q_k = 30$, $\mu_k = 1$, $c_k = 20$, $h_k = 0.5$, $\hat{b}_k =$

1. The estimated computing time is the result from running on the desktop PC machine with Pentium IV CPU 1.2G HZ, 2GB memory

10, $b_k = 5$, $B_k = 500$, $\omega = 1$ for $k = 1, 2$. These results show significant increases in the computing time when the number of phases increases from 1 to 8. For reasons of practical implementation, we choose to do our numerical experiments with the number of phases equal to 4 to assess the impact of the phase-type lead time.

Iter. No.	W_k	Av. Cost	Est. Time
4016	1	46.47	2 mins
2472	2	39.17	6 mins
1964	3	36.78	10 mins
1734	4	35.58	17 mins
1652	5	34.84	26 mins
1656	6	34.34	42 mins
1706	7	33.97	65 mins
1790	8	33.69	97 mins

Table 1. Average total cost rate for optimal transshipment policy with different numbers of phases

We do sensitivity tests on the parameters b_1 and $\hat{b}_1$. In these experiments we vary b_1 and $\hat{b}_1$ separately and observe the optimal average total cost rate under a general and a holdout transshipment policy. The optimal average total cost rate and the number of iterations for each value of b_1 or $\hat{b}_1$ when the number of phases is equal to 4 are shown in Table 2.

b_1	$\hat{b}_1$	Iter. No.	Av. Total Cost
1	10	1732	35.24
3	10	1730	35.44
5	10	1734	35.58
7	10	1738	35.68
5	2	1734	35.34
5	4	1734	35.41
5	6	1734	35.48
5	8	1734	35.53

Table 2. Average total cost rate for optimal transshipment policy with different values of b_1 and $\hat{b}_1$ when number of phases = 4

	$\hat{b}_1=10$			
I_2	$b_1=1$	$b_1=3$	$b_1=5$	$b_1=7$
0	35.43	35.52	35.61*	35.69*
1	35.40	35.51	35.61	35.71
2	35.37	35.51*	35.63	35.74
3	35.36	35.51	35.66	35.78
4	35.36*	35.54	35.70	35.84
5	35.37	35.57	35.76	35.92
6	35.39	35.62	35.83	36.01
7	35.42	35.69	35.91	36.11
8	35.47	35.76	36.01	36.22
9	35.53	35.85	36.12	36.35
10	35.60	35.94	36.23	36.48
11	35.67	36.05	36.36	36.62
12	35.76	36.16	36.49	36.76
13	35.85	36.28	36.62	36.91
14	35.95	36.39	36.75	37.06
15	36.04	36.51	36.89	37.20
16	36.14	36.63	37.03	37.35
17	36.24	36.75	37.16	37.49
18	36.33	36.87	37.28	37.62
19	36.42	36.98	37.41	37.75
20	36.51	37.08	37.52	37.88

Table 3. Average total cost rate under the holdout transshipment policy with different holdout thresholds and values of b_1 when number of phases = 4

The results in Table 3 show that, when $b_1 = 1$ and 3, the optimal holdout transshipment policy is the partial pooling policy with holdout threshold I_2 equal to 4 and 2 respectively. However, when $b_1 = 5$ and 7, the optimal holdout transshipment policy is complete pooling. The predicted improvements in the average total cost from the optimal holdout transshipment policy are 0.21%, 0.05%, 0.00% and 0.00% with respect to complete pooling ($I_2 = 0$) and 3.27%, 4.44%, 5.37% and 6.13% with respect to no pooling ($I_2 = M_2$) respectively. Further, the difference between the average total cost rates under the optimal holdout transshipment policy and the optimal general transshipment policy is 0.32%, 0.18%, 0.10% and 0.03% respectively.

	$b_1=5$			
I_2	$\hat{b}_1=2$	$\hat{b}_1=4$	$\hat{b}_1=6$	$\hat{b}_1=8$
0	35.46	35.50	35.54	35.57
1	35.43	35.48	35.52	35.57*
2	35.41	35.47*	35.52*	35.58
3	35.40*	35.47	35.54	35.60
4	35.41	35.49	35.56	35.63
5	35.42	35.51	35.60	35.68
6	35.45	35.56	35.65	35.74
7	35.49	35.61	35.72	35.82
8	35.55	35.68	35.80	35.91
9	35.62	35.76	35.89	36.01
10	35.70	35.85	35.99	36.12
11	35.79	35.95	36.10	36.23
12	35.88	36.06	36.22	36.36
13	35.99	36.17	36.34	36.49
14	36.10	36.29	36.46	36.62
15	36.21	36.41	36.59	36.74
16	36.32	36.53	36.71	36.87
17	36.43	36.65	36.83	37.01
18	36.54	36.76	36.96	37.13
19	36.64	36.87	37.07	37.24
20	36.74	36.97	37.18	37.36

