

Service differentiation in spare parts supply through dedicated stocks

E.M. Alvarez, M.C. van der Heijden, W.H.M. Zijm

Beta Working Paper series 373

BETA publicatie	WP 373 (working paper)
ISBN	
ISSN	
NUR	804
Eindhoven	February 2012

Service differentiation in spare parts supply through dedicated stocks

E.M. Alvarez, M.C. van der Heijden, and W.H.M. Zijm

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

Phone: +31-53-4893603; +31-53-4892852

E-mail: e.m.alvarez@utwente.nl; m.c.vanderheijden@utwente.nl; w.h.m.zijm@utwente.nl

Abstract

We investigate keeping dedicated stocks at customer sites in addition to stock kept at some central location as a tool for applying service differentiation in spare parts supply. We study the resulting two-echelon system in a multi-item setting, both under backordering and under emergency shipments assumptions (i.e. lost sales). In an extensive computational experiment, we show that dedicated stocks have significant added value: compared to an approach where all customers receive uniform service, we find average cost savings of 14% under backordering and 20% under emergency shipments. Furthermore, we find that dedicated stocks are comparable to critical level policies in terms of cost savings, while being much easier to implement in practice. Finally, we find further savings (20% under backordering, 23% under emergency shipments) by combining dedicated stocks and critical level policies in one aggregate differentiation strategy.

Key words: service differentiation, dedicated stocks, spare parts, multi-echelon systems, lost sales

1 Introduction

In the current business environment, suppliers of advanced capital goods, such as defense systems and chemical plants, increasingly provide their customers with service contracts that specify the services offered to that customer for system upkeep. Such contracts often contain quantified targets for key performance measures, such as a maximum response time in case of system failure. These so-called service level agreements may differ among customers to reflect the value each customer places on system availability. For instance, the maximum on-site response time may be 4 hours or next day. These varying service levels challenge suppliers to somehow incorporate such differentiation in their service processes.

In this paper, we consider dedicated customer stocks as a tool for handling differentiated service levels in spare parts supply. In this approach, a supplier keeps stock of certain items at premium customers' sites in addition to stock at some central location. Dedicated stocks are often used in practice, because it is a simple differentiation tool. For instance, we have seen such stocks used at a company specializing in luggage handling system at airports. Still, no research has yet been done on the savings possible with this approach: we expect the benefits from risk pooling to be smaller for this approach than for the case where all stock is kept centrally. We thus investigate if and when the approach has added value.

We evaluate the added value of dedicated stocks by comparing it to two approaches that have often been used in literature and practice, namely one-size-fits-all policies and critical level policies. One-size-fits-all policies provide all customers with uniform service irrespective of the individual service level agreements. As a result, they are usually excessively costly. Also, customers with standard contracts have no incentive to switch to premium contracts. In contrast, critical level policies reserve stock for only premium customers once the inventory level drops below a certain threshold. Requests that cannot be met by on-hand stock are either backordered or satisfied from another source (e.g. a production facility upstream in the supply chain). Although the policy can lead to large cost savings, there are barriers for implementing it in practice. For instance, service engineers responsible for speed of repair might use reserved stock for non-premium customers. Therefore, it is interesting to investigate whether a supplier can still obtain large savings with a simpler policy. In addition to comparing dedicated stock to critical level policies, we also examine the added value of combining both strategies in a single model. This combined policy enables us to judge which policy is best for which set of items.

In the remainder of the paper, we first give an overview of literature related to our research in Section 2. Here, we also state the main contributions of our research. We present our model in Section 3, and an optimization approach for this model in Section 4. Our optimization approach requires certain performance measures, such as waiting times, as input. We highlight how we find these performance measures in Section 5. We test the model in an extensive experiment (Section 6). Finally, we draw conclusions and discuss further research options in Section 7.

2 Literature overview

Our research contributes to the literature on service differentiation in spare parts supply of expensive slow movers. In general, the models for such systems are continuous-review models where demand arrives according to Poisson processes and one-for-one replenishment of items (i.e. base stock policies) is used.

In the area of service differentiation, most contributions consider critical level policies, a concept introduced by Veinott (1965). The optimality of this policy has been proven under various circumstances, such as under periodic and continuous review, both under backordering and lost sales assumptions (see Alvarez et al. (2010) for further details). Some recent contributions focusing on expensive slow movers are by Kranenburg and Van Houtum (2008) and Enders et al. (2008). Kranenburg and Van Houtum (2008) consider a multi-item single-location model with various demand classes and lost sales (emergency shipments are used for lost demand). The authors analyze the system through Markov chains and use an optimization approach based on decomposition and column generation, combined with local search. Enders et al. (2008) consider a single-item model with two demand classes. In addition to using critical level policies, the authors use different shipment modes when demand cannot be met from on-hand stock: non-premium demand is backordered, with premium demand being lost. In Alvarez et al. (2010), we also combine critical levels with differentiated shipment modes, but in a multi-item setting. In this model, demand that cannot be met from on-hand stock is either backordered or satisfied using an emergency shipment, which is similar to assuming a lost sale. The shipment option used depends on the customer class requesting the item and the item characteristics.

So far, all literature on critical level policies considers single-location models. In a multi-echelon setting, we find literature where the lowest echelon level consists of multiple locations that each have separate restrictions on performance (and as such are similar to a system where stock may be kept at customer sites). Under backordering, we find various contributions with recent ones from Wong et al. (2007) and Caggiano et al. (2008). Wong et al. (2007) consider a two-echelon system with a central depot and multiple local warehouses. Each warehouse as a service requirement in terms of a maximum mean waiting time for spares and the objective is to minimize system holding costs while meeting the service constraints per warehouse. The authors give exact and approximate approaches to

analyze the system. Furthermore, they present an optimization approach that is along similar lines as that of Kranenburg and Van Houtum (2008). Caggiano et al. (2008) consider a multi-item, multi-echelon system. They express the service restriction at each location in terms of channel fill rates, i.e. time-based fill rates where each location might have multiple fill rate restrictions (e.g. 90% of requests must be met instantaneously and 95% of requests must be met within 2 hours). Under lost sales, the contributions are limited, (see Bijvank and Vis (2011)). The analysis of such systems is complex, particularly of locations at higher echelons: there, the demand arrival process is often not Poisson and depends on the inventory states of locations at lower echelon levels. So far, literature considering multi-echelon systems with lost sales are either only accurate for limited problem instances (e.g. Andersson and Melchior, 2001) or under restrictive assumptions, e.g. that the transportation time from a central warehouse to a local stock point is at least the lead time to the central warehouse (e.g. Hill et al., 2007).

Our paper contributes to existing research in various ways. First, we extend the literature on two-echelon systems with lost sales. In our system, demand at a local stock point is only lost (and thus satisfied through an emergency shipment) if it cannot be met from stock at that stock point, from stock at the higher-level central location, or from items in the transportation pipeline between the two locations. We believe that such a system has not been considered in literature before, whereas it is a very reasonable model from a practical perspective: demand is only met through emergency shipments if no cheaper and faster alternatives are available. A companion paper (Alvarez and Van der Heijden, 2011) details how such a system can be accurately analyzed, whereas this paper gives an approach for optimizing such a system. Second, in an extensive experiment we investigate the added value of dedicated customer stocks by comparing it to alternative differentiation approaches. Such a comparison has not been done before, whereas it is very useful to know whether such an approach to differentiation can lead to large cost savings (and thus is a viable alternative to critical level policies). Finally, we consider a model where dedicated stocks and critical level policies are jointly used for differentiation to determine under what conditions each individual strategy (dedicated stocks, critical level policies) works best.

3 Model description

3.1 Outline

Consider a two-echelon network consisting of a warehouse supplying various items to multiple customers. We assume that all items are critical: any item failure causes a system failure. All customers belong to one of various customer classes, with each customer class having a distinct target service level in terms of a maximum time a customer of that class is willing to wait for spares.

To meet the various service requirements at minimal costs, the supplier can apply two differentiation strategies. First, he may keep *dedicated stocks (DS)* of certain items at a customer's facility next to stocks at the warehouse. Second, he may use a *critical level policy (CLP)*, where he concentrates all stocks at the warehouse and only uses warehouse stock to satisfy requests from a customer class if this stock exceeds a critical level for that customer class.

The supplier may opt to use the same differentiation strategy for all items. Furthermore, he may use a *combined (COMBO)* strategy, where the mode of differentiation (i.e. dedicated stocks or critical levels) can differ per item. In this case, only one differentiation strategy may be selected per item. The rationale behind COMBO is that an item's characteristics influence the added value of the individual differentiation strategies: we expect dedicated stocks to be most beneficial for inexpensive fast movers, since they are frequently requested and inexpensive to keep in stock. Conversely, for expensive slow movers it might be better to centralize stocks and differentiate through critical levels.

