Heuristics for multi-item two-echelon spare parts inventory control problem with batch ordering in the central warehouse

Research School for Operations Management and Logistics

Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse

Engin Topan, Z. Pelin Bayindir, Tarkan Tan

Beta Working Paper series 321

BETA publicatie	WP 321 (working paper)
ISBN	978-90-386-2346-7
ISSN	
NUR	804
Eindhoven	September 2010

Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse

Engin Topan[a,b], Z.Pelin Bayındır[a] and Tarkan Tan[c]
[a]Department of Industrial Engineering, Middle East Technical University, Ankara, Turkey
[b]Department of Industrial Engineering, Çankaya University, Ankara, Turkey
[c]School of Industrial Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

Abstract

We consider a multi-item two-echelon inventory system in which the central warehouse operates under a (Q, R) policy, and each local warehouse implements $(S-1, S)$ policy. The objective is to find the policy parameters minimizing expected system-wide inventory holding and fixed ordering costs subject to an aggregate mean response time constraint at each facility. We propose a Lagrangian heuristic that employs a column generation method and a greedy algorithm. We also consider variants of this heuristic, which are based on the sequential determination of policy parameters (first the order quantities and then the reorder points and basestock levels), as done frequently in practice. As opposed to the heuristics for multi-echelon inventory optimization problems in the literature, our heuristics guarantee feasibility. We propose a lower bound for the optimal expected total cost and show that this bound is asymptotically tight in the number of parts. In an extensive computational study, we test the performances of the heuristics against this lower bound. The performance of the Lagrangian heuristic is found to be extremely well, which improves even further as the number of parts increases. Finally, we show numerically that sequential determination of policy parameters performs less satisfactorily while the computational benefits are limited.

Keywords: Inventory, Two-Echelon, Multi-Item, Batch Ordering, Spare Parts, Lagrangean heuristics.

1 Introduction.

In this paper we consider a spare parts inventory control problem that we observe in two different capital goods manufacturers providing equipments and services for capital intensive markets. The first one of them is a leading manufacturer of industrial printing systems, whereas the other one is a leading supplier of advanced tools to the nanotechnology market. Each of these manufacturers produces an equipment that has a critical function for their customers, e.g., printing machine, electron microscope etc. Therefore, for the customers of these manufacturers, equipment breakdowns are of essential importance since these may lead to discontinuing a critical process at the customer site, resulting in high down-time costs, often on the order of thousands of euros per hour. In such an environment, customers are protected against down-time risks by service level agreements (SLAs). Under these agreements, it is the manufacturer's responsibility to keep spare parts that will satisfy service requirements, which are usually expressed as a target level on a certain service measure, such as fill rate, probability of no-stockout and response time. This makes the provision of the spare parts one of the most critical after-sales services for the manufacturers.

SLAs that are defined between the manufacturers and the customers can typically be classified into two groups: Under the "item approach", a target service level is defined for each individual part. It is widely considered in the inventory literature (Thonemann et al. 2002). Another approach is the "system approach", in which a target service level is defined for the demand weighted average of the relevant performance measure over all parts. Hence, the system approach defines an aggregate service measure. Although the number of end products that a typical manufacturer produces is quite limited, the number of spare parts associated with the products can be very large, often on the order of thousands or ten thousands. Since customers are primarily interested in their equipment or entire system being up and running, setting a target service level for each part does not make sense for them. Instead, they are interested in the availability of parts at an aggregate level. Since the system approach is based on the demand weighted average of the relevant performance over all parts, it enables holding more inventories for cheap parts while fewer for expensive parts. This brings substantial savings in inventory holding costs in comparison with the situation under the item approach (Thonemann et al. 2002). Hence, the system approach is more applicable and widely adopted in the SLAs for spare parts (Hopp et al. 1999, Al-Rifai and Rossetti 2007, Çağlar et al. 2004, Wong et al. 2007b), which is also the case for the manufacturers that we consider.

Since the manufacturers should supply spare parts for different customers at different locations, they operate an inventory distribution system that consists of a number of local warehouses at different locations and a central warehouse replenishing them. This type of distribution systems are prevalent in spare parts logistics (Cohen et al. 1997). For such a two-echelon distribution system, the optimal policy is not known. Hence, the usual practice is to operate under an appropriate inventory control policy depending on the demand and the cost structures of the system. Our experience with the two manufacturers is that most of the spare parts that the manufacturers provide are rarely used, slow moving items, a majority of which have a demand rate of less than 5 parts per year for the whole system. In addition, since ordering between the lower and the upper echelons is internal and automated, fixed ordering costs are insignificant at the lower echelon facilities. Hence, the batch sizes are low, often equivalent to one. Under this setting, it is reasonable to operate under a basestock policy, i.e., $(S-1, S)$ policy, at the lower echelon facilities. This is common and often justified in other applications of spare parts inventory control problems (Wong et al. 2007b, Hopp et al. 1999). However, at the central warehouse, parts move faster due to the accumulation of internal demands from local warehouses. Moreover, the central warehouse is typically fed by external suppliers, resulting in high fixed procurement/transportation costs. Consequently, it is a common practice to place orders in batches instead of individual units at the upper echelon. Furthermore, there are situations where batching decisions are motivated by aggregate performance targets on the order frequencies at the central warehouse or production smoothing requirements of a third-party supplier (Hopp et al. 1999, Al-Rifai and Rossetti 2007). Under these conditions, it is more reasonable for the central warehouse to operate under a reorder point, order quantity policy, i.e., (Q, R) policy. The two manufacturers that we mention indeed apply a (Q, R) policy at the central warehouse and a base-stock policy at the local warehouses.

For such an inventory distribution system, finding the optimal policy parameters of the inventory control policy minimizing the expected system-wide cost becomes a critical decision affecting the performance of the system in the medium and the long term. However, the problem of finding the optimal policy parameters is difficult even for the system with all facilities operating under basestock policy (Çağlar et al. 2004, Wong et al. 2007b). The reasons are as follows:

- Even a medium scale inventory system involves thousands of stock keeping units, for each of which the policy parameters should be optimized. Furthermore, under a system approach, the policy parameters for each part interacts with the others through constraints on an aggregate performance measure. This makes the resulting optimization problem very complex (Özer and Xiong 2008).

- The evaluation of the objective function and the constraints of such an optimization problem requires evaluating the probability distributions of the inventory levels, which are difficult to compute even in a single-item case.

Many efforts are devoted to propose heuristic procedures to find the policy parameters of multi-item multi-echelon systems under pure basestock policy (Çağlar et al. 2004, Wong et al. 2007b, Caggiano et al. 2007) and batch ordering policy (Hopp et al. 1999, Al-Rifai and Rossetti 2007). Almost all of these heuristics are based on approximate evaluation of the probability distributions of the inventory levels, hence, they do not guarantee feasible solutions with respect to constraints on service levels. To the best of our knowledge, the only heuristic that is based on an exact evaluation method is proposed by Wong et al. 2007b, which is developed for systems under basestock policy. For many of these heuristics, finding the policy parameters of a practical size problem becomes an issue. Similarly, finding an efficient and tractable benchmark solution, e.g., a tight lower bound on the optimal expected total cost for such problems, is difficult as well. This makes it hard to evaluate the performance of the heuristics (Çağlar et al. 2004, Al-Rifai and Rossetti 2007). As a result, there is a need for a solution procedure for finding the policy parameters of multi-item multi-echelon batch ordering systems guaranteeing feasibility and at the same time yielding satisfactory results in terms of both the relative errors and the computation time for practical-size systems. There is also a need for a generic procedure generating benchmark solutions that can be used to test the performance of the heuristics.

Finding the optimal policy parameters of a typical system under batch ordering is much more involved compared to the one in which each facility operates under a pure basestock policy since the reorder points and the order quantities at the central warehouse need to be determined simultaneously with the basestock levels at each local warehouse for each part, where the policy parameters of the parts interact with each other. A common practice is to follow a sequential approach, which assumes the dominance of the batching decisions over the others, and hence necessitates determining the batch sizes first, in most applications independent of the service level requirements, and then the other policy parameters. The method brings a significant computational saving and also results in very low percentage cost penalty in single-item single-echelon systems which is verified both empirically and theoretically by several researches (Zheng 1992, Axsäter 1996, Silver et al. 1998, Gallego 1998). Due to its excellent performance in single-item single-echelon systems, it is widely used also in general system settings (Hopp et al. 1997, Axsäter 1998, Hopp et al. 1999, Axsäter 2003) as well as in practical applications, e.g., the manufacturers considered in our paper adopt the sequential approach to find the policy parameters of their inventory control systems. Among these papers, Hopp et al. (1997) is the only paper that investigate the performance of the sequential approach in comparison with a simultaneous approach. They report that the performance of the sequential approach based heuristic developed for a multi-item single-echelon system varies depending on the problem setting. Therefore, the consequences of adopting the sequential approach in multi-item multi-echelon systems has not been fully addressed in the literature. In our paper, one of our objectives is to investigate the performance of the sequential approach in a multi-item two-echelon inventory control system.

In this paper, we consider a multi-item two-echelon spare parts inventory system consisting of a central warehouse operating under a continuous-review installation-stock (Q, R) policy and a number of local warehouses operating under a continuous-review installation-stock $(S-1, S)$ policy, all of which can serve external customers. The stocks at the local warehouses are replenished from the central warehouse, implying that the central warehouse has both internal and external demands to satisfy. Two demand types are not differentiated; they are served according to the FCFS rule. The stocks at the central warehouse are replenished from an external supplier. We assume that the external supplier has ample stock, and unsatisfied demand is backordered at all facilities. In order to maintain the service responsiveness, the aggregate mean response time, which is the demand weighted average of response times over parts, is considered. The system incurs inventory holding costs at all facilities and fixed ordering costs only at the central warehouse. The inventory at each location is reviewed continuously. Our objective is to find inventory control policy parameters minimizing the sum of inventory holding and fixed ordering costs subject to constraints on the aggregate mean response time at each facility.

This paper contributes to the literature on multi-item multi-echelon inventory control systems in the following directions: We propose four alternative heuristics to find the optimal policy parameters of a large, practical-size multi-item two-echelon inventory control problem with batch ordering at the central warehouse based on the exact evaluation of the probability distributions of the inventory levels. Hence, in contrast to the existing literature, our heuristics guarantee feasible solutions. The first heuristic, which we call the Lagrangian heuristic, is based on the simultaneous approach and relies on the integration of a a column generation method and a greedy algorithm. The other three heuristics are based on the sequential approach, in which first the order quantities are determined using a batch size heuristic, then the reorder points at the central warehouse and the basestock levels at the local warehouses are determined through the same method used for the Lagrangian heuristic, i.e., a column generation and a greedy algorithm. The latter three heuristics differ in the batch size heuristic used. We also propose a lower bound for the optimal expected total cost, which we show to be asymptotically tight in the number of parts. Considering the difficulties encountered in evaluating the performance of heuristics for different multi-item two-echelon inventory systems in the literature (Çağlar et al. 2004, Al-Rifai and Rossetti 2007), the lower bound that we propose also makes a significant contribution to the relevant literature.

By making use of the results of the computational study we develop several insights, some of which are summarized as follows: The lower bound for the optimal expected total cost is found to be quite tight, especially when the number of parts is high, e.g., the relative gap between the bound and the optimal expected total cost is less than 1% even when the number of parts is only 50. These results together with the asymptotic tightness of the lower bound with the number of parts motivates us using it in further numerical experiments with large number of parts as a benchmark solution. Based on the results of these further experiments, the Lagrangian heuristic performs quite well in terms of the relative difference between the expected total cost of the solution obtained by the heuristic and the lower bound. As the number of parts increases, the performance of the heuristic improves further, making the heuristic very promising for practical applications. The computational requirement of the heuristic is also quite tolerable. To be more specific, the experiment with 10000 parts and 12 warehouses reveals that the relative cost difference is 0.04%; problems of this size can be solved within 12 hours on an Intel 3 GHz processor with 3.5 GB RAM. We also show that some of the qualitative conclusions regarding the performance of the sequential approach in the single-item single-echelon literature (Zheng 1992, Axsäter 1996, Silver et al. 1998, Gallego 1998) do not hold for the multi-item two-echelon setting, which is more representative of practical situations: First, we empirically observe that the relative cost difference may reach up to 31.03%. Considering that this cost difference is fairly high compared to findings in the aforementioned papers on single-item single-echelon systems. The errors in practical applications are expected to be even higher considering that our sequential heuristics involve a column generation method, which is more sophisticated than the methods used in sequential approach applications in practice. Second, the computation times required for sequential heuristics are comparable to that of the Lagrangian heuristic, showing that the computational advantages of the sequential determination of policy parameters are limited in multi-item systems. While the main focus of this paper is spare parts inventory systems, our results apply to any inventory system with a similar cost and service level structure.

The outline of this paper is as follows: Section 2 provides a review of the literature relevant to our paper. In Section 3, we specify the problem environment and then present our model. Section 4 introduces the heuristics proposed. Also, we describe how we find the optimal solution and the lower bound for the optimal expected total cost, which are used as benchmark solutions to test the performances of the heuristics in the experiments. In Section 5, we study the asymptotic behavior of the lower bound. We also present theoretical results associated with the asymptotic performance. In Section 6, we report and discuss our computational results. Finally in Section 7, we draw some conclusions.

2 Literature Review.

Although there has been substantial research on multi-echelon spare parts inventory control systems with different settings (see e.g., Cachon 2001, Axsäter 2006, Caggiano et al. 2006, Simchi-Levi and Zhao 2007), our review involves the papers on continuous-review installation-stock policies. There are three main directions of research in the area: optimal policy characterization, policy evaluation and policy optimization (Simchi-Levi and Zhao 2007). Since the characterization of the optimal policy for our system is out of our scope, we review papers on policy evaluation and policy optimization.

2.1 Policy Evaluation.

The earlier papers focus on the problem of policy evaluation, i.e., they propose evaluation methods to determine relevant performance measures. Most of these evaluation methods rely on approximations. The METRIC (Sherbrooke 1968) and the two-moment approximation (Graves 1985) are the most well-known approaches of this kind, which are based on approximating the distributions of the number of outstanding orders at lower echelon facilities in two-echelon inventory systems operating under basestock policy. Although these approximations are developed for single-item systems, they are used in multi-item systems as well (Hopp et al. 1999, Çağlar et al. 2004, Wong et al. 2007b, Al-Rifai and Rosetti 2007, Caggiano et al. 2007). The main drawback of these approximations is that using them in a policy optimization problem for a multi-item system under service level constraints may result in infeasible solutions. For alternative approximation methods and additional references, see e.g., Axsäter (2006), Özer and Xiong (2008).

