Service differentiation through selective lateral transshipments

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

Beta Working Paper series 395

BETA publicatie	WP 395 (working paper)
ISBN	
ISSN	
NUR	804
Eindhoven	October 2012

Service differentiation through selective lateral transshipments

E.M. Alvarez[1], M.C. van der Heijden, I.M.H. Vliegen, and W.H.M. Zijm

University of Twente, School of Management and Governance, P.O. Box 217,

7500 AE Enschede, The Netherlands

Abstract

We consider a multi-item spare parts optimization problem with multiple warehouses and two customer classes, where lateral transshipments are used as a differentiation tool. Specifically, *premium* requests that cannot be met from stock at their preferred warehouse may be satisfied from stock at other warehouses (so-called lateral transshipments). We first derive approximations for the mean waiting time per class in a single-item model with selective lateral transshipments. Next, we embed our method in a multi-item model minimizing the holding costs and costs of lateral and emergency shipments from upstream locations in the network. Compared to the strategy where selective use of emergency shipments is the only differentiation option, we show that the addition of selective lateral transshipments can lead to significant further cost savings (14% on average). Adding the option of stock reservation for premium customers (so-called critical levels) appears to have little added value.

Key words: inventory, service differentiation, lateral transshipments, stock rationing, spare parts

1 Introduction

In the capital intensive industry, companies may fully rely on the performance of key equipment for their operations. Downtime of these systems, such as radar systems on frigates (Al Hanbali and Van der Heijden [2]) or lithography systems in the semiconductor industry (Kranenburg and Van Houtum [12]), can have severe consequences then. As the users of such systems are mainly interested in equipment use rather than maintenance, they tend to outsource maintenance and spare parts supply, with performance agreements formalized in service contracts by service level

[1] Corresponding author; phone: +31-53-4893603; e-mail adress: e.m.alvarez@utwente.nl

agreements (SLA's). Examples are a minimum system availability and a maximum response time to failures. Often, SLA's vary among customers to reflect the value placed on system uptime (see e.g. Jalil [10]). A key challenge for the supplier is to satisfy all SLA's at minimal costs.

Spare parts suppliers generally handle differentiated service levels in two extreme ways. The first is to provide all customers with a uniform service process – a so-called one-size-fits-all approach (see e.g. Cohen et al. [6]) – that is usually designed to meet the tightest service levels. Such an approach is costly and results in customers with a standard contract getting higher service than needed, possibly at the expense of service to customers with a premium contract. Also, customers with a standard contract have no incentive to switch to a premium contract. The other extreme is to design separate supply chains per customer segment, with high priority customers being served from stock points close to their facilities and lower priority customers served from locations that are further away. In this approach, the supplier cannot centralize stocks to benefit from risk pooling (Eppen and Schrage [8]), resulting in higher stock levels in the supply chain than needed.

The literature on differentiation has mainly focused on alternatives where stock is kept centrally for all customers, with differentiation occurring using so-called critical levels (Veinott [17]). Such an approach reserves some stock for premium customers once the inventory drops to a certain threshold, the critical level. Although this approach can lead to large savings in theory, there are practical drawbacks. For instance, the engineers who repair the system are often accountable for speed of repair. Therefore, they will not delay repair if a part is available. Case studies at various companies, such as a manufacturer of medical image processing equipment in the Netherlands, show that critical level policies are seldom used as a consequence.

An alternative approach for differentiation is the use of *selective emergency shipments* (Alvarez et al. [3]), where demand in out-of-stock settings may either be backordered or satisfied using emergency shipments from a location with infinite supply. In this paper, we extend the selective emergency shipment model by also allowing *lateral transshipments for premium customers.* Throughput the paper, we use the terms "lateral transshipments" and "transshipments" interchangeably. In practice, a warehouse that is out of stock can often obtain the needed item from a neighboring warehouse that still has the item on-hand (see e.g. Kranenburg and Van Houtum [12]). This is often both faster and cheaper than an emergency shipment. Pooling stocks in this way can also result in lower overall stock levels in the supply chain (Paterson et al. [14]).

Therefore, it could be beneficial to use such transshipments for meeting premium requests. We do not allow lateral transshipments for non-premium requests to avoid that such a transshipment depletes stock that could have been used for meeting a premium request arriving just a bit later. Demand that cannot be met from on-hand stock (either directly or through transshipments) is either backordered or satisfied using emergency shipments from a location with infinite supply.

In a multi-item setting where multiple warehouses each receive requests from two customer classes (a premium and a non-premium class), we investigate the benefits of using selective transshipments in addition to backordering and emergency shipments. Furthermore, as we found large savings when combining selective emergency shipments with critical level policies, we also combine selective emergency shipments *and* selective transshipments with critical level policies.

The paper is structured as follows: in Section 2, we give a literature overview and state our main contribution. We then describe the system in Section 3 and present an approach in Section 4 for analyzing this system for a single item when transshipments are used for premium requests. This analysis approach serves as a building block for solving the multi-item optimization problem that we address in Section 5, where we also present the solution approach. We then discuss extensions to a model where a critical level policy is combined with selective transshipments and emergency shipments in Section 6. In Section 7, we discuss our extensive computational experiment. Finally, we draw conclusions and indicate further research areas in Section 8.

2 Literature

Our research is related to literature on (i) service differentiation and (ii) lateral transshipments. In the *service differentiation* area, we find contributions on both a tactical and an operational level. Most papers on the *tactical level* consider single-location models, with the main differentiation tool being the critical level policy, a concept introduced by Veinott [17]. The optimality of this approach has been proven under various circumstances, both under backordering and lost sales. We refer to Teunter and Klein Haneveld [15] for a literature review. Alternatively, Alvarez et al. [3] use selective emergency shipments in a single-location setting as a differentiation tool. If a warehouse is out of stock, demand may either be backordered or met using an emergency shipment from a central stock point with infinite supply, depending on the customer class and item type. Their approach can lead to cost savings over a one-size-fits-all approach that are close to those using critical level policies and to considerably larger savings when combined with critical level policies. To our knowledge, the only paper considering a multi-location model is Alvarez et al. [4]. The authors consider a multi-item setting with various customers where dedicated stock may be kept at a customer's site in addition to stock kept at a central location. At an *operational level*, the amount of literature is much more limited and comprises a few multi-location models. Jalil [10] and Tiemessen et al. [16] consider single-item models with multiple warehouses and multiple customer classes, where a request can often be met from more than one warehouse. Differentiation occurs by not necessarily satisfying a low priority request from the nearest warehouse (or from any warehouse in the system) to reserve stock for premium requests.

The literature on *lateral transshipments* considers both models with backordering and models with emergency shipments. Models with *backordering* have initially been considered by Lee [13] and Axsäter [5], who consider a two-echelon setting consisting of a depot and various bases which are divided into transshipment pools. Axsäter uses an iterative analysis approach, where each base is analyzed separately over a number of iterations under the assumption that transshipment requests at each base arrive according to Poisson processes. This logic has often been used in other papers, e.g. Alfredsson and Verrijdt [1] and Van Wijk et al. [18]. Models with *emergency shipments* have initially been considered by Dada [7] and Alfredsson and Verrijdt [1], who analyze similar two-echelon models. Some recent contributions are Kranenburg and Van Houtum [12], who consider a model in which only a subset of warehouses can act as a transshipment source (so-called main warehouses), and Van Wijk et al. [18], who consider a model in which transshipment requests at a warehouse are only met if the stock level at that warehouse is above a so-called hold back level. We refer to Paterson et al. [14] for further details.

To our knowledge, lateral transshipments have not been used as a service differentiation tool before. In this paper, lateral transshipments may only be used to satisfy premium customer requests. This form of differentiation likely has significant added value: lateral transshipments are generally both faster and less expensive than emergency shipments. Hence, if there is added value to using selective emergency shipments for differentiation, it will likely be beneficial to use selective transshipments for this purpose as well. A complication, however, is that the feasibility of a lateral transshipments depends on the stock levels at other warehouses, whereas emergency shipments are always possible. Hence, we investigate under what conditions lateral transshipments are beneficial and how often each shipment option (i.e., lateral transshipments, emergency shipments, backordering) is used in a multi-item setting. To do so, we require single-item building blocks that have not been considered in literature so far, namely approaches to analyze the model when transshipments are used for premium customers only. We also investigate the added value of combining selective transshipments and emergency shipments with critical level policies. Our detailed contributions are:

1. We show how to analyze the system for a single item under lateral transshipments for premium customers. We also extend this approach for the combination with critical levels.
2. We develop an approach similar to Dantzig-Wolfe decomposition to optimize the overall multi-item model and show that this approach is fast and gives good quality solutions. Although such an approach has been used for solving similar problems before, its application is not straightforward for our problem, since we have a large number of control options.
3. In an extensive computational experiment, we show (i) that there is significant added value to using selective transshipments in addition to selective emergency shipments, especially in settings with slow moving items and (ii) that the combination of selective transshipments and selective emergency shipments is a good alternative to using critical level policies.

3 Model

3.1 Model outline

We consider a multi-item network of multiple local warehouses and a central depot with infinite supply. Each warehouse has its own customer base consisting of premium and non-premium customers. Per customer class, there is a maximum amount of time that customers of that class are willing to wait on average for parts. Naturally, the premium class has the strictest waiting time requirement. Direct requests at a warehouse (i.e., from its own customer base) are met from stock at the warehouse if possible, with a replenishment request being sent to the central depot (i.e., we consider a continuous-time, one-for-one replenishment policy).

A *premium* customer request that cannot be met from on-hand stock may be satisfied through a *lateral transshipment* from another warehouse. We consider a model where transshipments are only used for a subset of warehouses and items, with the selection of the most appropriate subset being part of the multi-item optimization problem (Section 5). If transshipments of a specific item are *not* allowed at a warehouse, that warehouse can neither request the item at another warehouse nor receive transshipment requests. Not allowing transshipments may be justified if a warehouse is far away from other warehouses or if an inexpensive fast moving item is considered (for which a lateral transshipment is relatively expensive). In contrast, if transshipments are allowed, the warehouse can both send and receive transshipment requests. On-hand stock is always used to satisfy an incoming transshipment request, i.e., no stock is reserved for direct requests. A warehouse issues transshipment requests to other warehouses in a predetermined order. Such an order is common in practice and will depend on shipment times and costs between warehouses.