Irrespective of the strategy, a one-for-one replenishment policy is used at all locations. Furthermore, we consider two settings for dealing with demand that cannot be met from on-hand stock: (1) we *backorder* demand at all locations, or (2) we satisfy demand using *emergency shipments* from an central depot with infinite supply (effectively a lost sales setting). In the emergency shipments setting, we assume that a shipment from the warehouse to any customer is faster than an emergency shipment. As a result, emergency shipments are only used if both the customer and the warehouse are out of stock, and there are no items in transit between warehouse and customer. In literature on differentiation, unmet demand is often satisfied through emergency shipments, since in practice suppliers will try to obtain a part for a customer as quickly as possible (see e.g.

Kranenburg and Van Houtum, 2008). However, in an earlier paper (Alvarez et al., 2010) we have shown that in some cases emergency shipments are excessively costly. Therefore, we also consider a backorder setting.

Overall, we focus on minimizing the system's holding costs and, if applicable, shipment costs under constraints on the mean aggregate waiting time per customer. Our decision variables are the item base stock levels at the various locations in the system, and the critical levels per customer class at the warehouse.

3.2 Assumptions and notation

3.2.1 Model assumptions

- Demand for parts at any customer occurs according to mutually independent Poisson processes.
- The shipment time from the warehouse to any customer is deterministic.
- The emergency shipment time from depot to customer is deterministic. This is most realistic, although it is no problem for the model to include variability (then we use the mean only).
- The regular shipment time to the warehouse is exponentially distributed. Although deterministic shipment times are generally more realistic, this assumption facilitates a performance evaluation based on Markov chain analysis. Also, inventory models for slow moving parts tend to be quite insensitive to lead time variability (Alfredsson and Verrijdt (1999)).
- We use priority backorder clearing under CLP with backordering: backorders from a certain class will only be cleared once all backorders have been cleared from classes with higher service requirements, and the stock level at the warehouse is at least the critical level for that class. In contrast, first-come-first-served clearing is used in the DS strategy.
- Possible emergency shipments are sent directly to the customers.

3.2.2 Notation

We keep stock of I items for K customers; index 0 refers to the warehouse, and indexes $1, \dots, K$ to the customers. Each customer belongs to one of J demand classes, with a class j customer ($j = 1, \dots, J$) willing to wait at most W_j^{\max} time units on average for any item. Without loss of generality, we assume that W_j^{\max} is

increasing in j (i.e., class 1 has the tightest waiting time constraint). We let $q(k)$ denote the class to which customer k belongs. Demand for item $i = 1, \dots, I$ from customer k occurs at rate m_{ik} , with $M_k = \sum_{i=1}^I m_{ik}$ denoting the total demand rate from customer k and $M_i^j = \sum_{k|q_k=j} m_{ik}$ the total demand rate for item i coming from class j customers. For each item i , the shipment time to the warehouse is denoted by T_{i0}^{reg} , the mean regular shipment time from the warehouse to customer k is denoted by T_{ik}^{reg} and the emergency shipment time from the central depot to customer k is denoted by T_{ik}^{em} ($> T_{ik}^{reg}$). Finally, for each item i we denote the unit holding costs per time unit at location k (i.e. including the warehouse) by h_{ik} and the additional costs of an emergency shipment compared to a regular replenishment at customer k by EC_{ik}^{em} . We only require the additional shipment costs, since each demand triggers either a regular or an emergency shipment.

For each item i , we have as decision variables (1) the base stock level S_{ik} at each location k ($k = 0, \dots, K$), and (2) the critical level $C_i(j)$ for class j customers at the warehouse. Note that $0 \leq C_i(j) \leq S_{i0} \forall j$, since we cannot reserve more items than we have in stock at the warehouse. As class 1 has the tightest waiting time restriction, we have that $C_i(1) = 0$. We use vectors $\mathbf{S}_i = [S_{i0}, \dots, S_{iK}]$ and $\mathbf{C}_i = [C_i(1), \dots, C_i(J)]$ to denote respectively the item i stock levels and critical levels in the system. We combine all variables for item i in an item policy $(\mathbf{S}_i, \mathbf{C}_i)$. For each item policy, we have as performance measures the expected waiting time $EW_{ik}(\mathbf{S}_i, \mathbf{C}_i)$ and the fraction of demand met through emergency shipments $\gamma_{ik}(\mathbf{S}_i, \mathbf{C}_i)$ for item i and customer k , and the total costs $TC_i(\mathbf{S}_i, \mathbf{C}_i)$ for item i . We now express optimization problem (P1) as follows:

$$\begin{aligned}
(P1) \quad & \min \sum_{i=1}^I TC_i(\mathbf{S}_i, \mathbf{C}_i) = \sum_{i=1}^I \sum_{k=0}^K h_{ik} S_{ik} + \sum_{i=1}^I \sum_{k=1}^K \gamma_{ik}(\mathbf{S}_i, \mathbf{C}_i) m_{ik} EC_{ik}^{em} \\
\text{s. t.} \quad & \sum_{i=1}^I \frac{m_{ik}}{M_k} EW_{ik}(\mathbf{S}_i, \mathbf{C}_i) \leq W_{q(k)}^{\max} \quad k = 1, \dots, K \quad (P1.1) \\
& S_{ik}, C_i(j) \in \mathbf{N}_0 \quad i = 1, \dots, I, k = 0, \dots, K, j = 1, \dots, J
\end{aligned}$$

As mentioned, our system costs consist of holding costs and, if applicable, additional emergency shipment costs. Under backordering, $\gamma_{ik}(\mathbf{S}_i, \mathbf{C}_i)$ will be 0,

and thus the total costs will only consist of holding costs then. Holding costs are computed over the stock in the entire system, including items in transit to the customers. However, the model can be adjusted to compute holding costs over on-hand stock only: we then subtract the average number of items in transit, which is a constant value, from the stock level. Each customer k has a restriction on the mean aggregate waiting time over all items, with m_{ik} / M_k being the fraction of item i waiting time that contributes to the aggregate waiting time. Note that the waiting time threshold $W_{q(k)}^{\max}$ depends on the customer's class.

4 Solution approach

We solve problem $(P1)$ by using an approach based on decomposition and column generation which closely resembles Dantzig-Wolfe decomposition. This approach has been used before to solve nonlinear integer spare parts optimization problems with multiple items and aggregate waiting time restrictions over all items (see e.g. Kranenburg and Van Houtum (2007, 2008), Wong et al. (2007), and Alvarez et al. (2010)). In the approach, we reformulate $(P1)$ to a linear integer programming problem and solve its LP-relaxation to obtain a lower bound. Then, we obtain a near-optimal integer solution by solving the integer problem itself.

This solution approach is suitable for all the strategies we consider in this paper (i.e. dedicated stocks and critical level policies, both under backordering and emergency shipments). As we show, we can use this solution approach for any kind of strategy, provided that we are able to determine the performance measures (e.g. expected waiting times) for that strategy. We elaborate on the computation of performance measures in Section 5.

Section 4.1 gives the reformulated variant of $(P1)$. Sections 4.2 and 4.3 detail how to find a lower bound and near-optimal integer solution respectively.

4.1 Reformulation to a linear problem

We obtain the linear variant of $(P1)$ by considering a set of item policies for each item. Our decision problem now becomes to select one item policy for each item such that the system costs are minimized while the waiting time restrictions per customer are still met.

Let b_i be shorthand notation for policy $(\mathbf{S}_i, \mathbf{C}_i)$, with B_i denoting the policy set considered for item i . Let binary variable x_{b_i} specify whether b_i is selected for item i or not ($x_{b_i} = 1$ or 0 respectively). The reformulated problem (P2) becomes:

$$\begin{aligned} \min \quad & \sum_{i=1}^I \sum_{b_i \in B_i} TC_i(b_i) x_{b_i} \\ \text{s.t.} \quad & \sum_{i=1}^I \sum_{b_i \in B_i} \frac{m_{ik}}{M_k} EW_{ik}(b_i) x_{b_i} \leq W_{q(k)}^{\max} \quad k = 1, \dots, K \quad (P2.1) \\ & \sum_{b_i \in B_i} x_{b_i} = 1 \quad i = 1, \dots, I \quad (P2.2) \\ & x_{b_i} \in \{0, 1\} \quad i = 1, \dots, I, b_i \in B_i \end{aligned}$$

4.2 Lower bound

To solve the LP-relaxation of (P2), we must determine what item policies to include in B_i for item i . We use a similar procedure as in earlier papers, i.e. we first construct an initial policy set for each item and we use this set to solve the LP-relaxation a first time. Subsequently, we use column generation to iteratively find unconsidered item policies that further improve the solution value. We proceed in this way until we cannot find any more relevant policies. A critical part for our specific model is to limit the number of item policies to be evaluated in the column generation step. This is important, since the computation time for policy evaluation may explode if we do not select relevant policies carefully. In Section 4.2.1, we show how we find an initial policy set. In Section 4.2.2, we give the column generation problem and the main steps in solving this problem. Finally, in Section 4.2.3 we give the formal column generation procedure.

4.2.1 Creating an initial set of policies for each item

As only criterion, our initial item policy set must result in a feasible solution to the LP-relaxation. In literature, such a policy set is typically found using a greedy approach: For each item, a policy is found by iteratively increasing stock until the waiting time for that item satisfies the tightest upper limit $\min_{q(k)} W_{q(k)}^{\max}$. This condition ensures that the resulting solution will be feasible (see e.g. Kranenburg and Van Houtum, 2008). As disadvantage, the stock levels found tend to be much

larger than in the optimal solution. Therefore, we look for a better initial solution by constructing a policy set over all items simultaneously.

