Università degli Studi di Roma Tre
Dipartimento di Informatica e Automazione
Via della Vasca Navale, 79 – 00146 Roma, Italy

Optimal stock allocation in single echelon inventory systems subject to a service constraint

Annalisa Cesaro[1], Dario Pacciarelli[1]

RT-DIA-184-2011 **Maggio 2011**

(1) Università degli Studi Roma Tre,
Via della Vasca Navale, 79
00146 Roma, Italy.

This work is partially supported by the Italian Ministry of Research, Grant number RBIP06BZW8, project FIRB "Advanced tracking system in intermodal freight transportation".

ABSTRACT

This paper addresses the problem of spare parts allocation in a single echelon inventory system with complete pooling characterized by expensive spares, long repairing time and a strict service constraint. The work is motivated by the practical problem faced by a large Italian logistics company that handles 17 warehouses supporting the daily activity of 38 civil airports spread over the Italian territory. The objective is the minimization of the total costs for inventory holding, lateral transshipments and emergency shipments. The spares allocation problem is formulated as a non convex integer program. The special structure of the problem allows to design an efficient branch and bound procedure. The lower bound is obtained by formulating a reduced problem with convex objective function, solvable at optimality very efficiently. Computational experiments, carried on practical data, show that this method solves at optimality many practical instances.

1 Introduction

Single echelon inventory systems are experiencing an increasing interest in practice, in particular for the management of expensive spare parts. In such a context, the supply chain involves at least three actors: equipment users, logistics companies and equipment suppliers. The users need spare parts to carry on their business without interruptions. Intermediate logistic companies are in charge of replenishing spare parts in the short term, by guarantying the contractual service level to the users at minimum cost. The suppliers are responsible for supplying new components and/or repaired items to the logistic companies.

This paper addresses the problem of spare parts allocation in a single echelon inventory system with complete pooling characterized by expensive spares, long repairing time and strict requirements of operational availability (i.e., the fraction of time during which all operational sites are working). Our work is motivated by a practical problem faced by a large Italian logistics company. The company handles 17 warehouses supporting the daily activity of 38 civil airports spread over the Italian territory. Each site is coupled with one regional warehouse. The base stock quantities are computed by the company through the VARIMETRIC algorithm of Sherbrooke (Sherbrooke 2004), based on a stiff hierarchic structure. However, in operation lateral transshipments take place between stocking points whenever there is an emergency requirement for parts, using couriers and overnight carriers to rapidly move parts. The company is therefore interested in determining the potential savings deriving from explicit inclusion of lateral transshipments in the model. To this aim, we propose and evaluate a new branch and bound procedure for stock level definition and spare parts allocation. The procedure exploits the particular cost structure of the maintenance supply chain under study and is very effective in this context. However, the method is general and we discuss the algorithm performance in a more general context.

The literature on inventory management deals with the analysis of many different models. Extensive surveys of inventory models with lateral transshipments can be found in [5, 15]. In these surveys, the existing approaches are classified according to a number of characteristics related to the inventory system, the ordering policy, the modeling of transshipments and other issues (e.g., the number of items, echelons, the number and role of locations, the unsatisfied demand, the timing of regular orders, the order policy, the type of transshipments, pooling, decision making).

Huiskonen [8] and Kennedy et al. [10] focus on review models for spare part management. They observe that the logistics of spare parts differs from those of other materials in several ways. Equipments may have remarkable costs, long repairing times and sporadic failures. The latter are difficult to forecast and may cause relevant financial effects, due to the economical implications of a lack of equipment at the operational sites. In such cases, continuous review policies are typically adopted to reduce both reaction time to stockouts and inventory levels [1, 7, 13].

We next present the most relevant references for our study. In spare part management the continuous review policies are typically adopted to reduce both reaction time to stockouts and inventory levels [1, 7, 13]. Several heuristic procedures can be found in the literature for allocating spares to warehouses in a single echelon context with complete pooling [2, 13, 19]. In such a context, Wong et al. [20] develop a solution procedure based on Lagrangian relaxation to obtain both a lower bound and an upper bound on the optimal total cost. A heuristic procedure for the base stock level determination is used also by Kranenburg et al. [12] in a single echelon context, where partial pooling is adopted. In a two echelon context, Koutanoglu et al. [14] consider a model to allocate stock levels at warehouses. They use approximate models for estimating the performance of a specific allocation of spare parts and exploit it to implicitly enumerate all possible allocations of spares.

In this paper, we describe a new branch and bound algorithm to compute the optimal base stock quantities that minimize the expected average total cost in a multi-location single echelon transshipment system. The algorithm makes use of a new lower bound based on Lagrangian relaxation. The computational experiments carried on practical instances show the effectiveness of the algorithm.

This paper is organized as follows. Relevant notation is described in Section 2. Section 4 describes the single echelon one-for-one ordering model with complete pooling and the spare parts allocation problem is formulated as a non-convex integer program. In Section 3 the Markov chain model used for computing transshipment costs and times. Heuristic and exact allocation algorithms are described in Section 6. Section 5 studies the mathematical structure of the optimization problem. Computational experiments are presented in Section 7, based on practical data from the airport maintenance context. Some conclusions follow in Section 8.

2 Notation

In order to formally define the problem, let us introduce the following notation. Let $A = \{1, 2, \dots, a\}$ be a set of operational sites (e.g., airports) where working equipments are located. We assume that operational sites are grouped on a regional basis, with a warehouse of spare parts for each region. Let $W = \{1, 2, \dots, w\}$ be the set of regional warehouses. Let s_i be the number of spare parts to allocate to each warehouse $i \in W$, $S = \sum_{i \in W} s_i$ be the total stock level and $s = (s_1, ..., s_w)$ be an allocation of spares to warehouses, i.e., the vector of decision variables. We denote with $MTTR$ the Mean Time To Return, i.e. the average replenishment time of the external supplier, with $MTBF$ the Mean Time Between Failures, with OS the Order and Ship time. Therefore, $MTTR+OS$ is the total time elapsed from the failure of a part to its replenishment in the warehouse. We denote with $MCMT$ the Mean Corrective Maintenance Time and with OA the Operational Availability of all the a sites. We also denote with MET the Mean Emergency Time, i.e., the mean time needed to replenish a part in emergency conditions, from the failure to the replenishment. In this paper we assume $MET = MTTR+OS$, which corresponds to the assumption that the order and ship time and the mean time to return from the supplier do not vary in emergency conditions and in the normal operations. This is a conservative choice since in practice MET is often smaller than $MTTR + OS$. We keep $MET = MTTR + OS$ since in our application the operational availability requirements are very strict and it is preferable to overestimate the blocking probability with respect to its underestimation. With this hypothesis, the service rate of a server at warehouse h is $\mu = \frac{1}{MTTR+OS}$. Let T_{hi} be the transfer time for a spare from warehouse h to warehouse i and $T_s(j, h)$ be the substitution time, i.e., the time needed to transfer a spare part to the site $j \in A$ from the warehouse $h \in W$ and to physically replace the failed item. Let λ_{jh} be the mean rate of failures from site j to warehouse h, let $\lambda_h = \sum_{j \in A} \lambda_{jh}$ be the arrival rate at warehouse h and let $\Lambda = \sum_{hinW} \lambda_h$. Note that $MTBF = \frac{1}{w} \sum_{i=1}^{w} \frac{1}{\lambda_i}$. Given an allocation s, the network blocking probability is the probability that a failure occurs at some site and no warehouse can satisfy the spare demand. We denote the network blocking probability as $P_B(S)$ since, as it will be shown in Section 3, it only depends on the total stock level S rather than on the particular allocation s.