A request that cannot be met from stock at the direct warehouse or through a lateral transshipment is either *backordered* or met using an *emergency shipment* from the central depot. Based on these shipment options, we consider the following three shipment strategies:

1. *Full backordering:* Premium and non-premium requests are backordered, with backorders cleared first-come-first-served. Premium backorders thus do not receive higher priority.
2. *Emergency shipments for premium customers only:* we backorder non-premium requests.
3. *Emergency shipments for all customers.*

We do not allow premium requests to be backordered when non-premium requests are met through emergency shipments. The shipment strategy may vary per item and warehouse. Alvarez et al. [3] have shown that the suitability of a shipment strategy depends on the characteristics of the item: full backordering is particularly beneficial for relatively inexpensive items with high demand rates, while emergency shipments are more suitable for premium requests for expensive slow moving items. As demand rates differ per warehouse, the shipment strategy may also vary among warehouses. The lateral and emergency shipment times do not have a specific distribution: we only use the mean shipment times in our model.

Figure 1 shows a single-item three-warehouse example where transshipments are only allowed among warehouses 1 and 2. The shipment strategies differ per warehouse: warehouse 2 uses full backordering, warehouse 1 uses emergency shipments for premium requests only, and warehouse 3 uses emergency shipments for all requests.

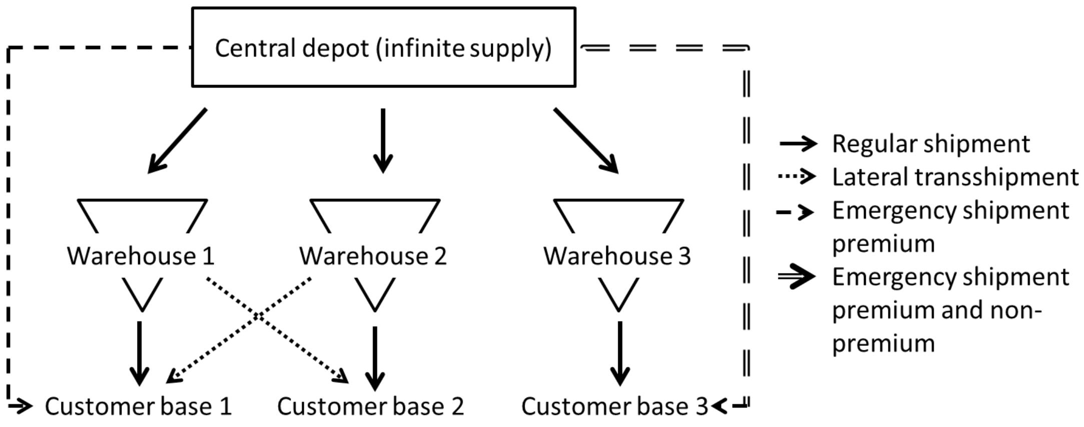

Figure 1 Example system with 3 warehouses

3.2 Key assumptions and notation

We make the following assumptions:

1. All direct requests arrive according to mutually independent Poisson processes.
2. The replenishment lead time to any warehouse is exponentially distributed. This assumption facilitates system analysis using continuous-time Markov chains. The system performance measures also tend to be insensitive to the lead time distribution, especially when emergency shipments are used for both classes (e.g. Alfredsson and Verrijdt [1], Alvarez et al. [3]).
3. The shipment time from any warehouse to a customer is negligible.
4. Lateral transshipments are faster than emergency shipments and also have lower shipment costs. As a result, they are preferred over emergency shipments.
5. Lateral and emergency shipments are sent directly to the customer and not via the warehouse.
6. Emergency shipment requests originate from the warehouse that needs the item: a second warehouse cannot request the item and then forward it to the warehouse who actually needs it.

We have K warehouses that each receive requests for I items from class 1 (premium) customers and class 2 (non-premium) customers. On average, class j customers ($j = 1,2$) are willing to wait at most W_j^{max} time units for spares. Direct requests for item i ($i = 1,..,I$) from class j customers at warehouse k ($k = 1, \dots, K$) occur at rate m_{ijk}, with $M_{jk} = \sum_{i=1}^{I} m_{ijk}$ denoting the total direct demand from class j customers arriving at warehouse k and $M_k = M_{1k} + M_{2k}$ denoting the total direct demand arriving at warehouse k. The mean replenishment lead time of item i to warehouse k is denoted by T_{ik}^{reg}, the emergency time by T_{ik}^{em}, and the lateral transshipment time from warehouse l by T_{ilk}^{lat} (with $T_{ilk}^{lat} \leq T_{ik}^{em} \leq T_{ik}^{reg}\ \forall l$). Warehouse k issues transshipment requests to other warehouses in the order specified by $\sigma_k = \{\sigma_k(1), \dots, \sigma_k(K-1)\}$, with $\sigma_k(n)$ being the n-th warehouse in the sequence. Note that σ_k is the same for all items, as the order will only depend on the shipment distances and costs among warehouses. Also, σ_k only indicates the order in which we try to find a transshipment source. Whether a warehouse can actually serve as a transshipment source for warehouse k also depends on the decision whether transshipments are allowed from that warehouse, and on the available stock at the time of a request. The holding cost parameter h_i denotes the item i unit costs per time unit, identical for all warehouses. Emergency and lateral shipments of item i to warehouse k occur at additional costs C_{ik}^{em} and C_{ilk}^{lat} over the costs of a regular replenishments, with l denoting the warehouse sourcing the item. We assume that $C_{ilk}^{lat} \leq C_{ik}^{em}, \forall l$, as this generally holds in practice.

We have three decision variables for each combination of item i and warehouse k, i.e., (1) the base stock level S_{ik}, (2) the lateral transshipment strategy L_{ik} denoting whether transshipments are allowed for that item and warehouse (L_{ik} then equals 1), and (3) the shipment strategy D_{ik} which denotes the highest customer class for which emergency shipments are used. In a setting with two customer classes, D_{ik} can take on three values: 0 (full backordering), 1 (emergency shipments for premium customers only), and 2 (emergency shipments for all customers). On a system level, the variables are denoted by vectors $\boldsymbol{S}_i = [S_{i1}, \dots, S_{iK}]$, $\boldsymbol{L}_i = [L_{i1}, \dots, L_{iK}]$ and $\boldsymbol{D}_i = [D_{i1}, \dots, D_{iK}]$. We aggregate all variables in an *item policy* $(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)$. As performance measures, we have $EW_{ijk}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)$, the expected class-$j$ waiting time for item i at warehouse k, and $TC_{ik}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)$, the total relevant costs for item i at warehouse k. The relevant costs consist of holding costs and extra costs for lateral and emergency shipments over regular replenishments.

4 Analysis

4.1 Approach

In this section, we focus on the special case where transshipments are allowed among all warehouses (i.e., $\boldsymbol{L}_i = [1, \dots, 1]$). The analysis under alternative values for $\boldsymbol{L}_i$ is straightforward: if $L_{ik} = 0$, warehouse k can be analyzed individually, as it does not send or receive transshipments of item i. An exact analysis with continuous-time Markov chains is intractable for more than 2 warehouses: we have to keep track of the inventory level at each warehouse separately to determine when transshipments are needed and where stocks are available. Solving such a Markov chain leads to very large computation times for systems with many warehouses. Therefore, we use a decomposition approach in which we analyze each warehouse separately and iteratively update the demand rates arising from lateral transshipments. Such an approach has led to accurate results for related models (Axsäter [5], Alfredsson and Verrijdt [1], and Van Wijk et al. [18]). A key approximation is that transshipment requests arrive according to *Poisson processes* with a known rate. Then, each warehouse can be analyzed separately, resulting in fill rate estimates. In turn, the fill rates at all warehouses determine the rate at which transshipment requests occur. Using a similar rationale, we develop an iterative procedure to analyze a system where lateral transshipments are only possible for a subset of all customers. We also assume that all warehouses operate independently of each other, which allows us to compute, amongst others, the fraction of demand met through transshipments as simple products of warehouse fill rates. Obviously, this assumption does not hold in reality. We include these dependencies to some extent by iteratively updating the transshipment rates among warehouses.

Section 4.2 gives further notation for computing $EW_{ijk}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)$ and $TC_{ik}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)$. Section 4.3 gives the main analysis steps, and Section 4.4 details the analysis of a warehouse. Section 4.5 evaluates the approach performance. We omit suffix i, as we consider a single item only.

4.2 Additional notation

We introduce the notation below, which applies for each warehouse k and customer class j (when applicable) The term 'demand at warehouse k' refers to the direct demand at that warehouse.

- e_k: the rate at which transshipment requests arrive.
- $\beta_{jk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$: the fill rate (i.e., the fraction of demand met directly from stock).
- $\alpha_{jlk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$: the fraction of demand met through transshipments from a warehouse l, with $\alpha_{2lk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}) = 0$ (transshipments are not used for non-premium customers). Also, $\alpha_{1lk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}) = 0$ when L_l or L_k equals 0: then, no transshipments are sourced from l.
- $\gamma_{jk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$: the fraction of demand met through emergency shipments, with $\gamma_{jk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}) = 0$ if $j > D_k$ (then, emergency shipments are not allowed for that class).
- $EBO_{jk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$: the mean backorder level, with $EBO_{jk}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}) = 0$ if $j \leq D_k$.

Using these performance measures, we find $EW_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ and $TC_k(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ as follows:

$$EW_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = EBO_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})/m_{jk} + \gamma_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})T_k^{em} + \sum_{l\in\sigma_k} \alpha_{jlk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})T_{lk}^{lat} \tag{1}$$

$$TC_k(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = hS_k + \sum_{l\in\sigma_k} \alpha_{1lk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})m_{1k}C_{lk}^{lat} + \sum_{j=1}^{2} \gamma_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})m_{jk}C_k^{em} \tag{2}$$

The first term of $EW_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ arises from backordering (using Little's formula), whereas the second and third term denote the waiting time arising from emergency and lateral transshipments. Note that $TC_k(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ consists of holding costs (which are computed over both on-hand stock and items in the pipeline), and the costs for using lateral and emergency shipments if applicable.

4.3 Main analysis steps

Our main analysis steps are:

1. **Initialization:** $e_k = 0, k = 1..K$, so we initially ignore lateral transshipments.
2. **Warehouse analysis:** Compute fill rates $\beta_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ and the expected number of backorders $EBO_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ for all warehouses k and classes j given the current value of e_k.
3. **Update** the transshipment rates e_k $\forall k$ given the current values of $\beta_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$
4. **Finish:** Stop if the change in e_k is smaller than some small ε $\forall k$. Otherwise, go to step 2.

