Int. J. Production Economics 157 (2014) 120–132

Contents lists available at ScienceDirect

Int. J. Production Economics

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

A vendor managed inventory model using continuous approximations for route length estimates and Markov chain modeling for cost estimates

Christian Larsen *, Marcel Turkensteen

CORAL—Cluster for Operations Research and Logistics, Aarhus University School of Business and Social Sciences, Department of Economics and Business, Fuglesangs Allé 4, DK 8210 Aarhus V, Denmark

ARTICLE INFO

Article history:
Received 26 September 2012
Accepted 1 August 2014
Available online 11 August 2014

Keywords:
Joint replenishment policies
Continuous approximation
Markov chain modeling

ABSTRACT

We consider a vendor who supplies goods to a set of geographically dispersed retailers and can monitor and control the inventory levels at the retailers. Such an arrangement is often called vendor managed inventory (VMI). The decisions in this set-up are the inventory levels at the warehouse and at the retailers and the routing along the retailers. Normally, the inventory levels at the vendor's warehouse and at the retailers are established by modeling the problem as a joint replenishment problem (JRP). Such a model ignores the differences in distances, number of retailers visited, and vehicle loads that may occur, in particular when these retailers are served on joint delivery trips. Some approaches that take routing and inventory decisions into account jointly, but these are so complex that only relatively small instances can be solved.

This paper develops a more detailed specification of the transport costs than other JRP models. In order to ensure that the complexity of the problem does not become overwhelming, we assume that the retailers are identical and uniformly distributed across an area, which can either be a two-dimensional area or a one-dimensional line structure. First of all, we decide to construct zones of retailers to be replenished jointly. The expected travel distances across a given number of retailers are estimated analytically using results from the field of continuous approximation for two-dimensional areas and using our own approximation for one-dimensional ones. Using a Markov chain approach we develop a model to minimize transport and inventory costs simultaneously. Our experiments show that our detailed specification gives more accurate set-up costs (and thereby different control policies) than regular JRP approaches when demand is infrequent and the transportation costs are low.

1. Introduction

Vendor managed inventory (VMI) refers to the situation in which a vendor is responsible for all decisions regarding inventories at its retailers, Chopra and Meindl (2007; p. 518). It implies that the vendor monitors the inventory levels and decides when and how much inventory to replenish at each of his retailers. In this way the vendor utilizes his transportation equipment better and provides cost and service efficient inventory control for his retailers. In this paper, we add the factor of geography, in the form of the locations of the retailers, into the analysis of VMI systems. In a joint delivery, a vehicle is dispatched for a delivery tour or route to a group of retailers. There are reasons for having fixed routes along groups of retailers: it is easy to plan fixed routes, and if demand occurs frequently, it is likely that each retailer on a tour needs replenishment. However, if demand is infrequent, only a fraction of the retailers may actually have had a demand and need replenishment when a vehicle is dispatched. It may then make sense only to visit these retailers on a delivery tour. We wish to include this routing flexibility into the process of setting inventory levels at the warehouse and at the retailers.

Thus, the routes and the resulting transportation costs are flexible, which poses large computational challenges. Due to the challenging nature of the problem, we develop an approach for a single item distribution system in this paper. A direction of future research is the extension to the multi-product setting based on our results.

In inventory studies, the problem of having joint deliveries to a set of retailers is known as a *joint replenishment problem* (*JRP*). One stream of research assumes deterministic demands, cf Viswanathan

* Corresponding author.
E-mail address: chl@econ.au.dk (C. Larsen).

http://dx.doi.org/10.1016/j.ijpe.2014.08.001
0925-5273/

(1996) and Wildeman et al. (1997). Another stream assumes that demands follow a stochastic process. As our paper accounts for stochastic demand, we therefore focus on the latter stream of research in this review of JRP studies. The focus in this type of research lies on analyzing good control policies such as the can order policy, analyzed first by Balinfy (1964) and subsequently by Silver (1981), Federgruen et al. (1984), Schultz and Johansen (1999), Melchiors (2002) and Johansen and Melchiors (2003). The main problem with the can order policy is to develop a valid mathematical model (for a good discussion, see Schultz and Johansen, 1999). Others have proposed policies where the coordination of replenishment decisions is secured by only being allowed to replenish at certain time points. This is accomplished by either having deterministic review intervals or a common stochastic review interval governed by the total demand since the last replenishment opportunity. Base-stock policies based on deterministic review intervals are developed by Atkins and Iyogun (1988). When letting the review interval be determined by the total demand since the last replenishment opportunity, one gets the *QS* policy proposed by Renberg and Planche (1969) and subsequently analyzed by Pantumsinchai (1992). Viswanathan (1997) formulated a $P(s,S)$ policy, where there is a common deterministic review interval and the replenishment decision of each item is governed by an (s, S) policy. He also proposed a stochastic variant of the $P(s,S)$ policy, which he called the $Q(s,S)$ policy, that has a stochastic review interval as described above. A mathematical model and an algorithm to compute optimal policy variables for this policy have been developed by Nielsen and Larsen (2005) and Larsen (2009) further extended this to the case of correlated demands. In finalizing this review on research in the stochastic JRP model we will also mention the work by Gürbuz et al. (2007). They study a distribution system with a single warehouse and multiple identical retailers and assume the fixed order cost is incurred irrespective of the number of retailers served.