Given an allocation s, let $\pi_{hi}(s)$ be the probability of the event: there are no spares in warehouse $h \in W$ and $i \in W$ is the closest warehouse with available spares (i.e., every warehouse l such that $T_{hl} < T_{hi}$, including the case $l = h$, is in stockout condition). Let $n = (n_1, \ldots, n_w, n_{w+1})$ be a vector representing the state of the network, in which n_i is the number of outstanding requests at warehouse $i \in W$, and n_{w+1} is the number of outstanding emergency requests to the external supplier. Let $p(n)$ be the probability that the warehouse network is in state n, and let c^h be the inventory holding cost for warehouse h, c^t_{ij} be the cost for a lateral transshipment from warehouse j to warehouse i, in stockout condition, and c^e be the emergency transshipment cost.

3 Multi-dimensional Markovian model

In this section, we compute the probability $p(n)$ of each state $n = (n_1, \ldots, n_w, n_{w+1})$ of the associated Markov chain for a given spare part allocation. In case of blocked network, the first repaired item returned by the external supplier is used for replacing a failed item in some operative site, if any. There are direct transitions among states just in case of a single arrival event (i.e. a request for a spare at some warehouse) or a single departure event (i.e. the replenishment of a repaired item by the external supplier). Let e^i be a vector with $w + 1$ elements, all equal to 0 but the element in position i that is equal to 1, and let $\psi(h, i)$ be equal to 1 if ($n_i < s_i$ and $n_l = s_l$ for each $l \in W$ such that $T_{lh} < T_{ih}$, included $l = h$) and be equal to 0 otherwise. With this notation, $n + e_i$ is the state of the Markov chain representing an arrival at the i-th warehouse (if $i \in W$ and $n_i < s_i$), due either to a failure in the i-th service region or to a re-forwarded request from some other warehouse h in stockout conditions, i.e., such that $\psi(h, i) = 1$. Similarly, $n - e_i$ is the state with a departure from the i-th warehouse (if $i \in W$ and $n_i > 0$). For the external supplier, $n + e_{w+1}$ represents a new emergency request (if $n_i = s_i$ for each $i \in W$) and $n - e_{w+1}$ represents the fulfillment of an emergency request (if $n_{w+1} > 0$). The transition rate $q(n, m)$ from state n towards state $m = n \pm e_i$ and $n \pm e_{w+1}$ is as follows.

- $q(n, n + e_i) = \lambda_i + \sum_{h \in W - \{i\}} \psi(h, i)\lambda_h$, for $i \in W$ and $n_i = 0, 1, \ldots, s_i - 1$;
- $q(n, n + e_{w+1}) = \sum_{i \in W} \lambda_i$, if $n_i = s_i \ \forall i \in W$;
- $q(n, n - e_i) = \sum_{i=1}^{w+1} n_i \mu$, for $i \in W$ and $n_i > 0$ and $n_{w+1} = 0$;
- $q(n, n - e_{w+1}) = \sum_{i=1}^{w+1} n_i \mu$, for $n_i = s_i \ \forall i \in W$ and $n_{w+1} \geq 1$.

Figure 1(a) shows the Markov chain for two warehouses, the first having two spares and the second having three available spares. Steady state probabilities can be computed for each state in the Markov chain by solving a linear system. Note that, to this aim the proposed Markov chain is equivalent to the one shown in figure 1(b), having a finite number of states. In the latter Markov chain, all states in which all warehouses are

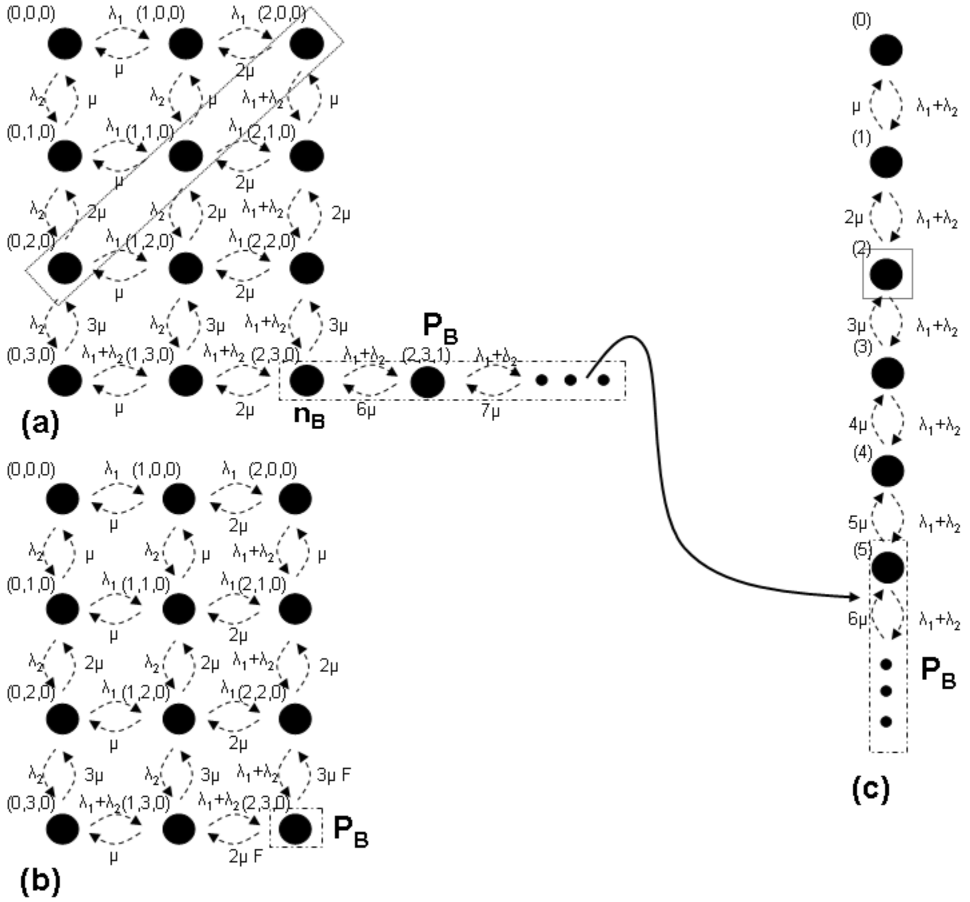

Figure 1: The infinite state space Markov chain for a system of two warehouses (a) and finite space state model (b) and the aggregated birth death model (c)

in stockout condition are grouped in one single state with different departure transition rates.