There exist exact evaluation methods for multi-echelon inventory systems as well (Axsäter 1998, 2006). Similar to approximations, the exact methods are primarily developed for single-item systems, but they are applicable to multi-item systems (Wong et al. 2007b, Topan et al. 2010). As opposed to the approximations, the exact methods guarantee feasible solutions when they are employed to find the optimal or near-optimal policy parameters. From the perspective of evaluation method, our work belongs to the latter group of research since we use an exact method to evaluate the inventory distribution system considered in this paper.

2.2 Policy Optimization.

Considering the three main directions of research in the relevant literature, our work belongs to the third -more recent- body of literature, which focuses on policy optimization and proposes search algorithms for multi-echelon spare parts inventory systems. Although finding the optimal policy parameters for such systems are difficult, there exist exact search algorithms proposed in single-item (Axsäter 1998) and multi-item settings (Topan et al. 2010). Nevertheless, they are tractable only for smaller problems and often used for benchmark purposes. Therefore, heuristics are common in the literature. Our focus is on those that are proposed for multi-item systems. As in our paper, in those papers that propose heuristics for multi-item multi-echelon inventory systems, (1) the system approach is common, i.e., performance targets are defined based on a system-related measure, (2) the heuristics are based on decomposing the problem by facilities and/or parts, predominantly by means of a Lagrangian relaxation, then applying an iterative procedure to combine the resulting subproblems. Among those papers that are most relevant to our work are Hopp et al. (1999) and Al-Rifai and Rosetti (2007). Hopp et al. (1999) consider a system which differs from ours in that it involves a target level on the aggregate ordering frequency rather than explicit part-specific fixed ordering costs at the upper echelon facility. To find the policy parameters of this system, they decompose the resulting problem by echelons using a Lagrangian relaxation, and then they use the two-moment approximation by Graves (1985) and a sequential approach based heuristic proposed by Hopp et al. (1997) to solve each subproblem. The performance of the heuristic is tested against two alternative lower bounds in a computational study. In terms of the solution procedure, our work differs from this paper in three aspects: First, we follow an exact evaluation method. Second, our heuristics are based on both the simultaneous and the sequential approaches. Third, to obtain the Lagrangian multipliers, we follow an exact search procedure while Hopp et al. (1999) use an iterative heuristic search procedure.

Al-Rifai and Rosetti (2007) consider a two-echelon spare parts inventory system consisting of one warehouse and multiple identical retailers, all of which operate under (Q, R) policy. Their system setting differs from ours in three aspects: First, a target level is considered for aggregate ordering frequency rather than fixed ordering costs for each facility. Second, the total expected number of backorders is set as the performance measure rather than the aggregate mean response times. Third, they consider only the identical retailer case, whereas our model allows for nonidentical local warehouses. Similar to Hopp et al. (1999), Al-Rifai and Rossetti (2007) propose a heuristic to find the policy parameters by decomposing the problem by echelons and then applying an iterative heuristic procedure to generate Lagrangian multipliers. Their heuristic relies on the normal approximation of the lead time demand distribution at retailers. From the solution procedure perspective, our work differs from theirs mainly in that the evaluation procedure to determine the relevant performance measures and the search procedure to obtain the Lagrangian multipliers in our paper are both exact.

As for the solution methodology, related papers to our work are Çağlar et al. (2004), Wong et al. (2007b) and Caggiano et al. (2007), where basestock policy is assumed for all locations. Wong et al. (2007b) propose four different heuristics to find the optimal basestock levels for a two-echelon pure basestock system. They report that the greedy heuristic combined with the decomposition and column generation (DCG) yields quite satisfactory results in their setting, but the heuristic is tractable for problems up to a size of 100 parts and 20 local warehouses. In this paper, we apply a similar procedure and obtain quite satisfactory results for our batch ordering problem. While implementing the method, (1) we employ an algorithm to solve the subproblems arising as a result of the decomposition in the entire procedure based on using lower and upper bounds on the optimal policy parameters, and (2) we also consider variants of this method that are based on the sequential approach. Consequently, while our problem is more complicated than Wong et al. (2007b) -as we consider (Q, R) policy in the upper echelon-, our heuristics solve yet larger-scale problems with up to 10000 parts. The other heuristics proposed by Wong et al. (2007b) are tractable for larger problems, but they yield less satisfactory results compared to the DCG, e.g., for the problem instances with 100 parts, the relative gap between the DCG and a lower bound proposed is 0.71%, while for the greedy heuristic which is tractable for large-scale problems, the maximum relative gap is 9.65%.

We note that almost none of the algorithms in the literature (except Wong et al. 2007b) guarantee feasible solutions since these algorithms rely on approximate evaluation of the objective function and the constraints. Furthermore, for some of the heuristics, it is not clear whether they are tractable for large, practical size problems (Hopp et al. 1999, Çağlar et al. 2004). The ones that are not known to be tractable for large-scale problems either encounter difficulties in evaluating the performance of their heuristics against an analytical solution or a bound (Al-Rifai and Rossetti 2007), or they are developed for systems under pure basestock policy (Caggiano et al. 2007, Wong et al. 2007b). Hence, our paper contributes to the vast literature on multi-item multi-echelon inventory optimization problems by proposing efficient heuristics based on an exact evaluation -hence guaranteeing feasible solutions- and also a tight lower bound for large-scale practical-size multi-item two-echelon inventory problems with batch ordering at the central warehouse

2.3 Sequential Approach Heuristics.

Another direction of research related to our paper is the development of sequential approach based heuristics and investigation of their performances. Zheng (1992) analyzes the performance of the EOQ with planned backorders formula in a sequential approach to obtain the order quantity in a single-item (Q, R) model. He reports that the EOQ with planned backorders performs well, resulting in a percentage cost penalty of less than 12.50% theoretically, while the numerical findings are found to perform better: In 80% of the numerical problem instances the percentage cost penalty is less than 1.00%, and the maximum percentage cost penalty is found to be 2.90%. Following this line of research, many other researches report that the sequential approach perform well in single-item single-echelon systems (Axsäter 1996, Silver et al. 1998, Gallego 1998). The only paper examining the performance of sequential approach in a multi-item setting is Hopp et al. (1997). They propose alternative heuristics for the problem, one of which is based on the sequential approach. The experiments reveal that the relative gap between the solution obtained by the heuristic based on the sequential approach and a lower bound may be large depending on the problem setting. They report that their sequential heuristic performs better when the target service levels are high and the order frequencies are low. The sequential approach is also common in multi-item two-echelon inventory control literature (Axsäter 1998, Hopp et al. 1999). To our knowledge, our paper is the first to evaluate the performance of the sequential approach against the simultaneous approach in multi-item two-echelon batch ordering systems operating under the system approach.

3 The model.

The two-echelon inventory system that we consider consists of a single central warehouse and a set, N, of local warehouses, each provides service for a set, I, of parts. The demand of part $i \in I$ at warehouse $n \in N \cup \{0\}$, where 0 denotes the central warehouse, is assumed to be Poisson with rate λ_{in}. This constitutes the external demands for the system, while the central warehouse should also meet the internal demands from local warehouses. The two demand types are served by the central warehouse on a FCFS basis. Hence, there is no customer differentiation. For simplicity, we consider a single-indenture model, implying that each part is managed at a product level, but not at the component level. Note that this is validated in many situations (Kim et al. 2009). All warehouses keep inventory and thereby incur inventory holding costs, while only the central warehouse incurs fixed ordering cost. Both types of costs are assumed to be part-specific. Under this cost structure, it is reasonable to assume that each local warehouse $n \in N$ employs a basestock policy with parameter S_{in}, and the central warehouse employs a batch ordering policy with reorder point R_i and order quantity Q_i for each part $i \in I$. The demand that cannot be satisfied from stock is backordered. Warehouses have no capacity restrictions.

The system operates as follows: When an external demand for a part i arrives at warehouse $n \in N \cup \{0\}$, it is immediately satisfied from stock if the part is available; otherwise, the demand is backordered. If the external demand is served by a local warehouse, a request for a part is placed at the central warehouse to replenish stocks. If the part is available at the central warehouse, this internal demand is satisfied within a transportation lead time of T_{in}, otherwise, the internal request is backordered. If the inventory position of the central warehouse drops to reorder point R_i, an order of size Q_i is placed at the outside supplier. The supplier lead time for part i is T_{i0}. We assume that all lead times are constant. The inventory positions are restricted to be nonnegative, implying that $R_i \geq -1$ and $S_{in} \geq 0$ for each part $i \in I$ and each warehouse $n \in N$. We note that this restriction is not essential for our analysis.

The problem is to find the policy parameters that will minimize the sum of expected inventory holding and fixed ordering costs subject to a constraint on the aggregate mean response time at each warehouse. In order to formulate the problem, we use the notation in Table 1. Accordingly, let $\Lambda_n = \sum_{i \in I} \lambda_{in}$ denote the total demand rate for warehouse $n \in N \cup \{0\}$. Then, by using the Little's law, the aggregate mean response time at local warehouse $n \in N$, $W_n(\vec{Q}, \vec{R}, \vec{S})$, can be expressed as a function of expected number of backorders for part $i \in I$, $E[B_{in}(Q_i, R_i, S_{in})]$.

$$W_n(\vec{Q}, \vec{R}, \vec{S}) = \sum_{i \in I} \frac{\lambda_{in}}{\Lambda_n} E[W_{in}(Q_i, R_i, S_{in})] = \sum_{i \in I} \frac{E[B_{in}(Q_i, R_i, S_{in})]}{\Lambda_n}.$$

Similarly, for the central warehouse, we have $W_0(\vec{Q}, \vec{R}) = \sum_{i \in I} \frac{E[B_{i0}(Q_i, R_i)]}{\Lambda_0}$. We can now formulate the problem as follows:

Problem P:

$$Min\ Z = \sum_{i \in I} \left[c_i h \left(E[I_{i0}(Q_i, R_i)] + \sum_{n \in N} E[I_{in}(Q_i, R_i, S_{in})] \right) + \frac{\lambda_{i0} K_i}{Q_i} \right] \quad (1)$$

$s.t.$

$$\sum_{i \in I} \frac{E[B_{i0}(Q_i, R_i)]}{\Lambda_0} \leq W_0^{\max}, \quad (2)$$

$$\sum_{i \in I} \frac{E[B_{in}(Q_i, R_i, S_{in})]}{\Lambda_n} \leq W_n^{\max}, \qquad \text{for } \forall\, n \in N, \quad (3)$$

$$Q_i \geq 1,\ R_i \geq -1,\ S_{in} \geq 0, \text{ and } Q_i,\ R_i,\ S_{in} \in Z, \qquad \text{for } \forall\, i \in I,\ \forall\, n \in N.$$

The objective function (1) minimizes the expected system-wide inventory holding and fixed ordering costs. Constraints (2) and (3) ensure that the aggregate mean response times at the central warehouse and local warehouses do not exceed target aggregate mean response times, $W_0^{\max}$ and $W_n^{\max}$, respectively. Alternatively, one could also model

Table 1: General Notation

i	Part index, $i \in I$
n	Warehouse index $n \in N \cup \{0\}$
c_i	Unit variable cost of part i
h	Inventory carrying charge
K_i	Fixed ordering cost of part i at the central warehouse
λ_{in}	Demand rate for part i at local warehouse $n \in N$
λ_{i0}^e	External demand rate for part i at the central warehouse
λ_{i0}	Demand rate (sum of internal and external) for part i at the central warehouse
Λ_n^e	Total external demand rate at the central warehouse
Λ_n	Total demand rate for warehouse $n \in N \cup \{0\}$
T_{i0}	Lead time for part i at the central warehouse from the outside supplier
T_{in}	Transportation lead time from the central warehouse to local warehouse $n \in N$ for part i
$W_n^{\max}$	Target aggregate mean response time at warehouse $n \in N \cup \{0\}$
R_i	Reorder point for part i at the central warehouse
Q_i	Order quantity for part i at the central warehouse
S_{in}	Basestock level for part i at local warehouse $n \in N$
$\vec{S}_i$	$[S_{i1}, S_{i2}, \ldots, S_{i\|N\|}]$ = Vector of basestock levels for part i
$\vec{S}$	$[\vec{S}_1, \vec{S}_2, \ldots, \vec{S}_{\|I\|}]$ = Vector of basestock levels
$\vec{Q}$	$[Q_1, Q_2, \ldots, Q_{\|I\|}]$ = Vector of order quantities
$\vec{R}$	$[R_1, R_2, \ldots, R_{\|I\|}]$ = Vector of reorder points
$I_{in}(Q_i, R_i, S_{in})$	On-hand inventory level for part i at warehouse $n \in N$ in the steady state
$I_{i0}(Q_i, R_i)$	On-hand inventory level for part i at the central warehouse in the steady state
$B_{in}(Q_i, R_i, S_{in})$	Backorder level for part i at warehouse $n \in N$ in the steady state
$B_{i0}(Q_i, R_i)$	Backorder level for part i at the central warehouse in the steady state
$W_{in}(Q_i, R_i, S_{in})$	Response time for part i at warehouse $n \in N$ in the steady state
$W_{i0}(Q_i, R_i)$	Response time for part i at the central warehouse in the steady state
$W_{i0}^e(Q_i, R_i)$	Response time for part i at the central warehouse (for external customers)
$W_n(\vec{Q}, \vec{R}, \vec{S})$	Aggregate mean response time at warehouse $n \in N$ in the steady state
$W_0(\vec{Q}, \vec{R})$	Aggregate mean response time at the central warehouse in the steady state
$W_0^e(\vec{Q}, \vec{R})$	Aggregate mean response time at the central warehouse (for external customers)

the situation in which only the external customers are incorporated in evaluating the performance of the central warehouse. In that case, the aggregate mean response time at the central warehouse is stated as follows: $W_0^e(\vec{Q}, \vec{R}) = \sum_{i \in I} \frac{\lambda_{i0}^e}{\Lambda_0^e} E[W_{i0}^e(Q_i, R_i)]$, where $\lambda_{i0}^e = \lambda_{i0} - \sum_{n \in N} \lambda_{in}$ is the external demand for part $i \in I$ and $\Lambda_0^e = \Lambda_0 - \sum_{n \in N} \Lambda_n$ is the total external demand, at the central warehouse. Since there is no differentiation between the external and the internal customers we simply have $W_0^e(\vec{Q}, \vec{R}) = W_0(\vec{Q}, \vec{R})$. Then, we obtain

$$W_0^e(\vec{Q}, \vec{R}) = \sum_{i \in I} \frac{\lambda_{i0}^e}{\Lambda_0^e} E[W_{i0}(Q_i, R_i)] = \sum_{i \in I} \frac{\lambda_{i0}^e}{\Lambda_0^e} \frac{E[B_{i0}(Q_i, R_i)]}{\lambda_{i0}},$$

which replaces constraint (2). In a way, this alternative model corresponds to weighing the individual aggregate mean response time values only with the rate of external customers.