In JRP studies, there is a fixed set-up cost for each delivery, meaning that each delivery to a fixed group of retailers, or a *zone*, costs about the same. The implicit assumption is that the same retailers are visited during each delivery and the cost of an average delivery tour is easily predictable. But especially for infrequent demand, the set-up costs may vary with the number of retailers visited and the transportation distance. Therefore, we study in more detail how they affect the transportation costs arising from the geographical locations of retailers and why it may make sense to adjust the retailers' inventory strategies to reduce transportation distances, even though this is not a strictly optimal pure inventory strategy.

Where JRP approaches tend to ignore routing considerations, the field of *inventory routing problems* (*IRPs*) specifically aims to simultaneously minimize routing and inventory costs. Actually, there is a wide range of IRP approaches; see the overview papers by Baita et al. (1998) and Andersson et al. (2010). Some studies allow for routing flexibility where, given a range of possible demand distributions, one should decide a time period to visit each retailer. However, approaches that have flexible routes are usually only able to solve relatively small instances, i.e. with few retailers and time periods, due to the combined complexity of inventory and routing. Dynamic programming approaches for this problem have been presented in Kleywegt et al. (2004) and Hvattum and Løkketangen (2009). In order to reduce the complexity of IRPs, many studies set a common inventory policy for fixed groups of retailers. Generally, retailers are replenished with a vehicle that travels along a fixed route. Such policies are collectively known as fixed-partition policies; see e.g. Anily and Federgruen (1990) and Viswanathan (1997). But these policies do not allow for flexible routing through zones and generally require deterministic demand. As alternatives, one can suggest fixed routes, as in fixed-partition policies or JRPs, and fixed delivery quantities to each retailer in each delivery. The challenge in flexible routing is that one should keep track of retailer's inventory levels, but the main challenge is that for many deliveries through a group of retailers, a new route has to be determined. However, the route lengths across such flexible routes may be clearly lower than when fixed routes are used, as is done in IRP approaches. The question is how large these benefits of having flexible routes are.

In this paper, we develop JRP inventory policies that address the randomness of demand and the flexibility of routing over a very long time horizon. For inventory modeling over an infinite horizon, a Markov modeling approach is generally applied. A straightforward way of explicitly incorporating geography is to apply a routing approach that determines a shortest route through all retailers needing replenishment in a given period, or alternatively, to use an IRP approach that also has the option to delay a delivery to a retailer or move it forward. The resulting Markov chain representation, however, could contain a huge number of states as it needs to keep track of the inventory position of each individual retailer. The complexity is compounded by the route computations in every state. Instead, we assume that retailers are identical and keep track of the inventory position of a single typical retailer. The probability can then be derived that a certain number of retailers m need a refill without needing to specify which retailers need replenishment and where they are located.

We base our estimation of the expected distance through m retailers on the assumption of uniformly distributed demand across the vendor's service area, meaning that the geographical locations of retailers are evenly spread across the service area and that the demand rates at the retailers are similar.

We consider two different types of service areas: a two-dimensional circular area (the *circular city*; actually, the shape is not so relevant, but for the sake of convenience we use this shape) and a one-dimensional area in which the retailers are assumed to be distributed on a line (the *linear city*); see Fig. 1. In both cases, we assume that the distribution of retailers over the area is uniform. The two-dimensional area is a standard representation of a service area; the linear city case can be found when cities and towns are located along a major traffic artery. The Danish region of East Jutland is an excellent example of this. Most of the major cities in that region are located along the E45 motorway as illustrated by the light distribution (a proxy for population density) in Fig. 2. Transport corridors frequently have such a structure as well; see Rodrigue et al. (2006, p. 84).

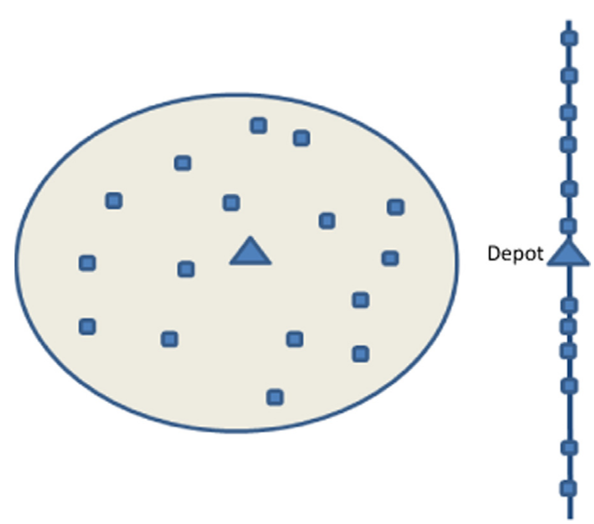

Fig. 1. Examples of the circular city (left) and linear city (right) *cases*.