Let n_B be the state in the original model in which $(n_B)_i = s_i$, $\forall i \in W$, and $(n_B)_{w+1} = 0$ (see Figure 1(a)). In the equivalent finite state model, the transition rates from state n_B after a replenishment is: $q(n_B, n_B - e_i) = (n_B)_i \mu F$, where $F = \frac{p(n_B)}{P_B(S)}$ (see Figure 1(b)). As proved in [3], quantities $p(n_B)$ and $P_B(S)$ can be easily computed by using the equivalence of the original model with the simple birth death model associated to a single queue with infinitely many servers (see figure 1(c)). Specifically, let $\rho = \frac{\sum_{i \in W} \lambda_i}{\mu}$ and consider a set W of warehouses, with total stock level S, in which the service process is exponentially distributed with average rate μ for each server and the demand flow to warehouse $i \in W$ is Poissonian with average rate λ_i. Then, the blocking probability of all warehouses is:

$$P_B(S) = 1 - \sum_{k=0}^{S-1} \frac{\rho^k}{k!} e^{-\rho}, \qquad (1)$$

while $p(n_B)$ is equal to $\frac{\rho^{N_B}}{N_B!} e^{-\rho}$, where $N_B = \sum_{i=1}^{w+1} (n_B)_i$.

By directly solving the finite state Markov chain model we can compute the steady state probability $p(n)$ of each state n and use them to compute the system cost and performance associated to a given allocation. Since the direct computation of state probabilities is expensive in terms of computational effort, we wish to carry out this computation for a limited number of allocations only. In [3] different heuristic methods for estimating costs and performance in such a context have been proposed and compared in terms of estimation accuracy, computation time and memory efforts with the exact Markov chain computation. In Section 6, we make use of the most effective of such methods to develop a heuristic procedure for spare allocation. On the other hand, we cannot use these heuristic methods to compute exact solutions. Therefore, in what follows we focus on the exact solution of the original Markov chain model.

The Markov chain model that we use is very similar to that of Wong et al. [18, 20]. However, while Wong et al. [18, 20] assume that the demand in state n_B is lost (i.e., the arrival rate in n_B is zero), we assume that this demand is backordered to an external supplier (i.e., we explicitly include the external supplier in the Markov chain). The resulting overall blocking probability in our case is strictly greater than in [18, 20] and therefore the results are more conservative. The latter property is specifically important for our application, in which operational availability requirements are very strict.

4 The problem

In our model, a logistic company aims to compute the stock level s_i of each warehouse $i \in W$ such that a minimum level of service is granted at the operational sites and the overall cost is minimum. Costs are related to inventory holding, transshipments and emergency shipments.

Given an allocation s of spares to warehouses, the model used to compute the level of service is a single, reparable item, single echelon, w-locations, continuous review, one-for-one replenishment policy inventory system, with lateral and emergency shipments, complete pooling and non-negligible transshipment times.

The Spares Allocation Problem is the problem of finding an allocation s which minimizes the overall cost for inventory holding, lateral and emergency shipments, subject to a constraint on the minimum operational availability of the system.

The contractual service level to grant is the operational availability OA of all operational sites for each item, computed as in [16].

$$OA = \frac{MTBF}{MTBF + MCMT}. \tag{2}$$

In equation (2), MCMT is the average time occurring from the failure of an item to its physical substitution. This is the substitution time if the spare is available at the regional warehouse. If no spares are locally available, the request is forwarded to the closest warehouse with available spares and MCMT increases by the deterministic transfer time between the two warehouses. When no warehouse has spares available, MCMT equals the substitution time plus the replenishment time from the external supplier. The MCMT can be therefore computed as follows:

$$\begin{aligned} MCMT = \ & \textstyle\sum_{h \in W} \sum_{j \in A} \frac{\lambda_{jh}}{\Lambda} T_s(j,h) + \\ & \textstyle\sum_{h \in W} \left(\frac{\lambda_h}{\Lambda} \sum_{i \in W} \pi_{hi}(s) T_{ih} \right) + \\ & P_B(S) MET \end{aligned} \tag{3}$$

We observe that the first term $\sum_{h \in W} \sum_{j=1}^{A} \lambda_{jh} T_s(j,h)$ of Equation (3) only depends on the failure process and on the distance between the sites and their respective regional warehouses. In other words, it does not depend on the specific spare parts management policy being used. Therefore, for sake of simplicity we assume it negligible in our model and omit its computation in the rest of this paper. As for the quantity $\pi_{hi}(s)$, we assume that a strict deterministic nearest chosen neighbor rule is adopted for sourcing a lateral transshipment, as in Kukreja [13]. Differently from [13], we compute the quantity $\pi_{hi}(s)$ by directly using its definition. In fact, we compute the probability $p(n)$ of each state $n = (n_1, \ldots, n_w, n_{w+1})$ by solving the associated Markov chain defined in Section 2 exactly and use these probabilities for computing the quantity $\pi_{hi}(s)$ as follows:

$$\pi_{hi}(s) = \sum_{n \in \Phi} p(n) \tag{4}$$

where $\Phi = \{n : n_i < s_i, n_l = s_l \forall l : T_{hl} < T_{hi}\}$.
We assume the Poisson distribution for the demand process, which is a typical assumption for modeling low demand processes [17]. We also use location dependent MTBF values. The replenishment time of the external supplier is a random variable, exponentially distributed, with known mean value MTTR. The capacity of the supplier repair shop is assumed to be infinite. It follows that also the number of replenishments from the external supplier follows the Poissonian distribution. These common assumptions make possible to use the Markovian analysis for modeling the multi-dimensional inventory system. Finally, we make the following assumptions.

1. Lateral transshipment is always more convenient than emergency shipment, i.e., the time and cost needed for a transshipment from warehouse i to warehouse j is always smaller than the time and cost required for an emergency shipment from warehouse

j:

$$\max_{i,j \in W}\{T_{ij}\} \quad < \quad MET \tag{5}$$

$$\max_{i,j \in W}\{c^t_{ij}\} \quad < \quad c^e \tag{6}$$

2. The cost for a lateral transshipment from warehouse i to warehouse j increases linearly with the transfer time T_{ij}, i.e.,

$$c^t_{ij} = \alpha T_{ij} \tag{7}$$

For the purpose of formulating the spare allocation problem, let L be the minimum operational availability level to be achieved by a feasible allocation. According to expression (2), the minimum operational availability constraint is as follows.

$$\frac{MTBF}{MTBF + \sum_{i \in W}[(\frac{\lambda_i}{\Lambda}\sum_{j \in W}\pi_{ij}(s)T_{ji})] + P_B(S)(MET)} \geq L.$$

This expression is equivalent to

$$MTBF \geq L[MTBF + \sum_{i \in W}[(\frac{\lambda_i}{\Lambda}\sum_{j \in W}\pi_{ij}(s)T_{ji})] + P_B(S)(MET)]$$

which, in turn, can be expressed as

$$\sum_{i \in W}[(\frac{\lambda_i}{\Lambda}\sum_{j \in W}\pi_{ij}(s)T_{ji})] + P_B(S)(MET) \leq \frac{(1-L)MTBF}{L}.$$

The latter inequality formulates the constraint by allowing a maximum waiting time $\frac{(1-L)MTBF}{L}$ to substitute all the failed items. In view of this expression, the Spares Allocation Problem P_0 can be formulated as the following integer program with non-convex objective function:

Problem P_0:

$$\begin{array}{ll} \min & \sum_{i=1}^{w} c^h s_i + \lambda_i \sum_{j \in W}\pi_{ij}(s)c^t_{ji} + \lambda_i P_B(S)c^e \\ s.t.: & \\ & \sum_{i=1}^{w}[\frac{\lambda_i}{\Lambda}\sum_{j \in W}\pi_{ij}(s)T_{ji}] + P_B(S)(MET) \leq \frac{(1-L)MTBF}{L} \end{array} \tag{8}$$

Let $f_1(S) = \sum_{i=1}^{w} c^h s_i$ be the total inventory holding cost, $f_2(s) = \left\{\sum_{i=1}^{w}\lambda_i \sum_{j \in W}\pi_{ij}(s)c^t_{ji}\right\}$ be the cost for lateral transshipments and $f_3(S) = \sum_{i=1}^{w}\lambda_i P_B(S)c^e$ be the cost for the emergency shipments. Similarly, for the waiting times we let $t_2(s) = \left\{\sum_{i=1}^{w}\frac{\lambda_i}{\Lambda}\sum_{j \in W}\pi_{ij}(s)T_{ji}\right\}$ be the waiting times due to lateral transshipments and $t_3(S) = P_B(S)(MET)$ be the emergency waiting times. Problem P_0 is very similar to that of Wong et al. [20]. The main difference is in the service level constraint. Due to the needs of our practical case study, we assume a single service level constraint for all locations, while in Wong et al. [18, 20] there is a constraint for each location.

5 Problem structure

In this section, we exploit the special structure of problem P_0. Let us first analyze the three functions $f_1(S)$, $f_2(s)$, $f_3(S)$. $f_1(S)$ is clearly linear and increasing with the total stock level S. As shown in [11], $f_3(S)$ is convex and decreasing with S for $S \geq \rho - 1$. Unfortunately, $f_2(s)$ is non-convex. The latter property can be easily shown by observing that $f_2(s)$ equals zero when $s = 0$ or when $s_i \to \infty$ for all $i = 1, \dots w$, since there are no transshipments in these two cases. On the other hand, $f_2(s) > 0$ otherwise, thus implying that $f_2(s)$ is non-convex. Similarly, it can be easily proved that $t_3(S)$ is convex with S for $S \geq \rho - 1$ and $t_2(s)$ is non-convex. We next show that the quantities $t_2(s) + t_3(S)$ and $f_2(s) + f_3(S)$ are decreasing. To this aim, let us consider an allocation s and a warehouse $i \in W$. Denote with $\hat{s}$ the allocation such that $\hat{s}_i = s_i + 1$ and $\hat{s}_h = s_h$ for all $h \in W$, $h \neq i$. Let also denote $S = \sum_{i=1}^{w} s_i$ and $\hat{S} = \sum_{i=1}^{w} \hat{s}_i = S + 1$.

Let us first observe that when passing from s to $\hat{s}$ the number of spares at each warehouse does not decrease and therefore the probability of an outstanding request at each warehouse cannot increase. Specifically, the following properties must hold.

- $P_B(S) > P_B(\hat{S})$.
- As for warehouse i, the probability that the aggregated arrival rate λ_i (without transshipments) is satisfied by the local stock increases when passing from s_i to s_i+1. Thus, the probability of re-forwarding the demand decreases, i.e. $\pi_{ij}(\hat{s}) < \pi_{ij}(s)$. It follows that $\sum_{j \in W} \pi_{ij}(s) T_{ji} + P_B(S)(MET) > \sum_{j \in W} \pi_{ij}(\hat{s}) T_{ji} + P_B(\hat{S})(MET)$.
- For what concerns warehouse $h \neq i$ the probability of re-forwarding the request to warehouse i increases and the probability of re-forwarding the request to a more far warehouse (or to the external supplier) decreases, i.e., $\sum_{j \in W} \pi_{hj}(s) T_{jh} + P_B(S)(MET) > \sum_{j \in W} \pi_{hj}(\hat{s}) T_{jh} + P_B(\hat{S})(MET)$.

In conclusion,

$$t_2(s) + t_3(S) - t_2(\hat{s}) - t_3(\hat{S}) =$$
$$= \sum_{i=1}^{w} \frac{\lambda_i}{\Lambda} \left[\sum_{j \in W} \left(\pi_{ij}(s) - \pi_{ij}(\hat{s}) \right) T_{ji} \right] + \left(P_B(S) - P_B(\hat{S}) \right) (MET) > 0.$$

Using assumption (6), a similar discussion for the costs leads to $f_2(s) + f_3(S) > f_2(\hat{s}) + f_3(\hat{S})$.

Given an upper bound UB on the optimum of problem P_0, an upper bound MAX on the total stock level S of an optimal solution can be efficiently computed by considering only the terms $f_1(S)$ and $f_3(S)$ of the objective function.

$$MAX = \min \{ S : f_1(S) + f_3(S) \geq UB \} \tag{9}$$

This value is quite close to the optimal stock level S^* when the transshipment cost $f_2(s^*)$ is small with respect to $f_1(S^*) + f_3(S^*)$. Similarly, a lower bound MIN on S^* can be efficiently computed by considering only the term $t_3(S)$, decreasing with S, in the constraint of the problem.

$$MIN = \min \left\{ S : t_3(s) \leq \frac{(1-L)MTBF}{L} \right\} \tag{10}$$

These bounds can be used to refine the formulation of Problem P_0, thus leading to the new formulation P_1.

Problem P_1:

$$\begin{aligned} \min \quad & f_1(S) + f_2(s) + f_3(S) \\ s.t.: \quad & \\ & t_2(s) + t_3(S) \le \tfrac{(1-L)MTBF}{L} \\ & MIN \le S \le MAX \end{aligned} \tag{11}$$

We observe that, if $MIN = \min\left\{S : t_3(s) \le \frac{(1-L)MTBF}{L}\right\} \ge \rho - 1$, then $f_3(S)$ and $t_3(S)$ are convex [11]. In the remaining part of this paper we assume that this condition is always satisfied. In fact, this condition holds for all the instances we tested and, more in general, it is a mild condition since, for example, even for $L \ge 0.5$ it can be easily checked that $MIN \ge \rho - 1$ for a wide range of ρ values (up to more than 100).

We next introduce a Lagrangian relaxation $P_2(\gamma)$ of problem P_1 by relaxing the waiting time constraint. We use the notation γ to denote the Lagrangian multiplier.

Problem $P_2(\gamma)$:

$$\begin{aligned} \min \quad & f_1(S) + f_2(s) + f_3(S) + \gamma\left(t_2(s) + t_3(S) - \tfrac{(1-L)MTBF}{L}\right) \\ s.t.: \quad & \\ & MIN \le S \le MAX \end{aligned} \tag{12}$$

It is well known that, for varying γ, $P_2(\gamma)$ is a concave, piecewise linear function. Calling breakpoint the values of $P_2(\gamma)$ in which the slope of $P_2(\gamma)$ changes, there is an optimal solution γ^* for the Lagrangian dual $\max\{P_2(\gamma) : \gamma \ge 0\}$ which is a breakpoint. If we let $\overline{s}$ be an optimal allocation for $P_2(\overline{\gamma})$, and $\overline{\gamma}$ is not a breakpoint, then the slope of $P_2(\gamma)$ in $\overline{\gamma}$ is [6]:

$$t_2(\overline{s}) + t_3(\overline{S}) - \tfrac{(1-L)MTBF}{L} \tag{13}$$

Theorem 1 *If $\overline{\gamma}$ is not a breakpoint there is a single optimal stock level for $P_2(\overline{\gamma})$.*

Proof. By contradiction, let us suppose that in $\overline{\gamma}$ there are two optimal allocations s and $\overline{s}$ with different stock levels S and $\overline{S}$, respectively. Therefore:

$$f_1(S) + f_2(s) + f_3(S) + \overline{\gamma}\left(t_2(s) + t_3(S) - \frac{(1-L)MTBF}{L}\right) =$$
$$= f_(\overline{S}) + f_2(\overline{s}) + f_3(\overline{S}) + \overline{\gamma}\left(t_2(\overline{s}) + t_3(\overline{S}) - \frac{(1-L)MTBF}{L}\right).$$

From equation (13) it follows that the constraint violation is the same for s and $\overline{s}$, i.e., $t_2(s) + t_3(S) = t_2(\overline{s}) + t_3(\overline{S})$. It follows from the proportionality assumption (7) that also $f_2(s) + f_3(S) = f_2(\overline{s}) + f_3(\overline{S})$ must hold. Hence, we obtain $f_1(S) = f_1(\overline{S})$, i.e., $c^h S = c^h \overline{S}$. This implies the thesis $S = \overline{S}$. □

Theorem 2 *If $\overline{\gamma}$ is a breakpoint and the slope of $P_2(\overline{\gamma})$ decreases from $t_2(s^1) + t_3(S^1) - \frac{(1-L)MTBF}{L}$ to $t_2(s^2) + t_3(S^2) - \frac{(1-L)MTBF}{L}$, then $S^2 > S^1$.*

Proof. At the breakpoint $\overline{\gamma}$ there are at least the two optimal solutions s^1 and s^2 for problem $P_2(\overline{\gamma})$, i.e.,

$$f_1(S^1) + f_2(s^1) + f_3(S^1) + \overline{\gamma}\left(t_2(s^1) + t_3(S^1) - \frac{(1-L)MTBF}{L}\right) =$$
$$= f_1(S^2) + f_2(s^2) + f_3(S^2) + \overline{\gamma}\left(t_2(s^2) + t_3(S^2) - \frac{(1-L)MTBF}{L}\right).$$

Since the slope of $P_2(\gamma)$ decreases, then $t_2(s^1) + t_3(S^1) > t_2(s^2) + t_3(S^2)$ and, from the proportionality assumption (7), also $f_2(s^1) + f_3(S^1) > f_2(s^2) + f_3(S^2)$ must hold. Hence, it follows that $f_1(S^1) < f_1(S^2)$, which implies the thesis $S^1 < S^2$. □

Theorem 3 *If the breakpoint γ^* is an optimal solution of the Lagrangian dual* $\max\{P_2(\gamma) : \gamma \geq 0\}$ *and the slope of $P_2(\gamma^*)$ decreases from $t_2(s^1) + t_3(S^1) - \frac{(1-L)MTBF}{L} \geq 0$ to $t_2(s^2) + t_3(S^2) - \frac{(1-L)MTBF}{L} \leq 0$, then:*

1. *s^2 is feasible for problem P_1;*
2. *either s^2 is optimal for P_1 or S^2 is greater than the optimal stock level for P_1.*

Proof. The feasibility of s^2 directly follows from $t_2(s^2) + t_3(S^2) - \frac{(1-L)MTBF}{L} \leq 0$. If s^2 is not optimal, let s^* be an optimal allocation and S^* be the corresponding stock level. From the optimality of S^* it follows that:

$$f_1(S^*) + f_2(s^*) + f_3(S^*) < f_1(S^2) + f_2(s^2) + f_3(S^2) \tag{14}$$

On the other hand at γ^* the objective function of the Lagrangian relaxation computed in s^* must be greater or equal than in s^2, i.e.,

$$f_1(S^*) + f_2(s^*) + f_3(S^*) + \gamma^*\left(t_2(s^*) + t_3(S^*) - \frac{(1-L)MTBF}{L}\right) \geq$$
$$= f_1(S^2) + f_2(s^2) + f_3(S^2) + \gamma^*\left(t_2(s^2) + t_3(S^2) - \frac{(1-L)MTBF}{L}\right).$$

Therefore, $t_2(s^*) + t_3(S^*) > t_2(s^2) + t_3(S^2)$ must hold. From assumption (7) it must hold also $f_2(s^*) + f_3(S^*) > f_2(s^2) + f_3(S^2)$. Therefore, from inequality (14), $f_1(S^*) < f_1(S^2)$, i.e., $S^* < S^2$. □

Despite the nice properties of $P_2(\gamma^*)$ shown in Theorem 3, the computation of $P_2(\gamma^*)$ requires the computation of quantity $f_2(s^*)+\gamma^* t_2(s^*)$, which is computationally expensive. In order to efficiently compute a lower bound to $P_2(\gamma^*)$, let us introduce problem $P_3(\gamma)$:

Problem $P_3(\gamma)$:

$$\begin{array}{ll} \min & f_1(S) + x + f_3(S) + \gamma\left(y + t_3(S) - \frac{(1-L)MTBF}{L}\right) \\ s.t.: & \\ & MIN \leq S \leq MAX \\ & x \leq f_2(s) \\ & y \leq t_2(s) \end{array} \tag{15}$$

Suitable values for x and y can be computed by exploiting the properties that $f_3(S)$ and $f_2(s)+f_3(S)$ are decreasing with S. Given any feasible allocation s and the corresponding $S=\sum_{i=1}^{w} s_i$, the following must hold:

$$\begin{array}{l} f_3(S) \leq f_3(MIN) \\ t_3(S) \leq t_3(MIN) \\ f_2(s) + f_3(S) \geq \min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{f_2(\overline{s})\} + f_3(MAX) \\ t_2(s) + t_3(S) \geq \min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{t_2(\overline{s})\} + t_3(MAX) \end{array} \tag{16}$$

Therefore, the values $x = \min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{f_2(\overline{s})\} + f_3(MAX) - f_3(MIN)$ and $y = \min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{t_2(\overline{s})\} + t_3(MAX) - t_3(MIN)$ guarantee that constraints $x \leq f_2(s)$ and $y \leq t_2(s)$ are satisfied by any allocation s such that $MIN \leq \sum_{i=1}^{w} s_i \leq MAX$. In what follows, we fix x and y to these values and omit the two latter constraints from the formulation of Problem $P_3(\gamma)$, which can be written as follows:

Problem $P_3(\gamma, MIN, MAX)$:

$$\begin{array}{ll} \min & f_1(S) + \min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{f_2(\overline{s})\} + f_3(MAX) - f_3(MIN) + f_3(S) + \\ & \gamma \left(\min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{t_2(\overline{s})\} + t_3(MAX) - t_3(MIN) + t_3(S) - \frac{(1-L)MTBF}{L} \right) \\ s.t.: & \\ & MIN \leq S \leq MAX \end{array} \tag{17}$$

Since $f_1(S)$ is linearly increasing while $f_3(S)$ and $t_3(S)$ are convex and decreasing with S, the objective function of Problem $P_3(\gamma, MIN, MAX)$ is convex for any given $\gamma \geq 0$. Therefore, given the values x and y, the optimal S can be efficiently computed by using a binary search approach in the interval $[MIN, MAX]$. We compute x as in [18]. The next theorem shows that computing y is not necessary in order to compute the maximum of $P_3(\gamma, MIN, MAX)$.

Theorem 4 *The value $\gamma = 0$ maximizes $P_3(\gamma, MIN, MAX)$.*

Proof. To prove the theorem it is sufficient to prove that quantity

$$\min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{t_2(\overline{s})\} + t_3(MAX) - t_3(MIN) + t_3(S) - \frac{(1-L)MTBF}{L}$$

is always non positive for $MIN \leq S \leq MAX$. This property follows by observing that $t_3(S) \leq t_3(MIN)$ and that a feasible solution exists for $S = MAX$, i.e.,

$$\min_{\overline{s}:\sum_{i=1}^{w} \overline{s}_i = MAX} \{t_2(\overline{s})\} + t_3(MAX) \leq \frac{(1-L)MTBF}{L}.$$

□

6 Solution procedure

In this section, a branch-and-bound algorithm for finding an optimal allocation of spares to warehouses is described. At each node of the enumeration tree the lower bound is computed by solving $P_3(0, MIN, MAX)$, where MIN and MAX are computed according to equations (10) and (9) at the root node and then updated by the branching rule. The heuristic algorithm described in subsection 6.1 provides an initial upper bound UB, then updated whenever a new feasible solution is found, and a sketch of the branch and bound procedure is presented in subsection 6.2.

6.1 Upper-bound computation

A simple upper bound to Problem P_0 is computed by distributing spare parts among warehouses with positive demand and by giving preference to warehouses with larger demand. In fact, simulation experiments carried out in [4] show that avoiding concentration of spares in few warehouses is an effective allocation policy. The heuristic

Procedure ISA
set $S = 0$ and $s_i = 0, \quad i = 1, \ldots, w$.
set $rhs = \frac{(1-L)MTBF}{L}$ and $k = 1$.
repeat
  repeat
    set $S = S + 1$, $s_k = s_k + 1$ and $k = k + 1$;
    if ($k = |\overline{W}|$ **then** set $k = 1$;
  until $\sum_{i=1}^{w} [\frac{\lambda_i}{\Lambda} \sum_{j \in W} \hat{\pi}_{ij}(s) T_{ji}] + P_B(S)(MET) \leq rhs$
  if $\sum_{i=1}^{w} [\frac{\lambda_i}{\Lambda} \sum_{j \in W} \pi_{ij}(s) T_{ji}] + P_B(S)(MET) > \frac{(1-L)MTBF}{L}$
  then $rhs = rhs - \sum_{i=1}^{w} \left[\frac{\lambda_i}{\Lambda} \sum_{j \in W} (\pi_{ij}(s) - \hat{\pi}_{ij}(s)) T_{ji}\right]$
until $\sum_{i=1}^{w} [\frac{\lambda_i}{\Lambda} \sum_{j \in W} \pi_{ij}(s) T_{ji}] + P_B(S)(MET) \leq \frac{(1-L)MTBF}{L}$
return s.

Figure 2: Pseudocode of the heuristic for Initial Spares Allocation

procedure ISA (Initial Spares Allocation), sketched in Figure 6.2, finds an allocation s, feasible for P_0, by greedily allocating one spare at a time to warehouses in the set $\overline{W} = \{i \in W : \lambda_i > 0\}$. Without loss of generality we assume that the warehouses are numbered for decreasing value of λ_i, i.e., $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_{|\overline{W}|}$. ISA terminates when the quantity $\sum_{i=1}^{w} [\frac{\lambda_i}{\Lambda} \sum_{j \in W} \pi_{ij}(s) T_{ji}] + P_B(S)(MET)$ becomes smaller than $\frac{(1-L)MTBF}{L}$. In order to speed up the computation of state probabilities $\pi_{ij}(s)$ at each step of the procedure, the heuristic computes approximate values $\hat{\pi}_{ij}(s)$, by using estimated values for the state probabilities of the associated Markov chain. Such values are computed with the fast multi-dimensional scaling down method described in [3]. The main idea of this method is to replace the original Markov chain with an equivalent one with a smaller number of states but with similar behavior in terms of operational availability. The reduced Markov chain is obtained by scaling the demand, the replenishment time and the stock level of each warehouse using a *scale factor* K. When

$$\sum_{i=1}^{w} [\frac{\lambda_i}{\Lambda} \sum_{j \in W} \hat{\pi}_{ij}(s) T_{ji}] + P_B(S)(MET) \leq \frac{(1-L)MTBF}{L} \tag{18}$$

in the reduced chain, the feasibility of allocation s is checked by solving the original Markov chain exactly. In case of a feasible solution, Procedure ISA stops and returns the feasible allocation s, otherwise the constraint (18) is strengthened by replacing the right-hand side $\frac{(1-L)MTBF}{L}$ with the smaller value

$$\frac{(1-L)MTBF}{L} - \sum_{i=1}^{w}\left[\frac{\lambda_i}{\Lambda}\sum_{j\in W}\left(\pi_{ij}(s)-\hat{\pi}_{ij}(s)\right)T_{ji}\right].$$

Procedure ISA then continues allocating one spare at a time and checking feasibility with the multi-dimensional scaling down method until a new apparently feasible solution is found. The procedure stops when the first feasible solution is found. In the next Section, we compare ISA with one of the most effective greedy algorithms for spare allocation, by Wong et al. [20]. The greedy heuristic of Wong et al. [20] consists of a first main step in which a feasible solution is generated by adding one unit of stock at a time. When a feasible solution is found, the algorithm tries to improve it by local search.

6.2 Branch and bound algorithm

Our BB (Branch and Bound) procedure maintains a queue $\mathcal{Q}$ of intervals $[MIN, MAX]$ for the stock level S, each corresponding to an instance of P_1.

Procedure ISA provides an initial solution $Bestsol$, from which the first upper bound UB on the optimum is derived. At the root node, $\mathcal{Q}$ is initialized with one open problem in which MIN and MAX are computed according to (10) and (9).

At each iteration of the BB procedure an open problem is removed from $\mathcal{Q}$ according to First In First Out rule and an optimal solution S^* to $P_3(0, MIN, MAX)$ is computed. If the lower bound $P_3(0, MIN, MAX) \geq UB$ the problem is closed. Otherwise, an allocation $s^* = \text{argmin}\{t_2(s) : \sum_{i=1}^{w} s_i = S^*\}$ is computed as in [18].

If $t_2(s^*)+t_3(S^*) \leq \frac{(1-L)MTBF}{L}$, then s^* is feasible for P_0 and, in view of assumption (7), it is also an optimal allocation for the restricted version of P_1 in which $S = S^*$. In this case, two new open problems are added to $\mathcal{Q}$ with $MIN \leq S \leq S^*-1$ and $S^*+1 \leq S \leq MAX$, respectively, and the upper bound UB is updated if $f_1(S^*) + f_2(s^*) + f_3(S^*) < UB$.

If $t_2(s^*) + t_3(S^*) > \frac{(1-L)MTBF}{L}$, then s^* is infeasible for P_1. Thus, for all values $MIN \leq S \leq S^*$ there is no feasible solution to P_0 and only the open problem with $S^* + 1 \leq S \leq MAX$ is added to $\mathcal{Q}$. The procedure terminates when $\mathcal{Q}$ is empty and the current allocation $Bestsol$ is an optimal solution to P_0.

7 Case study from the corrective airport maintenance context

In this section, we report on our computational experiments with the algorithms for spares allocation presented in section 6 applied to solve the practical problem from the airport maintenance context described in the introduction of the paper. The case study originates from the practical needs of an Italian logistics company supporting the activity of 38 civil airports spread over the Italian territory. The company manages the overall processes of purchasing, holding and replacing failed items, ensuring that the overall reliability of safety equipments is always within contractual limits. The aim of the company is therefore to grant the prescribed quality of service at minimum cost. While the company currently follows a two echelon policy for spare part management, the company managers

Procedure BB

Find an allocation $BestSol = ISA$
set $UB = f_1(BestSol) + f_2(BestSol) + f_3(BestSol)$
$MIN = \min\left\{S : t_3(s) \leq \frac{(1-L)MTBF}{L}\right\}$
$MAX = \min\{S : f_1(S) + f_3(S) \geq UB\}$
push $[MIN, MAX]$ in queue $\mathcal{Q}$
while $\mathcal{Q} \neq \emptyset$ **do**
pop $[x, y]$ from $\mathcal{Q}$
if $P_3(0, x, y) < UB$ **then**
let S^* be the optimal solution to $P_3(0, x, y)$
set $s^* = \text{argmin}\{t_2(s) : \sum_{i=1}^{w} s_i = S^*\}$
if $t_2(s^*) + t_3(S^*) \leq \frac{(1-L)MTBF}{L}$ **then**
if $f_1(S^*) + f_2(s^*) + f_3(S^*) < UB$ **then**
set $BestSol = s^*$ and $UB = f_1(S^*) + f_2(s^*) + f_3(S^*)$
set $y = \min\{y; \min\{S : f_1(S) + f_3(S) \geq UB\}\}$
end if
push $[x, S^* - 1]$ in $\mathcal{Q}$
push $[S^* + 1, y]$ in $\mathcal{Q}$
else
push $[S^* + 1, y]$ in $\mathcal{Q}$
end if
end if
end while
Return UB and $Bestsol$

Figure 3: Pseudocode of the BB algorithm

are interested in evaluating the potential benefits deriving from the adoption of a single echelon policy, which is generally acknowledged to achieve better performance in similar contexts. To this aim, the algorithms have been tested on a set of twelve instances from our case study, each based on the warehouse locations and demand rates of a particular item. However, in order to test the algorithms on a wider context than the real situation, we generated several scenarios by varying holding, transshipment and emergency costs of each item. For the replenishment time of an item we use the exponential distribution with average equal to three months for all items and scenarios while for the transshipment time and cost we use a deterministic value proportional to the distance between warehouses. Each pair holding-emergency cost defines a scenario for each of the twelve items. We consider 21 scenarios by choosing the cost of an item from the interval $[200, 1200]$ and fixing the emergency cost equal to 7000. The latter cost is a realistic one when emergency costs are the weighty ones, such as in contexts where safety is also involved. Fourteen additional scenarios are defined in order to analyze the influence of the emergency cost on the algorithms performance. In this second set of scenarios the item cost is fixed equal to 300 for each item while the emergency cost varies from 200 to 200000. Not all the emergency costs are realistic, since our purpose with the second set of scenarios is to investigate the effects of increasing emergency costs.

In total we obtain 420 instances. In Table 1 we summarize the values of the main parameters used in the computational experiment.

Parameter name	Unit	Values
Warehouses with positive demand		2, 3, 4, 5, 6, 7, 8, 9
Number of installed items		3, 5, 8, 9, 10, 11, 16, 18
Average MTBF	hours	16000, 17000, 26000, 38000, 61000, 79000, 81000, 94000, 101000, 132000, 191000, 200000
Holding cost	euros	$200, 250, 300, \ldots, 1000, 1150, 1200$
Emergency cost	euros	200, 300, ..., 700, 1000, 2000, ..., 5000, 7000, 50000, 100000, 200000
Min-Max average transshipment lead time	hours	$[5, 37.5]$
Emergency lead time	hours	2160

Table 1: Parameter values for the computational experiment

Specifically, we show the number of regional warehouses, which have some site with a failure process, which is directly coupled with it. Then, we show the range in which the number of items present in each single site varies in our computational experiments, the average MTBFs, the holding costs, the emergency costs, the minimum and the maximum transshipment lead time and the emergency lead time used in our test cases.

Table 2 reports on the comparison among the ISA heuristic and the heuristic proposed by Wong et al. [20] for all instances and scenarios. Since the latter makes use of a local search after the greedy phase, we compare four cases: ISA, the stand-alone greedy of [20] and both greedy algorithms followed by the local search described in [20].

We use the following notation: HEU is the cost generated by the considered heuristic algorithm, BB is the proven optimum, S_{HEU} is the number of spares allocated by HEU and S_{BB} is the number of spares allocated in the optimal solution found by the branch and bound algorithm. Column two in Table 2 reports the computation time (in seconds) of the four heuristics, column three shows the relative error $\frac{HEU-BB}{BB}$ of each heuristic and column four shows the relative difference in the number of allocated spares with respect to the branch and bound solution. Each row in table shows the average results obtained for all item by varying the costs over the 35 scenarios. We observe that the solutions found by ISA are slightly better with respect to the solutions found by the initial greedy of Wong et al. [20], and the computation time is slightly smaller. After the local search both algorithms have the same performance. However, the resulting values are still quite far from the optimum.

Method	Computation time	$\frac{HEU-BB}{BB}$	$\frac{S_{HEU}-S_{BB}}{S_{BB}}$
ISA	0.17	0.58	-0.36
ISA with local search	0.97	0.45	-0.14
Wong first step	0.25	0.69	- 0.33
Wong with local search	1.01	0.45	-0.13

Table 2: Performance of ISA, Wong et al. [20] first step, Wong et al. [20] approach with local search and ISA with local search algorithms for all the items and the scenarios

In tables 3, 4 and 5 we show the results obtained for the 12 instances and the 35 scenarios. Table 3 shows the solutions and the computation time (in seconds) of ISA and BB and the relative error of ISA, computed as the average over all instances of the values $\frac{ISA-BB}{BB}$. Each row in table shows the average results obtained for an item by varying the costs over the 35 scenarios. We observe that the BB algorithm is able to find the optimal solution within less than 100 seconds of computation for eight of the twelve items, while the optimum is found within approximately 30 minutes for other three items. ISA always finds a feasible solution within less than one second. The ISA value turns out to be the optimal solution for 72 out of 420 instances and the average error over the 35 scenarios varies in the range $[0.10, 0.96]$. These experiments show that ISA provides good solutions within short computation time, even if it is worth using the exact algorithm to find better solutions.

