European Journal of Operational Research 235 (2014) 334–338

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Short Communication

On two-echelon inventory systems with Poisson demand and lost sales

Elisa Alvarez *, Matthieu van der Heijden [1]

University of Twente, School of Management and Governance, P.O. Box 217, 7500 AE Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 15 May 2013
Accepted 22 December 2013
Available online 4 January 2014

Keywords:
Inventory
Two-echelon system
Spare parts
Lost sales
Approximation

ABSTRACT

We consider a two-echelon, continuous review inventory system under Poisson demand and a one-for-one replenishment policy. Demand is lost if no items are available at the local warehouse, the central depot, or in the pipeline in between. We give a simple, fast and accurate approach to approximate the service levels in this system. In contrast to other methods, we do not need an iterative analysis scheme. Our method works very well for a broad set of cases, with deviations to simulation below 0.1% on average and below 0.36% for 95% of all test instances.

1. Introduction

For advanced capital goods such as medical systems or defense systems, downtime may have serious consequences in terms of production loss and safety. To maintain such systems during their life cycle, the availability of spare parts to replace failed components is crucial. The related inventory holding costs can be huge, making the optional positioning of spare parts in the supply chain an important matter. Local stock points close to the installed base guarantee fast delivery, while a single central location facilitates low safety stocks because of the risk pooling effect. This leads to the need to analyze the performance of a single-item, two-echelon inventory model, consisting of various local warehouses that are replenished from a central depot.

For (expensive) slow moving parts, it is common to model demand by a Poisson distribution and to apply a one-for-one replenishment policy, see e.g. Muckstadt (2005). Also, since system downtime should be minimized, emergency procedures are typically used in out-of-stock situations. That is, a part is supplied from another source (e.g. the manufacturer) at additional costs. This demand can thus be considered as lost to the two-echelon system. The common approach in the literature is to use an emergency shipment if the local warehouse is out of stock. In this paper, we apply an emergency procedure only if the local warehouse is out of stock *and* no item is available at the central depot or in the pipeline between the depot and warehouse. The reason is that we focus on a setting where the spare part may be kept in stock both at the customer sites and at a forward stocking location serving all customers in a certain region. As the lead times between the forward stocking location and the customer sites are typically short, an emergency shipment from a (remote) external supplier does not make sense if a much cheaper regular shipment is feasible.

We found several related methods in the literature for the performance evaluation of single-item two-echelon models with lost sales, as we will explain in Section 2. For our variant, we develop a new method to analyze the performance of the inventory system in Section 3. We validate our method in Section 4 through a numerical experiment. We finish with our conclusions in Section 5.

2. Literature

We consider literature on single-item, two-echelon supply chains, consisting of multiple local warehouses facing Poisson distributed demand, and a central depot. Each location uses a continuous review, one-for-one replenishment policy. Demand that cannot be met from stock is served using an emergency shipment from an external source with ample supply and is thus lost to the system. A one-for-one replenishment policy is not always optimal under lost sales, since holding costs can be reduced by not requesting a replenishment directly after a demand occurs Hill (1999). The cost benefits of a sophisticated policy are limited to 1–2% Bijvank and Vis (2011).

It is well known that the analysis of such lost-sales inventory systems is more complex than its equivalent under full backordering Bijvank and Vis (2011). In particular, the analysis of the central depot is complex, since (i) the order process is *not* Poisson, and (ii) the order arrival rate depends on the inventory states of the warehouses: warehouses only generate replenishment orders if they

* Corresponding author.
E-mail addresses: e.m.alvarez@alumnus.utwente.nl (E. Alvarez), m.c.vanderheijden@utwente.nl (M. van der Heijden).
[1] Address: School of Management and Governance, P.O. Box 217, 7500 AE Enschede, The Netherlands. Tel.: +31 53 4892852.

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.12.031

have stock on hand. In the literature, solutions exist for specific cases. Muckstadt and Thomas (1980) consider a model where demand at any warehouse is met using an emergency shipment from the depot if the warehouse is out of stock and the depot still has stock. If neither location has stock, an emergency shipment from an external facility is used (with the demand thus lost to the system). They approximate the demand rate at the depot by a Poisson distribution with a constant rate.

In Andersson and Melchiors (2001), demand at a warehouse is lost if the warehouse is out of stock, even if the depot still has stock on hand. They approximate the arrival process at the depot by a Poisson distribution with a rate dependent on the warehouse fill rates and evaluate performance using an iterative approach: first, they find the mean waiting time for replenishment orders at the depot from the warehouse fill rates. Then, they find the warehouse fill rates given the mean waiting time at the depot. This iterative approach is reasonably accurate, but often does not converge when a lot of stock is kept at the depot, with little stock kept locally. Such a setting is very common for expensive slow movers to benefit from risk pooling. Seifbarghy and Jokar (2006) consider a similar approach for a model under an (R,Q) policy, i.e., with orders of size Q placed when the inventory position reaches a level R. The authors only consider identical warehouses and limit their experiments to settings with high service levels (fill rates of 90% or more).