We discuss Step 2 in more detail in Section 4.4. We update e_k in step 3 as follows: let e_{kl} denote the rate at which transshipment requests arrive at warehouse k from warehouse l. If $k = \sigma_l(n)$ for any positive integer n, k receives transshipment requests from l when l and all warehouses $\sigma_l(1)$ up to $\sigma_l(n-1)$ are out of stock. Assumption of independence among warehouses, we find:

$$e_{kl} = m_{1l}\left(1 - \beta_l(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})\right)\prod_{x=0}^{n-1}\left(1 - \beta_{\sigma_l(x)}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})\right) \tag{3}$$

$$e_k = \sum_{l|k\in\sigma_l} e_{kl} \tag{4}$$

We obtain $\alpha_{1kl}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ and, if applicable, $\gamma_{jl}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ from equations (5) and (6) respectively. We find α_{1kl} by multiplying the fraction of premium demand at l forwarded to k (i.e., e_{kl}/m_{1l}) by the probability that this demand can be met from on-hand stock at k (i.e., $\beta_k(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$).

$$\alpha_{1kl}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \beta_k(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) e_{kl}/m_{1l} \tag{5}$$

$$\beta_l(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) + \gamma_{jl}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) + \sum_{k\in\sigma_l} \alpha_{jkl}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) \; = 1 \quad j \le D_k \tag{6}$$

4.4 Detailed analysis of a single warehouse

We find the fill rate and the expected number of backorders per class from the distribution of the number of outstanding orders at the warehouse by using a continuous time Markov chain. For simplicity, we drop index k and denote $\mu = 1/T^{reg}$. Let $\lambda = m_1 + m_2 + e$ denote the demand rate including transshipment requests when the warehouse has stock on-hand, and $\theta(D)$ the demand rate under shipment strategy D when the warehouse is out of stock. We find for $\theta(D)$:

- $\theta(2) = 0$: all demand is lost (i.e., met through lateral or emergency shipments).
- $\theta(1) = m_2$: premium demand is met through lateral or emergency shipments.
- $\theta(0) = \pi_1 m_1 + m_2$: premium requests are backordered when the item cannot be obtained elsewhere in the system, which coincides with all warehouses in σ being out of stock. Hence, the probability π_1 of a premium backorder equals $\prod_{l\in\sigma}\big(1-\beta_l(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})\big)$.

Figure 2 shows the Markov chain. At S or more outstanding items, the arrival rate becomes $\theta(D)$.

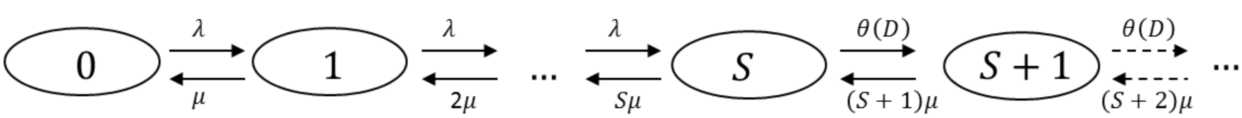

Figure 2 Markov chain of the number of outstanding orders at warehouse under shipment strategy *D*

Under full emergency shipments ($D = 2$), the Markov chain simplifies to an Erlang loss system with S servers. Using the notation $\rho = \lambda/\mu$, we thus have (see amongst others Gross et al. [9]):

$$\beta_1(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \beta_2(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = 1 - \frac{\rho^S/S!}{\sum_{w=0}^{S}\rho^w/w!} \tag{7}$$

Under (partial) backordering, we solve balance equations to find the steady-state probabilities p_n of n outstanding orders. With ρ as before, and ρ_1 equal to $\theta(D)/\mu$, we find:

$$p_0 = \left\{\sum_{w=0}^{S}\frac{\rho^w}{w!} + \left(\frac{\lambda}{\theta(D)}\right)^S\left(e^{\rho_1} - \sum_{w=0}^{S}\frac{\rho_1^w}{w!}\right)\right\}^{-1} \tag{8}$$

$$p_n = \rho^{\min\{n,S\}} {\rho_1}^{[n-S]^+} \frac{1}{n!} p_0 \tag{9}$$

$$\beta_1(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \beta_2(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \sum_{n=0}^{S-1} p_n \tag{10}$$

$$EBO(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \left(\frac{\lambda}{\theta(D)}\right)^S p_0 \left\{\rho_1 \left(e^{\rho_1} - \sum_{n=0}^{S-1} \frac{\rho_1^n}{n!}\right) - S\left(e^{\rho_1} - \sum_{n=0}^{S} \frac{\rho_1^n}{n!}\right)\right\} \tag{11}$$

Under partial backordering $(D = 1)$, (9) denotes the non-premium mean backorder level $EBO_2(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$. Under full backordering $(D = 0)$, (9) denotes the total mean backorder level. As premium and non-premium backorders occur at rates $\pi_1 m_1$ and m_2 respectively, we have:

$$EBO_1(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \frac{EBO(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})\pi_1 m_1}{\pi_1 m_1 + m_2} \tag{12}$$

$$EBO_{2k}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \frac{EBO(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})m_2}{\pi_1 m_1 + m_2} \tag{13}$$

4.5 Quality of the analysis approach

We compare our method to simulation for three performance measures: $\alpha_{1k}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \sum_{l\in\sigma_k} \alpha_{1lk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$, and $\beta_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ and $EW_{jk}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$, $(j = 1,2)$. We test 32 problem instances with either 6 or 18 warehouses and transshipments allowed at all warehouses (i.e., $L_k = 1\ \forall k$). For the simulation, we used a replication/deletion approach with at least 0.3 million requests for both premium and non-premium customers (average values are 1 million premium and 5 million non-premium requests). Table 1 shows the relative deviation of our method to simulation. The method is very accurate for slow movers and systems with many warehouses. We thus expect the approach to be accurate for practical instances. In systems with 6 warehouses and low stock levels (resulting in fill rates below 50%), the estimate of the transshipment fraction $\alpha_{1k}(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ can be poor. Still, in practice it will not occur that stocks of fast movers are low: these items have a high contribution to the overall waiting time. Therefore, waiting times for these items should be low. The maximum computation time for an instance is 12 milliseconds. Clearly, our approach is accurate and requires little computation time. As a result, it will be a suitable building block for solving multi-item problem of the next section. We refer to Appendix A for more details.

Settings			**Average deviation**				**Maximum deviation**			
M_k	K	S_k	β_k	α_k	EW_{1k}	EW_{2k}	β_k	α_k	EW_{1k}	EW_{2k}

0.05	6	1	0%	1%	1%	0%	0%	2%	5%	0%
		2	0%	0%	0%	0%	0%	1%	1%	0%
	18	1	0%	0%	0%	0%	0%	0%	0%	0%
		2	0%	0%	0%	0%	0%	1%	1%	1%
0.5	6	4	2%	8%	4%	1%	6%	20%	9%	5%
		8	0%	1%	1%	0%	0%	1%	1%	1%
	18	4	0%	1%	1%	0%	2%	3%	4%	1%
		8	0%	1%	1%	1%	0%	1%	1%	1%

Table 1 Relative deviations of the analysis approach to simulation.

5 Problem description and optimization

Problem $(P1)$ minimizes the total system costs $TC(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D})$ under restrictions on the mean aggregate waiting times per customer class *and* warehouse. A high waiting time at one warehouse thus cannot be compensated by a low waiting time at another warehouse, although such a variant (e.g. if a customer can be serviced from multiple warehouses) can be analyzed in a similar way.

$$(P1) \qquad \min TC(\boldsymbol{S},\boldsymbol{L},\boldsymbol{D}) = \sum_{i=1}^{I}\sum_{k=1}^{K} TC_{ik}(\boldsymbol{S}_i,\boldsymbol{L}_i,\boldsymbol{D}_i)$$

$$\text{s.t.} \qquad \sum_{i=1}^{I}\frac{m_{ijk}}{M_{jk}} EW_{ijk}(\boldsymbol{S}_i,\boldsymbol{L}_i,\boldsymbol{D}_i) \leq W_j^{max} \qquad j=1,2, k=1,\dots,K \qquad (14)$$

$$S_{ik} \in N_0, L_{ik} \in \{0,1\}, D_{ik} \in \{0,1,2\} \qquad (15)$$

We use an approach similar to Dantzig-Wolfe decomposition to solve $(P1)$. First, we reformulate the non-linear problem $(P1)$ to a linear problem by focusing on *item policies* as decision variables. The reformulated problem becomes to select one item policy for each item such that the system costs are minimized with the waiting time requirements still being met. Let B_i denote the set of item policies for item i, with $b_{ir} = \left(\boldsymbol{S}_i(r),\boldsymbol{L}_i(r),\boldsymbol{D}_i(r)\right)$ denoting a single item policy in B_i (i.e., $b_{ir} \in B_i$, with $r = 1,2,\dots,|B_i|$). Furthermore, let $x_{b_{ir}}$ be a binary variable indicating whether b_i is selected for item i (then, $x_{b_{ir}} = 1$). We then find linear problem $(P2)$:

$$(P2) \qquad \min \sum_{i=1}^{I}\sum_{k=1}^{K}\sum_{r=1}^{|B_i|} TC_{ik}(b_{ir}) x_{b_{ir}}$$

s.t. $$\sum_{i=1}^{I}\sum_{r=1}^{|B_i|}\frac{m_{ijk}}{M_{jk}}EW_{ijk}(b_{ir})x_{b_{ir}} \leq W_j^{max} \qquad j=1,2,k=1,\dots,K \qquad (16)$$

$$\sum_{r=1}^{|B_i|} x_{b_{ir}} = 1 \qquad i=1,\dots,I \qquad (17)$$

$$x_{b_{ir}} \in \{0,1\} \qquad i=1,\dots,I,r=1,\dots,|B_i| \qquad (18)$$

If B_i contains all item policies, $(P2)$ and $(P1)$ are equivalent and have the same optimal solution. Also, we find a lower bound on the costs by solving the LP-relaxation of (P2). Our challenge is the selection of policies to include in B_i for each item i, which is far from trivial: each policy $b_{ir} \in B_i$ refers to a *multi*-location problem. As we will show, an exact decomposition into single location problems is not possible under lateral transshipments. So we face a large set of relevant item policies. Furthermore, policy evaluation may take significant time when transshipments are allowed. The careful selection of item policies is thus crucial: we should select the minimal number of policies such that we still find a (near-) optimal solution to $(P2)$ and its LP-relaxation.

5.1 Solving the LP-relaxation

First, we first construct an initial set of item policies. Subsequently, we iteratively add policies to the policy set using column generation until no further interesting policies can be found.