Our approach is similar to the greedy approach by Wong et al. (2007) to find an integer solution (see Section 3.3 of their paper). We limit ourselves to policies without critical levels, irrespective of the strategy (DS or CLP, backordering or emergency shipments) being considered. Iteratively, we add stock at the item-location combination resulting in the largest decrease in the so-called *distance to the feasible region* per euro additional costs. We shall define this distance measure below. If options exist that are both closer to the feasible region *and* less expensive, we select the option closest to the feasible region, since our focus is to find a feasible solution. The procedure stops once the resulting customer waiting times satisfy the overall targets. Note that this approach is also suitable under CLP, provided that we only add stock at the warehouse.

We define the distance $Dist_i(\mathbf{S}_i)$ of a solution to the feasible region as the amount by which the customer waiting times exceed the thresholds, i.e.

$$Dist_i(\mathbf{S}_i) = \sum_{n=1}^K \left[\frac{m_{in}}{M_n} EW_{in}(\mathbf{S}_i) - W_{q(n)}^{\max} \right]^+ \quad \text{with } [a]^+ = \max\{0, a\} \quad (1)$$

At the end of the procedure, we have one item policy per item consisting of the system stock levels found. However, note that we have analyzed various other item policies as well during the procedure (i.e. in each iteration we changed one stock level value). By also adding these policies to the initial policy set, we limit the number of additional policies that must be found through column generation. We realize that some of the added policies might be poor options. Therefore, when looking for an integer solution later on, we first remove those poor policies from our policy set before optimizing the integer problem. Section 4.3 gives further details on the criteria we use to remove these policies.

4.2.2 The column generation problem

Through column generation, we iteratively look for unconsidered item policies that have negative reduced costs. Such policies further improve the solution value of the LP-relaxation. In each iteration, we find per item i the policy with minimum reduced costs and we add this policy to B_i if these reduced costs are negative. We proceed in this way until we cannot find any policy with negative reduced costs. We use the shadow prices found when solving the LP-relaxation as

input to obtain a policy's reduced costs. The reduced costs $RED(b_i)$ for policy b_i is given by equation (2), with $u_k \leq 0$ and $v_i \geq 0$ denoting the shadow prices associated with restrictions (P2.1) and (P2.2) respectively.

$$RED_i(b_i) = RED_i(\mathbf{S}_i, \mathbf{C}_i) = TC_i(\mathbf{S}_i, \mathbf{C}_i) - \sum_{k=1}^K \frac{u_k m_{ik}}{M_k} EW_{ik}(\mathbf{S}_i, \mathbf{C}_i) - v_i \quad (2)$$

For three strategies, we refer to earlier literature for the column generation procedures. Specifically, for DS under backordering we refer to Wong et al. (2007), for CLP under backordering we refer to Alvarez et al. (2010), and for CLP under emergency shipments we refer to Kranenburg and Van Houtum (2008). In the remainder of this section, we give the column generation procedure for DS under emergency shipments. For combination of DS and CLP (i.e. the COMBO strategies), we simply use both the procedures for DS and CLP. As we only consider DS, we omit the vector \mathbf{C}_i in the remainder of this section.

A complication under DS with emergency shipments is that the service level (i.e. fill rates and waiting times) at a customer does not only depend on the stock level at that customer and at the warehouse, but also on the stock levels at all other customers. In particular, if stock is increased at a customer k while all other customer stock levels remain unchanged, the service level at customer k will improve at the expense of the service levels at all other customers: the warehouse then sees a relatively large arrival rate from customer k due to replenishment requests, and must satisfy these additional demands using the same amount of stock. As a result, a smaller fraction of replenishment requests from other customers can be met from warehouse stock. Irrespective of this complication, we can find the policy with minimum reduced costs from the following observations:

- **Observation 1:** We can find an upper bound $S_{i_0}^{\max}$ on S_{i_0} . Let $RED_i^*(S_{i_0})$ denote the minimal reduced costs for a given S_{i_0} over the stock levels at all customer locations S_{ik} ($k \geq 1$). Equation (2) shows that $RED_i^*(S_{i_0})$ will be at least $h_{i_0} S_{i_0} - v_i$ since the total cost include the holding costs for the stock at the warehouse, and $u_k \leq 0$. We ignore the cost elements related to the various customers (e.g. the waiting times), as these depend on the other stock levels as well. We find $S_{i_0}^{\max}$ as follows: (i) We determine an upper bound RED_i^{UB} on the reduced costs by setting all customer stocks S_{ik} to zero and finding the

value of S_{i0} leading to the lowest reduced costs; given that $S_{ik} = 0$ ($k \geq 1$), the reduced costs are convex in S_{i0} , see Kranenburg and Van Houtum (2007). (ii) We find S_{i0}^{\max} as the smallest S_{i0} for which $h_{i0}S_{i0} - v_i$ exceeds $\min\{RED_i^{UB}, 0\}$. It is sufficient for the lower bound on the reduced costs $h_{i0}S_{i0} - v_i$ to be nonnegative, as we focus on item policies with negative reduced costs.

- **Observation 2:** We can find a rough upper bound S_{ik}^{MAX} on S_{ik} ($k \geq 1$). An increase of S_{ik} can only benefit the service level at customer k . Hence, we find S_{ik}^{MAX} once the additional holding costs of increasing S_{ik} outweigh the maximum reduction in that customer's emergency shipment and waiting time costs, i.e. S_{ik}^{MAX} is the smallest S_{ik} for which h_{ik} exceeds $\gamma_{ik}(\mathbf{S}_i)m_{ik}EC_{ik}^{em} - \frac{u_k m_{ik}}{M_k}EW_{ik}(\mathbf{S}_i)$. To ensure that S_{ik}^{MAX} is sufficiently large, we require upper bounds on $\gamma_{ik}(\mathbf{S}_i)$ and $EW_{ik}(\mathbf{S}_i)$. We find such bounds from the special case where demand at customer k can only be met from on-hand stock at that customer (i.e. that customer has no access to warehouse stock). Then, we have the worst-case scenario in terms of service level. The resulting system can be analyzed as an Erlang-loss system with S_{ik} servers.
- **Observation 3:** For a given value of S_{i0} , we can now find a tighter upper bound on S_{ik} ($k \geq 1$), denoted by $S_{ik}^{\max}(S_{i0})$. As in observation 2, we find $S_{ik}^{\max}(S_{i0})$ when the holding costs of increasing S_{ik} exceed the emergency shipment and waiting time costs of customer k . Compared to S_{ik}^{MAX} (in Observation 1), we now use more accurate values for $\gamma_{ik}(\mathbf{S}_i)$ and $EW_{ik}(\mathbf{S}_i)$, which we find by also considering the stock kept at other locations in the system. Specifically, we set all other customer stocks S_{in} $n \neq k$ to their rough upper bounds S_{in}^{MAX} and then we determine $\gamma_{ik}(\mathbf{S}_i)$ and $EW_{ik}(\mathbf{S}_i)$. As the service level at customer k is lowest when the stock levels at other customers are large, the values for $\gamma_{ik}(\mathbf{S}_i)$ and $EW_{ik}(\mathbf{S}_i)$ will still be sufficiently large.

In addition to these observations, we empirically find that the optimal value of S_{ik} for a given S_{i0} , denoted by $\hat{S}_{ik}(S_{i0})$, always lies between two thresholds $S'_{ik}(S_{i0})$

and $S''_{ik}(S_{i0})$. We find $S'_{ik}(S_{i0})$ as the value of $S_{ik} \in \{0, \dots, S_{ik}^{\max}(S_{i0})\}$ that minimizes $RED(\mathbf{S}_i)$ when all other customer stocks S_{in} $n \neq k$ are set to their upper bounds $S_{in}^{\max}(S_{i0})$. Similarly, we find $S''_{ik}(S_{i0})$ as the optimal S_{ik} when all other customer stocks are at their lower bounds $S_{in}^{\min}(S_{i0})=0$. As we look for optimal values of S_{ik} in two extreme cases (the remaining customer stocks are either at their maximum or at their minimum), we expect the true optimum to lie between these values. Note that $S'_{ik}(S_{i0})$ and $S''_{ik}(S_{i0})$ in fact give us new bounds on $\hat{S}_{ik}(S_{i0})$. We can thus repeat the mentioned steps (i.e. we can find new values for $S'_{ik}(S_{i0})$ and $S''_{ik}(S_{i0})$) by updating $S_{ik}^{\min}(S_{i0})$ and $S_{ik}^{\max}(S_{i0})$. We proceed in this way until the bounds stabilize (either because the values for $S_{ik}^{\min}(S_{i0})$ and $S_{ik}^{\max}(S_{i0})$ no longer change or because $S'_{ik}(S_{i0})=S''_{ik}(S_{i0})$ for all customers k).

