The selective use of emergency shipments for service-contract differentiation

Citation for published version (APA):

Alvarez, E. M., Heijden, van der, M. C., & Zijm, W. H. M. (2010). *The selective use of emergency shipments for service-contract differentiation*. (BETA publicatie : working papers; Vol. 322). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/2010

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

The selective use of emergency shipments for service-contract differentiation

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

Beta Working Paper series 322

BETA publicatie	WP 322 (working paper)
ISBN	978-90-386-2354-2
ISSN	
NUR	804
Eindhoven	September 2010

The selective use of emergency shipments for service-contract differentiation

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

University of Twente, School of Management and Governance

Abstract: Suppliers of capital goods increasingly offer performance-based service contracts with customer-specific service levels. We use selective emergency shipments of spare parts to differentiate logistic performance: We apply emergency shipments in out-of-stock situations for combinations of parts and customer classes that yield service levels close to the class-specific targets. We develop two heuristics to solve this problem. An extensive numerical experiment reveals average cost savings of 4.4% compared to the one-size-fits-all approach that is often used in practice. It is best to combine our policy with critical levels, which yields an average cost saving of 13.9%.

Key words: Inventory, customer differentiation, emergency shipments, service contracts, critical levels

1. Introduction

To service advanced capital goods - such as defense systems or medical systems - suppliers increasingly offer performance-based service contracts to their customers. This particularly applies in business situations where system downtime can have very serious consequences. For example, downtime of military equipment can lead to failed missions and downtime of medical equipment creates delay in the diagnosis and treatment of patients. Therefore, service contracts for such capital goods typically contain quantified targets for key performance measures such as a maximum response time in case of a system failure or a minimum system availability. Penalty regimes may apply if the supplier is unable to meet the target service levels. Because users typically value downtime differently, the service levels in contracts

may differ between customer groups. For example, the maximum on-site response time may be 4 hours or the next day.

Because it is hard to handle these different service levels, many companies use the so-called one-size-fits-all approach, in which a uniform logistics fulfillment process is used irrespective of the contractual service level agreements (cf. Cohen et al [3]). This approach can result in excessive costs if a supplier uses the premium service level to design the fulfillment process. Also, standard customers have no incentive to switch to premium contracts. Therefore, it seems better for the service provider to differentiate the logistics fulfillment process such that the actual service levels reflect the contractual agreements. While this can be accomplished in several ways (e.g. by prioritizing the assignment of service engineers to service calls or varying preventive maintenance frequencies), the emphasis in the literature is on differentiation in *spare parts supply*. One approach is to design separate supply chains per customer segment, such as stocking parts close to the customer site for premium customers and supplying non-premium customers from a central location with longer lead times (cf. Deshpande et al. [5]). A drawback of this separation is that the supplier can take less advantage of risk pooling (Eppen and Schrage [6]).

In the literature, a common approach for service differentiation is the use of *critical level policies* that reserve spare parts for premium customers once the inventory level drops below a certain threshold. Then, demand from non-premium customers is either backordered or satisfied from a secondary source that is usually assumed to have infinite supply (e.g. a central stock point upstream in the supply chain). Although shown to be effective and efficient, there are barriers for implementation in practice. First, customers may have access to stock information. In that case, suppliers are reluctant to refuse a spare part to a non-premium customer.

Second, service engineers responsible for system repair are usually unwilling to postpone their work waiting for a spare part when they are primarily accountable for the speed of repair and they know that the part is actually in stock.

These drawbacks prompted us to investigate the *selective use of emergency shipments* as an alternative. That is, a supplier uses on-hand stock to meet demand first-come-first-served. In case of a stock-out, he can request an emergency shipment from a secondary source (e.g. a central depot). As emergency shipments are both faster and more expensive than regular replenishments, the supplier can select combinations of customer segments and item types for which he applies emergency shipments. As main advantage, this approach is easier to implement in practice than critical level policies, while still giving the option to apply differentiation. We will show that our approach leads to clear savings over using simple one-size-fits all strategies. Furthermore, we will show that the *combination* of selective emergency shipments and critical level policies is clearly most efficient and effective.

The remainder of this paper is structured as follows. In Section 2, we give an overview of the literature and we state our contribution. Next, we introduce the optimization problem in Section 3 and we outline our solution approach in Section 4. It will become clear that we need to analyze various single-item models as building blocks. In Section 5, we analyze these models for the special case of two customer classes. We describe the results of an extensive numerical experiment in Section 6. In Section 7 we present our conclusions and suggest directions for further research.

2. Literature overview

Our research is related to two main literature streams: service differentiation and the use of emergency shipments in spare parts networks. In the service

differentiation stream, the focus particularly lies on the use of critical level policies that have been introduced by Veinott [19]. The optimality of this policy under periodic review has been proven by Topkis [17], for both backordering and lost sales. Under continuous review, the optimality has been proven assuming Poisson demand and either exponential or Erlang lead times, both for lost sales (Ha [9], [11]) and backordering (Ha [10], De Véricourt et al. [20], Gayon [7]).

Several approaches have been developed to find (near-) optimal base stock levels and critical levels. For *fast movers*, the focus is on continuous demand distributions where unmet requests are usually backordered. When a replenishment order arrives, it appears to be optimal to clear non-premium backorders if and only if the inventory level is above the critical level for premium demand (Ha [10]). Unfortunately, the mathematical analysis of such a model is intractable, since we must keep track of the non-premium backorders. Therefore, heuristic are often used, see e.g. Möllering and Thonemann ([15]) for two customer classes, and Arslan et al. [2] for an arbitrary number of classes. For *slow movers*, as are common in service logistics, the focus is on Poisson demand and one-for-one replenishment (Dekker et al. [4]). Our work shows most similarity to Kranenburg and Van Houtum [15], who analyze critical level policies in a multi-item model with the objective to minimize spare part holding and shipment costs under waiting time restrictions per customer class. Emergency shipments are used if a request cannot be met by on-hand stock. The authors find a lower bound on the minimum costs using decomposition and column generation. Next, they use a heuristic to obtain a near-optimal solution.

The second relevant literature stream is the use of emergency shipments in spare parts networks, which is in some cases combined with lateral transshipments between local warehouses at the same echelon level. Most authors consider a

single- or two-echelon model in which there is some central location with ample supply (see e.g. Muckstadt and Thomas [16], Hausman and Erkip [12]). Alfredsson and Verrijdt [1] combine lateral shipments with emergency shipments for a two-echelon single-item network. Other recent contributions in this area include Van Utterbeeck et al. [18] and Wong et al. [21]. We have not yet found any literature in which emergency shipments are used for customer differentiation.