5.1.1 Constructing an initial policy set

An initial policy set should lead to a feasible solution to the *integer* problem $(P2)$. One option to find such a set is to select a policy per item i such that $EW_{ijk}(b_{ir}) \leq W_j^{max}$ for each class j and warehouse k, guaranteeing $\sum_{i=1}^{I}\frac{m_{ijk}}{M_{jk}}EW_{ijk}(b_{ir}) \leq W_j^{max}$. As that option may lead to relatively large stock levels, we instead look for a policy over all items simultaneously. We use a ''biggest-bang-for-the-buck" algorithm, where we satisfy all unmet demand using emergency shipments, i.e., $L_{ik}=0$ and $D_{ik}=2$. This is justified since we only need a reasonable feasible solution as starting point for optimization. In each step of our algorithm, we increase the stock level S_{ik} by one unit at the item-warehouse combination (i,k) that leads to the greatest added value. We continue until all waiting time restrictions are met. To choose an option (i,k), we compute the decrease in waiting time relative to the extra investment needed. We find $\Delta W(\boldsymbol{S}_i+U_{ik})$, the decrease in waiting times for a unit stock increase at (i,k) (denoted by $\boldsymbol{S}_i+U_{ik}$), as follows:

$$\Delta W(\boldsymbol{S}_i + U_{ik}) = \sum_{j=1}^{2}\sum_{k=1}^{K}\left\{\left(\sum_{i=1}^{I}\frac{m_{ijk}}{M_{jk}}EW_{ijk}\big(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i\big) - W_j^{max}\right)^+ - \left(\sum_{i=1}^{I}\frac{m_{ijk}}{M_{jk}}EW_{ijk}(\boldsymbol{S}_i + U_{ik}, \boldsymbol{L}_i, \boldsymbol{D}_i) - W_j^{max}\right)^+\right\} \tag{19}$$

Here $[a]^+ = \max\{0, a\}$, which ensures that we only consider waiting time reductions above their respective thresholds. The extra investment $\Delta TC(\boldsymbol{S}_i + U_{ik}) = TC(\boldsymbol{S}_i + U_{ik}) - TC(\boldsymbol{S}_i)$ follows from (2). Note that options (i, k) may exist where both waiting times *and* costs decrease: a stock increase may lead to lower waiting times and fewer transshipments and emergency shipments (resulting in lower shipment costs). Then, we select the option with the largest $\Delta W(\boldsymbol{S}_i + U_{ik})$ among the options with lower costs. Otherwise, we select the option with the largest $\Delta W(\boldsymbol{S}_i + U_{ik})/\Delta TC(\boldsymbol{S}_i + U_{ik})$. During the procedure, we may find dominated policies that have both higher costs and higher waiting times at all warehouses than at least one other policy. We remove these policies before proceeding. Note that the initial policy set might contain more than one policy for each item: we expect that having many policies in the initial policy set reduces the amount of time needed for generating additional policies through column generation.

5.1.2 Finding additional policies through column generation

Column generation focuses on finding unconsidered item policies with negative reduced costs. Per item, we iteratively add the policy with minimal reduced costs to the policy set if these costs are negative. We stop once we cannot find further policies with negative reduced costs. Let u_{jk} (≤ 0) and v_i (≥ 0) denote the shadow price values for constraints (16) and (17) respectively, resulting from solving (P2) for a given set of item policies. The reduced costs $Z_i(b_i)$ for a policy b_i are now found as follows, with suffix r (i.e., the policy index) omitted for simplicity:

$$Z_i(b_i) = Z_i(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i) = \sum_{k=1}^{K}\left\{TC_{ik}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i) - \sum_{j=1}^{2}u_{jk}\frac{m_{ijk}}{M_{jk}}EW_{ijk}(\boldsymbol{S}_i, \boldsymbol{L}_i, \boldsymbol{D}_i)\right\} - v_i \tag{20}$$

It is far from trivial to find the item policy with the lowest reduced costs for an item i. If transshipments are allowed at a warehouse k, the performance at that warehouse depends on the rates at which it sends and receives transshipment requests. Hence, the optimal values for S_{ik} and

D_{ik} depend on the values of the decision variables at the other warehouses where transshipments are allowed. As a result, we can only guarantee optimality if all warehouses are jointly optimized.

For problems of realistic size, however, optimization over all warehouses jointly requires too much time: even for small instances with 10 items and 4 warehouses, the computation times can amount to three days. Instead, we opt for an approximate disaggregation of the overall problem into K single warehouse problems. Specifically, we can optimize the decision variable values at a warehouse k separately, if the decision variable values at the other warehouses are given. Clearly, the choice of variable values at warehouse k will influence the optimal values at other warehouses. Therefore, we iteratively optimize each warehouse until convergence occurs.

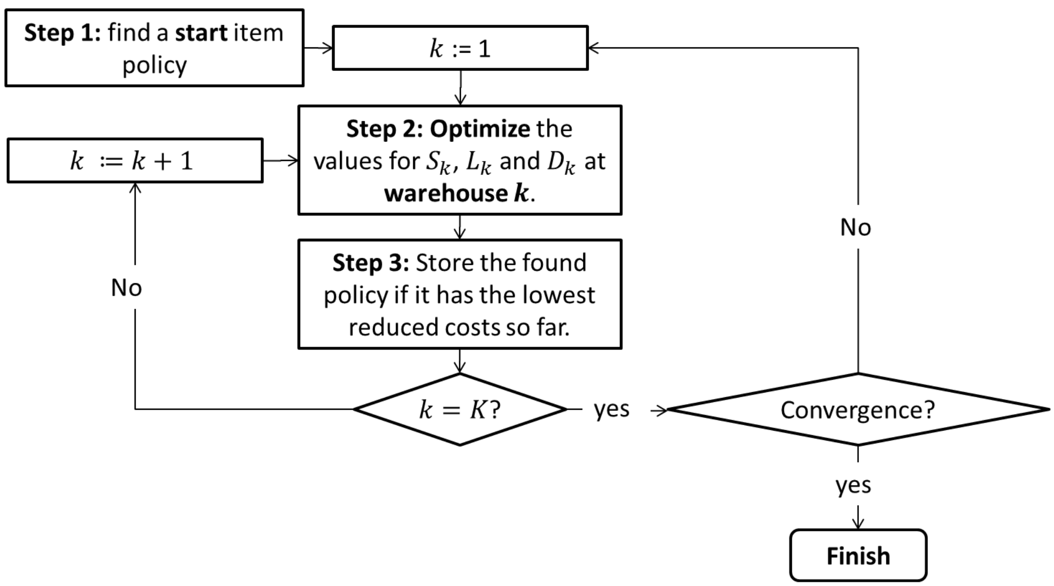

Figure 3 Column generation approach to find a near-optimal item policy for a particular item.

Error! Reference source not found. shows the main column generation steps for a single item. We omit suffix i in the figure and the rest of the section. First, we construct a start (i.e., initial) item policy. This policy serves as input for optimizing the decision variables at warehouse 1 a first time (i.e., the decision variable values for warehouses $l > 1$ serve as input for optimizing the values for warehouse 1). Then, we iteratively optimize decision variable values at a warehouse k, with the variable values of warehouses $n \neq k$ fixed to their most recent values. Each time we find a new item policy, we verify whether it has the lowest reduced costs so far and store it if this is the case. In an iteration, all warehouses in the system are considered. Convergence occurs when the decision variable values for all warehouses remain unchanged from one iteration to the next. We now give details on steps 1 and 2, with $(\boldsymbol{S}^*, \boldsymbol{L}^*, \boldsymbol{D}^*)$ being the best item policy found overall.

Step 1: finding a start item policy for the column generation procedure.

We can find a start policy in two extreme ways: either we allow transshipments at all warehouses (i.e., $L_k = 1\ \forall k$) or we do not allow them at any warehouse ($L_k = 0\ \forall k$). The advantage of the second option is that we can easily find good values for the remaining decision variables S_k and D_k, since each warehouse can be optimized separately. On the other hand, the first option will likely result in a more suitable start policy: we expect it to be easiest to move from a policy where transshipments are allowed at all warehouses to one where transshipments are only allowed at a subset of warehouses. In contrast, a move from a policy where transshipments are not used to one where transshipments are allowed can only occur if it is beneficial to transshipment among two or more warehouse (transshipments will not occur if they are only allowed at one warehouse).

These arguments prompt us to combine the options to find a start policy: first, we set $L_k = 0$ and optimize values for S_k and $D_k\ \forall k$. Then, we set $L_k = 1\ \forall k$ to obtain the start policy. In this way, we easily find values for S_k and D_k, while still obtaining a start policy where transshipments are allowed among all warehouses. Note that the values found for S_k and D_k result in a valid item policy both when $L_k = 0$ and when $L_k = 1$. Therefore, we analyze the system under both settings and store the policy with the lowest reduced costs $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ as the best policy so far $(\boldsymbol{S}^*, \boldsymbol{L}^*, \boldsymbol{D}^*)$.

Given that $L_k = 0$, we first optimize S_k for each value of $D_k \in \{0,1,2\}$ separately. Subsequently, we select the combination (S_k, D_k) leading to the lowest value for $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$. Given a value for D_k, we start with $S_k = 0$. We then iteratively increase S_k by one unit until a further increase has no benefit. Each time we increase S_k, we store the combination (S_k, D_k) if it leads to the lowest value for $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ so far (denoted by $Z^{min}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$). A further increase of S_k has no benefit once $h(\sum_{n=1}^{K} S_n + 1) - v \geq Z^{min}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$. Then, the minimal reduced costs for $S_k + 1$ (consisting of the system holding costs minus the item shadow price) already exceed the best reduced costs found so far. Note that the actual costs for $S_k + 1$ will be larger than that minimum value, as we ignore the shipment and waiting time costs.

Step 2: optimizing decision variable values at warehouse $\boldsymbol{k}$.

Our aim is to find the values for S_k, L_k and D_k that minimize the reduced costs $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ in the entire system. We do so, because the decision variable values at warehouse k influence the service levels at all warehouses. This influence can be significant: in particular, if stock is mainly (or even only) kept at warehouse k, the value of L_k is crucial, since it influences whether other warehouses have access to this stock.

We first optimize the decision variable values at k for each value of L_k separately. We then select the combination (S_k, L_k, D_k) that minimizes $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$. Note that when $L_k = 0$, the optimal values for S_k and D_k are the same as those found when looking for the start item policy (step 1), as the warehouse is not influenced by transshipment requests from other warehouses. When $L_k = 1$, we use the same approach as described in step 1 to find optimal values for S_k and D_k.