Overall, our column generation procedure works as follows: We increase S_{i0} from zero up to S_{i0}^{\max} with step size 1. In each step, we first compute $S'_{ik}(S_{i0})$ and $S''_{ik}(S_{i0})$ for each customer k . Then, we look for the combination of customer stock levels that has minimum reduced costs, given that $S'_{ik}(S_{i0}) \leq S_{ik} \leq S''_{ik}(S_{i0})$.

4.2.3 The formal steps the column generation procedure

Full column generation procedure

1. Find S_{i0}^{\max} from observation 1.
2. For each customer k , find a rough upper bound S_{ik}^{MAX} on the optimal stock level (observation 2).
3. For each $S_{i0} \in \{0, \dots, S_{i0}^{\max}\}$ do:
 - a. Find a tighter upper bound $S_{ik}^{\max}(S_{i0})$ on the optimal stock level for customer k (see observation 3).
 - b. Find thresholds $S'_{ik}(S_{i0})$ and $S''_{ik}(S_{i0})$ for each customer k .
 - c. Find the customer stock combination $[S_{i1}, \dots, S_{iK}]$ that minimizes $RED(\mathbf{S}_i)$, with $S'_{ik}(S_{i0}) \leq S_{ik} \leq S''_{ik}(S_{i0})$.

- d. If the solution is the best so far, store it. Also store the related reduced costs as RED_i^* . If $h_{i0}(S_{i0} + 1) - v_i$ (i.e. the lower bound on the reduced cost for $S_{i0} + 1$) exceeds RED_i^* , exit the procedure.

Next, we give further details on steps 1 through 3b.

Step 1. Finding S_{i0}^{\max} .

1. Determine an upper bound RED_i^{UB} on the reduced costs.
 - a. Set all customer stocks S_{ik} to zero ($k \geq 1$).
 - b. Find the S_{i0} that minimizes $RED_i(\mathbf{S}_i)$. Set RED_i^{UB} to this value.
2. Find S_{i0}^{\max} as the smallest S_{i0} for which $h_{i0}S_{i0} - v_i$ exceeds $\min\{RED_i^{UB}, 0\}$.

Step 2. Finding S_{ik}^{MAX} for each customer k .

1. Consider an Erlang Loss system with S_{ik} servers and replenishment rate $\mu_{ik} = 1/(T_{ik}^{reg} + T_{i0}^{reg})$. Our performance measures now only depend on S_{ik} : we find $\gamma_{ik}(S_{ik})$ as the probability of all servers being occupied, with $EW_{ik}(S_{ik})$ being equal to $T_{ik}^{em} \gamma_{ik}(S_{ik})$.
2. Find $S_{ik}^{\max}(S_{i0})$ as the smallest value of S_{ik} for which h_{ik} exceeds $\gamma_{ik}(S_{ik})m_{ik}EC_{ik}^{em} - \frac{u_k m_{ik}}{M_k} EW_{ik}(S_{ik})$. From that moment, the reduced costs cannot improve further.

Step 3a. Finding $S_{ik}^{\max}(S_{i0})$ customer k for any S_{i0} .

1. Set all other customer stocks S_{in} $n \neq k$ to S_{in}^{MAX} .
2. Find $S_{ik}^{\max}(S_{i0})$ as the smallest S_{ik} for which h_{ik} exceeds $\gamma_{ik}(\mathbf{S}_i)m_{ik}EC_{ik}^{em} - \frac{u_k m_{ik}}{M_k} EW_{ik}(\mathbf{S}_i)$.

Step 3b. Finding $S'_{ik}(S_{i0})$ and $S''_{ik}(S_{i0})$.

1. Set all customer lower bounds $S_{ik}^{\min}(S_{i0})$ ($k \geq 1$) to 0.
2. Find $S'_{ik}(S_{i0})$ for each customer k .

- a. Set all other customer stocks $S_{in} \ n \neq k$ to $S_{in}^{\max}(S_{i0})$.
 - b. Find $S'_{ik}(S_{i0})$ as the $S_{in} \in \{S_{ik}^{\min}(S_{i0}), \dots, S_{ik}^{\max}(S_{i0})\}$ that minimizes $RED(\mathbf{S}_i)$.
3. Find $S''_{ik}(S_{i0})$ for each customer k .
 - a. Set all other customer stocks $S_{in} \ n \neq k$ to $S_{in}^{\min}(S_{i0})$.
 - b. Find $S''_{ik}(S_{i0})$ as the $S_{in} \in \{S_{ik}^{\min}(S_{i0}), \dots, S_{ik}^{\max}(S_{i0})\}$ that minimizes $RED(\mathbf{S}_i)$.
 4. Exit if (i) $S'_{ik}(S_{i0}) = S''_{ik}(S_{i0})$ for all customers k , or (ii) neither $S'_{ik}(S_{i0})$ nor $S''_{ik}(S_{i0})$ has changed compared to the previous iteration for any customer. Otherwise, set $S_{ik}^{\min}(S_{i0})$ to $\min\{S'_{ik}(S_{i0}), S''_{ik}(S_{i0})\}$ and $S_{ik}^{\max}(S_{i0})$ to $\max\{S'_{ik}(S_{i0}), S''_{ik}(S_{i0})\}$ and proceed to step 2.

4.3 Near-optimal integer solution

The optimal solution to the LP-relaxation might be fractional, i.e. it might be that a combination of item policies has been selected for certain items. Therefore, we also require an approach to find a near-optimal integer solution. We obtain such a solution by solving the *integer* problem ($P2$) using a limited set of item policies. This approach worked well in an earlier paper (cf. Alvarez et al., 2010), where it outperformed a greedy heuristic in terms of solution quality.

In Alvarez et al. (2010), we solve the integer problem using the set of item policies generated when solving the LP-relaxation of ($P2$). In this paper, this policy set is a starting point, as it might contain many policies: when constructing our initial policy set (Section 4.2.1), we included all found policies in B_i . We also added additional policies during column generation. Such a large policy set is not an issue when solving an LP-relaxation, but computation times might explode when solving the integer problem. Therefore, we eliminate all *dominated* item policies from the LP-relaxation set before solving the integer problem. Dominated policies have both higher costs and higher waiting times than at least one other policy in the policy set. As a result, such policies will never be chosen and they can thus be eliminated from the policy set without sacrificing solution quality.

5 Evaluation of an item policy

Our solution approach requires as input the performance measures for each item policy. In this section, we specify how we can find such performance measures

using Markov chain analysis. For most item policies, the evaluation procedure is not new and has already been described in earlier papers. However, we needed a new procedure for evaluating a policy under dedicated stocks (DS) with emergency shipments. We describe this procedure in detail in a companion paper (Alvarez and Van der Heijden, 2011). In this paper, we limit ourselves to globally specifying the main evaluation steps. In Sections 5.1 and 5.2 we describe the evaluation under a critical level policy (CLP) and under DS respectively. For simplicity, we limit ourselves to two customer classes in this section and the remainder of the paper. However, in many cases the analysis can be extended to more classes. As we only have one relevant critical level in a two-class model (i.e. $C_i(2)$), we denote this critical level by C_i for simplicity.

5.1 Evaluation under a critical level policy (CLP)

Under CLP, we only keep stock at the warehouse. As a result, we only need to analyze the warehouse to obtain the performance measures needed. Indeed, under emergency shipments we find performance measures from the distribution of the *pipeline* to the warehouse: if we have more than $S_{i0} - C_i$ items in the pipeline, demands from both customer classes are satisfied from warehouse stock. Otherwise, only demand from premium customers is met from on-hand stock if possible, with non-premium demand being lost. Kranenburg and Van Houtum (2008) further detail how an item policy can be analyzed under lost sales.

In contrast, under backordering we can no longer analyze the system using only the pipeline to the warehouse: once we have $S_{i0} - C_i$ or more items in the pipeline, we can simultaneously have stock on-hand and backorders for non-premium customers. As a result, we need a two-dimensional state space (w, z) to analyze the warehouse, with w the number in the pipeline and z the number of class 2 backorders. From these values, we are able to compute the number of class 1 backorders as well. Further details are given in Alvarez et al. (2010).

Under backordering, it is difficult to analyze the system for more than two demand classes, as each additional class adds an additional dimension to the Markov chain. In contrast, such an extension is straightforward under emergency shipments, as shown by Kranenburg and Van Houtum (2008).

5.2 Evaluation under dedicated stocks (DS)

Under DS, we can also keep stock at the customers. As a result, we must analyze a two-echelon system to find the needed performance measures: Naturally, we must analyze the availability of stock at each customer separately. Furthermore, we must analyze the availability of stock at the warehouse, as this availability influences the lead time for replenishment orders to each customer.

Under full backordering, Wong et al. (2007) analyze such a system using the following steps: first, they analyze the warehouse to obtain per customer the distribution of items outstanding at the warehouse (i.e. the items that still need to be shipped to that customer). Using this distribution, the authors determine the distribution of the pipeline to each customer, which consists of the items outstanding for that customer at the warehouse and the additional demand that occurs during the shipment time from warehouse to customer. Finally, the authors use these pipeline distributions to determine the expected number of backorders at each customer (resulting in an expected waiting time through Little's Law).