The expected inventory levels and backorder levels in the central and local warehouses in problem P are established by following the exact evaluation method by Topan et al.

(2010), which is based on the disaggregation of the backorders at the central warehouse (Axsäter 2006).

4 Solution Procedures.

This section introduces the heuristics proposed in this paper. We also describe how we develop the optimal solution and the lower bound for our problem, which are used as benchmark solutions in the experiments. First, the exact solution procedure, through which we obtain the optimal solution, is explained in Section 4.1. Then, the column generation algorithm, through which we obtain the Lagrangian dual bound, is described in Section 4.2. Finally, the Lagrangian heuristic and the sequential heuristics are introduced in Sections 4.3 and 4.4, respectively.

4.1 Exact Solution Procedure: Branch-and-Price Algorithm.

In order to search for an exact solution, we employ a branch-and-price algorithm introduced by Topan et al. (2010). It is a variant of branch-and-bound algorithm in which a column generation method is applied to obtain a lower bound for each subproblem (node) of the branch-and-bound tree (Barnhart et al. 1998). The algorithm also employs a greedy algorithm to obtain a global upper bound to tighten the bounding scheme further. Depending on the lower and the upper bounds, a node is either fathomed or explored further by branching. The procedure is repeated until all nodes are fathomed. The exact solution procedure is tractable only for problems with limited number of parts and warehouses, e.g., in our experiments we have been able to solve problems of size up to 40 parts and 4 local warehouses. Hence, in order to be able to solve larger problems, we resort to heuristic methods.

4.2 Obtaining the Lagrangean dual bound for the problem

To obtain the Lagrangian dual bound for our problem, we use a column generation method. The method relies on an alternative formulation of the original problem P, which is known as the master problem (Lübbecke and Desrosiers 2002). The master problem simply corresponds to listing all set of feasible policies for each part $i \in I$ and then selecting exactly one of them. Since the column generation procedure works with the principle of generating only the policies (or as the name suggests columns) that improve the overall solution, it is not necessary to generate all set of columns, instead, one can continue with a restricted set. The method is widely used for solving various integer programming problems (Lübbecke and Desrosiers 2002, Guignard 2003). Before giving the details of the algorithm, we first introduce our notation. Let L denote the set of columns, i.e., control policy parameters $(Q_i, R_i, \vec{S}_i)$, for each part i, and x_{il} indicate whether column $l \in L$ is selected for part i or not. Let $C_{il} = c_i h E[I_{i0}(Q_i^l, R_i^l)] + c_i h \sum_{n \in N} E[I_{in}(Q_i^l, R_i^l, S_{in}^l)] + \frac{\lambda_{i0} K_i}{Q_i^l}$ be the expected total inventory holding and fixed ordering costs associated with column $l \in L$ for part i. Similarly let $A_{il0} = \frac{E[B_{i0}(Q_i^l, R_i^l)]}{\Lambda_0}$ and $A_{iln} = \frac{E[B_{in}(Q_i^l, R_i^l, S_{in}^l)]}{\Lambda_n}$ be the relevant terms for constraints (2) and (3) associated with column $l \in L$ for part i for each warehouse $n \in N$, respectively. Then the master problem (MP) is formulated as follows:

Problem MP:

$$
\begin{aligned}
& Min\ Z = \sum_{i \in I} \sum_{l \in L} C_{il} x_{il} \\
& s.t. \\
& \sum_{i \in I} \sum_{l \in L} A_{iln} x_{il} \le W_n^{\max}, \quad \text{for } \forall\, n \in N \cup \{0\}, \quad (\alpha_n) \quad (4) \\
& \sum_{l \in L} x_{il} = 1, \quad \text{for } \forall\, i \in I, \quad (\beta_i) \quad (5) \\
& x_{il} = 0/1, \quad \text{for } \forall\, i \in I,\ \forall\, l \in L.
\end{aligned}
$$

The solution of the LP-relaxation of problem MP $(LPMP)$ provides a lower bound on the optimal objective function value of MP and hence on that of P. This bound corresponds to the Lagrangian dual bound obtained through the Lagrangian relaxation of the constraints of problem P (Guignard 2003). In order to solve problem $LPMP$, we follow a column generation method by generating only the columns that improve the objective function value of $LPMP$. This step requires solving an integer programming problem known as the column generation (CG) or pricing problem. Letting $C_i(Q_i, R_i, \vec{S}_i) = c_i h E[I_{i0}(Q_i, R_i)] + c_i h \sum_{n \in N} E[I_{in}(Q_i, R_i, S_{in})] + \frac{\lambda_{i0} K_i}{Q_i}$, $A_{i0} = \frac{E[B_{i0}(Q_i, R_i)]}{\Lambda_0}$ and $A_{in} = \frac{E[B_{in}(Q_i, R_i, S_{in})]}{\Lambda_n}$ for $i \in I$ and $n \in N$, our pricing problem is stated as follows:

Problem CG:

$$
\begin{aligned}
& Min \sum_{i \in I} \left(C_i(Q_i, R_i, \vec{S}_i) - \sum_{n \in N \cup \{0\}} \alpha_n A_{in} - \beta_i \right) \\
& s.t. \\
& Q_i \ge 1,\ R_i \ge -1,\ \vec{S}_i \ge \vec{0}, \ \text{ and } \ Q_i,\ R_i,\ S_{in} \in Z, \quad \text{for } \forall\, i \in I,\ \forall\, n \in N,
\end{aligned}
$$

where $\alpha_n \le 0$ for $n \in N \cup \{0\}$ and β_i unrestricted in sign for $i \in I$ are the dual variables (Lagrangian multipliers) that are obtained from the solution of the problem MP. In an iterative procedure, CG provides the columns that are required for the solution of $LPMP$, whereas $LPMP$ provides the dual variables required for the solution of CG. In order to solve the problem CG, we decompose it into $|I|$ subproblems since the problem is decomposable by parts. Let $\theta_n = \frac{-\alpha_n}{\Lambda_n}$ for each $n \in N \cup \{0\}$ and $\vec{\theta} = [\theta_1, \theta_2, \ldots, \theta_{|N|}]$. Then the subproblem for part $i \in I$ for a given value of $\vec{\theta}$ is given as follows:

$SP_i(\vec{\theta})$:

$$\begin{aligned} Min \;\; G\left(Q_i, R_i, \vec{S}_i\right) &= c_i h \left(E[I_{i0}(Q_i, R_i)] + \sum_{n \in N} E[I_{in}(Q_i, R_i, S_{in})] \right) + \frac{\lambda_{i0} K_i}{Q_i} \\ &+ \theta_0 E[B_{i0}(Q_i, R_i)] + \sum_{n \in N} \theta_n E[B_{in}(Q_i, R_i, S_{in})] \\ & s.t. \;\; Q_i \geq 1, \, R_i \geq -1, \, \vec{S}_i \geq \vec{0}, \;\text{ and }\; Q_i, \, R_i, \, S_{in} \in Z, \qquad \text{for } \forall \, n \in N. \end{aligned}$$

Each time we solve $SP_i(\vec{\theta})$, we generate a column for part $i \in I$, i.e., $(Q_i^l, R_i^l, \vec{S}_i^l)$, that is required for solving *MP*. The procedure is repeated until none of the subproblems $SP_i(\vec{\theta})$ yields a negative optimal objective function value, confirming the optimality of the solution of *MP* (Lübbecke and Desrosiers 2002). The column generation method is known to converge to a solution provided that a nondegenerate basic feasible solution exists for the master problem (Dantzig, 1963). One can easily obtain a nondegenerate basic feasible solution for our problem by following Dantzig (1963). Therefore, our column generation algorithm converges as well.

Note that subproblem $SP_i(\vec{\theta})$ corresponds to the two-echelon version of the single-echelon (Q, R) model investigated deeply in the literature (Zheng 1992, Axsäter 1996, Silver et al. 1998, Gallego 1998), if θ_n is interpreted as the unit backorder cost per unit time at warehouse $n \in N \cup \{0\}$. To solve each of these subproblems we use an algorithm proposed for solving single-item two-echelon batching problems by Topan et al. (2010). The algorithm involves two nested loops: the outer loop searches for the optimal Q_i, and the inner loop searches for the optimal R_i for a given value of Q_i. Within these nested loops, an innermost subroutine optimizes S_{in} for given values of Q_i and R_i. In order to reduce the search space, we use upper bounds Q_i^{UB} and R_i^{UB}, lower bounds Q_i^{LB}, R_i^{LB} are proposed for the optimal values of Q_i and R_i by Topan et al. (2010). For a given value of Q_i, R_i^{UB} (R_i^{LB}) is obtained by optimizing R_i for $S_{in} = 0$ ($S_{in} \to \infty$) for all $n \in N$. Similarly, we obtain Q_i^{LB} by optimizing Q_i for $R_i \to \infty$ and $S_{in} \to \infty$ for all $n \in N$. This gives $Q_i^{LB} = \min\left\{Q_i : (Q_i + 1)\, Q_i \geq \frac{2K_i\lambda_{i0}}{c_i h}\right\}$, which corresponds to the discrete version of the EOQ formula. Finally, we use $Q_i^{UB} = \sqrt{\frac{2K_i\lambda_{i0} + (c_i h + p_i)\lambda_{i0} T_{i0}}{H_i}}$ as an upper bound, where $H_i = \frac{c_i h p_i}{c_i h + p_i}$ and $p_i = \theta_0 + \sum_{n \in N} \theta_n \frac{\lambda_{in}}{\lambda_{i0}}$.

4.3 Lagrangean (Greedy) Heuristic.

In this paper, we use the greedy algorithm to (1) solve the problem P heuristically by combining it with the column generation method introduced, which we call the overall procedure as the Lagrangian heuristic, (2) obtain alternative heuristics for P by integrating it with the sequential heuristics, (3) obtain an upper bound in the branch-and-price algorithm as it is mentioned in Section 4.1.

The greedy algorithm is a simple search algorithm that can be used to generate a feasible solution from an integer but infeasible (dual) solution. The method is known to perform quite well in multi-item two-echelon inventory control problems (Cohen et al. 1990, Wong et al. 2005, 2006, 2007a, 2007b). The main idea of the greedy algorithm is as follows: Starting with an infeasible solution, at each iteration, the algorithm iterates to a solution that is as close to the feasible region as possible while incurring as low additional cost as possible. This procedure is repeated until a feasible solution is obtained. Since the initial dual solution may yield fractional variables, this may require rounding fractional variables down to make sure that the new solution satisfies constraints (2) and (3) before iterating the greedy algorithm.

The greedy algorithm and a greedy move can formally be described as follows: Let $\vec{Q}$, $\vec{R}$ and $\vec{S}$ be vectors of order quantities, reorder points and basestock levels, respectively, and let $\omega(\vec{Q},\vec{R},\vec{S})$ denote the maximum constraint violation for given values of $\vec{Q}$, $\vec{R}$ and $\vec{S}$, i.e.,

$$\omega(\vec{Q},\vec{R},\vec{S}) = \max_{n \in N \cup \{0\}} \left\{ \left(W_n(\vec{Q},\vec{R},\vec{S}) - W_n^{\max} \right)^+ \right\},$$

and $Z(\vec{Q},\vec{R},\vec{S})$ be the corresponding objective function value of $(\vec{Q},\vec{R},\vec{S})$. Then, the neighborhood of $(\vec{Q},\vec{R},\vec{S})$, $V(\vec{Q},\vec{R},\vec{S})$, is defined as the set of all vectors $[\vec{Q},\vec{R},\vec{S}] + \varepsilon$, where ε is a vector in which exactly one of the entries is one and the rest are zero. Then, the greedy algorithm searches for the solution $(\vec{Q}',\vec{R}',\vec{S}') \in V(\vec{Q},\vec{R},\vec{S})$ that yields the maximum $r(\vec{Q}',\vec{R}',\vec{S}') = \frac{\omega(\vec{Q}',\vec{R}',\vec{S}') - \omega(\vec{Q},\vec{R},\vec{S})}{Z(\vec{Q}',\vec{R}',\vec{S}') - Z(\vec{Q},\vec{R},\vec{S})}$ ratio. The greedy algorithm converges finitely by nature.

"Lagrangian heuristic" is a generic name given to heuristics that first employ a Lagrangian relaxation to find a good -but often infeasible- relaxed solution, and then an algorithm to transform this relaxed solution into a feasible solution (Guignard 2003). In our paper, the Lagrangian heuristic simply corresponds to the entire procedure in which the column generation (to obtain the Lagrangian dual solution) and the greedy algorithm (to obtain a feasible solution starting from the Lagrangian dual solution) are integrated. Note that since the Lagrangian heuristic is based on determining the order quantities and the reorder points simultaneously, it is a simultaneous approach heuristic. An overview of the Lagrangian heuristic is given in Figure 1.

Figure 1: The Flowchart of the Lagrangian Heuristic.

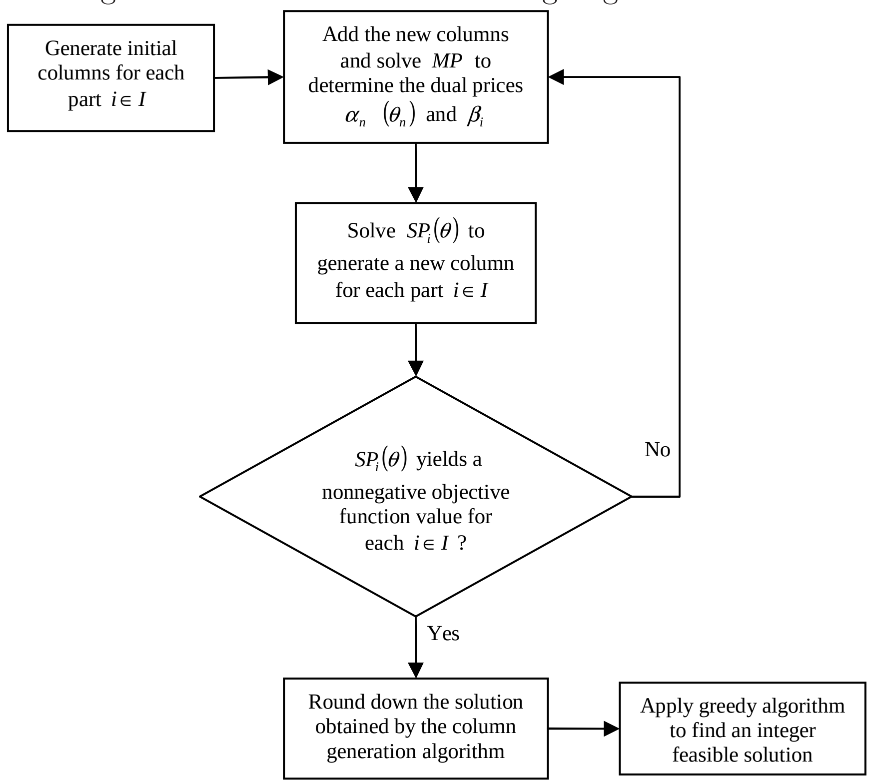

4.4 Sequential Approach Based Heuristics.

Similar to the Lagrangian heuristic, the sequential heuristics rely on the integration of the column generation and the greedy algorithm. However, in contrast to the Lagrangian heuristic, the order quantities at the central warehouse are determined offline. The sequential heuristics iterate as follows: First, the order quantities are determined through a batch size heuristic. Then, given the order quantities, the remaining policy parameters, i.e., the reorder points at the central warehouse and the basestock levels at the local warehouses, are determined by using the entire procedure developed for the Lagrangian heuristic in Section 4.2. This results in changes in the overall procedure: Q_i is discarded from problem $SP_i(\vec{\theta})$ for each $i \in I$, hence the outer loop of the algorithm proposed to solve $SP_i(\vec{\theta})$ is eliminated. This also brings a computational advantage to the sequential heuristics over the Lagrangian heuristic. An overview of the sequential heuristics is given in Figure 2.

To implement the sequential approach, we consider three alternatives for setting the order quantities:

- the EOQ formula, i.e., $Q_i = \sqrt{\frac{2\lambda_{i0}K_i}{c_i h}}$,
- the EOQ with planned backorders (EOQ^B) formula, i.e., $Q_i = \sqrt{\frac{2\lambda_{i0}K_i(c_i h+p_i)}{(c_i h)p_i}}$ (Zheng 1992, Gallego 1998), where p_i is the shortage cost defined per unit short of part $i \in I$ per unit time and obtained as in Section 4.2,
- an alternative batch size heuristic Q^{LU} based on the lower and upper bounds, Q_i^{LB} and Q_i^{UB}, proposed by Topan et al. (2010) for the single-item two-echelon batch ordering problem $SP_i(\vec{\theta})$ in Section 4.2. The heuristic is similar to the batch size heuristic proposed by Gallego (1998) for the single-echelon (Q, R) model. However, when Gallego's batch size heuristic is directly used in our model, i.e., $Q_i = min\left(\sqrt{2}Q_i^{LB}, \sqrt{Q_i^{LB} \cdot Q_i^{UB}}\right)$, the optimal order quantities are overestimated. Hence, we adopt it in our model by using the harmonic mean of Q_i^{LB} and Q_i^{UB} instead of using a geometric mean, which is less than or equal to the latter. In this way, we achieve better results. Accordingly, the order quantities are found from $Q_i = min\left(\sqrt{2}Q_i^{LB}, \frac{2Q_i^{LB}Q_i^{UB}}{Q_i^{LB}+Q_i^{UB}}\right)$.

In this manner, we obtain three alternative sequential heuristics, S_1, S_2 and S_3.

The batch size heuristics differ depending on how the service level requirements are taken into account in determining the order quantities. In S_1, the order quantities are determined independent of the service level requirements. This is the case in many practical applications, e.g., the manufacturers considered in our paper determine the order quantities using the EOQ. However, S_2 and S_3 incorporate the service level requirements by means of a part-specific shortage cost, p_i, for each part $i \in I$. In order to obtain each part-specific shortage cost, p_i, first, we apply the entire procedure in Figure 2 by using the EOQ, and then under the solution obtained, we compute the probability of no stockouts, γ_i, for each part $i \in I$. Then, by substituting γ_i in the newsboy ratio $\gamma_i = \frac{p_i}{c_i h+p_i}$, we determine p_i. Finally, the entire procedure iterates once more to obtain the solution of the corresponding sequential heuristic. Therefore, while the corresponding procedure iterates once in S_1, it iterates twice in S_2 and S_3; first to find the part part-specific shortage costs, second to obtain the overall solution. Since the greedy algorithm converges finitely and the column generation algorithm is guaranteed to converge to a solution, all our heuristics guarantee convergence. The heuristics proposed in this paper are summarized in Table 2.

Figure 2: The Flowchart of the Sequential Heuristics.

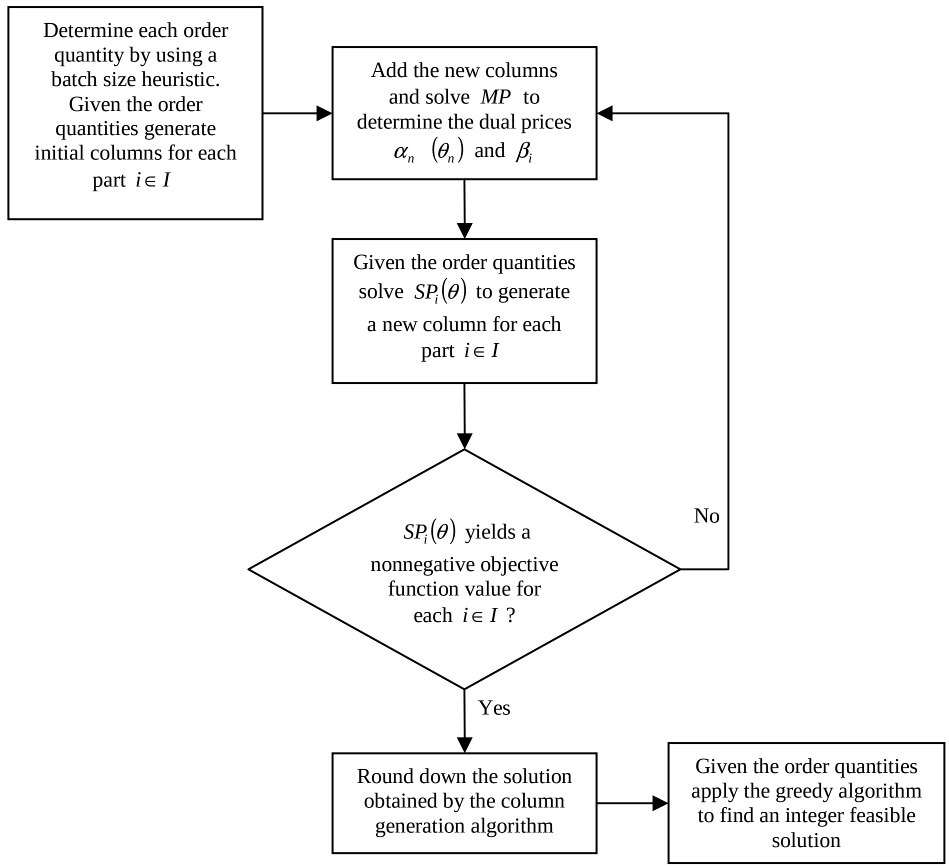

Table 2: Heuristics proposed.

Solution Approach	Heuristics
Simultaneous Approach • All the control parameters are determined simultaneously.	• The order quantities and reorder points at the central warehouse, basestock levels at local warehouses are obtained by using the column generation and the greedy algorithms.
Sequential Approach • Order quantities are predetermined by using a batch size heuristic. o EOQ o EOQ^B o Q^{LU}	• Given order quantities, reorder points at the central warehouse, and basestock levels at local warehouses are obtained by using the column generation and the greedy algorithms. o S_1 (uses EOQ) o S_2 (uses EOQ^B) o S_3 (uses Q^{LU})

5 Asymptotic analysis of the Lagrangean dual bound.

In this section we study the asymptotic behaviour of the Lagrangian dual bound for our problem and show that the Lagrangian dual bound is asymptotically tight in the number of parts. The analysis relies on the probabilistic analysis of combinatorial problems (Kellerer et al. 2004). Accordingly, we assume that for each $i \in I$ and $n \in N \cup \{0\}$, c_i, K_i, λ_{in} and $W_n^{\max}$ are independent and identically distributed random variables drawn from a uniform distribution $U[\underline{c}, \overline{c}]$, $U[\underline{K}, \overline{K}]$, $U[\underline{\lambda}, \overline{\lambda}]$ and $U[\underline{W}, \overline{W}]$, respectively. We further assume that $\underline{c}, \underline{W} > 0$ and $\overline{K}, \overline{c}, \overline{\lambda} < \infty$, implying that

- the fixed ordering cost is strictly finite for each part $i \in I$, i.e., $K_i < \infty$,
- the unit holding cost is strictly positive and finite for each part $i \in I$, i.e., $0 < c_i h < \infty$,
- the target aggregate mean response time at each warehouse $n \in N \cup \{0\}$ is strictly positive, i.e., $W_n^{\max} > 0$,
- the average lead time demand for each part $i \in I$ at each warehouse $n \in N \cup \{0\}$ is finite, i.e., $\lambda_{in} T_{in} < \infty$.

Note that these four assumptions are practically nonrestrictive, but necessary for our model to be stable and the problems to have finite solutions.

Through Theorem 1, we first show that the optimal objective function value of MP, z^{MP}, increases at least linearly with the number of parts. Then, in Theorem 2, we show that the gap between the optimal objective function value of MP, z^{MP}, and its LP-relaxation, z^{LPMP}, grows only with an order of the number of local warehouses, meaning that this gap is independent of the number of parts. Finally, in Theorem 3, we combine the results of Theorem 1 and 2 and show that for a given number of local warehouses, as the number of parts increases the relative gap between z^{MP} and z^{LPMP} with respect to z^{MP} approaches to zero since the absolute gap between z^{MP} and z^{LPMP} grows faster than z^{MP}. Hence, this shows that the Lagrangian dual bound for problem P is asymptotically tight in the number of parts. Under the assumptions given above, we show that the following propositions hold for every realization of random parameters c_i, K_i, λ_{in} and $W_n^{\max}$ for each $i \in I$ and each $n \in N$.

The following lemma is necessary for the proof of Theorem 1. It shows that for any part $i \in I$, the cost associated with each column generated through the column generation algorithm is bounded below by the optimal objective function value of the EOQ model with unit backorder cost of θ_0, i.e., $z_i^{EOQ(\theta_0)} = \sqrt{\frac{2K_i \, \lambda_{i0} \, c_i h \, \theta_0}{c_i h + \theta_0}}$ (Gallego 1998).

Lemma 1. *For a given value of θ_0, $C_{il} \geq z_i^{EOQ(\theta_0)}$ for each $i \in I$ and $l \in L$.*

Proof. Proof is provided in the appendix. □

Theorem 1. *The optimal objective function value of MP, z^{MP}, is in $\mathbf{\Omega}(|I|)$, i.e., z^{MP} is asymptotically bounded below by a function in the order of $|I|$ with probability 1.*

Proof. Proof is provided in the appendix. □

The following two lemmas are used in the proof of Theorem 2.

Lemma 2.

***a**) The column generation method yields finite solutions (columns).*

***b**) The total cost associated with each column generated by the column generation method is finite and increases only with the order of $|N|$.*

Proof. Proofs of part (a) and (b) are provided in the appendix. □

Lemma 3. *The optimal solution of $LPMP$ contains at most $|N|$ non-integer variables.*

Proof. PROOF. Proof is provided in the appendix. □

Theorem 2. *The gap between the optimal objective function value of MP, z^{MP}, and its LP-relaxation, z^{LPMP}, is in $\boldsymbol{O}(|N|^2)$, meaning that the gap is asymptotically bounded above by a function of $|N|^2$.*

Proof. Proof is provided in the appendix. □

Theorem 3. *For a given number of local warehouses $|N|$, the Lagrangian dual bound for problem P is asymptotically tight in the number of parts $|I|$ with probability 1.*

Proof. Proof is provided in the appendix. □

Theorem 3 shows that the Lagrangian dual bound can be used as a benchmark solution for problem P with large number of parts. Considering that the size of the problems in practice grows especially with the number of parts (compared to the number of warehouses), this also shows that the corresponding bound can be used as a benchmark solution for practical problems.

6 Computational Study.

In this section, we conduct an extensive computational study to further explore the performances of the heuristics and the Lagrangian dual bound developed in our paper. First, the performance of the Lagrangian dual bound is tested against the optimal solution for small-size problems to see how reasonable it is to employ the Lagrangian dual bound as a benchmark solution. Then, the performances of the heuristics, i.e., S_1, S_2, S_3 and the Lagrangian heuristic are evaluated relative to the Lagrangian dual bound for larger problems, where this bound yields better results. In our analysis, the expected total cost corresponding to each solution is considered as the performance criterion. The performance of the Lagrangian dual bound is mainly evaluated in terms of the percentage dual gap with the optimal expected total cost, $PGAP$. However, we also consider the absolute dual gap, GAP. Similarly, the performances of the heuristics are mainly evaluated in terms of the percentage cost difference between the solution obtained by the heuristic and the Lagrangian dual bound, PCD, but we also consider the absolute cost difference between the solution and the bound, ACD. Let z^* be the optimal objective function value, z_{LD} be the objective function value of the Lagrangian dual solution, and z be the objective function value of any solution to be tested, then the GAP and the $PGAP$ are computed as $GAP = |z_{LD} - z^*|$ and $PGAP = \frac{|z_{LD}-z^*|}{z^*}$, whereas the PCD and the ACD are calculated as $ACD = |z - z_{LD}|$ and $PCD = \frac{|z-z_{LD}|}{z_{LD}}$.

6.1 Experimental Design.

We consider the following six system parameters as the experimental factors: (i) number of parts, $|I|$, (ii) number of local warehouses, $|N|$, (iii) demand rates, λ_{in}, (iv) unit variable costs, c_i, (v) fixed ordering costs, K_i, and (vi) target aggregate mean response times at the warehouses, $W_n^{\max}$. Since lead time, T_{in}, contributes to the model in the form of lead time demand, $\lambda_{in}T_{in}$, we do not consider it as a distinct factor. This also means that we do not distinguish the effect of demand rate from that of lead time demand. Using these factors, we conduct a full factorial experiment to investigate the overall performance of the heuristics and the Lagrangian dual bound and perform an analysis of variance (ANOVA) to investigate (i) the individual effect of each factor on the performance of the heuristics and the Lagrangian dual bound and (ii) the interactions between factors.

To generate the problem instances, we first generate a base case setting. Then, based on this base case setting, we build the testbeds for the experiments. For the base case setting, the following parameters are set identical; lead time at the central warehouse, T_{i0}, across all parts, the target aggregate mean response times at the warehouses, W_n^{max}, across all warehouses, the lead times at the local warehouses, T_{in}, across all parts and local warehouses. We assume that the unit variable costs, c_i, and the fixed ordering costs, K_i, are nonidentical across all parts, the demand rates, λ_{in}, are nonidentical across all parts and warehouses. The fixed ordering cost of each part is generated from a uniform distribution. To represent skewnesses of the demand rates and the unit variable costs across the population of parts, we follow an approach similar to the one described in Thonemann et al. (2002). Following this approach, the demand rates are generated through a two-step procedure: First, a part-specific average demand rate is generated randomly for each part, then by multiplying it with a second random number representing the demand intensity at each warehouse part-specific and location-dependent rates are obtained, whereas the unit variable costs are generated in one step since they are only part-specific. To obtain the part-specific average demand rate for any part $i \in I$, say ν_i, we first randomly generate a continuous number, $u_d \sim U[0,1]$, representing the percentile of part $i \in I$ with respect to demand. Next, we obtain ν_i from $\nu_i(\lambda) = \frac{\lambda}{\rho_d} u_d^{\frac{1-\rho_d}{\rho_d}}$, where ρ_d is the demand skewness parameter, and λ is the average demand rate of all parts. Similarly, the unit variable cost, c_i, for any part $i \in I$ is generated from $c_i(c) = \frac{c}{\rho_c} u_c^{\frac{1-\rho_c}{\rho_c}}$, where ρ_c is the cost skewness parameter, c is the average unit variable cost of parts, and $u_c \sim U[0,1]$ is the percentile of part $i \in I$ with respect to unit variable costs. In this way, we obtain the part-specific average demand rate, ν_i, and the unit variable cost, c_i, for each part $i \in I$. Finally, by multiplying ν_i with a second random number generated for each location from U[0,2], we obtain the part-specific location-dependent demand rate, λ_{in}, at each warehouse $n \in N$. To obtain the part-specific location-dependent demand rate λ_{i0} for each part $i \in I$ at the central warehouse, we first generate the corresponding external demand rate, λ_{i0}^e, the same way we generate λ_{in}. After obtaining λ_{i0}^e and λ_{in} for all $n \in N$, λ_{i0} is obtained from $\lambda_{i0} = \lambda_{i0}^e + \sum_{n \in N} \lambda_{in}$. For any given part, this ensures the differences in the demand rates among warehouses. However, the demand of each part relative to that of the others remains identical at each warehouse. Note that this corresponds to a practical situation where each warehouse serves a market with a similar demand structure. We refer to this case as the symmetric demand case. However, in different geographical regions or markets, the demand of spare parts relative to each other may differ. In order to represent the demand asymmetry across warehouses, the second multiplier is generated from U[0,2] for each part $i \in I$ and each warehouse $n \in N \cup \{0\}$. We call this second case the asymmetric demand case. Based on the data available for the spare parts systems considered in our paper, the demand rate (unit variable cost) skewness is approximated as 20%/80% (20%/90%), i.e., $\rho_d = 0.139$ ($\rho_c = 0.097$), meaning that 20% of the parts represent approximately 80% (90%) of the total demand rate (cost) of parts. Table 3 summarizes the base setting used in our paper.

For the full factorial analysis, we consider 3 levels of the average demand rates: average unit variable costs, average fixed ordering costs and target aggregate mean response times. To generate the problem instances for the experiments, we first generate the base case setting, then we multiply the value of each parameter in the base case setting by the multiplier associated with each level in Table 4. Furthermore, to avoid explosion of the number of problem instances, we consider 2 levels of the number of parts and the number of local warehouses. For the first set of experiments, we consider small-size problems; the number of parts is set to 4 and 8, and the number of local warehouses is set to 2 and 4. In the second set of experiments, in which we experiment with larger problems, the number of parts is set to 100 and 500, and the number of local warehouses is set to 4 and 9. Based on these, 20 random problem instances are generated for each of the 243

Table 3: Base case setting for the experiments.

Factors	λ_{in} (units/day)	c_i ($/unit)	K_i ($/order)	$W_n^{\max}$ (day)	$W_n^{\max}$ (day)	h (per year)	T_{i0} (day)	T_{in} (day)
Values	$\lambda = 0.015$	$c = 3000$	U[50, 100]	0.3	0.3	0.25	10	1

Table 4: Multipliers for the average demand rates, average unit variable cost of parts, average fixed ordering costs and target aggregate mean response times.

Parameters	**Number of Levels**	**Level Multipliers**
λ_{in}	3	1/3, 1, 10/3
c_i	3	1/3, 1, 10/3
K_i	3	1/3, 1, 10/3
$W_n^{\max}$	3	1/3, 1, 3

$(2^4 \times 3^2)$ different settings, resulting in a total of 6480 problem instances for each set of experiments.

In addition to the full factorial analysis, we also carry out sensitivity analysis to precisely observe the effect of each factor on the performance of the heuristics and the Lagrangian dual bound. The problem instances for the sensitivity analysis are generated by using the base case setting in Table 3 in a similar way that the testbeds for the factorial analysis are generated. The results that we present here are based on those of the experiments with problem instances with symmetric demand structure. We also experiment with problem instances with asymmetric demand structure. We report the results of the latter only when there is an inconsistency between these two settings. In our experiments, we consider the cases in which (1) only external customers, (2) both type of customers are incorporated in evaluating the performance of the central warehouse. The experiments do not reveal any significant difference between the results of the two cases (in the symmetric demand case both models are the same). Therefore we only present the results for the former case, which is more common. In all experiments, the inventory carrying charge is taken as 25% annual. The algorithms are coded in C++ and the experiments are run on an Intel 3 GHz processor with 3.5 GB RAM. In the remainder of this section the results of the experiments are presented and discussed.

6.2 Performance of the Lagrangian Dual Bound.

A summary of the results regarding the factorial experiment to test the performance of the Lagrangian dual bound is given in Table 5. The main findings are as follows:

- As depicted in Table 5, both the average and the maximum $PGAP$ are high, however, both improve when the number of parts is larger.
- Table 5 also indicates that the results are sensitive to the factors considered. According to the ANOVA results, all the parameters are found to be significant at 0.05 significance level. The results also show that the parameters highly interact. The most significant interaction effects are the interactions between the number of parts and the average demand rate, the number of parts and target aggregate mean response time and the average demand rate and target aggregate mean response time, each having a p-value of 0.000.

The effect of the number of parts on the performance of the Lagrangian dual bound deserves further attention since it is used as a benchmark solution in the second part of experiments, in which we experiment with larger number of parts. Therefore, we carry out a sensitivity analysis to observe the effect of number of parts more deeply. We examine 9 cases with $|I|$ = 10, 15, 20, 25, 30, 35, 40, 45, 50, in each of which $|N| = 2$. For each case, we generate 5 random problem instances using the base case setting in Table 3. Figure 3 shows the results of the sensitivity analysis. Each point in the figure represents the average of $PGAP$s for 5 problem instances. As shown in this figure, the performance of the Lagrangian dual bound improves with the number of parts. Note that this result is consistent with Theorem 3, in which the Lagrangian dual bound is shown to be asymptotically tight in the number of parts. This result together with Theorem 3 suggests that the Lagrangian dual bound can confidently be used as a benchmark solution in the experiments with larger problems, which will be the case in the remaining of this chapter.

6.3 Performance of the Lagrangian Heuristic.

The results of the experiments are summarized in Table 6. Based on the results, we make the following observations:

- As shown in Table 6, the performance of the Lagrangian heuristic is quite satisfactory. The average PCD obtained by the Lagrangian heuristic is less than 1%. This result is even better for problem instances with large number of parts. When the

Table 5: Effect of the parameters on the performance of the Lagrangian dual bound.

Parameters		GAP		PGAP	
		Avg.	Max.	Avg.	Max.
All instances		524.6	26662.5	3.84%	53.87%
Number of Parts	4	556.5	26662.5	4.48%	53.87%
	8	492.7	13918.0	3.21%	32.20%
Number of LWHs	2	456.1	19033.5	4.06%	53.87%
	4	592.2	26662.5	3.63%	46.16%
Average Demand Rate (units/day)	0.005	517.6	19033.5	5.17%	53.87%
	0.015	472.3	13918.0	3.42%	31.68%
	0.05	610.1	26662.5	2.37%	21.17%
Average Unit Cost ($/unit)	1000	106.4	1907.4	2.73%	27.59%
	3000	305.3	10331.5	3.76%	39.39%
	10000	1147.8	26662.5	5.02%	53.87%
Average Fixed Ordering Cost ($/order)	25	521.5	19021.2	5.07%	53.87%
	75	569.0	24202.5	3.86%	40.88%
	250	486.5	26662.5	2.71%	27.59%
Target Aggregate Mean Response Time (day)	0.1	509.4	10919.8	2.82%	22.24%
	0.3	839.6	26662.5	4.49%	29.02%
	0.9	308.9	13637.6	4.26%	53.87%

Figure 3: Effect of number of parts on the performance of the Lagrangian dual bound.

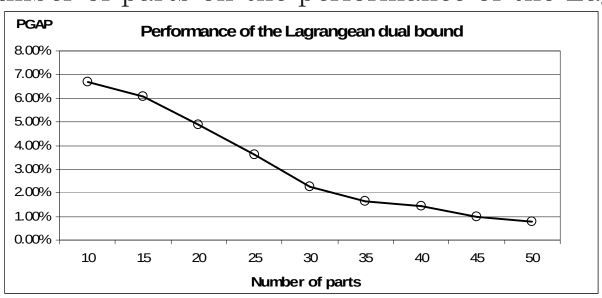

number of parts is 500, the PCD obtained by the Lagrangian heuristic is less than 1% for all of the 3240 problem instances studied.

- The ANOVA results reveal that the main effects of all the parameters except the number of local warehouses are significant at 0.05 significance level, each having a p-value of 0.000. On the other hand, the effect of the number of local warehouses on the performance of the Lagrangian heuristic is insignificant in terms of the PCD, but significant in terms of the ACD. The most significant interaction effects are the interactions of the number of parts with the average demand rate, the average unit variable cost, the average fixed ordering cost and the target aggregate mean response time and those of the target aggregate mean response time with the average demand rate and the number of warehouses. Each of these interactions is significant with a p-value of 0.000.

Table 6: Effect of the parameters on the performance of the Lagrangian heuristic.

Parameters		ACD		PCD	
		Avg.	Max.	Avg.	Max.
All instances		2686.7	94984.1	0.61%	10.33%
Number of Parts	100	3206.3	94984.1	1.08%	10.33%
	500	2167.1	39792.7	0.14%	0.93%
Number of LWHs	4	2063.2	37882.6	0.61%	8.16%
	9	3310.1	94984.1	0.61%	10.33%
Average Demand Rate (units/day)	0.005	2910.3	51784.2	0.83%	8.16%
	0.015	2785.2	94984.1	0.62%	10.33%
	0.05	2364.6	57048.9	0.38%	4.94%
Average Unit Cost ($/unit)	1000	586.0	7065.2	0.48%	8.16%
	3000	1699.1	17733.2	0.60%	7.53%
	10000	5774.9	94984.1	0.75%	10.33%
Average Fixed Ordering Cost ($/order)	25	2713.2	94984.1	0.74%	10.33%
	75	2678.9	50501.8	0.62%	7.61%
	250	2668.0	57048.9	0.47%	8.16%
Target Aggregate Mean Response Time (day)	0.1	4859.1	94984.1	0.79%	10.33%
	0.3	2420.7	35131.3	0.63%	7.82%
	0.9	780.2	9147.9	0.41%	7.45%

To identify the effect of parameters more precisely, we perform a sensitivity analysis for each factor. Figure 4 shows the results of the sensitivity analysis for the average demand rate, average unit variable cost, average fixed ordering cost and target aggregate mean response time, whereas Figure 5 illustrates the results of the analysis for the number of parts. As shown in Figure 4, we consider 10 levels for each of the average demand rate (λ = 0.005, 0.010,..., 0.050), the average unit variable cost (c = 1000, 2000,..., 10000), the average fixed ordering cost (K = 25, 50,..., 250) and the target aggregate mean response time (W_n = 0.1, 0.2,..., 1.0). We also consider two different values of the number of parts ($|I|$ = 50 and 250), abbreviated as NumPart in Figure 4, to explore the interactions between the effect of the number of parts and the effects of the corresponding four parameters. As shown in Figure 5, for the analysis of the effect of the number of parts, we consider 7 levels with $|I|$ = 50, 100, 250, 500, 1000, 3000 and 5000. To see the effect of the number of local warehouses in the same figure, we also consider four different values of the number of warehouses ($|N|$ = 3, 6, 9 and 12), abbreviated as NumWare in Figure 5. We randomly generate 10 problem instances for each level of the parameters, using the base case setting in Table 3. Hence, each point in the figures represents the average of PCDs for 10 problem instances. The main observations drawn from the sensitivity analysis are given as follows:

- Figure 4 indicates that the Lagrangian heuristic offers a better performance for problem instances with
  - high average demand rate,

– low average unit variable cost,

– high average fixed ordering cost,

The results are similar for ACDs and in line with the results of the factorial analysis presented in Table 6. We interpret these three observations as follows: Each problem instance conforming the conditions given above corresponds to a situation where the optimal inventory policy parameters, i.e., Q_i, R_i and S_{in}, are high. Hence, this shows that when the value of optimal policy parameters are high, both the PCD and the ACD obtained by the Lagrangian heuristic decrease. This observation is in line with Wong et al. (2007a) and (2007b).

Figure 4: Sensitivity Analysis: The effects of parameters on the performance of the Lagrangian heuristic - average demand rates, average unit variable cost of parts, average fixed ordering costs and target aggregate mean response times.

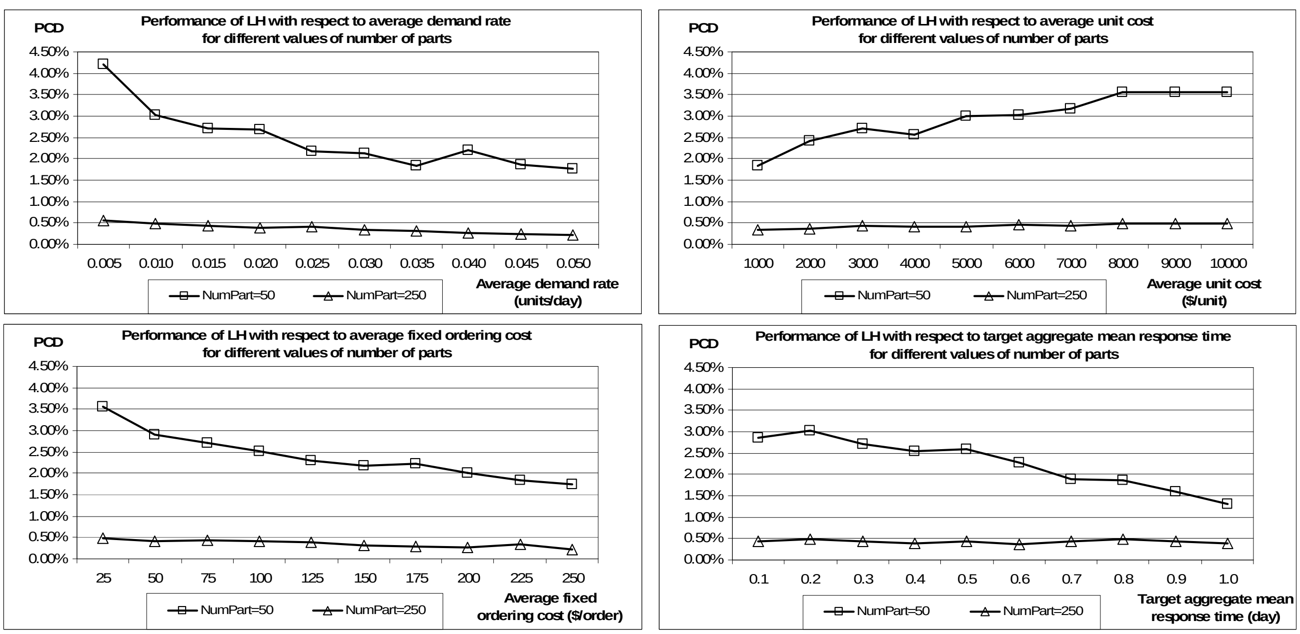

- The results of the factorial analysis in Table 6 shows that the PCD obtained by the Lagrangian heuristic decreases with the number of parts. Figure 5 further shows that as the number of parts increases, independent of the number of local warehouses, the PCD obtained by the Lagrangian heuristic approaches to zero. Note that this observation is similar to our findings regarding the asymptotic behavior of the Lagrangian dual solution. Intuitively, since the solution that the Lagrangian heuristic yields and the Lagrangian dual solution are the primal and dual solutions obtained through Lagrangian relaxation, respectively, it makes sense to have similar

Figure 5: Sensitivity Analysis: The effects of parameters on the performance of the Lagrangian heuristic - number of parts, number of local warehouses.