Table 4. Average total cost rate under the holdout transshipment policy with different holdout thresholds and values of $\hat{b}_1$ when number of phases = 4

Similarly, when $\hat{b}_1 = 2, 4, 6$, and 8 , in Table 4 the optimal holdout transshipment policy is the partial pooling policy with holdout threshold I_2 equal to $3, 2, 2$ and 1 respectively. The predicted improvements in the average total cost from the optimal holdout transshipment policy are 0.16% , 0.09% , 0.04% and 0.02% with respect to complete pooling ($I_2 = 0$) and 3.77% , 4.24% , 4.65% and 5.02% with respect to no pooling ($I_2 = M_2$) respectively. Further, the difference between the average total cost rates under the optimal holdout transshipment policy and the optimal general transshipment policy is 0.18% , 0.15% , 0.14% and 0.11% respectively.

The results in Tables 3 and 4 illustrate properties one would expect of an optimal holdout policy. That is, when the expected stockout or backorder cost increases at location 1, the optimal holdout threshold I_2 at location 2 is lowered to give more transshipment support. The evidence is clear from the movement of the optimal holdout threshold in the tables.

3.2 Optimal transshipment decisions

In this section, we investigate how optimal transshipment decisions depend on the state of the

replenishment order process. Due to the memoryless property of the exponential distribution, the same transshipment decisions are taken when there is no outstanding order at location k as when there are W_k phases remaining in the lead time of an outstanding order at location k . Hence, when considering the state of the replenishment order process, it is only necessary to consider situations where there are outstanding orders at both locations (i.e. $1 \leq w_k \leq W_k$ for $k = 1$ and 2). We are also keen to know how the optimal holdout transshipment policy compares to the optimal general transshipment policy.

Figures 1 to 8 plot the optimal transshipment decision as a function of the inventory levels at the two locations when $b_1 = 5$ and $\hat{b}_1 = 10$ for different states of the replenishment order process. Grey shading represents the decision to use transshipment to meet demand at location 1 and black shading represents the decision to meet demand at location 1 from local stock or by backorders. So for example from Figure 1, when $w_1 = 2$, $w_2 = 4$ and the inventory level at location 1 is between 0 and -8 , transshipment is only optimal if the inventory level at location 2 is greater than 6 . We only draw the optimal decisions within a range where $-30 \leq i_1 \leq 0$ and $1 \leq i_2 \leq 10$. For the examples we consider, it is optimal to use transshipment outside this range if and only if $i_1 \leq 0$ and $i_2 > 0$.

Figures 1 to 4 show the optimal transshipment decisions when there are 2 phases remaining in the lead time of the outstanding order at location 1. Each plot represents a different number of remaining phases in the lead time at location 2. We see that complete pooling is optimal when there is only 1 phase remaining in the lead time of the outstanding order at location 2. In other cases ($w_2 \geq 2$), there is some evidence of partial pooling. We note that the optimal level of transshipment support provided by location 2 increases: as the inventory level at location 2 increases; as the inventory level at location 1 decreases; and as the number of phases remaining in the lead time of the outstanding order at location 2 decreases. None of these features are captured in a holdout transshipment policy. For this example, the optimal holdout policy is complete pooling ($I_2 = 0$).

Similarly, Figures 5 to 8 plot the optimal transshipment policy when there are 4 phases remaining in the lead time of the outstanding order at location 2. We see that, as the number of phases remaining in the lead time at location 1 decreases, location 2 provides less transshipment support. Again this feature is not captured in a holdout transshipment policy.

The examples show that the optimal transshipment policy depends on the inventory levels at

the two locations and on the times to the next replenishments of the two locations. In the examples we have considered, we note that the optimal transshipment policy changes gradually as the inventory level at location 1 falls from 0. As most inventory systems are managed to provide high levels of customer service, we expect instances of negative inventory levels to be relatively rare. This suggests that the difference between the average total cost rates for the optimal transshipment policy and the optimal holdout transshipment policy is most likely due to the fact that the holdout threshold cannot vary with the times to replenishment at the two locations.