Under emergency shipments, the main analysis steps are similar to the backordering case. However, the analysis of the warehouse is now far from trivial: If the warehouse is out of stock, it will only receive item requests from customers that either still have stock on-hand or have stock in transit from the warehouse (demand from other customers is met through emergency shipments). As a result, the distribution of outstanding items at the warehouse depends on the availability of stock in the entire system. Under backordering, this complication does not exist, as each customer demand triggers a demand at the warehouse. In a companion paper (Alvarez and Van der Heijden, 2011), we detail how such a system can still be efficiently analyzed.

Note that a customer's demand class has no influence on the analysis of the system, so an extension to more demand classes is straightforward.

6 Computational experiment

In this section, we describe our computational experiment. We give the objectives in Section 6.1, the experiment design in Section 6.2, and the results in Section 6.3.

6.1 Experiment objectives

First, we investigate the performance of our optimization approach for the DS strategy with emergency shipments in terms of solution quality – expressed as a relative gap to the lower bound – and computation time. Second, we evaluate the added value of using dedicated customer stocks for differentiation. We do so by comparing results under DS to those under a one-size-fits-all approach and under CLP. Finally, we investigate the added value of the COMBO strategy where the differentiation mode (i.e. dedicated stocks or critical levels) may differ per item.

6.2 Experiment design

Figure 1 shows the parameter values we used for our problem instances.

Parameter	Values
1. Number of items I	20; 100
2. Number of customers K	8; 16
3. Percentage premium customers (% of K)	12.5; 25
4. $(W_1^{\max}; W_2^{\max})$ (days)	$(\frac{1}{12}; \frac{1}{3}); (\frac{1}{12}; \frac{2}{3})$
5. Intervals for demand rates m_{ik} (per day)	$[0.002 - 0.05]; [0.002 - 0.25]$
6. Intervals for holding costs h_i (per day)	$[10 - 100]; [10 - 1000]$
7. T_{i0}^{reg} (days)	5; 15
8. T_{ik}^{reg} (days)	$\frac{1}{48}; \frac{1}{16}$
9. T_{ik}^{em} (as a % of T_{i0}^{reg})	10; 20
10. EC_{ik}^{em} (per shipment)	1000

Figure 1 Parameter settings in problem instances

Parameters 2 and 3 give the number of premium customers in the system. Note that T_{ik}^{reg} has very small values compared to the other shipment times: under CLP, the mean waiting time for each customer will at least be T_{ik}^{reg} . Hence, we can only find solutions under CLP if T_{ik}^{reg} is smaller than W_1^{\max} .

Except for the demand rates and holding costs, the parameter values are the same for all items, and if applicable all customers, in a problem instance. The demand rates m_{ik} and holding costs h_i are randomly drawn from uniform distributions on

the specified intervals. We use the same holding cost value at all locations in a problem instance. Therefore, we omit suffix k in h_i .

For each combination of parameters 1 through 9, we create 3 samples of demand rates and holding costs, thereby ensuring that our results are not influenced by the specific values of one sample. In total, we have 2304 instances: 3 (samples) $\times 2^8$ (parameters 1..8) = 768 instances for each of the strategies (i) backordering, (ii) emergency shipments with T_{ik}^{em} as 10% of T_{i0}^{reg} , (iii) emergency shipments with T_{ik}^{em} as 20% of T_{i0}^{reg} .

6.3 Results

We discuss the results for each experiment objective in a separate section.

6.3.1 Performance of the optimization approach

We determine the solution quality of the optimization approach by comparing the integer solutions found to their lower bounds. The solution quality is expressed as a gap to the lower bound, i.e. $(TC_{IP} - TC_{LB})/TC_{LB}$ with TC_{LB} and TC_{IP} respectively denoting the lower bound and integer solution value.

Table 1 shows the solution quality and computation times of the approach. The approach works very well in terms of solution quality: the average gap is 0.3% and the maximum gap is 2.2%. Also, the maximum gap decreases greatly as the number of items increases. The procedure will thus give good solutions for practical instances with many items. The average computation time of an instance is 23 minutes, with a maximum of 940 minutes. However, most instances (91%) require fewer than 60 minutes computation time, with only 9 instances (out of 1536) having a computation time larger than 360 minutes. The bulk of the computation time lies in the column generation step, where sometimes many item policies need to be evaluated. Furthermore, the computation time mainly depends on the number of items and customers, and the shipment time to the warehouse.

Parameter	Values	Relative gap to lower bound		Computation time (mins)	
		Average	Maximum	Average	Maximum
Num. Items I	20	0.4%	2.2%	10	404
	100	0.1%	0.2%	36	940
Num. Customers K	8	0.2%	2.2%	3	28
	16	0.3%	2.0%	43	940
T_{i0}^{reg}	5	0.2%	2.2%	4	36
	15	0.3%	2.0%	42	940
Grand Total		0.3%	2.2%	23	940

Table 1 Performance of the optimization approach for DS with emergency shipments

6.3.2 The added value of dedicated stocks (DS)

We determine the added value of DS by comparing the solutions under DS to those under a one-size-fits-all (OSFA) approach and under critical level policies (CLP). Under OSFA, stock may only be kept at the warehouse, with no stock reserved for premium customers. Possible backorders at the warehouse are cleared first-come-first-served. We compute the relative cost saving of DS and CLP over OSFA, expressed as $(TC_{OSFA} - TC_{DS})/TC_{OSFA}$. Here, TC_{OSFA} denotes the costs under OSFA and TC_{DS} the costs under a differentiation strategy. Figure 2 shows the overall savings of DS over OSFA, both under backordering and under emergency shipments (columns BO and ES respectively). Under backordering, the average savings are 14% with a maximum of 36%. Under emergency shipments, the savings are even larger (20% on average and a maximum of 63%).

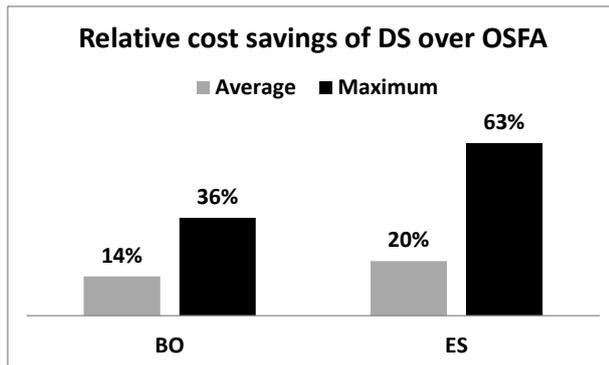


Figure 2 Overall cost savings of dedicated stocks (DS) over one-size-fits-all (OSFA)

Table 2 shows the parameters of greatest influence on the savings for the emergency shipment cases. Except for the holding cost values, the trends are similar for the backordering cases. Dedicated stocks are particularly beneficial if we have relatively few premium customers. Then, we only need to keep stock at those few premium customers to effectively differentiate. As the (relative)

number of premium customers grows, we either need to keep stock at more customers (which increases holding costs), or we find that it is more beneficial to pool stock centrally – which benefits all customers – instead of keeping dedicated stock. Either way, the savings compared to OSFA decrease. Dedicated stocks are also beneficial when the shipment time from the warehouse to the customers is large: As this time increases, it becomes more interesting to keep stock at the customer instead of centrally. Indeed, in practical settings – where shipment times to customers are often larger than the times we tested – dedicated stocks might even be necessary for meeting the premium customer requirements.

Parameter	Values	Relative savings over OSFA	
		Average	Maximum
Num. Customers K	8	23%	63%
	16	17%	59%
Percentage premium customers (%)	0.125	24%	63%
	0.25	16%	51%
T_{ik}^{reg}	0.5	15%	52%
	1.5	25%	63%
$(W_1^{max}; W_2^{max})$	(2; 8)	17%	58%
	(2; 16)	23%	63%
Holding cost interval	[10 - 100]	9%	37%
	[10 - 1000]	31%	63%
Grand Total		20%	63%

Table 2 Parameter-specific savings of DS over OSFA under emergency shipments

In contrast to the backordering cases, where the holding cost values have very little influence on the savings, we see that the savings under emergency shipments clearly increase as the holding costs increase. This is caused by the height of the holding costs relative to the emergency shipment costs: when holding costs are low, we keep stock both to meet waiting time requirements and to limit expensive emergency shipments. As a result, it is most beneficial to keep stock at the warehouse, where it reduces the number of emergency shipments needed for *all* customers. Dedicated stocks are not beneficial in this setting, since they benefit only those customers where stock is kept. In contrast, when holding costs are high, emergency shipments are relatively inexpensive and we focus less on avoiding them. As a result, we use dedicated stocks to minimize the gap between actual and required waiting times, especially for non-premium customers (see Figure 3 for the aggregate non-premium waiting times when $W_2^{max} = \frac{2}{3}$ days).