All papers mentioned above use rather simple approximations for the analysis of the central depot, ignoring the fact that demand is not Poisson. Also, they ignore that the demand rate at the depot depends on the inventory levels at all warehouses. The only exception to the latter is Hill, Seifbarghy, and Smith (2007) who assume that (i) each local warehouse may have at most one outstanding order, and (ii) the shipment time from depot to warehouse is at least the central depot lead time. The second assumption is particularly restrictive, since upstream lead times tend to exceed downstream ones.

In this paper, we apply an accurate analysis of the order arrival rates and the pipeline at the central depot. We can handle non-identical local warehouses and allow any number of outstanding orders. As mentioned in the introduction, our notion of lost sales differs from literature, as we satisfy demand using the two-echelon network if an item is available in local stock, central stock or in transit in between. We find very accurate approximations for the service levels.

3. Model and analysis

3.1. Notation and assumptions

Fig. 1 illustrates our two-echelon model. We use index $k = 0$ for the depot and indices $k = 1,\ldots,K$ for the local warehouses. Demand arriving at warehouse k is served through the two-echelon network (the *regular channel*) if an item is available at the warehouse, at the depot, or in the transport pipeline in between. Note that an item in the transport pipeline can only be available if it has not yet been assigned to backordered demand at the warehouse. If no such item is available, we use an external *emergency channel* with infinite capacity at additional costs, and the demand is lost to the network. We use the following notation and assumptions:

- The stock at location $k = 0,\ldots,K$ is controlled using an (S_k-1,S_k) installation stock policy.
- Demand at local warehouse k occurs according to a Poisson process with rate m_k.
- The replenishment lead time to the depot has an exponential distribution with mean L_0. This assumption facilitates an analysis using Markov chains. Also, the performance of such lost sales models is not very sensitive to the lead time distribution, see Alfredsson and Verrijdt (1999).

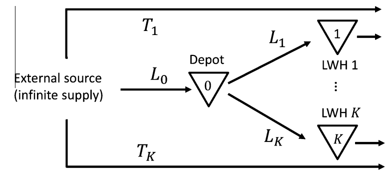

Fig. 1. A graphical representation of the supply system.

- The transportation time from the depot to warehouse k is deterministic and equal to L_k.
- The emergency shipment time to local warehouse k has an arbitrary probability distribution function with mean T_k.

Our performance indicators are (i) the fraction of warehouse k demand that is satisfied through the regular channel α_k, and (ii) the mean waiting time for demand satisfied through the regular channel $E[W_k]$ $(k = 1,\ldots,K)$. These performance indicators enable us to evaluate the mean downtime waiting for the spare part $DTWP_k$ at local warehouse k as $\alpha_k E[W_k] + (1-\alpha_k)T_k$.

3.2. Analysis

We find $E[W_k]$ from Little's Law, i.e. $E[W_k] = E[BO_k]/\alpha_k m_k$, with BO_k denoting the number of items backordered at warehouse k. In turn, we find both $E[BO_k]$ and α_k from the distribution of the number of backorders *at the depot* destined for local warehouse k, which we denote by BO_0^k. This is the critical and novel part of our analysis. Depot backorders occur when a warehouse sends a replenishment request to the depot when the latter is out of stock. The distribution of BO_0^k allows us to determine the distribution of the number of outstanding orders (i.e. the pipeline) at warehouse k, which in turn allows us to determine the distribution of BO_k. From the distribution of the depot backorders BO_0^k we also directly obtain α_k: once the depot is out of stock, at most S_k additional requests can still be met through the regular channel (either directly from warehouse stock or from items in the transport pipeline). Hence, once we have S_k depot backorders for warehouse k, further demand at that warehouse is lost. We thus find α_k as follows:

$$\alpha_k = \Pr\{BO_0^k < S_k\} = 1 - \Pr\{BO_0^k = S_k\} \tag{1}$$

We first show how to find the distribution of BO_0^k. Then, we show how to find $E[BO_k]$. For the analysis, we also need PI_k, the number of outstanding orders at each location $(k = 0,\ldots,K)$.

3.2.1. Distribution of BO_0^k, the number of backorders at the central depot for warehouse k

We condition on the number of outstanding orders at the depot PI_0. Under full backordering, we can easily disaggregate the depot backorders over the warehouses: Given $y_0 > S_0$ outstanding orders, the conditional probability $\Pr\{BO_0^k = x_k | PI_0 = y_0\}$ of having x_k backorders at the depot for warehouse k follows a binomial distribution with $y_0 - S_0$ trials and "success probability" $q_k = m_k/\sum_{h=1}^K m_h$ Graves (1985). However, this does not apply for our model. Also, the arrival rates of replenishment orders at the depot are state-dependent: the conditional probability of warehouse k being out of stock given $PI_0 = y_0$ increases in y_0, and so the arrival rate of replenishment orders at the depot decreases in y_0.

We now first show how to compute $\Pr\{BO_0^k = x_k | PI_0 = y_0\}$ exactly. Using these conditional probabilities, we compute the state-dependent arrival rates at the depot and find the distribution of PI_0. As an approximation in this second step, we assume that these arrival rates have a Poisson distribution irrespective of the value of PI_0. Finally, we find the distribution of BO_0^k as follows: