European Journal of Operational Research 237 (2014) 474–486

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Inventory sharing via circular unidirectional chaining

Dina Smirnov [a,*,1], Yigal Gerchak [b]

[a] Faculty of Industrial Engineering and Management, Technion – Israel Institute of Technology, 32000 Haifa, Israel
[b] Department of Industrial Engineering, Tel-Aviv University, 69978 Tel-Aviv, Israel

ARTICLE INFO

Article history:
Received 29 October 2012
Accepted 9 February 2014
Available online 20 February 2014

Keywords:
Inventory pooling
Supply chain management
Lateral transshipments

ABSTRACT

This research studies the performance of *circular unidirectional chaining* – a "lean" configuration of lateral inventory sharing among retailers or warehouses – and compares its performance to that of no pooling and complete pooling in terms of expected costs and optimal order quantities. Each retailer faces uncertain demand, and we wish to minimize procurement, shortage and transshipment costs. In a circular unidirectional chain all retailers are connected in a closed loop, so that each retailer can cooperate with exactly two others as follows: receive units (if needed\available) from the left "neighbor" and send units (if needed\available) to the right, and a retailer who receives units from one neighbor is not allowed to send any units to its other neighbor. If the chain consists of at least three nodes and demands across nodes are i.i.d., its performance turns out to be independent of the number of nodes. The optimal stocking is therefore solved *analytically*. Analytical comparative statics with respect to cost parameters and demand distributions are provided. We also examine thoroughly the cases of uniform demand distribution (analytically) and normal demand distribution (numerically). In the uniform case with free transshipment, a unidirectional chain can save up to 1/3 of the expected cost of separate newsvendors caused by uncertainty. For three nodes, the advantage of complete pooling over unidirectional chaining does not exceed 19%.

1. Introduction

Horizontal Inventory pooling is a powerful cost reducing tool in supply chain management (Eppen, 1979). Horizontal pooling is performed among retailers or warehouses located at the same echelon of a centrally-managed[2] supply chain, in which each retailer faces uncertain demand. The purpose of the current study is to provide an extended analysis of limited pooling of inventory through a circular unidirectional chain. We also compare the performance of the circular unidirectional chain to that of other (extreme) configurations of inventory sharing.

Throughout the study we focus on a single-period problem, a setting appropriate for perishable products, products with seasonal demand and products with long lead times. We conduct virtual inventory pooling, often referred to as *lateral transshipments* in the literature. At the beginning of the period, and before actual demand is revealed, each location (viewed in this work as a node in a graph) places and (instantaneously) receives an order from the upper level in the supply chain. After demand at each node is observed and satisfied locally to the extent possible, units are transferred (instantaneously) from nodes with excess inventory to nodes with excess demand. Lateral transshipments thus serve as emergency supply sources. Each item transshipped between two locations saves a unit shortage cost, but incurs a unit transshipment cost. The problem of determining the optimal replenishment and transshipment quantities has been widely addressed in the literature (e.g., Krishnan & Rao, 1965; Tagaras, 1989). An extensive review on the topic of lateral transshipments is available in Paterson, Kiesmüller, Teunter, and Glazebrook (2011).

If one wishes to pool inventories, he must decide on the sharing configuration of the system. That is, for every node, we need to define the nodes it can cooperate with. A directed link between two nodes represents their possibility to share inventory unidirectionally (thus, two links are required so that both nodes can send to and receive from each other). We assume that the unit production (procurement), shortage and transshipment costs are identical across nodes. Our main purpose is comparing the performances of different inventory sharing configurations and obtaining general insights into the solution properties rather than developing ways

* Corresponding author. Tel.: +972 544906081.
E-mail addresses: dinna@tx.technion.ac.il (D. Smirnov), yigal@post.tau.ac.il (Y. Gerchak).

[1] This work was done as a part of the author's M.Sc. thesis at Tel-Aviv University.
[2] For an overview on decentralized inventory systems with lateral collaboration, see Summerfield and Dror (2012) and references therein.

http://dx.doi.org/10.1016/j.ejor.2014.02.019
0377-2217/

to solve (large) instances of a particular configuration. We thus find it useful to consider symmetric configurations with i.i.d. demands. Additionally, if demands are not i.i.d., there is no analytical solution for the sharing policies in question. Consequently, in most of the parts of this study, unless explicitly stated, we assume that demands at all nodes are i.i.d. However, later we computationally analyze a system with different demand distributions as well.

There are two extreme inventory sharing configurations. The first is *no inventory pooling* (*NP*), in which each location operates alone. No transshipment is performed in such a system. The other extreme configuration is *complete inventory pooling* (*CP*), where each retailer is allowed to share inventory with all others (bidirectionally), with no limitations but often at transshipment costs (usually, simply referred to as "pooling"). It has been shown that when the pooled system is managed optimally, inventory pooling always reduces the expected cost of the system (e.g., Eppen, 1979; Gerchak & He, 2003) because the aggregate inventory problem constitutes a relaxation of the collection of separate inventory problems. Moreover, pooling could be beneficial even at the presence of transshipment cost (e.g., Krishnan & Rao, 1965).

From the pooling perspective, complete pooling is clearly the most effective arrangement (Tagaras, 1989). However, complete pooling has three main drawbacks. First, it requires constructing and maintaining many links between the nodes. The existence of each link may be associated with a (high) fixed cost, which is independent of the amount transshipped along the link. This cost (referred to as the "establishment cost" henceforth) is incurred for establishing and maintaining the infrastructure that would allow inventory sharing between a pair of nodes. In particular, the fixed cost usually consists of (but is not limited to) the cost of drivers, trucks, fixed payments to a third party for delivery, information systems, any kind of administration associated with pooling, etc. Second, systems under complete pooling tend to have high computational complexity, especially if a variable transshipment cost is present. The third reason concerns multi-period systems (a possible extension of our setting). From a practical standpoint, complete pooling may be difficult to operate despite its theoretical effectiveness. In complete pooling, the actual links selected for transshipment may vary from one period to another (according to the demand realization) – possibly confusing for the retailers. In this paper we present a simpler approach with fewer possibilities of selecting the "from-to" pairs of nodes.

One moderate sharing configuration which constitutes partial inventory pooling and requires far fewer links is a *circular unidirectional chain*. This configuration has been found to be very effective in reducing shortage costs in inventory systems (e.g., a recent study by Lien, Iravani, Smilowitz, & Tzur, 2011) and capturing much of the benefits of total flexibility in a manufacturing environment (e.g., Jordan & Graves, 1995). We define a "circular unidirectional chain" as a configuration where all retailers are connected unidirectionally in a closed loop, such that each retailer cooperates with two other retailers (the neighboring nodes). Items may be sent to only one neighboring node, and may be received only from *the other* neighboring node. In our chaining setting, a node which receives units due to a shortage is not allowed to further transfer units to other nodes. As such, our rules of chaining are more restrictive than in previous transshipment literature (e.g., Lien et al., 2011). Note that the issue of which node will actually transship to which does not arise here, as the symmetric shortage and transshipment costs do not affect transshipment decisions. From a practical point of view, such policy is much simpler to operate.

Ignoring the costs of establishing links and of transshipment, circular unidirectional chaining with no-pooling level inventories is clearly superior to no-pooling for any demand realization (and thus in expectation). That is also true when the chain and the system without pooling use their respective optimal quantities. Further, circular unidirectional chaining, even if it orders optimally, is clearly inferior to complete pooling in terms of expected costs. The extent of this gap is an issue explored here.

Our research is thus motivated by the tradeoff between the possible benefits of consolidation of demands and the high (and certain) establishing cost of each link, as well as the high computational complexity associated with complete pooling at the presence of transshipment costs. We aim to determine whether a restrictive and "minimalistic" circular unidirectional chaining configuration is effective enough in reducing shortages. Bidirectional chaining, although more effective than unidirectional, is considerably more complex to manage and optimize (Kalikhman, 2011). Although we do not explicitly consider the savings achieved by pooling versus the establishment cost, this fixed cost could easily be incorporated at the end. If one has information about the establishment cost of each link and the cost of no pooling, he can easily obtain the net savings from creating a circular unidirectional chain, and compare the cost of a circular unidirectional chain to that of complete pooling when both systems include the corresponding establishment cost, according to the number of links in each.

An issue of independent interest is the relative magnitudes of the optimal order quantities in different sharing policies. It is not obvious that the optimal order quantity always decreases with inventory pooling (e.g., Gerchak & Mossman, 1992; Yang & Schrage, 2009).

The paper makes two main contributions: first, it provides an extended analytical approach to circular unidirectional chaining, including the optimality condition, and offers various insights on this sharing configuration, which, besides being very lean and economical in terms of the number of links, is easier to solve than complete pooling with positive unit transshipment cost (see Section 2). Further, circular unidirectional chaining can serve as an approximation of complete pooling in terms of expected costs. Such approximation is easier to compute than exact expected cost of complete pooling, primarily because it is lean. We find combinations of the problem parameters for which the approximation is good (i.e., the unidirectional chain is close to complete pooling), as well as those for which it is not. Note that a symmetric system is of interest, since it serves as an approximation of symmetric complete pooling. Also, a symmetric system allows exploring the effect of number of retailers, without a need to define their order along the chain.

A recent work by Lien et al. (2011) has shown that a configuration of chaining (either uni or bi-directional) with the requirement that each retailer uses *its own* units for transshipment (but is allowed to receive units at the same time), is superior to other configurations suggested in the literature, such as grouping configurations. Unlike Lien et al. (2011), who focus on the design of transshipment networks and assume that the order-up-to level can be found by using an infinitesimal perturbation analysis procedure suggested by Herer, Tzur, and Yucesan (2006), we provide an analytical study of circular unidirectional chaining, including the optimality condition and comparative statics for the case of general demand distribution, and show that our model of unidirectional chaining can be viewed as an extension of the multi-location newsvendor model. Further, we provide a comparison between the policy of unidirectional chaining and those of no pooling and complete pooling (which were not included in the study of Lien et al. (2011)). Moreover, Lien et al. (2011) make transshipment decisions before satisfying demand locally.

Another relevant study is that of Kranenburg and van Houtum (2009), who investigated an inventory system where only some of the locations are allowed to serve as sources of lateral transshipments, while receiving items is allowed by all locations. The base-stock levels are found through a greedy heuristic, and it is shown numerically that such partial pooling achieves much of the benefits