The contribution of our paper is fourfold: First, we present a *new approach to service differentiation* in spare parts supply using selective emergency shipments. Second, we develop *two efficient and effective heuristics* to determine near-optimal base stock levels and shipment strategies. Third, we show the *added value of selective emergency shipments* compared to both one-size-fits-all policies and critical level policies. Finally, we show the *added value of combining selective emergency shipments and critical level policies for service differentiation*.

3. Model

We first give an outline of our model. Next, we discuss the validity of our selection of shipment policies (Section 3.2). In Section 3.3, we present our model assumptions and notation. We give the formal optimization problem in Section 3.4.

3.1 Model outline

We consider a local warehouse supplying multiple types of spare parts to multiple customer groups, and a central depot with ample supply that replenishes the local warehouse. We assume that all items are critical, i.e., any item failure causes a system failure. Each customer group has a distinct target service level, defined as a maximum on the mean waiting time for spares. Such a restriction reflects the downtime caused by lack of spares (see also Kranenburg and Van Houtum [14]).

The local warehouse fills all customer requests on a first-come-first-served basis, regardless of the customer's class. If the warehouse is out of stock, it may either backorder the demand, or request an emergency shipment from the central depot. Service differentiation is accomplished by only using emergency shipments for customer classes with tight waiting time restrictions. We expect that this is particularly advantageous for expensive slow movers that usually have relatively low fill rates, which makes the difference between regular and emergency shipment time crucial. In some cases, we may find that it is better to apply emergency shipments for all customer classes, avoiding stocks as much as possible. For cheap fast movers, we expect that it will be better to keep large inventories avoiding expensive emergency shipments and to backorder demand for all customer classes. Hence, the choice for backordering or emergency shipments depends both on the item characteristics and the service requirements of the demand classes.

Our objective is to minimize the total relevant system costs, consisting of holding and shipment costs, under restrictions on the mean aggregate waiting per class. Our decision variables are the shipment mode (regular, emergency) and the stock levels for each combination of item and customer class.

3.2 Selection of shipment policies

In fact, we only consider a limited number of shipment policies: *if* we apply emergency shipments, we always do so if we are out-of-stock. However, if there are plenty of items in the pipeline, the waiting time for a backorder may well be comparable to – or even lower than – the emergency shipment time. Then, it is better to backorder the item instead of using an expensive emergency shipment. The backorder waiting time for a customer depends both on the number of items in the pipeline and the number of earlier backorders that must be cleared before the

customer's backorder. So, the decision when to use emergency shipments for an item should ideally not only depend on the customer class, but also on the number of items in the pipeline and the composition of the backorder queue. In principle, we can include such more advanced policies in our model. Still, we ignore this refinement in this paper to keep notation transparent and the number of item policies within reasonable limits. In the end, the basic research question is whether and when it makes sense to apply selective emergency shipments for service differentiation compared to critical level policies and the "one-size-fits-all" approach.

3.3 Assumptions and notation

3.3.1 Main assumptions

1. *Demand for each item occurs according to a Poisson process.*
2. *An $(S-1, S)$ base stock policy is applied for all items.* In practice, spares often tend to be expensive slow movers. Therefore, holding costs usually dominate ordering costs and hence the optimal ordering quantity is usually 1.
3. *Regular shipment times from depot to warehouse are exponentially distributed.*
We use this assumption to facilitate Markov chain analysis. Alfredsson and Verrijdt [1] show that inventory models for slow moving spare parts tend to be quite insensitive to lead time variability.
4. *The shipment time from the local warehouse to the customer is negligible.*
5. *An emergency shipment is directly shipped from central depot to customer.*

3.3.2 Notation

The local warehouse carries stock of stock-keeping unit (SKU) $i = 1, 2, \dots, I$. We denote the corresponding mean replenishment lead time from depot to local warehouse by T_i^{reg} , the mean emergency shipment time by T_i^{em} , the holding costs

per time unit by h_i and the additional costs for an emergency shipment instead of a normal replenishment by c_i^{em} . The latter cost factor is sufficient, since each request triggers either a normal replenishment or an emergency shipment. The target service level for customer class $j = 1, 2, \dots, J$ is a maximum on the mean waiting time for spares, denoted by $W_j^{\max}$. We order the J classes according to an increasing target service level $W_j^{\max}$. Demand from class j for SKU i occurs at rate $m_{ij} (> 0)$.

$M_{\bullet j} = \sum_{i=1}^I m_{ij}$ and $M_{i\bullet} = \sum_{j=1}^J m_{ij}$ denote the total class j demand and the total demand for SKU i , respectively. The decision variables for SKU i are: (1) the base stock level S_i (2) the *shipment strategy* D_i , denoting the highest customer class index for which emergency shipments are used in a stock-out situation. Because it does not make sense to use emergency shipments for low priority customers only, D_i is an integer between 0 and J . We combine S_i and D_i to an *item policy* (S_i, D_i) . The relevant performance indicators are the mean waiting time $EW_{ij}(S_i, D_i)$ of class j for SKU i and the fill rate $\beta_i(S_i, D_i)$ for SKU i .

3.4 Formal optimization problem

We express the formal optimization problem (P) as follows:

$$\begin{aligned}
 (P1) \quad & \min_{S_i, D_i} \sum_{i=1}^I \left\{ h_i S_i + c_i^{em} \sum_{j=1}^J m_{ij} (1 - \beta_i(S_i, D_i)) I_{\{1..D_i\}}(j) \right\} \\
 \text{s.t.} \quad & \sum_{i=1}^I \frac{m_{ij}}{M_{\bullet j}} EW_{ij}(S_i, D_i) \leq W_j^{\max} \quad j = 1..J \quad (P1.1) \\
 & S_i \in N_0, D_i \in \{0, 1, \dots, J\} \quad i = 1..I \text{ and } j = 1..J \quad (P1.2)
 \end{aligned}$$

We minimize the sum of holding costs and emergency shipment costs under the restriction that the weighted mean waiting time for class j (with the demand rates as weights) does not exceed the target $W_j^{\max}$ for that class. We calculate the holding costs over the total stock S_i . If needed, we can easily replace this by the on-hand stock (Kranenburg and Van Houtum [14]). Emergency shipment costs arise if a request cannot be filled from stock on-hand and an emergency shipment is used under shipment strategy D_i . We denote the latter by the indicator function $I_{\{1..D_i\}}(j)$ that equals 1 if $j \in \{1..D_i\}$ and 0 otherwise. The index set $\{1..D_i\}$ is empty if $D_i = 0$.

To include critical level policies in our model, we replace the fill rate by $\beta_{ij}(S_i, D_i, c_i)$ and the mean waiting time by $EW_{ij}(S_i, D_i, c_i)$, where c_i is the critical level for item i . Note that the fill rate also depends on the customer class j now.

4. Solution approach

Problem (P) is a nonlinear integer problem that we cannot decompose into separate single-item problems because of the aggregate waiting time restriction (P1.1). We therefore use an approach similar to Dantzig-Wolfe decomposition, i.e. we reformulate (P) to a linear integer programming problem, where binary decision variables specify whether a specific item policy is selected for SKU i or not. We obtain a lower bound on the system costs by solving the LP-relaxation of the reformulated problem. To determine the set of relevant item policies (S_i, D_i) to consider, we use column generation, see Section 4.1 for details. In the optimal solution to the LP-relaxation, a linear combination of two item policies may be selected for an SKU. In that case, we also require a method to obtain a near-optimal integer solution from the lower bound. We describe two heuristics in Section 4.2

4.1 Using decomposition and column generation to find a lower bound

In Section 4.1.1, we first elaborate on reformulating (P) to a linear problem.

Next, we describe in Section 4.1.2 how we solve the LP-relaxation.

4.1.1 Problem reformulation

Let us use b_i as a shorthand notation for item policy (S_i, D_i) . The binary variable x_{b_i} denotes whether policy b_i is selected for SKU i ($x_{b_i} = 1$) or not ($x_{b_i} = 0$).

Let B_i denote the set of item policies that we consider for SKU i . We then obtain the integer program (P2):

$$(P2) \quad \min \sum_{i=1}^I \sum_{b_i \in B_i} TC_i(b_i) x_{b_i}$$

$$\text{s.t.} \quad \sum_{i=1}^I \sum_{b_i \in B_i} \frac{m_{ij}}{M_{\bullet j}} EW_{ij}(b_i) x_{b_i} \leq W_j^{\max} \quad j = 1..J \quad (P2.1)$$

$$\sum_{b_i \in B_i} x_{b_i} = 1 \quad i = 1..I \quad (P2.2)$$

$$x_{b_i} \in \{0, 1\} \quad \forall b_i \in B_i, i = 1..I$$

Here, $TC_i(b_i)$ is a shorthand notation for the total relevant costs related to SKU

$$i, \text{ so } TC_i(b_i) = TC_i(S_i, D_i) = h_i S_i + c_i^{em} \sum_{j=1}^J m_{ij} (1 - \beta_i(S_i, D_i)) L_{\{1..D_i\}}(j)$$

The optimal solution to the LP-relaxation may consist of linear combinations of item policies for at most J SKUs: (P2) has $I + J$ restrictions, so at most $I + J$ decision variables have nonzero values. Since we must select at least I item policies, we have at most J SKUs for which additional item policies can be selected. This is convenient, since the number of customer classes J is typically small.

4.1.2 Solving the LP-relaxation

We must specify which policies to include in B_i for each SKU i . Initially, we choose $B_i = \{b_i\}$, where b_i is an item policy guaranteeing that all waiting time

restrictions are satisfied. We then find $x_{b_i} = 1$ for all i as the only feasible solution to the LP-relaxation. A simple choice for the initial item policy is $D_i = 0$ (backorder all demand) combined with the smallest stock level S_i such that $EW_{ij}(b_i) \leq W_1^{\max}$ for all classes $j = 1 \dots J$. Since $W_1^{\max} < W_j^{\max}$ for $j > 1$, all waiting time restrictions are satisfied.

Next, we use *column generation* to find any item policies that improve the objective value if added (Gilmore and Gomory [8]). Those policies have negative reduced costs. If such policies exist for SKU i , we include the one with minimum reduced costs to B_i . Given an optimal solution to the LP-relaxation, let u_j and v_i be the shadow prices for constraints (P2.1) and (P2.2), respectively. Note that $u_j \leq 0$ and $v_i \geq 0$. Then the column generation problem for SKU i boils down to:

$$(SUB(i)) \quad \min_{b_i} RED(b_i) = TC_i(b_i) - \sum_{j=1}^J u_j \frac{m_{ij}}{M_{\bullet j}} EW_{ij}(b_i) - v_i$$

For each shipment strategy D_i , we determine the S_i that minimizes $RED(b_i) = RED(S_i, D_i)$. Let b_i^d denote the resulting policy for $D_i = d$. We select the policy with the lowest reduced costs from the set $\{b_i^d \mid d = 0 \dots J\}$. A complication in determining the optimal S_i for a given shipment strategy is that the objective function is *only* convex if we use emergency shipments for all demand classes (Kranenburg and Van Houtum [13]). If emergency shipments are used for class 1 requests only, we can find a function for the expected waiting time for class 1 as shown in Figure 4.1. The figure shows the expected waiting time for an item with a high demand rate. For such items, the waiting time function might initially be concave in the base stock level: keeping little stock of these items has a small impact on the waiting time. As the stock level increases, the marginal decrease in the waiting time becomes larger.

Figure 4.1 An example expected waiting time function for class 1.

We expect the objective function to become convex from some base stock level S_i^{conv} onwards: for high base stock levels, the costs increase almost linearly, whereas the waiting time decreases convexly. This suggests to select the base stock level with smallest reduced costs from the set of (i) all stock levels $S_i \leq S_i^{conv}$, plus (ii) all stock levels $S_i > S_i^{conv}$ until the reduced costs no longer decrease.

4.2 Methods for finding a near-optimal integer solution

Below, we discuss two methods for finding a near-optimal integer solution from the solution of the LP-relaxation: (1) we solve the integer problem (P2) using all policies generated when solving the LP-relaxation; (2) we use the (non-integer) solution of the LP relaxation as a starting point of a local search algorithm.

4.2.1 Method 1: use of Integer Programming (IP)

We empirically found that we usually generate only 4-7 item policies for each SKU in the LP-relaxation. Therefore, we should be able to solve the corresponding IP problem with a commercial solver (we used CPLEX) for most problems of realistic size. However, the set of policies for certain items – particularly those with large demand rates – might be completely unrelated. For instance, we have found policies (0,1) (i.e. keep no stock and use emergency shipments for premium customers only),

(9,0) and (10,0) (i.e. high base stock levels and backordering for all classes). If we use IP, we must choose between an inexpensive policy with high waiting times and an expensive policy with low waiting times, while a better, but unconsidered, policy might exist. Therefore, we tested our approach by generating many additional columns (item policies) bridging the gap between distinct scenarios as shown above and including them in the IP problem¹. We found that these additional columns significantly increase the computation time, while they improve the final solution only marginally: the average gap with the lower bound drops from 0.041 to 0.038, and the maximum gap drops from 0.259 to 0.258. Therefore, we conclude that it is sufficient to use the columns as generated when solving the LP relaxation only.

4.2.2 Method 2: use a local search algorithm

Simply rounding the LP relaxation solution to integer values may give a poor solution to the corresponding IP problem. We therefore develop a local search algorithm to find a near-optimal solution for the original problem. First, we round the fractional variables in the LP-relaxation solution to obtain an infeasible, low cost solution. We then use local search to move to a feasible solution with minimum cost increase². To this end, we need (i) an initial solution, (ii) a neighborhood structure, (iii) a criterion to select the “best” neighbor.

A simple *initial solution* is the least expensive item policy b_i for each SKU i with $x_{b_i} > 0$. This solution is infeasible, since such policies have high waiting times.

Our *neighborhood* must contain solutions that have lower waiting times than the given solution and require little additional cost. To obtain a feasible solution, the

¹ This was a mid-sized experiment of 100 problem instances with 25, 100 and 400 SKUs.

² Another approach is to round the variables to obtain a costly feasible solution and apply local search to find alternatives with lower costs. There is no reason a priori to prefer one method over the other.

waiting times must decrease for those demand classes j for which the aggregate waiting time is higher than $W_j^{\max}$. Let the term *distance to the feasible region* denote the total amount by which the aggregate waiting times exceed the targets. Each solution in the neighborhood has exactly one SKU for which the item policy is modified such that the solution is “closer” to the feasible region according to the distance measure above. We can reduce this distance either by using emergency shipments for more demand classes (increase D_i), or by increasing the stock levels S_i . If we increase D_i , we combine this with all values of S_i smaller than or equal to the current stock level insofar the policy is closer to the feasible region. Similarly, when we increase S_i , we combine this with all values of D_i smaller than or equal to the current value insofar the policy is closer to the feasible region. This gives us several neighbors for each SKU i , and hence many neighbors in the neighborhood.

We use as *selection criterion* the weighted reduction in waiting time per euro. This is the quotient of the weighted reduction in distance to the feasible region compared to solution $\mathbf{b}$ ($M_{\mathbf{b}'}$) and the additional investment needed ($N_{\mathbf{b}'}$), where:

$$M_{\mathbf{b}'} = \sum_{j=1}^J \frac{1}{W_j^{\max}} \left\{ \left[\sum_{i=1}^I \frac{m_{ij}}{M_{\bullet j}} EW_{ij}(b_i) - W_j^{\max} \right]^+ - \left[\sum_{i=1}^I \frac{m_{ij}}{M_{\bullet j}} EW_{ij}(b_i') - W_j^{\max} \right]^+ \right\}, \quad [a]^+ = \max\{0, a\}$$

$$N_{\mathbf{b}'} = \sum_{i=1}^I TC_i(b_i') - TC_i(b_i)$$

Note that the inverse of $W_j^{\max}$ acts as a weight for class j in $M_{\mathbf{b}'}$: in general, we expect that a larger marginal investment is needed to reduce a small waiting time by the amount Δ compared to reducing a large waiting time by that same amount. We therefore assign a high weight to demand classes with tight waiting time restrictions.

We may find neighbors with *lower* costs than $\mathbf{b}$ (and thus a *negative* value for $M_{\mathbf{b}'} / N_{\mathbf{b}'}$). From the neighbors with a negative value for $M_{\mathbf{b}'} / N_{\mathbf{b}'}$, we then select the one with the largest $M_{\mathbf{b}'} / N_{\mathbf{b}'}$. Otherwise, we choose the neighbor with the largest $M_{\mathbf{b}'} / N_{\mathbf{b}'}$ from the entire neighborhood.

5. Performance evaluation of single-item policies with two demand classes

We use continuous-time Markov chain analysis to find the average waiting time and fill rate per customer class for each policy. Both performance indicators are functions of the state probabilities. For simplicity, we limit ourselves to two demand classes. We thus consider three shipment strategies: emergency shipments for both classes ($D_i = 2$), backorder class 2 demand only ($D_i = 1$), or backorder all demand ($D_i = 0$). The backorder clearing mechanism depends on the shipment strategy and whether critical levels are used. In subsections 5.1 and 5.2, we analyze the strategies for the selective emergency shipment model and the critical level model, respectively. For simplicity, we omit the item index i . We denote the normal replenishment rate by $\mu = 1/T^{reg}$ and the critical level by c .

5.1 Building blocks for the selective emergency shipment model

In this case, the fill rate is identical for all customer classes. We use Little's Law to compute waiting times when requests for a class are backordered.

Emergency shipments for both classes ($D_i = 2$)

The state k represents the number of items in the pipeline. We find the state probabilities from the $M | M | n | n$ loss system with $n = S$ (Kranenburg and van Houtum [13]):

$$p_k = \frac{(M/\mu)^k / k!}{\sum_{l=0}^S (M/\mu)^l / l!}, \quad k = 0..S$$

Hence, the fill rate and the average waiting time per class equal:

$$\beta(S, D) = p_S$$

$$EW_j(S, D) = \beta(S, D) T^{em}$$

Backorder class 2 demand only ($D_i = 1$)

Again, we represent the state by the number of items in the pipeline k . The number of class 2 backorders then equals $[k - S]^+$. Figure 5.1 displays the corresponding Markov chain, which yields the state probabilities p_k and the performance indicators as given below.

Figure 5.1 The transition diagram of the Markov chain

$$p_0 = \left\{ \sum_{k=0}^S \frac{1}{k!} \left(\frac{M}{\mu} \right)^k + \left(\frac{M}{m_2} \right)^S \left(e^{\frac{m_2}{\mu}} - \sum_{k=0}^S \frac{1}{k!} \left(\frac{m_2}{\mu} \right)^k \right) \right\}^{-1}$$

$$p_k = \frac{1}{k!} \left(\frac{M}{\mu} \right)^k p_0 \quad k = 1, 2, \dots, S$$

$$p_k = \frac{1}{k!} \left(\frac{M}{\mu} \right)^S \left(\frac{m_2}{\mu} \right)^{k-S} p_0 \quad k \geq S+1$$

$$\beta(S, D) = 1 - \sum_{k=0}^{S-1} p_k$$

$$EW_1(S, D) = \beta(S, D) T^{em}$$

We can derive the following expression for the expected backorders of class 2:

$$EBO_2(S, D) = \left(\frac{M}{m_2} \right)^S p_0 \left\{ \frac{m_2}{\mu} \left(e^{\frac{m_2}{\mu}} - \sum_{k=0}^{S-1} \frac{1}{k!} \left(\frac{m_2}{\mu} \right)^k \right) - S \left(e^{\frac{m_2}{\mu}} - \sum_{k=0}^S \frac{1}{k!} \left(\frac{m_2}{\mu} \right)^k \right) \right\}$$

Backorder demand from all classes ($D_i = 0$)

For this policy, we do not need the fill rate, since there are no emergency shipments. In this case, it is best to use priority clearing: class 1 backorders are cleared before class 2 backorders, even if a class 2 customer has been waiting longer. As a result, the number of items in the pipeline is not sufficient to describe the system state, since the waiting time depends on the number of backorders per class. Therefore, we describe the system state as (k, l) , with k =number of items in the pipeline and l =number of class 2 backorders. Then, the number of class 1 backorders equals $\left[[k - S]^+ - l \right]^+$. Figure 5.2 depicts the corresponding Markov chain. Replenishment flows usually go from state (k, l) to state $(k-1, l)$, since we clear class 1 backorders first. Hence, the pipeline decreases while the number of class 2 backorders remains the same. Replenishment flows from state (k, l) to state $(k-1, l-1)$ only take place if the system contains class 2 backorders only in (k, l) .

Figure 5.2 The transition diagram of the Markov chain.

We could not find analytical expressions for the state probabilities, so we solved the balance equations numerically, assuming that the probability of having more than $k^{\max}$ items in the pipeline is negligible. We find $k^{\max}$ by aggregating all

customer demand into a single class and analyzing the resulting model exactly. Note that this model gives the same distribution for the number in the pipeline as the two-class model, since the demand and replenishment rates are the same. We find $k^{\max}$ such that $\sum_{k=k^{\max}+1}^{\infty} p_k \leq \varepsilon$, where ε denotes some small number, say $\varepsilon = 10^{-8}$. Then, the

expected backorders $EBO_j(S, D)$ are given by:

$$EBO_1(S, D) = \sum_{k=S+1}^{k^{\max}} \sum_{l=0}^{k-S} (k-S-l) p_{kl}$$

$$EBO_2(S, D) = \sum_{k=S+1}^{k^{\max}} \sum_{l=0}^{k-S} l \cdot p_{kl}$$

5.2 Building blocks for the critical level model

For the case of a critical level and emergency shipments for both classes ($D_i = 0$), we refer to Kranenburg and Van Houtum [14]. Below, we give the performance analysis for critical level policies under (partial) backordering. We use Little's law to find the average waiting time from the expected backorders.

Critical level and backorder class 2 demand only ($D_i = 1$)

When the on-hand stock is at most c_i , it is optimal to increase the on-hand stock to c_i before clearing class 2 backorders (cf. Ha [10]). Since we cannot derive the number of class 2 backorders from the number in the pipeline, we include the class 2 backorders in the state definition. **Error! Reference source not found.** displays the Markov chain. Note that we have no states with class 1 backorders: once there is no more stock on-hand, the pipeline is only increased further by class 2 requests.

Figure 5.3 The transition diagram of the Markov chain

We solve the balance equations numerically assuming that the probability of more than $k^{\max}$ items in the pipeline is negligible. We determine $k^{\max}$ as shown in Section 5.1, although we overestimate $k^{\max}$ then (we assume that all demand is backordered). Then, we find for the class 2 backorders and the class 1 fill rate:

$$EBO_2(S, D, c) = \sum_{k=S-c+1}^{k^{\max}} \sum_{l=0}^{k-S+c} l \cdot p_{kl}$$

$$\beta_1(S, D, c) = \left(1 - \sum_{k=0}^{S+1} p_{kl} \right)$$

Critical level and backorder demands from all classes ($D_i = 0$)

We only need expressions for the mean waiting time per class, since there are no emergency shipments. The optimal backorder clearing mechanism is to clear class 2 backorders only if all class 1 backorders (if present) have been cleared *and* the on-hand stock is at least c_i (Ha [10]). Then we find a Markov chain similar to Figure 5.2. The main difference is that the Markov chain now branches out once $S_i - c_i$ items are in the pipeline. We can thus use a similar approach to that of Section 5.3. Once we have calculated the state probabilities for a certain value of $S_i - c_i$, we can immediately find the performance characteristics for all other policies

with the same value of $S_i - c_i$, since the state probabilities are derived from the same set of equations. This significantly reduces the computational burden when generating and analyzing item policies. We find the following the expressions for the expected backorders $EBO_j(S, D, c)$:

$$EBO_1(S, D, c) = \sum_{k=S-c+1}^{k^{\max}} \sum_{l=0}^{k-S+c} \max(0, k - S - l) p_{kl}$$

$$EBO_2(S, D, c) = \sum_{k=S-c+1}^{k^{\max}} \sum_{l=0}^{k-S+c} l \cdot p_{kl}$$

6. Computational experiment

We conducted a numerical experiment, for which we state the objectives in Section 6.1. Section 6.2 covers the problem instances and Section 6.3 the results.

6.1 Objectives

The objectives of the experiment are: (i) to evaluate the two methods for obtaining a near-optimal solution (i.e. IP and local search) in terms of solution quality and computation time, (ii) to find out whether and when selective emergency shipments are efficient and effective for service differentiation, (iii) to determine how well selective emergency shipments perform compared to critical level policies, (iv) to find the added value of combining selective emergency shipments and critical level policies into an aggregate policy.

6.2 Experiment design

Table 6.1 shows the parameters we varied in our experiment and their values. We base our parameter values on Kranenburg and Van Houtum [14], who use values based on observations from practice. In all instances, we choose a C_i^{em} of 1000 as cost normalization. For each combination of parameters 1, 3, 4, 6 and 7, we generate 4 random sets of instances as follows: per SKU, the demand rate is randomly drawn

from a uniform distribution on the specified interval. The holding costs are selected from the holding cost interval such that the correlation between demand rates and holding costs is -0.8. We do so, because in practice high demand items tend to have low (holding) costs and vice versa. Note that the holding cost range (i.e. the quotient between the upper and lower bound) remains the same in all cases, namely 999. Such a large range is common in practice. With the exception of the item demand rates and holding costs, we choose the same parameter values for all SKUs in a problem instance. In total, we test 3456 problem instances: we have $3 \times 2 \times 3 \times 4 \times 3 \times 4 = 864$ parameter combinations and we test each combination with 4 samples of demand rates and holding costs.

Parameter		Values
1	Number of SKUs I	25, 100, 400
2	Daily demand rate per SKU $M_{i\bullet}$	$U[0,0.1]$, $U[0,0.5]$
3	Fractions of class demand per item $m_{i1}/M_{i\bullet}$; $m_{i2}/M_{i\bullet}$	$(0.2;0.8)$, $(0.5;0.5)$, $(0.8;0.2)$
4	$(T_i^{reg}; T_i^{em})$ (in days)	$(4;1)$, $(8;1)$, $(8;2)$, $(16;2)$
5	Item holding cost interval (per unit per day)	$[0.02,19.98]$, $[0.2,199.8]$, $[2,1998]$
6	c_i^{em} (per shipment)	1000
7	Target service levels $(W_1^{\max}; W_2^{\max})$ (in hours)	$(0.5,2)$, $(0.5,4)$, $(3,12)$, $(3,24)$

Table 6.1 Parameter values of the tested instances

6.3 Results

We evaluate the performance of the two methods for obtaining near-optimal integer solutions in Section 6.3.1. In Section 6.3.2, we investigate whether and when the emergency shipment strategy has added value over one-size-fits-all strategies. In Section 6.3.3, we compare the emergency shipment policy to the critical level policy and we determine the added value of combining both policies.

6.3.1 Performance of the heuristics

To evaluate the solution quality, we compare the integer solutions constructed to the lower bounds found by solving the LP-relaxation. For each heuristic, we express the solution quality in terms of a relative gap to the lower bound, defined as $100 \cdot (TC_H - TC_{LB}) / TC_{LB}$. Here, TC_H denotes the solution value for the heuristic (Integer Programming or Local Search). Table 6.2 displays the average and maximum gap for different numbers of SKUs, which is the parameter with the largest impact on the gaps. Overall, both solutions perform well in terms of solution quality: the average gap to the lower bound is well below 1%. Integer programming gives the best results: the gap to the lower bound is generally less than half that of local search. We observe that the gap clearly decreases with the number of SKUs. This result is relevant, because practical instances typically contain hundreds of items.

Parameter	Values	gap integer programming (%)		gap local search (%)	
		average	maximum	average	maximum
Number of SKUs	25	0.25	2.16	0.51	3.81
	100	0.02	0.11	0.05	0.35
	400	0.00	0.02	0.00	0.08
Overall		0.09	2.16	0.19	3.81

Table 6.2 Gap of integer programming and local search to the lower bound

Table 6.3 shows the heuristic calculation times on a Dell optiplex 760 computer with Intel quad core, 2.83 GHz processor. Both methods require little computation time, although the run times of IP increase substantially with the problem size.

Parameter	Values	computation time IP (sec)		computation time LS (sec)	
		average	maximum	average	maximum
Number of SKUs	25	0.10	0.64	0.01	0.06
	100	0.14	0.73	0.03	0.73
	400	1.48	10.34	0.50	2.66
Overall		0.58	10.34	0.18	2.66

Table 6.3 Calculation times for the problem instances (IP=Integer Programming, LS = Local Search)

6.3.2 The added value of using emergency shipments

We compare our selective emergency shipment policy to two one-size-fits-all policies: one with emergency shipments (OSFA ES), and one where we can choose for every item whether we backorder or use emergency shipments (OSFA BO+ES). In the latter policy, we differentiate between items, but *not* between customer classes. Backorder clearing is done first-come-first-served. Both policies aggregate all customers into a single demand class with maximum waiting time $W_1^{\max}$.

We use OSFA ES as a benchmark, since it has often been used in literature to measure the added value of critical level policies (see, e.g. Kranenburg and Van Houtum [14]). Table 6.4 shows the average and maximum percentage savings of OSFA BO+ES and our policy to OSFA ES. We also show the results for different holding cost intervals, because this parameter has most influence on the savings.

Parameter	Values	savings over OSFA ES (%)			
		OSFA BO+ES		Selective em shipments	
		average	maximum	average	maximum
holding cost range interval	[0.02, 19.98]	7.9	38.9	11.7	45.8
	[0.2 - 199.8]	0.2	2.7	1.2	12.6
	[2 - 1998]	0.1	1.9	0.3	3.0
Overall		<u>2.7</u>	38.9	<u>4.4</u>	<u>45.8</u>

Table 6.4 Savings of selective emergency shipments over OSFA BO+ES and the division of items over shipment strategies

OSFA BO+ES leads to average savings of 2.7% over OSFA ES. Selective emergency shipments yield additional average savings of 1.7%, resulting in average savings of 4.4% over OSFA ES with a maximum of 45.8%. The savings are largest when the holding costs are low, because emergency shipments are less appealing then: it is cheaper to keep high stocks and to reserve emergency shipments for

requests from premium customers. We also obtain large savings when waiting time restrictions for class 1 demand are loose (7.4% on average).

Figure 6.1 displays on the left the fraction of items assigned to each shipment strategy for OSFA BO+ES and our policy (SES). The figure to the right shows the division of items over shipment strategies per holding cost interval for our policy.

Figure 6.1 Fraction of items per shipment strategy for selective emergency shipments and OSFA BO + ES (left) and fraction of items per holding cost interval (right)

The left side of Figure 6.1 shows that the fraction of items for which all demand is backordered ($D_i = 0$) is similar for both policies. We make the same observation for distinct parameter settings. Hence, we clearly use the selective shipment strategy ($D_i = 1$) to limit using expensive emergency shipments for both classes ($D_i = 2$). On average, $D_i = 1$ for 20% of the items. This fraction increases to more than 40% when the holding costs are low (see figure on the right).

To analyze the types of items for which differentiation is used most frequently, we determine the average holding costs and item demand rates associated with each strategy, see Figure 6.2. It appears that the differentiation strategy is generally used for expensive slow movers: little or no stock is kept for this type of item. The shipment strategy is thus essential for meeting waiting time restrictions: we use emergency shipments for premium customers, but backorder class 2 requests.

Figure 6.2 Average holding costs and demand rates per shipment strategy over all problem instances

6.3.3 Comparison to the critical level policy and a combined policy

We compare selective emergency shipments to a well-known critical level policy with emergency shipments (referred to as CLP ES) and we investigate the added value of combining both policies (referred to as CLP+SES). Table 6.5 shows the relative savings of the policies compared to OSFA ES.

Parameter	Values	savings over OSFA ES (%)		
		Em ship policies	CLP ES	CLP + SES
Treg - Tem (days)	4-1	7.6	5.8	16.1
	8-1	4.5	8.5	14.3
	8-2	4.1	7.5	13.1
	16-2	1.5	9.7	12.1
holding cost range interval	[0.02, 19.98]	11.7	0.5	16.2
	[0.2 - 199.8]	1.2	8.3	10.5
	[2 - 1998]	0.3	14.8	15.0
W1-max - W2-max (hours)	0.5 - 2	1.4	4.4	6.2
	0.5 - 4	1.4	7.6	10.4
	3 - 12	7.4	8.4	16.7
	3 - 24	7.5	11.1	22.4
Overall		4.4	7.9	13.9

Table 6.5 Savings of different policies over OSFA ES

Critical level policies generally perform better than selective emergency shipments, with average savings of 7.9%. This is caused by the mode of differentiation: critical level policies reserve stock for premium customers, who thus

often obtain a part right away. In contrast, customers need to wait at least T_i^{em} time units for an emergency shipment. The selective emergency shipment policy is also less sensitive to the waiting time restrictions than the critical level policy: the waiting time restriction for class 1 is usually dominant. Then, increasing $W_2^{\max}$ has little impact on the solutions found.

Selective emergency shipments outperform critical level policies when regular shipment times are short, holding costs are low, and waiting time restrictions for class 1 requests are loose. Then it is viable to meet (a part of) the requests through regular replenishment instead of expensive emergency shipments. Indeed, the fraction of items for which D_i is 0 or 1 is relatively high then (for the holding costs we can see this in Figure 6.1). Note that selective emergency shipments do not outperform CLP ES for the given waiting time restrictions, but this happens if we further increase $W_1^{\max}$. Under the mentioned conditions, the base stock levels with CLP ES tend to be high in order to avoid expensive emergency shipments.

It is obvious that the combined policy yields the best results, but the size of the additional gain is surprising: it is more than the combined savings of the individual policies. The key reason is that under CLP ES, the actual average waiting time tends to be considerably below the target *for class 2 customers*. The waiting time for class 1 customers is usually the bottleneck. If we include selective emergency shipments, we are able to push the actual performance of low priority customers closer to the target (0.04% instead of 29% in the experiments with 100 SKUs).

7. Conclusions and directions for further research

Based on the research, we draw the following key conclusions.

1. ***Both methods for finding a near-optimal solution yield an average gap to the lower bound of less than 1% and require little computation time.*** Greedy approaches are not always necessary to find a near-optimal solution: integer programming with the columns from the LP-relaxation works well and is simple.
2. ***Selective emergency shipments may have significant added value compared to one-size-fits-all approaches.*** The average savings are 4.4% compared to a one-size-fits-all strategy with emergency shipments.
3. ***Differentiation through emergency shipments is most useful for expensive slow-movers.*** It has most impact if little or no stock is kept.
4. ***The selective emergency shipment strategy outperforms critical level policies when item holding costs are low, regular shipment times are short, or class 1 target waiting times are loose.*** Then, emergency shipments are very expensive compared to regular replenishments.
5. ***We can achieve large savings by combining critical levels and selective emergency shipments (13.9% on average).*** The combined policy enables us to effectively differentiate in spare parts supply.

We see the following opportunities for further research.

1. ***Extend the model to more than two demand classes.*** It is not straightforward to do so under priority backorder clearing, because we then obtain a multi-dimensional Markov chain with the number of classes as dimension. Analysis is simple under FCFS backorder clearing, but the performance will be clearly suboptimal then. Further research is thus needed in this area.
2. ***Use better shipment strategies.*** We expect further savings if the decision when to use emergency shipments also depends on the system state (see Section 3.2).

3. **Consider a multi-echelon system.** We need to stock items somewhere to meet emergency shipment requests, resulting in costs that we have not considered yet.
4. **Include lateral transshipments between local warehouses.** This flexibility option may yield another interesting source of customer differentiation.

References

- [1] Alfredsson, P., and Verrijdt, J. (1999). Modeling emergency supply flexibility in a two-echelon inventory system. *Management science*, 45, 1416 – 1431
- [2] Arslan, H., Graves, S.C., and Roemer, T.A. (2007). A single-product inventory model for multiple demand classes. *Management Science*, 53, 1486–1500
- [3] Cohen, M.A., Agrawal, N. and Agrawal, V. (2006). Winning in the aftermarket. *Harvard Business Review*, 84, 129 – 138
- [4] Dekker, R., Hill, R.M., Kleijn, M.J., and Teunter, R.H. (2002). On the $(S-1, S)$ lost sales inventory model with priority demand classes, *Naval Research Logistics*, 49, 593 – 610
- [5] Deshpande, V., Cohen, M.A., and Donohue, K. (2003). A threshold inventory rationing policy for service-differentiated demand classes. *Management Science*, 49, 683-703
- [6] Eppen, G., and Schrage, L., (1981). Centralized ordering policies in a multi-warehouse system with lead times and random demand. In: *Multi-level Production/Inventory Control Systems: Theory and Practice*, North-Holland, L.B. Schwarz, Editor, 51–67
- [7] Gayon, J.-P., De Véricourt, F., Karaesmen, F. and Dallery, Y. (2005). Stock rationing in a multi-class make-to-stock queue with information on the production

status, Working paper, Ecole Centrale Paris, France, Duke University, USA, and Koç University, Turkey