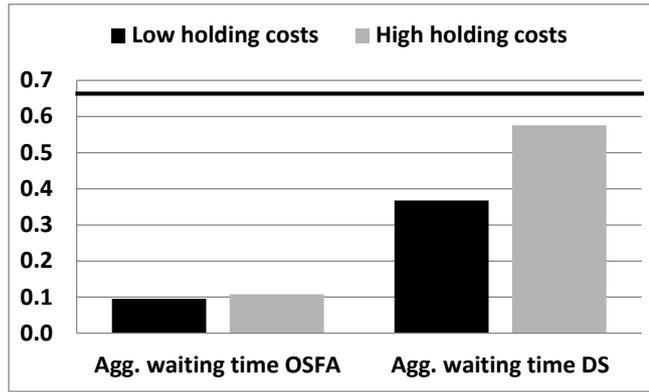


Figure 3 Aggregate non-premium waiting times under OSFA and DS (waiting threshold = 2/3 days)

Further analysis of the solutions shows that we mainly keep stock at premium customers' sites: on average, we keep 61% of items in stock at premium sites, with only 1% kept in stock at non-premium sites. However, we often only keep 1 unit of an item at any customer. Overall, we keep dedicated stocks of fast-moving items that are slightly cheaper than average (see Figure 4).

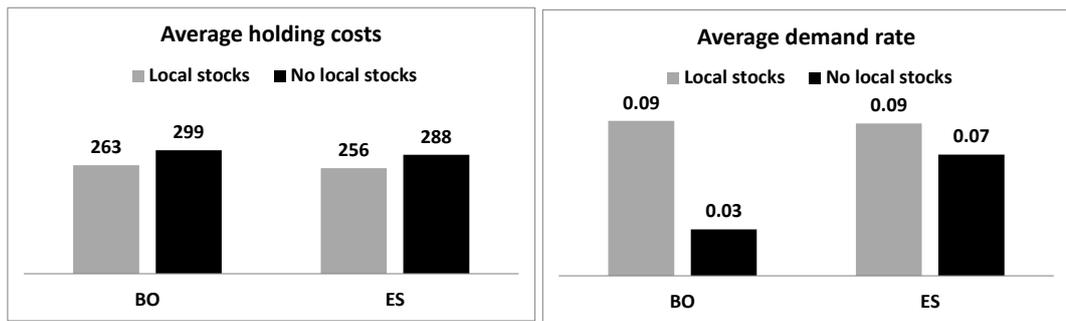


Figure 4 The item characteristics of items that are kept at customers vs. items that are not

Table 3 shows the savings of both DS and CLP over OSFA. Note that the savings under DS are close to those under CLP, particularly when emergency shipments are used. As we expect DS to also be easier to implement in practice, we consider it a very viable alternative to CLP.

Parameter	Values	Backordering		Em. Shipments	
		DS	CLP	DS	CLP
Num. Customers K	8	17%	19%	23%	21%
	16	12%	19%	17%	21%
Percentage premium customers (%)	0.125	18%	20%	24%	23%
	0.25	11%	18%	16%	20%
T_{ik}^{reg}	0.5	10%	17%	15%	19%
	1.5	19%	21%	25%	23%
Grand Total		14%	19%	20%	21%

Table 3 Relative savings of DS and CLP over OSFA

6.3.3 The combined policy

Figure 5 summarizes the overall savings of the various differentiation strategies over OSFA, including the COMBO strategy where the mode of differentiation (dedicated stocks, critical level policies) may vary per item. Note that the savings under COMBO are not much larger than those under DS or CLP.

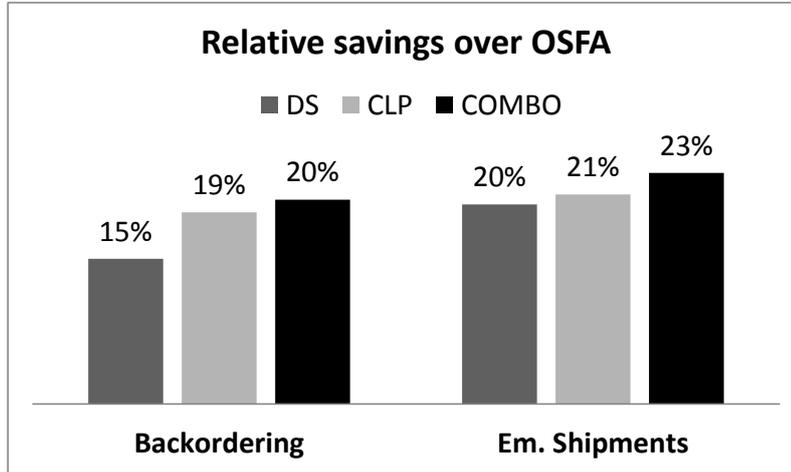


Figure 5 Overall savings over OSFA of the various differentiation strategies

Even though the additional savings under COMBO are not very large, we do see that most solutions (68% of the instances) contain a mix of both dedicated stocks and critical level policies. Still, the fraction of items per strategy (DS or CLP) depends greatly on the parameter values, see Table 4.

Parameter	Values	Percentage of items per strategy			
		Backordering		Em. Shipments	
		DS	CLP	DS	CLP
Num. Customers K	8	29%	38%	42%	40%
	16	21%	42%	26%	60%
Percentage premium customers (%)	0.125	25%	32%	40%	41%
	0.25	23%	45%	28%	58%
T_{ik}^{reg}	0.5	10%	41%	21%	61%
	1.5	36%	41%	43%	43%
(W_1^{max}, W_2^{max})	[2 - 8]	23%	29%	31%	48%
	[2 - 16]	24%	53%	35%	54%
Maximum daily demand rate	0.05	21%	47%	26%	56%
	0.25	26%	35%	39%	47%
Holding cost interval	[10 - 100]	23%	41%	29%	45%
	[10 - 1000]	24%	41%	35%	56%
Grand Total		24%	41%	33%	52%

Table 4 Average percentage of items per differentiation strategy in the COMBO approach

Finally, for each strategy (DS, CLP) we analyzed the item holding costs of the items assigned to that strategy. We find that DS is mainly used for inexpensive

items, with CLP being used for expensive items. Figure 6 shows the average holding costs per strategy over the instances with holding cost interval [10 – 1000]. We find similar figures for other intervals. We were unable to draw clear conclusions on the item demand rates per strategy.

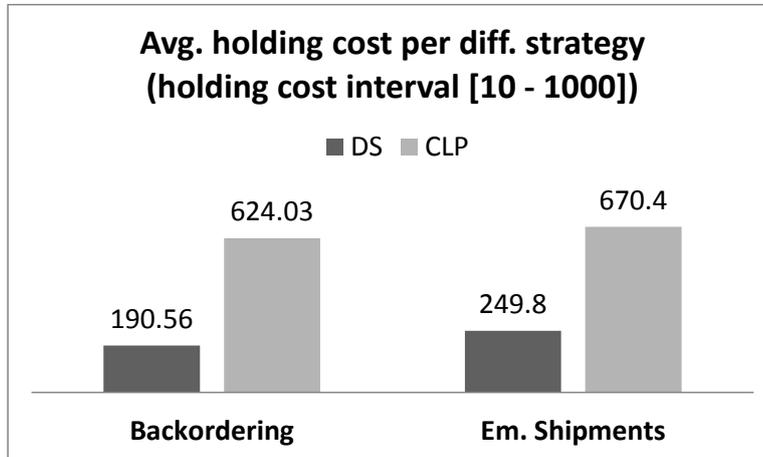


Figure 6 The average item holding costs per differentiation strategy

7 Conclusions and further research

7.1 Conclusions

- **The optimization approach of Section 4 gives near-optimal solutions under dedicated stocks with emergency shipments:** the average and maximum gap to the lower bound are 0.3% and 2.2% respectively. The approach works particularly well (maximum gap of 0.2%) when there are many items in a problem instance.
- **Dedicated stocks have significant added value.** Under backordering, DS leads to average savings of 14% compared to an approach where no differentiation is used (i.e. OSFA), with a maximum of 36%. Under emergency shipments, the average and maximum savings even amount to 20% and 63% respectively. Furthermore, the savings obtained under DS are comparable to those under CLP (who has average savings of 19% and 21% under backordering and emergency shipments respectively).
- **Dedicated stocks are very beneficial, if not necessary, when the shipment time to customers is large.** As shipment times to customers increase, it might no longer be possible to only keep stock centrally if customers have high service requirements. So far, this fact has been largely ignored in literature on

critical level policies, where the shipment time to customers is assumed to be negligible as it is often much smaller than the shipment time to the warehouse.