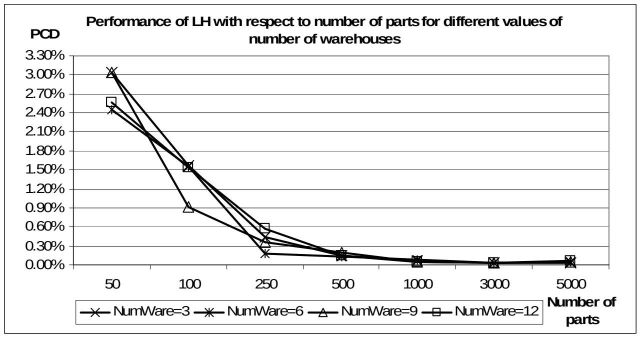

results for both the Lagrangian heuristic and the Lagrangian dual bound. We argue that these two results has some connection with the multi-item approach. Under the multi-item approach, which makes risk pooling possible among parts, as the number of parts increases, the benefits of risk pooling increases. This will also increase the number of alternative near-optimal solutions. In this situation, it is more likely to find a feasible solution that is close to the optimum by using an appropriate heuristic method, e.g., the Lagrangian heuristic. In a similar way, one can find a lower bound for our problem, e.g., the Lagrangian dual bound, by using an appropriate relaxation method, e.g., the Lagrangian relaxation. Similar results exist in the literature for other combinatorial problems as well, e.g., there exists greedy algorithms guaranteeing asymptotically optimal solutions for multidimensional knapsack problem and generalized assignment problem (Rinnooy Kan et al. 1993, Romeijn and Morales 2000); the relative gap between the LP relaxation of the knapsack problem and its optimal solution approaches to zero as the number of items increases (Kellerer et al. 2004).

- Figure 4 and Table 6 also show that the Lagrangian heuristic offers a better performance for problem instances with long target aggregate mean response time. This finding is also in line with our previous observation. Intuitively, long target aggregate mean response times yield loose constraints for the problem P. Under the multi-item approach, such loose constraints increase the risk pooling among parts. This increases the number of alternative near-optimal solutions. Hence, under long target aggregate mean response times, it is more likely to find a feasible solution that is close to the optimal solution by using the Lagrangian heuristic.

- Figures 4 and 5 also indicate that the effect of other parameters vanishes as the number of parts increases. Hence, as the problem gets larger the Lagrangian heuristic becomes more robust.

- In general, we observe that the Lagrangian heuristic yields better results in problem instances with asymmetric demand, and this corresponds to practical situations in which each warehouse serves a distinct market with a different demand structure. This result is in line with Wong et al. (2007b), reporting that the greedy heuristic performs better with asymmetric demand instances.

6.4 Performance of the Sequential Heuristics.

The results of the experiments to evaluate the performance of the sequential heuristics S_1, S_2, and S_3 are summarized in Table 7. Accordingly, we make the following observations:

- The average and worst case performances of S_2 and S_3 are better than those of S_1, indicating that the sequential approach performs better when service level requirements are taken into account in calculating the order quantities. However, the results indicate that neither of the methods dominates the others. S_3 outperforms S_1 and S_2 in 95.73% and 70.73% of all problem instances, respectively, and S_2 outperforms S_1 in 92.25% of all problem instances.

- The average and the maximum PCDs that the sequential approach yield are higher compared to the corresponding results of single-echelon systems operating under (Q, R) policy (Zheng 1992, Axsäter 1996, Silver et al. 1998, Gallego 1998): The PCD obtained by S_2, which uses EOQ^B to determine order quantities, can be as high as 21.75% in our experiments, while the maximum PCD obtained by the EOQ^B in single-echelon systems is 2.90% (Zheng 1992), empirically, and 11.80%, theoretically (Axsäter 1996). Similarly, the average and the maximum PCDs obtained by S_1, which uses EOQ, are 5.22% and 31.03%, respectively, which are fairly high for a batch size heuristic commonly used in practical applications. In contrast to findings for single-item models, our results are comparable with (Hopp et al. 1997), who report that the relative gap between the solution that their sequential heuristic yields and a lower bound may reach 15.37% in multi-item single-echelon inventory systems. Therefore, we conclude that in multi-item systems, which is a more realistic setting for spare parts, the performance of the sequential approach is not as satisfactory as in single-item systems. Although sequential approach is commonly used in practice, the results show that the approach results in very high errors. Note that all sequential heuristics considered in our study utilizes a column generation and a greedy algorithm to determine the reorder points at the central warehouse and the basestock levels at the local warehouses. However, in practical applications these parameters are typically determined by using simpler methods, which may yield much higher errors than the PCDs presented here. All these indicate that some of the conjectures in the literature about the performance of the sequential approach are misleading, and using the sequential approach in practical applications may not be the best option.

Table 7: Effect of the parameters on the performance of the sequential heuristics.

Parameters		S1				S2				S3			
		ACD		PCD		ACD		PCD		ACD		PCD	
		Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.	Avg.	Max.
All instances		35302.23	666547	5.22%	31.03%	16969.71	248755.6	2.54%	21.75%	14244.54	261775.6	2.33%	19.60%
Number of Parts	100	12974.6	146829.9	5.60%	31.03%	7946.908	100419.5	3.14%	21.75%	7014.1	101433.1	2.87%	19.60%
	500	57629.87	666547	4.83%	28.34%	25992.5	248755.6	1.95%	13.59%	21475.0	261775.6	1.78%	13.15%
Number of LWHs	4	26613.65	418397.8	4.83%	25.21%	13448.6	149855.9	2.50%	15.03%	10834.6	159583.2	2.19%	11.12%
	9	43990.82	666547	5.61%	31.03%	20490.81	248755.6	2.59%	21.75%	17654.4	261775.6	2.47%	19.60%
Average Demand Rate (units/day)	0.005	12927.39	116452.5	3.16%	11.79%	9693.469	99016.23	2.30%	11.64%	7068.1	66833.0	1.85%	13.49%
	0.015	27500.98	260993.2	4.68%	14.88%	15825.71	115150.5	2.56%	10.92%	11730.4	101433.1	2.12%	10.60%
	0.05	65478.33	666547	7.81%	31.03%	25389.94	248755.6	2.78%	21.75%	23935.1	261775.6	3.01%	19.60%
Average Unit Cost ($/unit)	1000	10777.22	83248.66	4.77%	29.24%	5057.16	77763.86	2.45%	21.75%	3537.1	26444.2	1.89%	10.95%
	3000	28273.51	226439	5.62%	31.03%	12975.09	174922	2.65%	15.03%	10306.5	78532.7	2.44%	19.60%
	10000	66855.97	666547	5.26%	31.02%	32876.86	248755.6	2.54%	11.61%	28890.1	261775.6	2.65%	18.66%
Average Fixed Ordering Cost ($/order)	25	22737.91	259462.3	5.31%	31.03%	12526.99	123329.4	2.54%	11.64%	11311.0	131554.9	2.65%	19.60%
	75	35135.03	435269.8	5.62%	31.02%	16906.63	170287.3	2.67%	14.69%	14656.4	198428.6	2.48%	18.34%
	250	48033.77	666547	4.72%	29.24%	21475.5	248755.6	2.42%	21.75%	16766.2	261775.6	1.86%	10.99%
Target Aggregate Mean Response Time (day)	0.1	30656.15	403908.2	2.59%	11.79%	22185.97	248755.6	2.03%	11.64%	13809.0	122389.8	1.51%	13.49%
	0.3	32858.9	492391.8	4.04%	11.18%	16274.28	160121.3	2.26%	11.17%	12532.1	157118.3	1.83%	11.81%
	0.9	42391.65	666547	9.02%	31.03%	12448.88	196347.9	3.35%	21.75%	16392.6	261775.6	3.65%	19.60%

- The Lagrangian heuristic generally performs better than the sequential heuristics. When the number of parts is 100, the Lagrangian heuristic outperforms S_2 (S_3) in 95.77% (94.51%) of the problem instances, whereas when the number of parts is 500, the Lagrangian heuristic outperforms S_2 (S_3) in 99.88% (99.48%) of the problem instances. This shows that as the number of parts increases, the Lagrangian heuristic, which is based on the simultaneous approach, becomes much more dominant over the sequential heuristics. Intuitively, under the multi-item approach, determining the order quantity of each part independent of the other parts' parameters benefits less from the opportunities of risk-pooling among parts. Therefore, for a system operating under the multi-item approach, (1) as the number of parts increases, or (2) as we switch from single-item to multi-item setting, the performance of the sequential heuristics deteriorates relative to that of the simultaneous approach.

In addition to the observations given above, the ANOVA results indicate the followings:

- The effects of all the parameters considered in the factorial analysis are significant at 0.05 significance level, each having a p-value less than 0.005. The most significant observation is that the sequential heuristics generally perform better in problem instances with

  – high average fixed ordering cost,
  – short target aggregate mean response time,
  – low average demand rate.

  These findings are in line with the literature (Zheng 1992, Hopp et al. 1997). However, even with those problem instances that are in favor of the sequential heuristics, the PCD obtained by the sequential approach is still higher compared to that of using the Lagrangian heuristic. This can be seen also in Table 7.

- Almost all the interactions between the factors are significant at 0.05 level. Because of the interactions between the factors, we also analyze and interpret the effects of parameters considering the interactions. The most critical observation is that the performance of the sequential heuristics relative to that of the Lagrangian heuristics is best when the average demand rate and the average fixed cost are low, the average unit variable cost is high, and the target aggregate mean response time is short. We note that this corresponds to instances in which the performance of the Lagrangian heuristic performs worse. The results show that even in this setting, the Lagrangian heuristic outperforms the sequential heuristics in most of instances. Under this specific setting, when the number of parts is 100, the Lagrangian heuristic outperforms S_3 (S_2) in 70.0% (65.0%) of the problem instances, whereas when the number of parts is 500, the Lagrangian heuristic outperforms S_3 (S_2) in 95.5% (100.0%) of the problem instances. Again, this shows that as the number of parts increases, the Lagrangian heuristic becomes much more dominant over the sequential heuristics even for the problem instances less favorable for the Lagrangian heuristic.

We further perform a sensitivity analysis to identify the effect of parameters more precisely. We conduct the experiments based on the testbed used for the sensitivity analysis of the Lagrangian heuristic. The main observations drawn from the sensitivity analysis are given as follows:

- Figures 6 and 7 illustrate the result of analysis regarding the effect of the number of parts on the performance of the heuristics based on the PCDs and ACDs, respectively. As shown in Figure 6, similar to the findings regarding the performance of the Lagrangian heuristic, as the number of parts increases, the PCD obtained by the sequential heuristics decreases. However, Figure 6 also shows that neither of them converges to zero, although that of the Lagrangian heuristic does, e.g., the average PCD obtained by the sequential heuristic with EOQ is 4.38% even the number of parts is 5000. Figure 7 shows that the ACD obtained by the sequential heuristics increases with the number of parts quite faster than that of the Lagrangian heuristic.

Figure 6: Performance of the heuristics with respect to number of parts and number of parts - PCD.

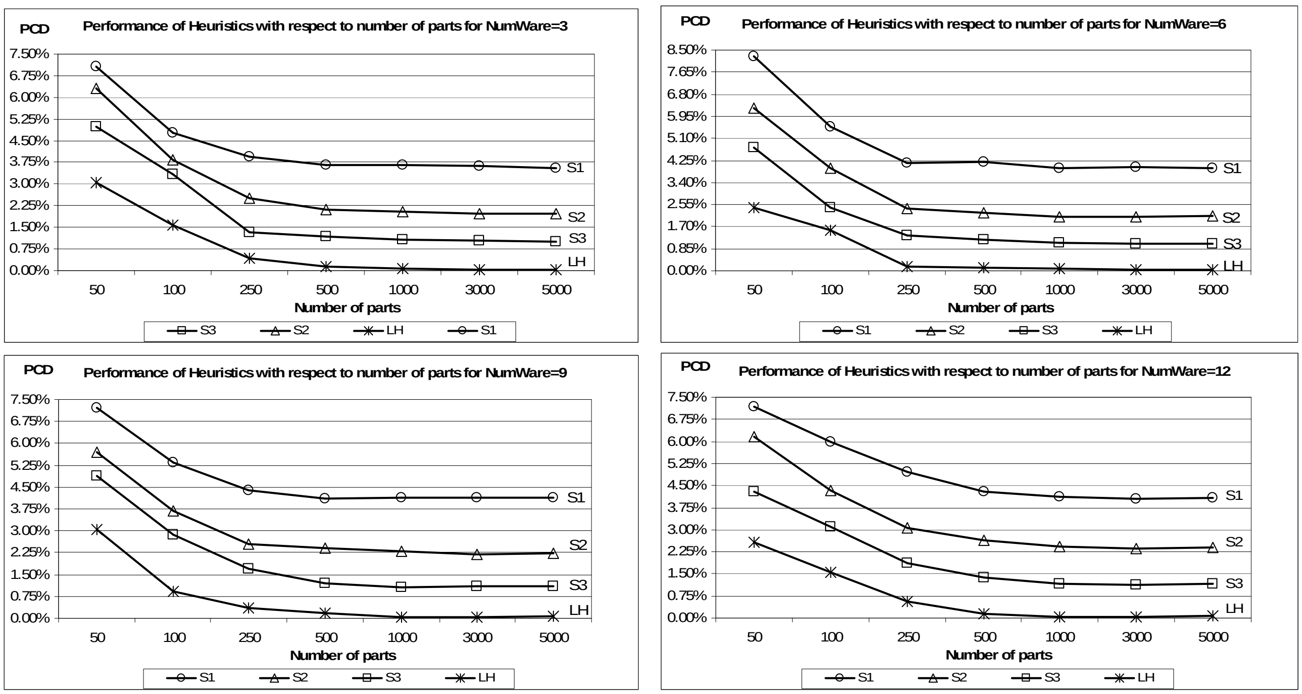

Figure 7: Performance of the heuristics with respect to number of parts and number of parts - ACD.