- [8] Gilmore, P.C., and Gomory, R.E. (1961). A linear programming approach to the cutting-stock problem. *Operations research*, 9, 849 – 859
- [9] Ha, A.Y. (1997a). Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science*, 43, 1093 – 1103
- [10] Ha, A.Y. (1997b). Stock-rationing policy for a make-to-stock production system with two priority classes and backordering. *Naval Research Logistics*, 44, 457 – 472
- [11] Ha, A.Y. (2000). Stock rationing in an M/Ek/1 make-to-stock queue. *Management Science*, 46, 77 – 87
- [12] Hausman, W.H., and Erkip, N.K. (1994). Multi-echelon vs. single-echelon inventory control policies for low demand items. *Management science*, 40, 597 – 602
- [13] 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
- [14] 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
- [15] Möllering, K.T., and Thonemann, U.W., (2010). An optimal constant level rationing policy under service level constraints. *OR Spectrum*, 32, 319-341
- [16] Muckstadt, J.A., and Thomas, L.J. (1980). Are multi-echelon inventory methods worth implementing in systems with low-demand-rate items? *Management Science*, 26, 483 – 494

- [17] Topkis, D.M. (1968), Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes, *Management Science*, 15, 160 – 176
- [18] Van Utterbeeck, F., Wong, H., Van Oudheusden, D., and Cattrysse, D. (2009). The effects of resupply flexibility on the design of service parts supply systems. *Transportation Research Part E: Logistics and Transportation Review*, 45, 72 – 85
- [19] Veinott, A.F. (1965). Optimal Policy in a Dynamic, Single Product, Nonstationary Inventory Model with Several Demand Classes. *Operations Research*, 13, 761 – 778
- [20] De Véricourt, F., Karaesmen, F., and Dallery, Y. (2002). Optimal stock allocation for a capacitated supply system. *Management Science*, 48, 1486 – 1501
- [21] Wong, H., Van Oudheusden, D., and Cattrysse, D. (2007). Two-echelon multi-item spare parts systems with emergency supply flexibility and waiting time constraints. *IIE Transactions*, 39, 1045 - 1057

Working Papers Beta 2009 - 2010

nr.	Year	Title	Author(s)
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 Teaching Retail Operations in Business and Engineering Schools	Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo.
317	2010	Design for Availability: Creating Value for Manufacturers and Customers	Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak.
316	2010	Transforming Process Models: executable rewrite rules versus a formalized Java program	Pieter van Gorp, Rik Eshuis.
315	2010		Bob Walrave, Kim E. van Oorschot, A. Georges L. Romme
314	2010	Ambidexterity and getting trapped in the suppression of exploration: a simulation model	

313 2010	A Dynamic Programming Approach to Multi-Objective Time-Dependent Capacitated Single Vehicle Routing Problems with Time Windows	S. Dabia, T. van Woensel, A.G. de Kok
312 2010	Tales of a So(u)rcerer: Optimal Sourcing Decisions Under Alternative Capacitated Suppliers and General Cost Structures	Osman Alp, Tarkan Tan
311 2010	In-store replenishment procedures for perishable inventory in a retail environment with handling costs and storage constraints	R.A.C.M. Broekmeulen, C.H.M. Bakx
310 2010	The state of the art of innovation-driven business models in the financial services industry	E. Lüftenegger, S. Angelov, E. van der Linden, P. Grefen
309 2010	Design of Complex Architectures Using a Three Dimension Approach: the CrossWork Case	R. Seguel, P. Grefen, R. Eshuis
308 2010	Effect of carbon emission regulations on transport mode selection in supply chains	K.M.R. Hoen, T. Tan, J.C. Fransoo, G.J. van Houtum
307 2010	Interaction between intelligent agent strategies for real-time transportation planning	Martijn Mes, Matthieu van der Heijden, Peter Schuur
306 2010	Internal Slackening Scoring Methods	Marco Slikker, Peter Borm, René van den Brink
305 2010	Vehicle Routing with Traffic Congestion and Drivers' Driving and Working Rules	A.L. Kok, E.W. Hans, J.M.J. Schutten, W.H.M. Zijm
304 2010	Practical extensions to the level of repair analysis	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
303 2010	Ocean Container Transport: An Underestimated and Critical Link in Global Supply Chain Performance	Jan C. Fransoo, Chung-Yee Lee
302 2010	Capacity reservation and utilization for a manufacturer with uncertain capacity and demand	Y. Boulaksil; J.C. Fransoo; T. Tan
300 2009	Spare parts inventory pooling games	F.J.P. Karsten; M. Slikker; G.J. van Houtum
299 2009	Capacity flexibility allocation in an outsourced supply chain with reservation	Y. Boulaksil, M. Grunow, J.C. Fransoo
298 2010	An optimal 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
297 2009	Responding to the Lehman Wave: Sales Forecasting and Supply Management during the Credit Crisis	Robert Peels, Maximiliano Udenio, Jan C. Fransoo, Marcel Wolfs, Tom Hendriks
296 2009	An exact approach for relating recovering surgical patient workload to the master surgical schedule	Peter T. Vanberkel, Richard J. Boucherie, 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 Programming Approach	Ton G. de Kok, Gabriella Muratore
275	2009	Inventory redistribution for fashion products under demand parameter update	G.P. Kiesmuller, S. Minner
274	2009	Comparing Markov chains: Combining aggregation and precedence relations applied to sets of states	A. Busic, I.M.H. Vliegen, A. Scheller-Wolf

273	2009	Separate tools or tool kits: an exploratory study of engineers' preferences	I.M.H. Vliegen, P.A.M. Kleingeld, G.J. van Houtum
272	2009	An Exact Solution Procedure for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering	Engin Topan, Z. Pelin Bayindir, Tarkan Tan
271	2009	Distributed Decision Making in Combined Vehicle Routing and Break Scheduling	C.M. Meyer, H. Kopfer, A.L. Kok, M. Schutten
270	2009	Dynamic Programming Algorithm for the Vehicle Routing Problem with Time Windows and EC Social Legislation	A.L. Kok, C.M. Meyer, H. Kopfer, J.M.J. Schutten
269	2009	Similarity of Business Process Models: Metics and Evaluation	Remco Dijkman, Marlon Dumas, Boudewijn van Dongen, Reina Kaarik, Jan Mendling
267	2009	Vehicle routing under time-dependent travel times: the impact of congestion avoidance	A.L. Kok, E.W. Hans, J.M.J. Schutten
266	2009	Restricted dynamic programming: a flexible framework for solving realistic VRPs	J. Gromicho; J.J. van Hoorn; A.L. Kok; J.M.J. Schutten;

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