We propose a dynamic holdout transshipment policy in which the threshold depends on the state of the replenishment order process. The policy has $W_1 W_2$ different threshold values each corresponding to a different combination of (w_1, w_2) where $1 \leq w_k \leq W_k$. The case $w_k = 0$ is like $w_k = W_k$ due to the memory-less property of the exponential distribution. The holdout threshold for given (w_1, w_2) is defined to be the largest inventory level for which it is optimal for location 2 to refuse a transshipment request when the inventory level at location 1 is zero ($i_1 = 0$). Figure 9 plots the estimated threshold values for different states of the replenishment order process when $b_1 = 5$ and $\hat{b}_1 = 10$. The size of the bubbles represents the size of the holdout threshold which is also shown by the numbers next to the bubbles. For example, when $w_1 = 2$ and $w_2 = 3$, the threshold is 3 and the policy would use transshipment when location 1 has no stock and the inventory level at location 2 is greater than 3. The value of the threshold is non-increasing in the time remaining in the lead time at location 1 and non-decreasing in the time remaining in the lead time at location 2.

b_1	$\hat{b}_1$	Iter. No.	Av. Total Cost
1	10	1720	35.24
3	10	1718	35.44
5	10	1728	35.58
7	10	1740	35.68
5	2	1666	35.34
5	4	1692	35.41
5	6	1708	35.48
5	8	1708	35.53

Table 5. Average total cost rate for dynamic holdout transshipment policy with different values of b_1 and $\hat{b}_1$ when number of phases equals to 4

The results in Table 5 show that, in all the cases we consider, the dynamic holdout transshipment policy is as good as the optimal transshipment policy in terms of the average total cost rate. These results demonstrate the importance of considering the times to replenishment of the locations and provide strong support for the dynamic holdout transshipment policy proposed.

Fig. 1. Optimal transshipment decision when $b_1 = 5, \hat{b}_1 = 10, w_1 = 2, w_2 = 4$

Fig. 2. Optimal transshipment decision when $b_1 = 5, \hat{b}_1 = 10, w_1 = 2, w_2 = 3$

Fig. 3. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 2, w_2 = 2$

Fig. 5. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 4, w_2 = 4$

Fig. 4. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 2, w_2 = 1$

Fig. 6. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 3, w_2 = 4$

Fig. 7. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 2, w_2 = 4$

Fig. 8. Optimal transshipment decision
when $b_1 = 5, \hat{b}_1 = 10, w_1 = 1, w_2 = 4$

Fig. 9. Thresholds for a dynamic holdout policy when $b_1 = 5, \hat{b}_1 = 10$

4. CONCLUSIONS

In order to investigate the transshipment decision in a two-location inventory system with constant replenishment lead time, we have developed two **SMDP** models with phase-type exponential lead time. Numerical experiments with the models demonstrate interesting properties of the optimal

transshipment decisions, although, compared to complete pooling and no pooling, the total cost savings from selective transshipment in the examples we considered were not very significant.

Interestingly the numerical experiments demonstrate that the level of transshipment support provided by location 2 can vary significantly with

the state of the replenishment order process. We observed that it is optimal for location 2 to provide more transshipment support as the number of phases remaining in the lead time at location 2 decreases and as the number of phases remaining in the lead time at location 1 increases. These are intuitively appealing properties.

We argue that the difference between the optimal cost rates of the general and holdout transshipment policies is due to the use of a fixed holdout threshold which does not depend on the times until replenishment of the locations. We propose a dynamic holdout transshipment policy in which the threshold is a function of the state of the replenishment order process. The policy is specified by defining a threshold for each pair of values (w_1, w_2) satisfying $1 \leq w_k \leq W_k$. Determining an optimal set of thresholds for such a policy would be computationally difficult due to the large number of variables. However, we propose a method to estimate the set of threshold values from the optimal general transshipment policy.

Compared to general transshipment policies, the proposed dynamic holdout transshipment policy is easy to implement and to explain. These properties make it an attractive policy to use in practice. Further, we have demonstrated that, for the examples we considered, the cost rate for the proposed dynamic holdout transshipment policy is close to the optimal cost rate for the general transshipment policy. For situations in which the replenishment lead times are constant, the phases of the lead time can be interpreted as follows. Let L_k denote the lead time at location k . Divide the lead time into W_k intervals each of length $\frac{L_k}{W_k}$ and associate each interval with a phase of the lead time. Hence, when time until replenishment at location k in the fixed lead time model is between $\frac{(w_k - 1)L_k}{W_k}$ and $\frac{w_k L_k}{W_k}$ ($k = 1, 2$), we would use the threshold corresponding to (w_1, w_2) .