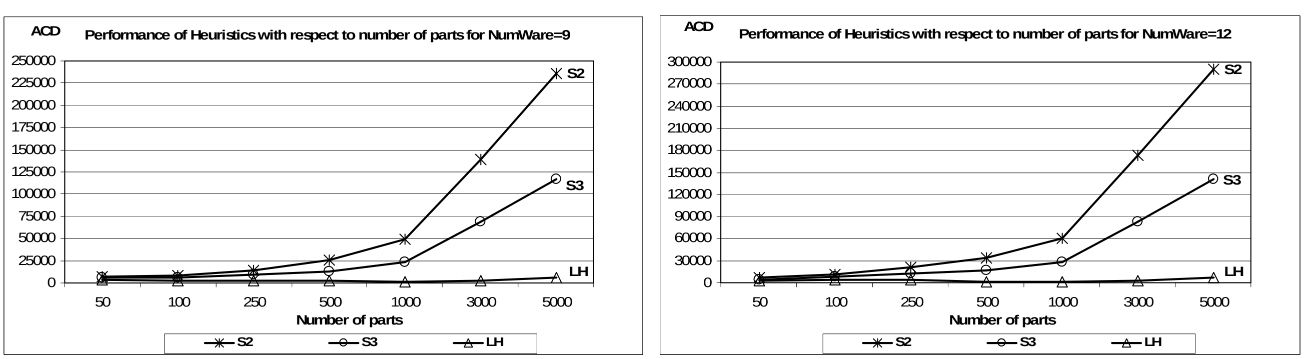

6.5 Computational requirements of the solution procedures and experiments with Practical-Size Problems.

In this section, we test the performance of the heuristics in terms of the computational requirements. Figure 8 illustrates the computational requirements of the heuristics that have run to test the effect of the number of parts on the heuristics in Sections 6.3 and 6.4. The figure shows only the results for $|N| = 9$ and 12. As shown in the figure the average CPU time required by heuristics increases with the number of parts. This is rather intuitive. The figure also shows that the average CPU time for the Lagrangian heuristic is higher than those for the sequential heuristics, but they are comparable even when $|I| = 5000$. This indicates the computational savings that can be obtained by sequential heuristics are quite limited. Recall that by using the sequential heuristics, the computational time required to solve the problem P is reduced by eliminating the outer loop in the single-item two-echelon subroutine in the column generation algorithm. However, the experiments reveal that for small-size problems, the bulk of the computational effort is devoted to the column generation algorithm, whereas for problems with larger number of parts, the most of the computational time is spent by the greedy algorithm. Thus, elimination of the subroutine, which makes sequential heuristics faster, contributes less to computational savings of the overall method for large-scale problems.

Figure 8: Computational requirements of the heuristics with respect to number of parts and number of local warehouses

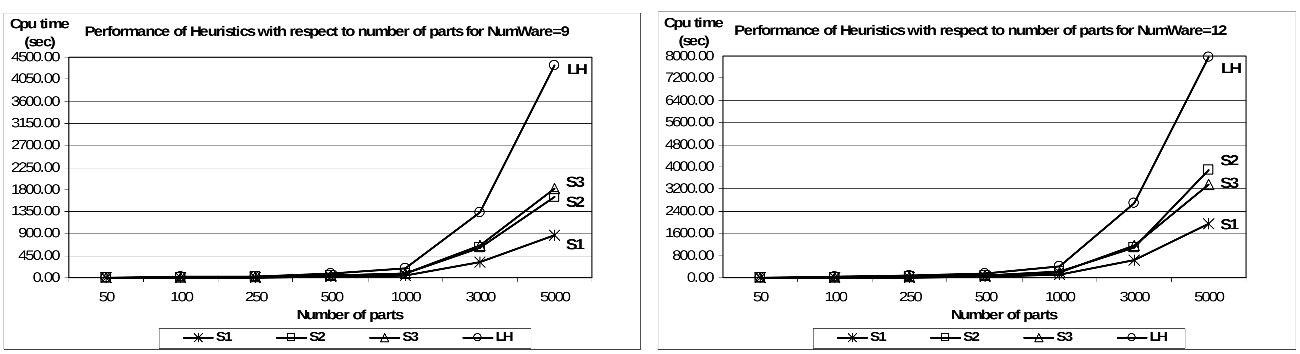

Furthermore, we experiment with larger problem instances to further explore the performance of the Lagrangian heuristic in practical-size problems. The number of parts is set to 10000, and the number of local warehouses is set to 12. The problem instances are generated by using the base case setting in Table 3. We consider both the symmetric and the asymmetric problem instances. The results of the experiment show that the average PCDs obtained by the Lagrangian heuristic are 0.09% and 0.04%, and the average CPU times required by the Lagrangian heuristic are 15.55 and 11.94 hours for the symmetric and the asymmetric problem instances, respectively.

Although the papers in the literature deal with systems under different settings, a comparison with them is still possible to a certain extent. Hopp et al. (1999) who consider a multi-item two-echelon batch ordering system similar to ours report that their heuristic can solve problem instances with 1263 parts and 2 regional facilities. They also show that for relatively small problem instances (with up to 10 parts and 5 regional facilities) the relative gap between the expected total cost of the solution obtained by the heuristic and the lower bound that they propose is less than 5%. Compared to Hopp et al. (1999), our heuristics seem to be significantly better both in terms of relative errors and the computational efficiency. When we compare our results with Wong et al. (2007b), who apply a column generation and decomposition method similar to the one in our paper for a multi-item two-echelon system under pure basestock policy, our results seem to be comparable in terms of relative errors. Furthermore, although our problem is more difficult than theirs, our results are significantly better in terms of the computational efficiency, e.g., while they can solve problems problems up to a size of 100 parts, we can solve problems with up to 10000 parts. Note that when we fix the order quantities at 1, our model reduces to their system, hence our heuristics are applicable in their setting as well. We also compare our results with Caggiano et al. (2007), who propose heuristics to solve large-scale multi-item multi-echelon systems under basestock policy. Although their system is very different from our system (our model involves batch ordering decision, while their model is applicable to more than two-echelon), one can see that our results are slightly better than theirs both in terms of relative errors and computational efficiency. For example, while they can solve problems with 27175 part-location combinations in almost 21 hours, in our paper we can solve 130000 part-location combinations (10000 parts, 12 local warehouses and 1 central warehouse) within 12 hours. As evident from these comparisons, our paper contributes to the relevant literature by proposing an efficient and tractable heuristic for large-scale spare parts inventory problems. As opposed to these papers, a comparison with Al-Rifai and Rossetti (2007) is not possible since the performance of their heuristic is not compared against an analytical solution or a bound, since such a solution or a bound is not available for large-scale problems. Çağlar et al. (2004) encounter a similar problem. They report that the lower bound that they use to test their heuristic is not tight, and hence, the relative gap between the their heuristics and the lower bound is used as a conservative estimate of the true relative error. As opposed to these studies, the Lagrangian dual bound for our problem is tractable and performs quite well for larger-scale problems. Hence, our findings regarding the performance of the Lagrangian dual bound also contribute to the relevant literature.

7 Conclusion.

In this paper, we consider a multi-item two-echelon inventory system in which the central warehouse operates under a (Q, R) policy, and each local warehouses implements an $(S-1, S)$ policy. Our objective is to find the policy parameters minimizing expected system-wide inventory holding and fixed ordering costs subject to an aggregate mean response time constraint at each facility. We propose four alternative heuristics to find the optimal parameters based on the exact evaluation of the probability distributions of the inventory levels in the problem. The first one, which we call the Lagrangian heuristic, is based on the simultaneous approach. The other three heuristics are based on the sequential approach and differ in the batch size heuristics used. We also propose a lower bound for the optimal total cost and analytically show that this bound is asymptotically tight in the number of parts. We conduct an extensive computational study to test the performance of the heuristics and the lower bound. Based on the results of the computational study, we provide several interesting results, some of which are summarized as follows: The Lagrangian heuristic performs quite well and provides an excellent performance as the number of parts increases. Furthermore, as the number of parts increases, the Lagrangian heuristic becomes robust and becomes insensitive to other parameters and whether the demand is symmetric or not. The heuristic is also quite tractable. It can be used to solve very large practical problems in reasonable computation time. This makes the Lagrangian heuristic very promising for practical applications. The performance of the heuristics that are based on the sequential approach are also satisfactory, but not as much as the Lagrangian heuristic. As the number of parts increases the performance of the sequential approach deteriorates compared to that of the Lagrangian heuristic. Furthermore, the computational advantage of the sequential heuristics is found to be limited compared to the Lagrangian heuristic. Hence, despite the fact that the sequential approach is widely used in practice and that its performance is experimentally verified in single-item problems, our paper shows that in a multi-item setting and under a multi-item approach, the performance of the sequential approach heuristics is inferior compared to the Lagrangian heuristic, which yields superior results even for large problems in reasonable time. These results show that some of the conjectures in the literature about the performance of the sequential approach are misleading. We also find that the performance of the sequential heuristics depends on the batch size heuristic used. The batch size heuristic that take service level requirements into account outperforms the

EOQ. This shows that if the sequential heuristics are used to solve batching problems in multi-item two-echelon inventory systems, it is better to take the service level requirements into account in calculating the order quantities. For this purpose, one can use the batch size heuristics that we consider in this paper.

To summarize, we contribute to the literature by proposing efficient and tractable heuristics to solve large, practical-size multi-item two-echelon inventory control problems with batch ordering at the central warehouse, one of which significantly outperforms the others. The comparisons based on our heuristics also makes a contribution to the literature in evaluating the performance of the sequential approach against the simultaneous approach in a multi-item multi-echelon setting. Our paper also contributes to the literature by proposing a tight and efficient lower bound on the optimal total cost for practical-size problems.

Although we consider systems with fixed ordering costs and aggregate mean response time constraints, our work can be extended to systems with target ordering frequency constraints and/or other service measures or backorder costs. We also note that although our main concern is spare parts inventory systems, our results are applicable to a wide range of inventory systems with a similar cost and service level structure as in this paper.

Since both the heuristics and the lower bound that we propose have a general framework, they can be adopted to more complex systems as well, e.g., multi-item, more than two-echelon systems, the ones with more complex control policies. In this situation, again, to obtain a lower bound, the column generation can be used to decompose the resulting problem into single-item problems, and the greedy algorithm can be employed to find a feasible solution using the lower bound obtained by the column generation. Similarly, the sequential approach can easily be extended to more complex systems using the batch size heuristics proposed in this paper. For all such extensions, the difficulty arises in solving the resulting single-item (multi-echelon) problems, just like the one in our paper. Since the structure of the corresponding systems will resemble ours, we expect that one can obtain results similar to ours, such as the asymptotic tightness of the Lagrangian dual bound, and the improvement of the Lagrangian heuristic's performance in the number of parts. Furthermore, it may also be interesting to investigate the issues raised in our paper in a more general system setting, e.g., exploring the performance of the sequential approach.

Since the Lagrangian heuristic and the Lagrangian dual bound yield quite satisfactory results with practical size problems, they can be used to provide several managerial insights about problems encountered in practice, e.g., cost-benefit analysis of opening up a new local warehouses, determining the optimal number of local warehouses, and cost-benefit analysis of increasing the service levels at the warehouses.

8 Appendix.

PROOF OF LEMMA 1. In problem $SP_i(\vec{\theta})$ let $\theta_n = 0$ for all $n \in N$, meaning that no penalty cost is incurred due to backorders at each local warehouse $n \in N$. This yields an optimal solution in which $S_{in} = 0$ for each $i \in I$ and $n \in N$ and the optimal expected inventory holding and backorder cost at each local warehouse is zero. In this situation, the subproblem $SP_i(\vec{\theta})$ is reduced to a single-echelon batch ordering problem with expected cost function $c_i h E[I_{i0}(Q_i, R_i)] + \theta_0 E[B_{i0}(Q_i, R_i)] + \frac{\lambda_{i0} K_i}{Q_i}$ (Zheng 1992, Gallego 1998). Note that for a given value of θ_0, the solution of this single-echelon batch ordering problem is a lower bound to the optimal objective function value of $SP_i(\vec{\theta})$. Furthermore, it is known that the optimal objective function value of the EOQ model with backorders, i.e., $z_i^{EOQ(\theta_0)}$, is a lower bound to the solution of the single-echelon batch ordering problem whose cost function is given above (Gallego 1998). Combining these two arguments we establish $C_{il} \geq z_i^{EOQ(\theta_0)}$ for each $i \in I$ and $l \in L$.

□

PROOF OF THEOREM 1. Using the result of Lemma 1 in constraint (5), we have $\sum\limits_{l \in L} C_{il} x_{il} \geq z_i^{EOQ(\theta_0)}$ for each $i \in I$. Then, by summing up these expressions over all $i \in I$, we obtain $\sum\limits_{i \in I} \sum\limits_{l \in L} C_{il} x_{il} \geq \sum\limits_{i \in I} z_i^{EOQ(\theta_0)}$. Let C_{ave} be the average of $z_i^{EOQ(\theta_0)}$ over all $i \in I$, hence, defined as $C_{ave}(I) = \frac{1}{|I|} \sum_{i \in I} z_i^{EOQ(\theta_0)}$, then, we simply have $\sum\limits_{i \in I} \sum\limits_{l \in L} C_{il} x_{il} \geq |I| \cdot C_{ave}(I)$. Since this holds for all feasible solutions for problem MP, we also have $z^{MP} \geq |I| \cdot C_{ave}(I)$. Furthermore, since we define each c_i, K_i, λ_{i0}, θ_0 as an independent and identically distributed random variable, $z_i^{EOQ(\theta_0)}$ is an independent and identically distributed random variable for each $i \in I$ as well. From the convergence of random variables, $C_{ave}(I)$ converges to a constant with probability 1 as $|I|$ goes to infinity. This suffices to show that z^{MP} is asymptotically bounded below by a function in the order of $|I|$ with probability 1.

□

PROOF OF LEMMA 2.