- **We find relatively small additional gains under the combined strategy (COMBO) compared to DS or CLP.** The practical relevance of this observation is that dedicated stocks indeed have significant added value, as we do not find much greater savings by adding critical levels.
- **Under the combined strategy, we keep dedicated stocks of inexpensive items, while using critical level policies for expensive items.**

7.2 Further research options

We also discuss the feasibility and relevance of various research options:

- **More efficient optimization approaches.** The optimization approach of Section 4 currently requires a lot of computation time, particularly the column generation step. The main cause of this computation time is the fact that many item policies need to be analyzed during column generation. Further research is thus necessary to limit the number of item policies that must be analyzed.
- **More sophisticated shipment strategies.** At present, we use backordering for all items and demand classes, or we always use emergency shipments. By also distinguishing the shipment mode per item (and possibly per class), further savings might be possible. This extension will not impact the methods for analyzing an item policy. However, the number of item policies to choose from then increases greatly, which complicates the optimization procedure of Section 4. In particular, further research is needed for efficiently finding item policies with negative reduced costs.
- **Combining dedicated stocks and critical level policies for a single item.** In this paper, we did not consider strategies where dedicated stocks and critical levels can jointly be used for the same item. However, we expect this combination to be of little added value: for a subset¹ of problem instances, we compared the solutions under the current COMBO approach (i.e. at most one differentiation option per item) with backordering to the approach where both differentiation options may be used per item. We concluded that the solutions rarely changed: the relative difference between the solutions of the two approaches was at most 0.6%.

¹ We considered one sample of demand rates and holding costs.

- **Alternative differentiation techniques:** In addition to dedicated stocks, we see the selective use of *lateral transshipments* between warehouses as an additional promising tool for applying differentiation in spare parts supply.

References

- Alfredsson, P. and Verrijdt, J. (1999). Modeling emergency supply flexibility in a two-echelon inventory system. *Management science*, 45(10), 1416–1431
- Alvarez, E.M., Van der Heijden, M.C., Zijm, W.H.M. (2010). The selective use of emergency shipments for service-contract differentiation. Beta working paper 322, http://cms.ieis.tue.nl/Beta/Files/WorkingPapers/wp_322.pdf
- Alvarez, E.M. and Van der Heijden, M.C. (2011). On two-echelon inventory systems with Poisson demand and lost sales. Beta working paper 366, http://cms.ieis.tue.nl/Beta/Files/WorkingPapers/wp_366.pdf
- Andersson, J. & Melchior, P. (2001), Two-echelon inventory model with lost sales, *International Journal of Production Economics*, 69(3), 307-315.
- Bijvank, M. & Vis, I.F.A. (2011), Lost-sales inventory theory: A review, *European Journal of Operational Research*, 215(1), 1-13.
- Caggiano, K.E., Jackson, P.L., Muckstadt, J.A., Rappold, J.A. (2007). Optimizing Service Parts Inventory in a Multiechelon, Multi-item Supply Chain with Time-based Customer Service-Level Agreements. *Operations Research*, 55, 303 – 318
- Enders, P., Adan, I., Scheller-Wolf, A., and Van Houtum, G-J. (2008). Inventory Rationing for a System with Heterogeneous Customer Classes. *Tepper School of Business*, paper 431, <http://repository.cmu.edu/tepper/431>
- Hill, R.M. (2007), Continuous-review, lost-sales inventory models with Poisson demand, a fixed lead time and no fixed order cost, *European Journal of Operational Research*, 176(2), 956-963.
- Kranenburg, A.A., and Van Houtum, G.J. (2007). Effect of commonality on spare parts provisioning costs for capital goods. *International Journal of Production Economics*, 108, 221 – 227.
- Kranenburg, A.A., and Van Houtum, G.J. (2008). Service differentiation in spare parts inventory management. *Journal of the Operational Research Society*, 59, 946 – 955
- Veinott, A.F. (1965). Optimal Policy in a Dynamic, Single Product, Nonstationary Inventory Model with Several Demand Classes. *Operations Research*, 13(5), pp. 761 – 778
- Wong, H., Kranenburg, B., V. Houtum, G.-J., Cattrysse, D. (2007). Efficient heuristics for two-echelon spare parts inventory systems with an aggregate mean waiting time constraint per local warehouse. *OR spectrum*, 29, 699-722

Working Papers Beta 2009 - 2012

nr.	Year	Title	Author(s)
373	2012	Service differentiation in spare parts supply through dedicated stocks	E.M. Alvarez, M.C. van der Heijden, W.H.M. Zijm
372	2012	Spare parts inventory pooling: how to share the benefits	Frank Karsten, Rob Basten
371	2012	Condition based spare parts supply	X.Lin, R.J.I. Basten, A.A. Kranenburg, G.J. van Houtum
370	2012	Using Simulation to Assess the Opportunities of Dynamic Waste Collection	Martijn Mes
369	2012	Aggregate overhaul and supply chain planning for rotables	J. Arts, S.D. Flapper, K. Vernooij
368	2012	Operating Room Rescheduling	J.T. van Essen, J.L. Hurink, W. Hartholt, B.J. van den Akker
367	2011	Switching Transport Modes to Meet Voluntary Carbon Emission Targets	Kristel M.R. Hoen, Tarkan Tan, Jan C. Fransoo, Geert-Jan van Houtum
366	2011	On two-echelon inventory systems with Poisson demand and lost sales	Elisa Alvarez, Matthieu van der Heijden
365	2011	Minimizing the Waiting Time for Emergency Surgery	J.T. van Essen, E.W. Hans, J.L. Hurink, A. Oversberg
364	2011	Vehicle Routing Problem with Stochastic Travel Times Including Soft Time Windows and Service Costs	Duygu Tas, Nico Dellaert, Tom van Woensel, Ton de Kok
363	2011	A New Approximate Evaluation Method for Two-Echelon Inventory Systems with Emergency Shipments	Erhun Özkan, Geert-Jan van Houtum, Yasemin Serin
362	2011	Approximating Multi-Objective Time-Dependent Optimization Problems	Said Dabia, El-Ghazali Talbi, Tom Van Woensel, Ton de Kok
361	2011	Branch and Cut and Price for the Time Dependent Vehicle Routing Problem with Time Window	Said Dabia, Stefan Röpke, Tom Van Woensel, Ton de Kok

360	2011	Analysis of an Assemble-to-Order System with Different Review Periods	A.G. Karaarslan, G.P. Kiesmüller, A.G. de Kok
359	2011	Interval Availability Analysis of a Two-Echelon, Multi-Item System	Ahmad Al Hanbali, Matthieu van der Heijden
358	2011	Carbon-Optimal and Carbon-Neutral Supply Chains	Felipe Caro, Charles J. Corbett, Tarkan Tan, Rob Zuidwijk
357	2011	Generic Planning and Control of Automated Material Handling Systems: Practical Requirements Versus Existing Theory	Sameh Haneyah, Henk Zijm, Marco Schutten, Peter Schuur
356	2011	Last time buy decisions for products sold under warranty	M. van der Heijden, B. Iskandar
355	2011	Spatial concentration and location dynamics in logistics: the case of a Dutch province	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
354	2011	Identification of Employment Concentration Areas	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
353	2011	BOMN 2.0 Execution Semantics Formalized as Graph Rewrite Rules: extended version	Pieter van Gorp, Remco Dijkman
352	2011	Resource pooling and cost allocation among independent service providers	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
351	2011	A Framework for Business Innovation Directions	E. Lüftenegger, S. Angelov, P. Grefen
350	2011	The Road to a Business Process Architecture: An Overview of Approaches and their Use	Remco Dijkman, Irene Vanderfeesten, Hajo A. Reijers
349	2011	Effect of carbon emission regulations on transport mode selection under stochastic demand	K.M.R. Hoen, T. Tan, J.C. Fransoo G.J. van Houtum
348	2011	An improved MIP-based combinatorial approach for a multi-skill workforce scheduling problem	Murat Firat, Cor Hurkens
347	2011	An approximate approach for the joint problem of level of repair analysis and spare parts stocking	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
346	2011	Joint optimization of level of repair analysis and	R.J.I. Basten, M.C. van der Heijden,

	spare parts stocks	J.M.J. Schutten
345 2011	Inventory control with manufacturing lead time flexibility	Ton G. de Kok
344 2011	Analysis of resource pooling games via a new extension of the Erlang loss function	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
343 2011	Vehicle refueling with limited resources	Murat Firat, C.A.J. Hurkens, Gerhard J. Woeginger
342 2011	Optimal Inventory Policies with Non-stationary Supply Disruptions and Advance Supply Information	Bilge Atasoy, Refik Güllü, TarkanTan
341 2011	Redundancy Optimization for Critical Components in High-Availability Capital Goods	Kurtulus Baris Öner, Alan Scheller-Wolf Geert-Jan van Houtum
339 2010	Analysis of a two-echelon inventory system with two supply modes	Joachim Arts, Gudrun Kiesmüller
338 2010	Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh	Murat Firat, Gerhard J. Woeginger
335 2010	Attaining stability in multi-skill workforce scheduling	Murat Firat, Cor Hurkens
334 2010	Flexible Heuristics Miner (FHM)	A.J.M.M. Weijters, J.T.S. Ribeiro
333 2010	An exact approach for relating recovering surgical patient workload to the master surgical schedule	P.T. Vanberkel, R.J. Boucherie, E.W. Hans, J.L. Hurink, W.A.M. van Lent, W.H. van Harten
332 2010	Efficiency evaluation for pooling resources in health care	Peter T. Vanberkel, Richard J. Boucherie, Erwin W. Hans, Johann L. Hurink, Nelly Litvak
331 2010	The Effect of Workload Constraints in Mathematical Programming Models for Production Planning	M.M. Jansen, A.G. de Kok, I.J.B.F. Adan
330 2010	Using pipeline information in a multi-echelon spare parts inventory system	Christian Howard, Ingrid Reijnen, Johan Marklund, Tarkan Tan
329 2010	Reducing costs of repairable spare parts supply systems via dynamic scheduling	H.G.H. Tiemessen, G.J. van Houtum