The modelling approach considered here offers benefits over other approaches such as decomposition approximation models. Firstly, rather than modelling each location as an independent location, we model the two-location system as a whole system. Hence, it is possible to capture all the interactions which occur between the two locations. Such an approach provides us with a powerful tool to improve our modelling of a system with strong interactions. Secondly, using numerical experiments, we can monitor how the optimal decision depends on the state of the process. Hence, we benefit from insights on optimal transshipment decisions for the **SMDP** model. For example, the use of the **SMDP** model to pre-

dict suitable holdout thresholds. These properties of the **SMDP** models prove **SMDP** modelling to be an effective approach for our problem domain.

However, the **SMDP** modelling approach has its weaknesses. Firstly, we need a large number of states and actions to model the problem. As a consequence of this, the solution method is computationally burdensome. It is unlikely that the method could be extended to more than two locations due to the exponential increase in the number of states and actions. Secondly, the **SMDP** approach does not provide explicit expressions for the average total cost rate and other performance measures such as direct fill rate and backorder fill rate.

In summary, the **SMDP** modelling technique provides a new angle of viewing general transshipment policies for the two-location inventory system with unidirectional transshipment. In turn this has led to the development of a dynamic holdout transshipment policy that is easy to implement yet potentially very efficient in terms of cost.

REFERENCES

- T. W. Archibald. Modelling replenishment and transshipment decisions in periodic review multilocation inventory systems. *Journal of the Operational Research Society*, 58 (7): 948–956, 2007.
- T. W. Archibald, S. A. E. Sassen, and L. C. Thomas. An optimal policy for a two depot inventory problem with stock transfer. *Management Science*, 43: 173–183, 1997.
- S. Axsäter. Modelling emergency lateral transshipments in inventory systems. *Management Science*, 36: 1329–1338, 1990.
- S. Axsäter. Evaluation of unidirectional transshipments and substitutions in inventory systems. *European Journal of Operational Research*, 149 (2): 438–447, 2003.
- L. P. Bertrand and J. H. Bookbinder. Stock redistribution in two-echelon logistics systems. *Journal of the Operational Research Society*, 49: 966–975, 1998.
- J. Grahovac and A. Chakravarty. Sharing and lateral transshipment of inventory in a supply chain with expensive low-demand items. *Management Science*, 47 (4): 579–594, 2001.
- D. Gross. Centralized inventory control in multilocation systems. In H. E. Scarf, D. M. Gilford, and M. W. Shelly, editors, *Multistage inventory models*

- and techniques, pages 47–84. Stanford University Press, Stanford, 1963.
- X. X. Hu, I. Duenyas, and R. Kapuscinski. Optimal joint inventory and transshipment control under uncertain capacity. *Operations Research*, 56 (4): 881–897, 2008.
- U. S. Karmarkar. The multilocation multiperiod inventory problem — bounds and approximations. *Management Science*, 33: 86–94, 1987.
- K. S. Krishnan and V. R. K. Rao. Inventory control in N warehouses. *Journal of Industrial Engineering*, 16: 212–215, 1965.
- E. Kutanoglu. Insights into inventory sharing in service parts logistics systems with time-based service levels. *Computers and Industrial Engineering*, 54 (3): 341–358, 2008.
- E. Kutanoglu and M. Mahajan. An inventory sharing and allocation method for a multi-location service parts logistics network with time-based service levels. *European Journal of Operational Research*, 194 (3): 728–742, 2009.
- H. L. Lee. A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science*, 33: 1302–1316, 1987.
- F. Olsson. An inventory model with unidirectional lateral transshipments. *European Journal of Operational Research*, 200 (3): 725–732, 2010.
- R.W. Seifert, U. W. Thonemann, and M. A. Sieke. Integrating direct and indirect sales channels under decentralized decision-making. *International Journal of Production Economics*, 103 (1): 209–229, 2006.
- H. C. Tijms. *Stochastic models: an algorithmic approach*. John Wiley, Chichester, 1995.
- H. C. Tijms. *A first course in stochastic models*. John Wiley, Chichester, 2003.
- J. Zhang. *Holdout transshipment policy in two-location inventory systems*. PhD thesis, Business School, University of Edinburgh, 2008.
- H. Zhao, J. K. Ryan, and V. Deshpande. Optimal dynamic production and inventory transshipment policies for a two-location make-to-stock system. *Operations Research*, 56 (2): 400–410, 2008.
- L. Zou, M. Dresner, and R. Windle. A two-location inventory model with transshipments in a competitive environment. *International Journal of Production Economics*, 125 (2): 235–250, 2010.