Proof of part (a): The proof is rather intuitive. Since $K_i < \infty$, $c_i h > 0$, $W_n^{\max} > 0$ and $\lambda_{in} T_{in} < \infty$, $SP_i(\vec{\theta})$ is guaranteed to yield finite solutions, e.g., for each $n \in N \cup \{0\}$, $W_n^{\max} > 0$ implies $\theta_n < \infty$ since θ_n has correspondence with $W_n^{\max}$ through α_n, the dual price for the relevant constraint (4), and this is necessary to have finite R_i and S_{in}. Hence, our column generation method yields finite columns.

Proof of part (b): Provided that the objective function parameters in $SP_i(\vec{\theta})$ are finite, e.g., $K_i < \infty$, $c_i h < \infty$, $W_n^{\max} > 0$, i.e., $\theta_n < \infty$, the optimal objective function value of $SP_i(\vec{\theta})$ is finite. Hence, the total cost associated with each column generated by the column generation method is finite. The proof of the second part relies on that due to the decomposition of problem CG into parts, the size of each subproblem $SP_i(\vec{\theta})$ grows only with the order of $|N|$.

□

PROOF OF LEMMA 3. The proof is based on an alternative formulation of problem MP obtained by substituting the equality constraints (5) in (4): Arbitrarily, we select x_{i1} and then substitute $x_{i1} = 1 - \sum_{l \in L'} x_{il}$ in (4) for each $i \in I$, where $L' = L - \{1\}$. In this way, we establish the alternative formulation AP.

Problem AP:

$$
\begin{aligned}
&Min\ Z = \sum_{i \in I} \sum_{l \in L'} C_{il} x_{il} \\
&s.t. \\
&\sum_{i \in I} \sum_{l \in L'} (A_{iln} - A_{i1n}) x_{il} \le W_n^{\max} - A_{i1n}, \qquad \text{for } \forall\, n \in N \cup \{0\}, \\
&x_{il} = 0/1, \qquad \text{for } \forall\, i \in I,\ \forall\, l \in L',
\end{aligned}
$$

Note that in the optimal solution of the LP-relaxation of AP ($LPAP$), there exist $|N|$ basic variables. Furthermore, a variable has a fractional value only if it is a basic variable. Hence, the optimal solution of $LPAP$ contains at most $|N|$ variables with fractional values. Finally, since AP is exactly the same problem as MP, this result also holds for $LPMP$. □

PROOF OF THEOREM 2. Our proof consists of two parts. In the first part, we introduce a repair algorithm to generate an integer feasible solution to MP by adjusting only the fractional variables in the solution of $LPMP$. Then in the second part, by using the repair algorithm, we show that the gap between the expected cost obtained by the repair algorithm and the Lagrangian dual solution is asymptotically bounded above by a function of $|N|^2$.

The repair algorithm relies on the following observation: For each part $i \in I$, the $LPMP$ yields a solution that is a convex combination of columns generated by the column generation algorithm. Accordingly, the solution for each part $i \in I$ corresponds to either

- an integer solution corresponding to one of the columns generated by the algorithm (pure policy), i.e., one of the variables x_{il} in constraint (5) is 1 while the others are 0, or,

- a fractional solution that is a mixture of a set of columns generated by the algorithm (randomized policy), i.e., a set of variables x_{il} in constraint (5) have fractional values summing to 1 while others are 0.

Based on these observations, the following repair algorithm generates an integer solution for each part $i \in I$ whose solution is fractional so that the overall solution still remains feasible.
The Repair Algorithm:

1. For each part $i \in I$ having an integer solution, do nothing.

2. For each part $i \in I$ having a fractional solution, generate a new solution by taking the maximum value of the policy parameters defined by the columns that constitute the fractional solution, and replace the corresponding fractional solution with this new one. To be more specific, let $m \in I$ be any of those parts whose solution is fractional and $(Q_m, R_m, \vec{S}_m)$ be the corresponding fractional solution. Then, we replace the fractional solution $(Q_m, R_m, \vec{S}_m)$ with the solution $(\widetilde{Q}_m, \widetilde{R}_m, \vec{\widetilde{S}}_m) = \left(\underset{k\in\Gamma_m}{max} \left\{ Q_m^k \right\}, \underset{k\in\Gamma_m}{max} \left\{ R_m^k \right\}, \underset{k\in\Gamma_m}{max} \left\{ \vec{S}_m^k \right\} \right)$, where Γ_m is the set of (integer) columns that constitutes $(Q_m, R_m, \vec{S}_m)$.

The entire solution is feasible for MP, because (1) for each part $i \in I$, the solution generated by the repair algorithm satisfies integrality, (2) for each part $i \in I$, the new solution yields lower $E[B_{i0}(Q_i, R_i)]$ and $E[B_{in}(Q_i, R_i, S_{in})]$ values for all $n \in N$ than the fractional solution yields, meaning that just like the former columns, the new solutions are guaranteed to satisfy the constraints (4).

After introducing the repair algorithm and our notation, now we begin our proof: First, it is a direct consequence of Lemma 2 that the additional cost incurred by switching from $(Q_m, R_m, \vec{S}_m)$ to $(\widetilde{Q}_m, \widetilde{R}_m, \vec{\widetilde{S}}_m)$ is finite and bounded above by an order of N. Second, for any part $i \in I$, the repair algorithm requires increasing the value of at most $|N| + 2$ policy parameters, i.e., in our example one for Q_m, one for R_m and $|N|$ for $\vec{S}_m$. Combining these two results we find that for any part $i \in I$, the solution obtained through the repair algorithm has an additional cost bounded above by a finite value in the order of $|N|$. Furthermore, it follows from Lemma 2 that we need to reassign at most $|N|$ fractional variables to obtain a feasible integer solution for MP by using the solution of $LPMP$. Let the objective function value of this integer feasible solution be z^H. Then, by combining the arguments above, we establish that $z^H - z^{LP}$ is asymptotically bounded above by a function of $|N|^2$. Since z^H is an upper bound on z^{MP}, the result also holds for z^{MP}. Hence, $z^{MP} - z^{LP}$ is in $\mathbf{O}(|N|^2)$. This also proves that $z^H - z^{LP}$ grows only with an order of $|N|$. □

PROOF OF THEOREM 3. It follows from Theorem 1 and Theorem 2 that for any given value of $|N|$, $\lim_{|I|\to\infty} \frac{z^{MP}-z^{LPMP}}{z^{MP}} \to 0$ with probability 1. Since P and MP are identical problems, we have $z^P = z^{MP}$. Also, since the solution of $LPMP$ gives the Lagrangian dual solution of P, z^{LD}, we obtain $z^{LD} = z^{LPMP}$. Hence, for any given value of $|N|$, we have $\lim_{|I|\to\infty} \frac{z^P - z^{LD}}{z^P} \to 0$ with probability 1. This shows that z^{LD} is asymptotically tight in the number of parts (Anily and Federgruen, 1990).
□

References

[2007] Al-Rifai, M.H., Rossetti, M.D., (2007). "An efficient heuristic optimization algorithm for a two-echelon (R, Q) inventory system". *Int. J. Production Economics*, Vol. 109 (1-2), 195-213.

[1990] Anily, S., Federgruen, A. (1990). "A class of euclidean routing problems with general route cost functions", *Mathematics of Operations Research*, Vol.**15 (2)** 268-285.

[1996] Axsäter, S., (1996). "Using the Deterministic EOQ Formula in Stochastic Inventory Control", *Management Science*, Vol. 42 (6), 830-834.

[1998] Axsäter, S., (1998) "Evaluation of Installation Stock Based (R,Q) Policies for Two Level Inventory Systems with Poisson Demand", *Operations Research*, Vol.46, 135-145.

[2003] Axsäter, S., (2003). "Approximate Optimization of A Two-Level Distribution Inventory System", International Journal of Production Economics, Vol.81-82, 545-553.

[2006] Axsäter, S., (2006). *Inventory Control.* Springer, New York.

[2001] Cachon, G.P., (2001). "Exact Evaluation of Batch-Ordering Inventory Policies in Two-Echelon Supply Chains with Periodic Review", *Operations Research*, Vol. 49 (1), 79-98.

[2004] Çağlar, D., Li, C.L., and Simchi-Levi, D., (2004). "Two-echelon spare parts inventory system subject to a service constraint", *IIE Transactions*, Vol.36, 655-666.

[2006] Caggiano, K.E., Muckstadt, J.A., Rappold, J.A., (2006). “Integrated real-time capacity and inventory allocation for reparable service parts in a two-echelon supply system”, *Manufacturing and Service Operations Management*, Vol. 8, 292319.

[2007] Caggiano, K.E., Jackson, P.L, Muckstadt, J.A., Rappold, J.A., (2007). “Optimizing Service Parts Inventory in a Multiechelon, Multi-Item Supply Chain with Time-Based Customer Service-Level Agreements. *Operations Research* **55 (2)** 303318.

[1990] Cohen, M., Kamesam, P.V., Kleindorfer, P., Lee, H., and Tekerian, A., (1990). “Optimizer: IBM’s multi echelon inventory system for managing service logistics”, *Interfaces*, Vol. (20), 65-82.

[1997] Cohen, M.A, Zheng, Y.S., Agrawal, V., (1997). “Service parts logistics: a benchmark analysis”, *IIE Transactions*, Vol. 29, 627-639.

[1963] Dantzig, G.B., (1963). *Linear programming and extensions*. Princeton University Press, Princeton, NJ.

[1998] Gallego, G., (1998). “New bounds and heuristics for (Q,R) policies”, Management Science Vol. 44 (2), 219-233.

[1985] Graves, S.C., (1985). “A Multi-Echelon Inventory Model for a Repairable Item with One-for-One Replenishment”, *Management Science*, Vol.31, 1247-1256.

[2003] Guignard, M., (2003). “Lagrangean Relaxation”, *Sociedad de Estadistica e Investigacion Operativa Top*, Vol. 11 (2), 151-228.

[1997] Hopp, W.J., Spearman, M.L., Zhang, R.Q., (1997). “Easily Implementable Inventory Control Policies”, *Operations Researchs*, Vol.45, 327-340.

[1999] Hopp, W.J., Zhang, R.Q., Spearman, M.L., (1999). “An Easily Implementable Hierarchical Heuristic for a Two-Echelon Spare Parts Distribution System”, *IIE Transactions*, Vol.31, 977-988.

[2004] Kellerer, H., Pferschy, U., Pisinger, D., (2004). *Knapsack Problems*. Springer, Berlin.

[2009] Kim, S.H., Cohen, M.A. Netessine, S., (2009). “Reliability or Inventory? Analysis of Product Support Contracts in the Defense Industry”, Working Paper.

[2002] Lübbecke, M.E., and Desrosiers, J. (2002). “Selected Topics in Column Generation”, *Operations Research*, Vol.53 (6), 1007-1023.

[2008] Özer, Ö., and Xiong, H., (2008). “Stock positioning and performance estimation for distribution systems with service constraints”, *IIE Transactions*, Vol. 40 (12), 1141-1157.

[1993] Rinnooy Kan, A.H.G., Stougie, L., Vercellis, C., (1993). “A class of generalized greedy algorithms for the multi-knapsack problem”, *Discrete Applied Mathematics*, Vol. **42** 279-290.

[2000] Romeijn, H.E., Morales, D.R., (2000). “A class of greedy algorithms for the generalized assignment problem”, *Discrete Applied Mathematics*, Vol. **103** 209-235.

[1968] Sherbrooke, C.C., (1968). “A multi-echelon technique for recoverable item control”, *Operations Research*, Vol. 16, 122-141.

[1998] Silver, E.A., Pyke, D., and Peterson, R., (1998). Inventory Management and Production Planning and Scheduling, 3rd edition, Wiley, New York.

[2007] Simchi-Levi, D., and Zhao. Y., (2007). “Three Generic Methods for Evaluating Stochastic Multi-Echelon Inventory Systems”, *Working paper*.

[2002] Thonemann, U.W., Brown, A.O., and Hausman, W.H., (2002). “Easy Quantification of Improved Spare Parts Inventory Policies”, *Management Science*, Vol. 48 (9), 12131225.

[2009] Topan, E., Bayındır, Z.P., Tan, T., (2010). “An exact solution procedure for multi-item two-echelon spare parts inventory control problem with batch ordering in the central warehouse”, *Operations Research Letters*, doi:10.1016/j.orl.2010.05.006.

[2007] Van Houtum, G.J., Scheller-Wolf, A., and Yi, J., (2007). Optimal control of serial inventory systems with fixed replenishment intervals. Operations Research, 55(4), 674-687.

[2005] Wong, H., Van Houtum, G.J., Cattrysse, D., and Van Oudheusden, D., (2005). “Simple, efficient heuristics for multi-item multi-location spare parts systems with lateral transshipments and waiting time constraints”, *Journal of the Operational Research Society*, Vol. 56, 1419-1430.

[2006] Wong, H., Van Houtum, G.J., Cattrysse, D., and Van Oudheusden, D., (2006). “Multi-item spare parts systems with lateral transshipment and waiting time constraints”, *European Journal of Operational Research*, Vol. 171, 1071-1093.

[2007a] Wong, H., Van Oudheusden, D., and Cattrysse, D., (2007a). "Two-Echelon Multi-Item Spare Parts Systems with Emergency Supply Flexibility and waiting Time Constraints", *IIE Transactions*, Vol. 39 (2007), 1045-1057.

[2007b] Wong, H., Kranenburg, D., Van Houtum, G.J and Cattrysse, D., (2007b). "Efficient heuristics for two-echelon spare parts inventory systems with an aggregate mean waiting time constraint per local warehouse", *OR Spectrum*, Vol. 29, 699-722.

[1992] Zheng, Y.S., (1992). "On Properties of Stochastic Inventory Systems", *Management Science*, Vol.38 (1), 87-103.

Working Papers Beta 2009 - 2010

nr.	Year	Title	Author(s)
321	2010	Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse	
320	2010	Preventing or escaping the suppression mechanism: intervention conditions	B. Walrave, K. v. Oorschot, A.G.L. Romme
319	2010	Hospital admission planning to optimize major resources utilization under uncertainty	Nico Dellaert, Jully Jeunet.
318	2010	Minimal Protocol Adaptors for Interacting Services	R. Seguel, R. Eshuis, P. Grefen.
317	2010	Teaching Retail Operations in Business and Engineering Schools	Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo.
316	2010	Design for Availability: Creating Value for Manufacturers and Customers	Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak.
315	2010	Transforming Process Models: executable rewrite rules versus a formalized Java program	Pieter van Gorp, Rik Eshuis.
314	2010	Ambidexterity and getting trapped in the suppression of exploration: a simulation model	Bob Walrave, Kim E. van Oorschot, A. Georges L. Romme
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	R.J.I. Basten, M.C. van der Heijden,

		analysis	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 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