328	2010	Identification of Employment Concentration and Specialization Areas: Theory and Application	F.P. van den Heuvel, P.W. de Langen, K.H. van Donselaar, J.C. Fransoo
327	2010	A combinatorial approach to multi-skill workforce scheduling	Murat Firat, Cor Hurkens
326	2010	Stability in multi-skill workforce scheduling	Murat Firat, Cor Hurkens, Alexandre Laugier
325	2010	Maintenance spare parts planning and control: A framework for control and agenda for future research	M.A. Driessen, J.J. Arts, G.J. v. Houtum, W.D. Rustenburg, B. Huisman
324	2010	Near-optimal heuristics to set base stock levels in a two-echelon distribution network	R.J.I. Basten, G.J. van Houtum
323	2010	Inventory reduction in spare part networks by selective throughput time reduction	M.C. van der Heijden, E.M. Alvarez, J.M.J. Schutten
322	2010	The selective use of emergency shipments for service-contract differentiation	E.M. Alvarez, M.C. van der Heijden, W.H. Zijm
321	2010	Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse	B. Walrave, K. v. Oorschot, A.G.L. Romme
320	2010	Preventing or escaping the suppression mechanism: intervention conditions	Nico Dellaert, Jully Jeunet.
319	2010	Hospital admission planning to optimize major resources utilization under uncertainty	R. Seguel, R. Eshuis, P. Grefen.
318	2010	Minimal Protocol Adaptors for Interacting Services	Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo.
317	2010	Teaching Retail Operations in Business and Engineering Schools	Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak.
316	2010	Design for Availability: Creating Value for Manufacturers and Customers	Pieter van Gorp, Rik Eshuis.
		Transforming Process Models: executable rewrite rules versus a formalized Java program	Bob Walrave, Kim E. van Oorschot, A.

315	2010		Georges L. Romme
		Getting trapped in the suppression of exploration: A simulation model	
314	2010		S. Dabia, T. van Woensel, A.G. de Kok
		A Dynamic Programming Approach to Multi-Objective Time-Dependent Capacitated Single Vehicle Routing Problems with Time Windows	
313	2010		
	2010		
		Tales of a So(u)rcerer: Optimal Sourcing Decisions Under Alternative Capacitated Suppliers and General Cost Structures	
312	2010		Osman Alp, Tarkan Tan
		In-store replenishment procedures for perishable inventory in a retail environment with handling costs and storage constraints	
311	2010		R.A.C.M. Broekmeulen, C.H.M. Bakx
		The state of the art of innovation-driven business models in the financial services industry	
310	2010		E. Lüftenegger, S. Angelov, E. van der Linden, P. Grefen
		Design of Complex Architectures Using a Three Dimension Approach: the CrossWork Case	
309	2010		R. Seguel, P. Grefen, R. Eshuis
		Effect of carbon emission regulations on transport mode selection in supply chains	
308	2010		K.M.R. Hoen, T. Tan, J.C. Fransoo, G.J. van Houtum
		Interaction between intelligent agent strategies for real-time transportation planning	
307	2010		Martijn Mes, Matthieu van der Heijden, Peter Schuur
		Internal Slackening Scoring Methods	
306	2010		Marco Slikker, Peter Borm, René van den Brink
		Vehicle Routing with Traffic Congestion and Drivers' Driving and Working Rules	
305	2010		A.L. Kok, E.W. Hans, J.M.J. Schutten, W.H.M. Zijm
		Practical extensions to the level of repair analysis	
304	2010		R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
		Ocean Container Transport: An Underestimated and Critical Link in Global Supply Chain Performance	
303	2010		Jan C. Fransoo, Chung-Yee Lee
		Capacity reservation and utilization for a manufacturer with uncertain capacity and demand	
302	2010		Y. Boulaksil; J.C. Fransoo; T. Tan
		Spare parts inventory pooling games	
300	2009		F.J.P. Karsten; M. Slikker; G.J. van Houtum
		Capacity flexibility allocation in an outsourced supply chain with reservation	
299	2009		Y. Boulaksil, M. Grunow, J.C. Fransoo
		An optimal approach for the joint problem of level of repair analysis and spare parts stocking	
298	2010		R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
		Responding to the Lehman Wave: Sales Forecasting and Supply Management during the Credit Crisis	
297	2009		Robert Peels, Maximiliano Udenio, Jan C. Fransoo, Marcel Wolfs, Tom Hendriks
		An exact approach for relating recovering	
296	2009		Peter T. Vanberkel, Richard J. Boucherie,

	surgical patient workload to the master surgical schedule	Erwin W. Hans, Johann L. Hurink, Wineke A.M. van Lent, Wim H. van Harten
295 2009	An iterative method for the simultaneous optimization of repair decisions and spare parts stocks	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
294 2009	Fujaba hits the Wall(-e)	Pieter van Gorp, Ruben Jubeh, Bernhard Grusie, Anne Keller
293 2009	Implementation of a Healthcare Process in Four Different Workflow Systems	R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J.M. Bakker
292 2009	Business Process Model Repositories - Framework and Survey	Zhiqiang Yan, Remco Dijkman, Paul Grefen
291 2009	Efficient Optimization of the Dual-Index Policy Using Markov Chains	Joachim Arts, Marcel van Vuuren, Gudrun Kiesmuller
290 2009	Hierarchical Knowledge-Gradient for Sequential Sampling	Martijn R.K. Mes; Warren B. Powell; Peter I. Frazier
289 2009	Analyzing combined vehicle routing and break scheduling from a distributed decision making perspective	C.M. Meyer; A.L. Kok; H. Kopfer; J.M.J. Schutten
288 2009	Anticipation of lead time performance in Supply Chain Operations Planning	Michiel Jansen; Ton G. de Kok; Jan C. Fransoo
287 2009	Inventory Models with Lateral Transshipments: A Review	Colin Paterson; Gudrun Kiesmuller; Ruud Teunter; Kevin Glazebrook
286 2009	Efficiency evaluation for pooling resources in health care	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
285 2009	A Survey of Health Care Models that Encompass Multiple Departments	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
284 2009	Supporting Process Control in Business Collaborations	S. Angelov; K. Vidyasankar; J. Vonk; P. Grefen
283 2009	Inventory Control with Partial Batch Ordering	O. Alp; W.T. Huh; T. Tan
282 2009	Translating Safe Petri Nets to Statecharts in a Structure-Preserving Way	R. Eshuis
281 2009	The link between product data model and process model	J.J.C.L. Vogelaar; H.A. Reijers
280 2009	Inventory planning for spare parts networks with delivery time requirements	I.C. Reijnen; T. Tan; G.J. van Houtum
279 2009	Co-Evolution of Demand and Supply under Competition	B. Vermeulen; A.G. de Kok
278 2010	Toward Meso-level Product-Market Network Indices for Strategic Product Selection and (Re)Design Guidelines over the Product Life-Cycle	B. Vermeulen, A.G. de Kok
277 2009	An Efficient Method to Construct Minimal Protocol Adaptors	R. Seguel, R. Eshuis, P. Grefen
276 2009	Coordinating Supply Chains: a Bilevel	Ton G. de Kok, Gabriella Muratore

	<u>Programming Approach</u>	
275	2009 <u>Inventory redistribution for fashion products under demand parameter update</u>	G.P. Kiesmuller, S. Minner
274	2009 <u>Comparing Markov chains: Combining aggregation and precedence relations applied to sets of states</u>	A. Basic, I.M.H. Vliegen, A. Scheller-Wolf
273	2009 <u>Separate tools or tool kits: an exploratory study of engineers' preferences</u>	I.M.H. Vliegen, P.A.M. Kleingeld, G.J. van Houtum
272	2009 <u>An Exact Solution Procedure for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering</u>	Engin Topan, Z. Pelin Bayindir, Tarkan Tan
271	2009 <u>Distributed Decision Making in Combined Vehicle Routing and Break Scheduling</u>	C.M. Meyer, H. Kopfer, A.L. Kok, M. Schutten
270	2009 <u>Dynamic Programming Algorithm for the Vehicle Routing Problem with Time Windows and EC Social Legislation</u>	A.L. Kok, C.M. Meyer, H. Kopfer, J.M.J. Schutten
269	2009 <u>Similarity of Business Process Models: Metrics and Evaluation</u>	Remco Dijkman, Marlon Dumas, Boudewijn van Dongen, Reina Kaarik, Jan Mendling
267	2009 <u>Vehicle routing under time-dependent travel times: the impact of congestion avoidance</u>	A.L. Kok, E.W. Hans, J.M.J. Schutten
266	2009 <u>Restricted dynamic programming: a flexible framework for solving realistic VRPs</u>	J. Gromicho; J.J. van Hoorn; A.L. Kok; J.M.J. Schutten;

Working Papers published before 2009 see: <http://beta.ieis.tue.nl>