Given values for S_k, L_k and D_k, we have two options for estimating $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$. The first is to evaluate the system using the approach of Section 4. This option leads to the most accurate estimate of $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$, but is very time-consuming, in particular since we need to analyze the system for various combinations of S_k, L_k and D_k. The second option is to analyze only warehouse k (as in Section 4.4) and update the estimates of $\alpha_{jl}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ and $\gamma_{jl}(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ for the other warehouses l in the system through equations (5) and (6). Although the latter option is less accurate, we still use it since it is much faster (we only need to analyze one warehouse). Furthermore, it leads to sufficiently good solutions, as we will show in Section 5.1.3.

Once we have optimized the values of S_k, L_k and D_k, we use the approach of Section 4 to determine the actual value of $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ related to the newfound policy (as opposed to the approximate value found with the fast evaluation option). Using this actual value, we determine whether the new policy is the best so far (i.e. step 3) and store it if this is the case.

5.1.3 Quality of the obtained lower bound

We cannot guarantee that our column generation procedure always finds the item policy with the lowest reduced costs. Hence, we cannot ensure that we find the true optimal solution to the LP-relaxation of $(P2)$. Therefore, we compare the lower bound found with our column generation procedure to the lower bound when using an optimal column generation approach based on complete enumeration. As the latter approach is time-consuming, we only test small problem instances. We tested 128 problem instances, each with 5 or 10 items, and 2 or 4 warehouses. The remaining parameter values have been marked by an asterisk in Table 3 (Section 7.1).

I	***K***	**Relative deviation to true LB**	
		Average	Maximum
5	2	0.24%	2.23%
	4	0.13%	1.23%
10	2	0.24%	2.29%
	4	0.06%	0.46%

Table 2 Relative deviation to the true lower bound

From Table 2, we see that our approach indeed does not always find the true lower bound. Still, the deviation is at most 2.29%. Also, the deviations decrease in the number of warehouses and items, with the deviation being at most 0.46% for instances with 10 items and 4 warehouses. We thus expect the lower bound estimate to be accurate for larger instances that occur in practice.

5.2 Finding a near-optimal integer solution

The solution to the LP-relaxation might be fractional, with a combination of item policies being selected for certain items. Therefore, we need an approach to find a near-optimal solution to the integer problem $(P2)$. A simple option would be the intelligent rounding of the fractional values $x_{b_{ir}}$ of the LP-relaxation solution. However, such rounding will not be trivial, as we can have many items for which multiple policies are used: $(P2)$ has $2K + I$ constraints, leading to $2K + I$ item policies b_{ir} being basis variables (i.e., where $x_{b_{ir}} > 0$). For each item, at least one policy will be selected. We thus can have up to $2K$ items for which multiple policies are selected. Also, even if rounding is used to find a starting point for a local search procedure, the resulting solution is usually inferior to that obtained when solving the integer problem using linear optimization software such as CPLEX (see Alvarez et al. [3]). Therefore, we also solve $(P2)$ using CPLEX.

The policy set used for solving the LP-relaxation serves as a starting point for the integer problem policy set, as this set has worked well before (e.g. Alvarez et al. [3]). From the LP-relaxation set, we remove dominated policies (i.e., policies with both higher costs and waiting times than at least one other policy) and policies b_{ir} where $\frac{m_{ijk}}{M_{jk}} EW_{ijk}(b_{ir}) > W_j^{max}$ for at least one item i and warehouse k (the overall waiting time $\sum_{i=1}^{I} \sum_{r=1}^{|B_i|} \frac{m_{ijk}}{M_{jk}} EW_{ijk}(b_{ir}) x_{b_{ir}}$ also exceeds W_j^{max} then).

Still, computation times remain extensive under this smaller policy set and can amount to several hours. To decrease computation times, we consider two options, namely (i) further reducing the number of item policies per item or (ii) setting a limit on the time for CPLEX to find a solution. We choose for option (ii) because computation times remained large under option (i), irrespective of the criterion used for removing item policies (e.g. when reduced costs of a policy exceed a certain threshold). Also, the solutions found could be very poor, such as a gap to the lower bound of 14%. Option (ii) outperformed option (i) both in solution quality and computation times. The reason is that CPLEX often finds a good solution in the first few minutes, with improvements being minor from then on. Most time is spent on evaluating options that turn out to be infeasible. In an experiment with 80 problem instances – with 20 to 50 items and 10 to 20 warehouses – we considered time limits from 15 to 60 minutes. We found that a limit of 15 minutes already works well, with an average gap to the lower bound of 0.85%. Further improvements in quality were negligible under larger time limits (e.g., under 60 minutes the average gap reduced to 0.84%).

6 Extension to a model with critical levels

We now extend the model of Section 3 to include positive *critical levels*, i.e., where an amount of stock can be reserved for premium requests (either direct or transshipment requests). We let C_{ik} denote the critical level for item i at warehouse k, with $\boldsymbol{C}_i = [C_{i1}, \dots, C_{iK}]$ denoting the system critical levels. As before, premium requests may be met through transshipments when the direct warehouse is out of stock. However, warehouses with positive critical levels must always use emergency shipments to satisfy all (premium and non-premium) requests that cannot be satisfied through stock or transshipments. In other words: $D_{ik} = 2$ if $C_{ik} > 0$. We choose this model for its simplicity: as we shall see, the combination of critical levels with emergency shipments leads to a simple analysis of a warehouse. Also, it remains a realistic model: critical levels are mainly beneficial for expensive slow movers (see e.g. Alvarez et al. [3], Kranenburg and Van Houtum [11]). For such items, all unmet demand is generally satisfied through emergency shipments.

We can easily extend the analysis and optimization approaches for the combined model. In the analysis approach, the main steps and the computation of the transshipment rates (Section 4.3) remain the same. We must only be able to analyze a single warehouse under a critical level policy with emergency shipments. For the optimization procedure, we require a slight modification to the column generation method. Specifically, we must be able to optimize decision variable values – including the critical level – at a single warehouse. We discuss the warehouse analysis in Section 6.1 and the optimization in Section 6.2. For simplicity, suffixes i and k are omitted.

6.1 Warehouse analysis

Figure 4 shows the Markov chain of the number of outstanding orders, with g and μ as in Section 4.4. Non-premium demand is lost once $S - C$ orders or more are outstanding (equivalent to having at most C items on-hand). This Markov chain is similar to that of Kranenburg and Van Houtum [11] (the difference is that they do not consider transshipments, with e thus being zero).

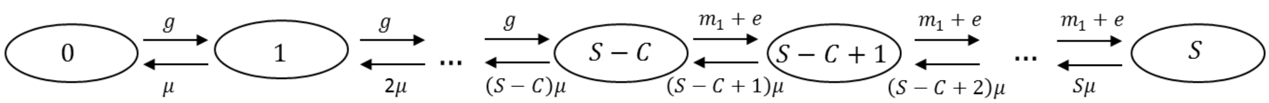

Figure 4 Markov chain of the pipeline at the local warehouse under a critical level policy with emergency shipments

The steady-state probabilities p_n and fill rate values β_j ($j = 1,2$) follow directly from the balance equations. We refer to Kranenburg and Van Houtum [11] for further details.

6.2 Warehouse optimization

Optimization occurs a similar way to Section 5.1.2, except that we must also optimize C for any combination (S, L, D) when $D = 2$. Given values for S and L, and $D = 2$, we find an optimal value for C as follows: starting with $C = 0$, we iteratively increase C by one unit, with the value leading to the lowest $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C})$ being stored. We keep increasing C until either (i) $C = S$ (we can reserve at most the base stock level) or (ii) $\beta_1(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C}) \geq 1 - \varepsilon$, with ε a specified tolerance. As C increases, the service level at premium customers improves (leading to lower reduced costs for those customers) at the expense of non-premium customers (for whom we find higher reduced costs). Overall, $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C})$ will thus first decrease and then increase. Still, we are unable to prove convexity of $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C})$ in C. However, once $\beta_1(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C})$ is close to 1, we can be certain that the reduced costs for premium customers will barely decrease further.

As before (see step 2 of Section 5.1.2), we have two options for estimating $Z(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D}, \boldsymbol{C})$ for given values of S, L, D and C, i.e., (i) a more accurate but time-consuming option of analyzing the entire system, and (ii) a faster but less accurate option of analyzing a single warehouse and only updating the values of $\alpha_j(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ and $\gamma_j(\boldsymbol{S}, \boldsymbol{L}, \boldsymbol{D})$ for the other warehouses. We use the first option for the model with critical levels only. This model serves as a benchmark for evaluating the model with lateral transshipments and emergency shipments as the only differentiation tools. The first evaluation option results in a stronger benchmark, as it generally gives better solutions.

7 Computational experiment

In an extensive computational experiment, we investigate (i) the performance of our optimization approach in terms of solution quality and computation time, (ii) the added value of the selective transshipment approach by comparing it to alternative differentiation approaches, and (iii) the suitability of the various shipment and transshipment strategies.

7.1 Experiment design

We construct 1024 problem instances, with T_{ilk}^{lat} always equal to 1 day and C_{ik}^{em} equal to 1000. Table 3 gives the other parameter values. The asterisks specify the values considered when evaluating the quality of our lower bound estimate (Section 5.1.3). Shipment times and costs are the same for all items and warehouses in a problem instance, with the lateral times and costs equal for any warehouse pair. Using a uniform distribution, the holding costs h_i are randomly drawn on the specified interval. Below, we detail how we obtain values for demand rates m_{ijk}.

	Parameter	Value
1	I	20, 50
2	K	10, 20
3	T_{ik}^{reg} (days)	8*, 16*
4	T_{ik}^{em} (days)	2*, 4*
5	$[W_1^{max}; W_2^{max}]$ (hours)	$[0.5; 2]$*, $[3; 24]$*
6	C_{ilk}^{lat}	100*, 500
7	Avg. M_{ik} – interval (p. day)	$[0.002; 0.05]$*,$[0.002; 0.5]$*
8	Avg. fraction premium $frac_p$	0.2*, 0.5
9	h_i – interval (p. day)	$[0.1; 10]$*,$[0.1; 50]$*

Table 3 Tested parameter values

Our demand rates m_{ijk} should differ among warehouses *and* items, with the *overall* fraction of premium demand in the system equal to $frac_p$. We find m_{ijk} in three steps: first, (1) we draw a value on the M_{ik}–interval (using a uniform distribution) to obtain the average demand rate for item i at one warehouse. By multiplying this value by K we find the total *system* demand rate M_i. Then, (2) we find the total premium demand in the system M_i^p by multiplying M_i by $frac_p$, with M_i^n denoting the remaining non-premium demand. Finally, (3) we disaggregate M_i^p and M_i^n over the warehouses to obtain m_{ijk}. Each warehouse is assigned a fraction of M_i^p and M_i^n (using a normal distribution), with normalization ensuring that $\sum_{k=1}^{K} m_{i1k} = M_i^p$ and $\sum_{k=1}^{K} m_{i2k} = M_i^n$.

The parameter values used by Kranenburg and Van Houtum [11][12] formed the basis for our values, as their values are based on practice. We consider items that have both high and low values, and high and low demand rates. The annual demand rates vary between 0.7 units and 183 units. In practice, an item's annual holding cost is a fraction (roughly 25%) of its value. In this study, we thus consider item values between 146 and 73000 euro's.

For simplicity, a warehouse k sends transshipment requests to other warehouses in the same order in all problem instances: $\sigma_k = \{k+1, k+2, \dots, K, 1, 2, \dots\}$. So, if warehouse k is out of stock, it first requests an item at warehouse $k+1$, then at warehouse $k+2$, etc.

For each combination of parameters in Table 3, we construct 2 sets of item demand rates and holding costs to ensure that our results are not dependent on the specific values of one sample. Combined with $2^9 = 512$ possible parameter combinations, we thus have 1024 instances in total.

7.2 Performance of the optimization procedure

Table 4 shows the solution quality – expressed as a relative gap to the lower bound estimate – and computation times of the optimization procedure. We used a Dell optiplex 760 with Intel quad core 2.83 GHz processor. Overall, the relative gap is 0.8% on average, with a maximum of 5.5%. The average and maximum gap decrease greatly as the number of items increases. We therefore expect the approach to work very well in realistic settings with many items. The average instance computation time is 12 minutes, with the maximum being 34 minutes. The computation time mainly increases with the number of items and warehouses in an instance.

Parameter	Values	Gap to lower bound estimate (%)		Computation time (min.)	
		Average	Maximum	Average	Maximum
I	20	1.3	5.5	7	21
	50	0.3	1.3	17	34

K	10	0.6	2.9	7	16
	20	1.0	5.5	16	34
Grand Total		**0.8**	**5.5**	**12**	**34**

Table 4 Solution quality and computation times of optimization procedure

7.3 Comparison to alternative differentiation approaches

We compare the selective transshipment model (ST_SES) to two alternatives: (i) a **selective emergency shipment model (SES)**, which is the special case of ST_SES with transshipments not allowed for any item or warehouse, and (ii) the **selective transshipment model with critical levels (CLP_ST_SES)** of Section 6. The added value of both ST_SES and CLP_ST_SES is expressed in terms of relative cost savings to SES, shown in Table 5. Notice that ST_SES has significant savings compared to SES: the average savings are 14% and can amount up to 34%. The savings are particularly large for instances with many slow moving items; for fast movers, lateral transshipments are not beneficial, as we will see in Section 7.4. Savings are also large when emergency shipment times are large and waiting times are not very strict, although the influence of these parameters is mainly large in settings with expensive slow movers.

Parameter	Values	Average savings over SES		Maximum savings over SES	
		ST_SES	CLP_ST_SES	ST_SES	CLP_ST_SES
T_{ik}^{em}	2	12%	12%	28%	28%
	4	17%	17%	34%	35%
$[W_1^{max}; W_2^{max}]$	$[0.5; 2]$	11%	11%	19%	19%
	$[3; 24]$	18%	19%	34%	35%
Max. M_{ik}	0.05	19%	20%	34%	35%
	0.5	9%	10%	20%	20%
Grand Total		**14%**	**15%**	**34%**	**35%**

Table 5 Relative savings of ST_SES over SES

The savings of CLP_ST_SES are similar to those of ST_SES. Clearly, there is little benefit to also allowing stock reservation for premium customers. The reason for this is that ST_SES is already able to differentiate very effectively: the aggregate waiting times per class j are close to their thresholds W_j^{max}. Adding critical levels therefore does not lead to extra savings.

7.4 Suitability of shipment and transshipment strategies

For each combination (L_{ik}, D_{ik}), Figure 5 shows the overall fraction of items and warehouses for which that combination is used. Clearly, lateral supply is very suitable for meeting premium requests: overall, transshipments are allowed at 96% of all item-warehouse combinations. For the remaining 4% of combinations where transshipments are not allowed, we always use full backordering. This is logical: if it is not beneficial to use lateral transshipments, it will also not be beneficial to use more expensive (and slower) emergency shipments. We can thus limit the combinations (L_{ik}, D_{ik}) that we should consider during optimization. The instances where lateral transshipments are not beneficial have many inexpensive fast moving items, high transshipment costs and loose waiting time restrictions, making transshipments expensive and unnecessary.

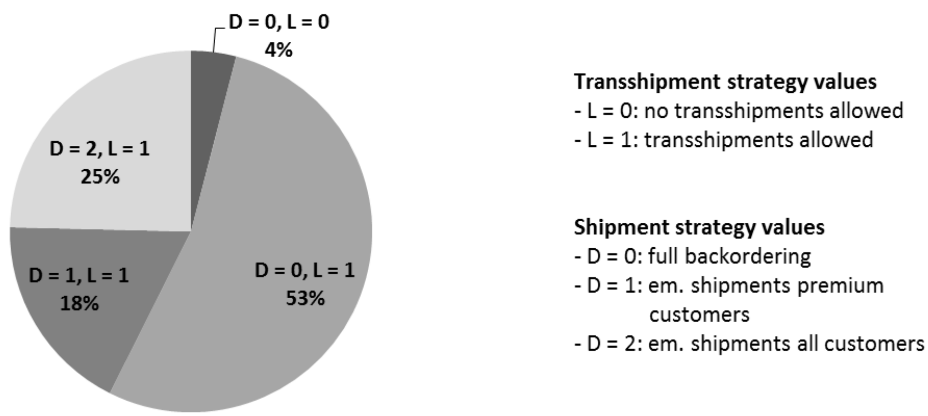

Figure 5 The fraction of items and warehouses using a particular (trans-)shipment combination

Overall, full backordering ($D = 0$) is the most frequently used shipment strategy (see Figure 5). Still, the added value of each shipment strategy depends heavily on the shipment times and type of item, as shown in Figure 6. Full backordering ($D = 0$) is especially beneficial when emergency shipments are slow relative to regular shipments, and when items are mostly cheap fast movers. Then, that strategy is used for roughly 85% of all items and warehouses. This coincides with earlier findings (Alvarez et al. [3]). Clearly, it is beneficial to consider backordering in addition to emergency shipments, even though it is common in both literature and business for emergency shipments to be the only shipment mode.

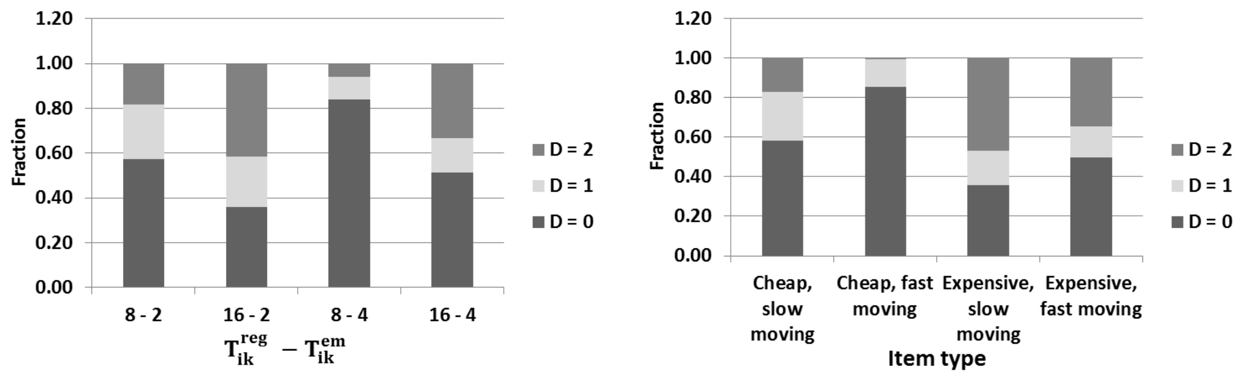

Figure 6 The influence of shipment times (left) and item type (right) on the use of various shipment strategies

Figure 7 shows for various problem instances how the strategies (L_{ik}, D_{ik}) are distributed over the items in each instance. We focus on instances with an M_{ik}-interval of $[0.002; 0.5]$ and a holding cost interval of $[0.1; 50]$; the results are similar for other parameter values. As expected, neither lateral transshipments nor emergency shipments are used for inexpensive fast movers, with both transshipments and (partial) emergency shipments used for expensive slow movers.

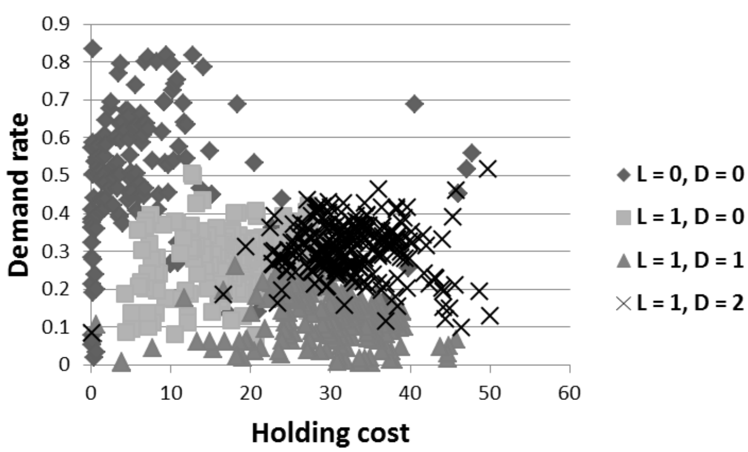

Figure 7 Item characteristics per (trans-)shipment strategy

8 Conclusions and further research

We considered a system with two customer classes where lateral transshipments and emergency shipments are both used selectively for service differentiation purposes. For a single-item setting, we developed an analysis approach when selective transshipments may only be used for premium requests. We also developed an approach similar to Dantzig-Wolfe decomposition to optimize the multi-item system under class-specific waiting time restrictions. Key conclusions are:

- Our **analysis approach is accurate and fast.**
- Our multi-item **solution approach** gives **near-optimal solutions in little computation time.**
- **Selective lateral transshipments lead to significant cost savings** when combined with selective emergency shipments. The savings are 14% on average and can amount to 34%. The savings can be particularly large (19% on average) if we have many expensive slow movers.
- Using **critical levels** besides selective (trans-)shipments does not lead to significant extra gains. Clearly, the combination of selective transshipments and emergency shipments is a good alternative to critical level policies, while being easier to implement in practice.
- **Backordering should also be considered as a shipment option in spare parts settings.** This is in contrast to the practice of always using emergency shipments for unmet demand.

From the findings in this paper and those in Alvarez et al. [3], our conjecture is that significant cost savings can be obtained by using any two differentiation tools jointly, such as critical levels with selective emergency shipments Alvarez et al. [3] or selective transshipments with selective emergency shipments. Combining three differentiation tools does not lead to additional benefits, but clearly "two out of three (options) is not bad". This flexibility to choose differentiation tools allows service providers to select those tools that are easiest to implement (with critical level policies possibly not being used in favor of options with fewer practical drawbacks).

We see two areas for further research:

1. **An extension to more than two customer classes**. If we have more than two customer classes, transshipments and emergency shipments might be used for a subset of customers (with the subset possibly varying per item). The analysis approach for such a system follows directly from that in Section 4. However, optimization will become much more difficult: we obtain additional variables (i.e., for what customer classes do we allow transshipments and emergency shipments) and additional constraints. Further research is thus needed to carefully select relevant item policies for the optimization problem.
2. **More flexible supply chains with warehouse clusters.** In this paper, lateral transshipments may occur among any pair of warehouses. In practice, however, warehouses may be clustered in regions (e.g. an EMEA region, a US region, and an Asia region), with transshipments only allowed among warehouses in the same region, i.e., lateral transshipments among regions are not allowed. Then, we must also decide how many regions there will be and how warehouses

are assigned to these regions. These additional decisions increase the complexity of the multi-item optimization problem, as we then have various configurations (and hence item policies) to choose from. Further research is thus needed to quickly select relevant item policies.

References

[1] Alfredsson, P., Verrijdt, J. (1999). Modeling emergency supply flexibility in a two-echelon inventory system. *Management science,* 45, 1416 – 1431

[2] Al Hanbali, A., Van der Heijden, M.C. (2011). Interval availability analysis of a two echelon, multi-item system. *Beta working paper series* 359. http://beta.ieis.tue.nl/node/1967

[3] Alvarez, E.M., Van der Heijden, M.C., Zijm, W.H.M. (2012a). The selective use of emergency shipments for service-contract differentiation. *International Journal of Production Economics,* http://dx.doi.org/10.1016/j.ijpe.2012.02.019

[4] Alvarez, E.M. Van der Heijden, M.C., Zijm, W.H.M. (2012b). Service differentiation in spare parts supply through dedicated stocks. *Beta working paper series* 373, http://beta.ieis.tue.nl/node/1996

[5] Axsäter, S. (1990). Modelling emergency lateral transshipments in inventory systems. *Management Science*, 36, 1329–1338

[6] Cohen, M., Agrawal, N., Agrawal, V., (2006). Winning in the aftermarket. *Harvard Business Review* 84, 129–138.

[7] Dada, M. (1992). A two-echelon inventory system with priority shipments. *Management Science* 38, 1140–1153.

[8] Eppen, G., 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.

[9] Gross, D., Shortle, J.F., Thompson, J.M., Harris, C.M. (2008). *Fundamentals of queuing theory,* 4^{th} edition,John Wiley & Sons, Hoboken, New Jersey.

[10] Jalil, M.N. (2011). *Customer information driven after sales service management: lessons from spare parts logistics,* PhD thesis, Erasmus Research Institute of Management (ERIM), http://repub.eur.nl/res/pub/22156/

[11] Kranenburg, A.A., Van Houtum, G.J.,(2008). Service differentiation in spare parts inventory management. *Journal of the Operational Research Society* 59, 946 – 955.

[12] Kranenburg, A., Van Houtum, G.J., (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research* 199, 908–921.

[13] Lee, H. L. (1987). A multi-echelon inventory model for repairable items with emergency lateral transshipments. *Management Science* 33, 1302–1316.

[14] Paterson, C., Kiesmüller, G., Teunter, R., Glazebrook, K., (2011). Inventory models with lateral transshipments: A review. *European Journal of Operational Research* 210, 125–136.

[15] Teunter, R. H., Klein Haneveld, W. K., (2008). Dynamic inventory rationing strategies for inventory systems with two demand classes, Poisson demand and backordering. *European Journal of Operational Research* 190, 156–178.

[16] Tiemessen, H.G.H., Fleischmann, M., Van Houtum, G.J., Van Nunen, J.A.E.E., Pratsini, E. (2012). Dynamic demand fulfillment in spare parts networks with multiple customer classes, *Beta working paper series* 377, http://beta.ieis.tue.nl/node/2002

[17] Veinott, A.F. (1965). Optimal Policy in a Dynamic, Single Product, Nonstationary Inventory Model with Several Demand Classes. *Operations Research* 13, 761 – 778

[18] Van Wijk, A.C.C., Adan, I.J.B.F., Van Houtum, G.J. (2012). Approximate evaluation of multi-location inventory models with lateral transshipments and hold back levels. *European Journal of Operational Research* 218, 624–635.

Appendix A: detailed performance analysis approach

We now give details on the comparison of our analysis approach to simulation from Section 4.5. Table A1 shows the parameter values tested. In all instances, the shipment strategies are spread evenly over the warehouses, i.e., one third of all warehouses uses full backordering, one third uses emergency shipments for premium customers only, etc. The demand rates and shipment times are the same at all warehouses, with a fraction $frac_p$ of demand coming from premium customers. We let large demand rates coincide with large stock levels.

Parameter	**Values**
K	6; 18
$[T_k^{reg}, T_{lk}^{lat}, T_k^{em}]$	[8,1,2]
Premium fraction $frac_p$	0.1; 0.2; 0.3; 0.5

M_k	0.05	0.5
S_k	1; 2	4; 8

Table A1 Parameter values considered for testing the analysis approach

Table A2 shows both the simulated and computed values for various performance measures. The values are averages over all warehouses (e.g. β_{avg} shows the average fill rate at a warehouse).

Settings					β_{avg}		α_{avg}		EW_{1-avg}		EW_{2-avg}	
Case	$\boldsymbol{K}$	$\boldsymbol{M_k}$	$\boldsymbol{frac_p}$	$\boldsymbol{S}$	**Sim**	**Analytic**	**Sim**	**Analytic**	**Sim**	**Analytic**	**Sim**	**Analytic**
1	6	0.05	0.1	1	0.68	0.68	0.32	0.32	0.32	0.32	1.14	1.14
2	6	0.05	0.1	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
3	6	0.05	0.2	1	0.67	0.67	0.33	0.33	0.33	0.33	1.16	1.16
4	6	0.05	0.2	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
5	6	0.05	0.3	1	0.66	0.66	0.33	0.34	0.34	0.34	1.18	1.18
6	6	0.05	0.3	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
7	6	0.05	0.5	1	0.65	0.65	0.34	0.35	0.37	0.36	1.21	1.22
8	6	0.05	0.5	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
9	6	0.5	0.1	4	0.50	0.50	0.48	0.48	0.52	0.51	1.23	1.23
10	6	0.5	0.1	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
11	6	0.5	0.2	4	0.49	0.48	0.48	0.50	0.56	0.54	1.22	1.22
12	6	0.5	0.2	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
13	6	0.5	0.3	4	0.47	0.46	0.48	0.51	0.59	0.56	1.20	1.21
14	6	0.5	0.3	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
15	6	0.5	0.5	4	0.44	0.42	0.47	0.55	0.67	0.62	1.19	1.22
16	6	0.5	0.5	8	0.96	0.96	0.04	0.04	0.04	0.04	0.05	0.05
17	18	0.05	0.1	1	0.68	0.68	0.32	0.32	0.32	0.32	1.14	1.14
18	18	0.05	0.1	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
19	18	0.05	0.2	1	0.67	0.67	0.33	0.33	0.33	0.33	1.16	1.16
20	18	0.05	0.2	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
21	18	0.05	0.3	1	0.66	0.66	0.34	0.34	0.34	0.34	1.18	1.18
22	18	0.05	0.3	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
23	18	0.05	0.5	1	0.65	0.65	0.35	0.35	0.35	0.35	1.22	1.22
24	18	0.05	0.5	2	0.94	0.94	0.06	0.06	0.06	0.06	0.15	0.15
25	18	0.5	0.1	4	0.50	0.50	0.50	0.50	0.50	0.50	1.23	1.23
26	18	0.5	0.1	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
27	18	0.5	0.2	4	0.48	0.48	0.52	0.52	0.52	0.52	1.23	1.22
28	18	0.5	0.2	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
29	18	0.5	0.3	4	0.46	0.46	0.54	0.54	0.55	0.54	1.23	1.22
30	18	0.5	0.3	8	0.96	0.96	0.04	0.04	0.04	0.04	0.06	0.06
31	18	0.5	0.5	4	0.40	0.40	0.59	0.60	0.62	0.60	1.25	1.25
32	18	0.5	0.5	8	0.96	0.96	0.04	0.04	0.04	0.04	0.05	0.05

Table A2 Detailed comparison between simulation and our analysis approach

Working Papers Beta 2009 - 2012

nr.	Year	Title	Author(s)
95	2012	Service differentiation through selective lateral transshipments	E.M. Alvarez, M.C. van der Heijden, I.M.H. Vliegen, W.H.M. Zijm
394	2012	A Generalized Simulation Model of an Integrated Emergency Post	Martijn Mes, Manon Bruens
393	2012	Business Process Technology and the Cloud: Defining a Business Process Cloud Platform	Vasil Stoitsev, Paul Grefen
392	2012	Vehicle Routing with Soft Time Windows and Stochastic Travel Times: A Column Generation And Branch-and-Price Solution Approach	D. Tas, M. Gendreau, N. Dellaert, T. van Woensel, A.G. de Kok
391	2012	Improve OR-Schedule to Reduce Number of Required Beds	J.T. v. Essen, J.M. Bosch, E.W. Hans, M. v. Houdenhoven, J.L. Hurink
390	2012	How does development lead time affect performance over the ramp-up lifecycle?	Andres Pufall, Jan C. Fransoo, Ad de Jong
389	2012	Evidence from the consumer electronics industry	Andreas Pufall, Jan C. Fransoo, Ad de Jong, Ton de Kok
388	2012	The Impact of Product Complexity on Ramp-Up Performance	Frank P.v.d. Heuvel, Peter W.de Langen, Karel H. v. Donselaar, Jan C. Fransoo
387	2012	Co-location synergies: specialized versus diverse logistics concentration areas	Frank P.v.d. Heuvel, Peter W.de Langen, Karel H. v.Donselaar, Jan C. Fransoo
386	2012	Proximity matters: Synergies through co-location of logistics establishments	Frank P. v.d.Heuvel, Peter W.de Langen, Karel H.v. Donselaar, Jan C. Fransoo
385	2012	Spatial concentration and location dynamics in logistics:the case of a Dutch province	Zhiqiang Yan, Remco Dijkman, Paul Grefen
384	2012	FNet: An Index for Advanced Business Process Querying	W.R. Dalinghaus, P.M.E. Van Gorp
383	2012	Defining Various Pathway Terms	Egon Lüftenegger, Paul Grefen, Caren Weisleder

382	2012	The Service Dominant Strategy Canvas: Defining and Visualizing a Service Dominant Strategy through the Traditional Strategic Lens	Stefano Fazi, Tom van Woensel, Jan C. Fransoo
381	2012	A Stochastic Variable Size Bin Packing Problem With Time Constraints	K. Sharypova, T. van Woensel, J.C. Fransoo
380	2012	Coordination and Analysis of Barge Container Hinterland Networks	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
379	2012	Proximity matters: Synergies through co-location of logistics establishments	Heidi Romero, Remco Dijkman, Paul Grefen, Arjan van Weele
378	2012	A literature review in process harmonization: a conceptual framework	S.W.A. Haneya, J.M.J. Schutten, P.C. Schuur, W.H.M. Zijm
377	2012	A Generic Material Flow Control Model for Two Different Industries	H.G.H. Tiemessen, M. Fleischmann, G.J. van Houtum, J.A.E.E. van Nunen, E. Pratsini
376	2012	Dynamic demand fulfillment in spare parts networks with multiple customer classes	K. Fikse, S.W.A. Haneyah, J.M.J. Schutten
375	2012	Improving the performance of sorter systems by scheduling inbound containers	Albert Douma, Martijn Mes
374	2012	Strategies for dynamic appointment making by container terminals	Pieter van Gorp, Marco Comuzzi
373	2012	MyPHRMachines: Lifelong Personal Health Records in the Cloud	E.M. Alvarez, M.C. van der Heijden, W.H.M. Zijm
372	2012	Service differentiation in spare parts supply through dedicated stocks	Frank Karsten, Rob Basten
371	2012	Spare parts inventory pooling: how to share the benefits	X.Lin, R.J.I. Basten, A.A. Kranenburg, G.J. van Houtum
370	2012	Condition based spare parts supply	Martijn Mes J. Arts, S.D. Flapper, K. Vernooij
369	2012	Using Simulation to Assess the Opportunities of Dynamic Waste Collection	J.T. van Essen, J.L. Hurink, W. Hartholt, B.J. van den Akker

368	2012	Aggregate overhaul and supply chain planning for rotables	
367	2011	Operating Room Rescheduling	Kristel M.R. Hoen, Tarkan Tan, Jan C. Fransoo, Geert-Jan van Houtum
366	2011	Switching Transport Modes to Meet Voluntary Carbon Emission Targets	Elisa Alvarez, Matthieu van der Heijden
365	2011	On two-echelon inventory systems with Poisson demand and lost sales	J.T. van Essen, E.W. Hans, J.L. Hurink, A. Oversberg
364	2011	Minimizing the Waiting Time for Emergency Surgery	Duygu Tas, Nico Dellaert, Tom van Woensel, Ton de Kok
363	2011	Vehicle Routing Problem with Stochastic Travel Times Including Soft Time Windows and Service Costs	Erhun Özkan, Geert-Jan van Houtum, Yasemin Serin
362	2011	A New Approximate Evaluation Method for Two-Echelon Inventory Systems with Emergency Shipments	Said Dabia, El-Ghazali Talbi, Tom Van Woensel, Ton de Kok
361	2011	Approximating Multi-Objective Time-Dependent Optimization Problems	Said Dabia, Stefan Röpke, Tom Van Woensel, Ton de Kok
360	2011	Branch and Cut and Price for the Time Dependent Vehicle Routing Problem with Time Window	A.G. Karaarslan, G.P. Kiesmüller, A.G. de Kok
359	2011	Analysis of an Assemble-to-Order System with Different Review Periods	Ahmad Al Hanbali, Matthieu van der Heijden
358	2011	Interval Availability Analysis of a Two-Echelon, Multi-Item System	Felipe Caro, Charles J. Corbett, Tarkan Tan, Rob Zuidwijk
357	2011	Carbon-Optimal and Carbon-Neutral Supply Chains	Sameh Haneyah, Henk Zijm, Marco Schutten, Peter Schuur
356	2011	Generic Planning and Control of Automated Material Handling Systems: Practical Requirements Versus Existing Theory	M. van der Heijden, B. Iskandar
355	2011	Last time buy decisions for products sold under warranty	Frank P. van den Heuvel, Peter W. de Langen, Karel H. van Donselaar, Jan C. Fransoo
			Frank P. van den Heuvel, Peter W. de

			Langen, Karel H. van Donselaar, Jan C. Fransoo
354	2011	Spatial concentration and location dynamics in logistics: the case of a Dutch provence	Pieter van Gorp, Remco Dijkman
353	2011	Identification of Employment Concentration Areas	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
352	2011	BOMN 2.0 Execution Semantics Formalized as Graph Rewrite Rules: extended version	E. Lüftenegger, S. Angelov, P. Grefen
351	2011	Resource pooling and cost allocation among independent service providers	Remco Dijkman, Irene Vanderfeesten, Hajo A. Reijers
350	2011	A Framework for Business Innovation Directions	K.M.R. Hoen, T. Tan, J.C. Fransoo G.J. van Houtum
349	2011	The Road to a Business Process Architecture: An Overview of Approaches and their Use	Murat Firat, Cor Hurkens
348	2011	Effect of carbon emission regulations on transport mode selection under stochastic demand	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
347	2011	An improved MIP-based combinatorial approach for a multi-skill workforce scheduling problem	R.J.I. Basten, M.C. van der Heijden, J.M.J. Schutten
346	2011	An approximate approach for the joint problem of level of repair analysis and spare parts stocking	Ton G. de Kok
345	2011	Joint optimization of level of repair analysis and spare parts stocks	Frank Karsten, Marco Slikker, Geert-Jan van Houtum
344	2011	Inventory control with manufacturing lead time flexibility	Murat Firat, C.A.J. Hurkens, Gerhard J. Woeginger
343	2011	Analysis of resource pooling games via a new extenstion of the Erlang loss function	Bilge Atasoy, Refik Güllü, TarkanTan
342	2011	Vehicle refueling with limited resources	Kurtulus Baris Öner, Alan Scheller-Wolf Geert-Jan van Houtum
341	2011	Optimal Inventory Policies with Non-stationary Supply Disruptions and Advance Supply Information	Joachim Arts, Gudrun Kiesmüller
		Redundancy Optimization for Critical Components in High-Availability Capital Goods	Murat Firat, Gerhard J. Woeginger
339	2010	Analysis of a two-echelon inventory system with two supply modes	

338 2010		Murat Firat, Cor Hurkens
335 2010	Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh	
		A.J.M.M. Weijters, J.T.S. Ribeiro
334 2010	Attaining stability in multi-skill workforce scheduling	P.T. Vanberkel, R.J. Boucherie, E.W. Hans, J.L. Hurink, W.A.M. van Lent, W.H. van Harten
333 2010	Flexible Heuristics Miner (FHM)	
		Peter T. Vanberkel, Richard J. Boucherie, Erwin W. Hans, Johann L. Hurink, Nelly Litvak
332 2010	An exact approach for relating recovering surgical patient workload to the master surgical schedule	M.M. Jansen, A.G. de Kok, I.J.B.F. Adan
331 2010	Efficiency evaluation for pooling resources in health care	Christian Howard, Ingrid Reijnen, Johan Marklund, Tarkan Tan
330 2010	The Effect of Workload Constraints in Mathematical Programming Models for Production Planning	H.G.H. Tiemessen, G.J. van Houtum
329 2010	Using pipeline information in a multi-echelon spare parts inventory system	F.P. van den Heuvel, P.W. de Langen, K.H. van Donselaar, J.C. Fransoo
328 2010	Reducing costs of repairable spare parts supply systems via dynamic scheduling	
		Murat Firat, Cor Hurkens
327 2010	Identification of Employment Concentration and Specialization Areas: Theory and Application	
		Murat Firat, Cor Hurkens, Alexandre Laugier
326 2010	A combinatorial approach to multi-skill workforce scheduling	M.A. Driessen, J.J. Arts, G.J. v. Houtum, W.D. Rustenburg, B. Huisman
325 2010	Stability in multi-skill workforce scheduling	
		R.J.I. Basten, G.J. van Houtum
324 2010	Maintenance spare parts planning and control: A framework for control and agenda for future research	
		M.C. van der Heijden, E.M. Alvarez, J.M.J. Schutten
323 2010	Near-optimal heuristics to set base stock levels in a two-echelon distribution network	
		E.M. Alvarez, M.C. van der Heijden, W.H. Zijm
	Inventory reduction in spare part networks by	

		selective throughput time reduction	
322	2010		
321	2010	The selective use of emergency shipments for service-contract differentiation	B. Walrave, K. v. Oorschot, A.G.L. Romme
			Nico Dellaert, Jully Jeunet.
320	2010	Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse	
			R. Seguel, R. Eshuis, P. Grefen.
319	2010	Preventing or escaping the suppression mechanism: intervention conditions	Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo.
318	2010	Hospital admission planning to optimize major resources utilization under uncertainty	Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak.
317	2010	Minimal Protocol Adaptors for Interacting Services	Pieter van Gorp, Rik Eshuis.
316	2010	Teaching Retail Operations in Business and Engineering Schools	Bob Walrave, Kim E. van Oorschot, A. Georges L. Romme
315	2010	Design for Availability: Creating Value for Manufacturers and Customers	S. Dabia, T. van Woensel, A.G. de Kok
314	2010	Transforming Process Models: executable rewrite rules versus a formalized Java program	
313	2010	Getting trapped in the suppression of exploration: A simulation model	
	2010	A Dynamic Programming Approach to Multi-Objective Time-Dependent Capacitated Single Vehicle Routing Problems with Time Windows	
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 Hendrikx
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	P.T. Vanberkel; R.J. Boucherie; E.W.

		health care	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