Item	# warehouses	BB cost values			ISA	$\frac{ISA-BB}{BB}$	Computation time	
		holding	transshipment	emergency	value		ISA	BB
1	2	1607.96	107.98	15.95	1915.79	0.11	0.15	0.82
2	3	1232.31	360.45	25.13	3167.76	0.96	0.17	0.94
3	4	1709.10	456.66	8.09	3086.06	0.43	0.15	1.40
4	4	2221.00	517.70	41.56	2794.14	0.05	0.15	2.00
5	4	1474.52	563.24	22.42	3101.16	0.60	0.15	1.39
6	5	1345.48	668.23	37.11	2436.62	0.20	0.16	1.80
7	6	1904.55	869.60	70.29	3060.88	0.10	0.19	30.80
8	6	1450.23	914.24	60.45	3588.61	0.50	0.16	68.09
9	7	1512.84	1095.64	22.60	2694.90	0.10	0.15	1743.62
10	7	1528.90	1086.95	28.10	2980.42	0.12	0.16	1867.27
11	8	2068.41	1350.92	15.43	3788.17	0.10	0.17	1047.60
12	9	1984.05	1300.86	38.01	4036.73	0.23	0.14	4494.06

Table 3: Performance of ISA and BB algorithms for the 12 items

In Table 4 we analyze the performance of ISA and BB for varying the holding cost of the items. Each row in table reports the average results over the 12 items for one of the first 21 scenarios. We also show the three components of the optimum cost, i.e., holding, transshipment and emergency cost. It can be observed that the transshipment cost is often comparable with the holding cost, and therefore it cannot be neglected in the solution of the problem. For a holding cost lower than 700 the number of spares allocated by ISA is smaller than the optimal value, while for holding costs higher than 350 ISA always allocates a number of spares optimal or slightly greater than the optimal one. In fact, we observe that the number of spares allocated by ISA does not depend on the spare holding cost and therefore the number of spares allocated by ISA is always the same for all scenarios. A consequence of this behavior is that the gap between ISA and the optimum is influenced by the holding cost. When the holding cost of an item increases from 200 to 500 the error decreases from 40% down to 5%. For larger holding costs the error increases regularly until 13%. Smaller errors are attained when the number of spares allocated by ISA is close to the optimal one and the holding cost is small. In these cases the error only depends on the warehouses to which spares are allocated. As the holding cost increases, the cost related to the different allocation becomes more and more relevant and the relative gap between ISA and BB increases.

Holding cost	Emerg. cost	BB cost values			ISA value	$\frac{ISA-BB}{BB}$	# spares	
		holding	transshipment	emergency			BB	ISA
200	7000	1050.00	225.57	2.40	1769.05	0.40	5.2	2.5
250	7000	1229.17	302.33	2.71	1898.21	0.30	4.9	2.5
300	7000	1250	511.14	3.38	2027.38	0.16	4.1	2.5
350	7000	1341.67	608.98	12.51	2156.55	0.10	3.8	2.5
400	7000	1433.33	702.55	13.23	2285.71	0.10	3.5	2.5
450	7000	1537.50	770.17	15.87	2414.88	0.10	3.4	2.5
500	7000	1541.67	914.79	29.99	2544.05	0.05	3.0	2.5
550	7000	1604.17	985.84	45.09	2673.21	0.05	2.9	2.5
600	7000	1700.00	1030.90	49.58	2852.38	0.05	2.8	2.5
650	7000	1733.33	1125.55	56.92	3112.70	0.06	2.5	2.5
700	7000	1866.67	1125.55	56.92	3260.71	0.07	2.5	2.5
750	7000	1937.50	1182.59	62.38	3389.88	0.07	2.5	2.5
800	7000	2066.67	1182.59	62.38	3539.05	0.07	2.5	2.5
850	7000	2195.83	1182.59	62.38	3668.21	0.07	2.5	2.5
900	7000	2325.00	1182.59	62.38	3877.38	0.08	2.5	2.5
1000	7000	2454.17	1182.59	62.38	4006.75	0.08	2.5	2.5
1050	7000	2583.33	1182.59	62.38	4164.88	0.08	2.5	2.5
1100	7000	2712.50	1182.59	62.38	4394.05	0.11	2.5	2.5
1150	7000	2841.67	1182.59	62.38	4623.21	0.13	2.5	2.5
1200	7000	2875	1182.59	62.38	4652.38	0.13	2.5	2.5

Table 4: Performance of ISA and BB for different holding costs

In Table 5 we analyze the performance of ISA and BB for varying the item emergency costs. Each row in table shows the average results over the 12 instances for one of the 15 scenarios. Similarly to the previous scenario, the number of spares allocated by ISA is the same for all scenarios since this value does not depend on the spare emergency cost. The gap between ISA and the optimum depends therefore on the emergency cost, even if the emergency cost has a smaller influence on the error of ISA with respect to the holding cost. When the emergency cost varies in the range $[200 - 2000]$, the error remains the same. In order to observe relevant errors, the emergency cost must increase up to more than 50000.

As a concluding observation, computational experiments show that the overall behavior of ISA is acceptable as an initial solution for subsequent optimization. In general, the performance of ISA depends on the specific holding and emergency costs being considered and, therefore, it may be quite erratic. BB algorithm seems to be more promising, since it finds the proven optimum within acceptable computation time for all tested instances.

Holding cost	Emerg. cost	BB cost values			ISA value	$\frac{ISA-BB}{BB}$	# spares	
		holding	transshipment	emergency			BB	ISA
300	200	1200	556.27	0.16	1966.78	0.12	4	2.5
300	300	1200	556.27	0.31	1967.67	0.12	4	2.5
300	400	1200	556.27	0.47	1968.56	0.12	4	2.5
300	500	1200	556.27	0.63	1969.46	0.12	4	2.5
300	600	1200	556.27	0.78	1970.35	0.12	4	2.5
300	700	1200	556.27	0.94	1971.24	0.12	4	2.5
300	1000	1200	556.27	1.57	1973.91	0.12	4	2.5
300	2000	1200	556.27	3.15	1982.82	0.12	4	2.5
300	3000	1200	556.27	4.72	1991.73	0.13	4	2.5
300	4000	1200	556.27	6.30	2000.64	0.14	4	2.5
300	5000	1250	511.14	2.41	2009.56	0.14	4.17	2.5
300	7000	1250	511.14	3.38	2027.38	0.15	4.17	2.5
300	50000	1250	511.14	24.12	2410.56	0.35	4.17	2.5
300	100000	1325	442.93	38.76	2856.13	0.60	4.42	2.5
300	200000	1400	406.72	15.40	3747.26	1.05	4.67	2.5

Table 5: Performance of ISA and BB for different emergency costs

8 Conclusions

In this paper we propose and evaluate a solution methodology for optimizing inventory stock allocation of repairable spare parts in a single echelon, w-locations system, where lateral and emergency shipments occur in response to stockouts. We model our problem as a non-convex integer program and develop a new heuristic and a new branch and bound algorithm for allocating the spare parts optimally. Both algorithms are evaluated by using practical data from the Italian airport corrective maintenance context. Computational experiments demonstrate that the branch and bound technique is able to optimally solve almost all tested instances within reasonable computation time. The heuristic algorithm finds sub-optimal solutions within very limited computation time, thus being a promising approach for finding feasible solutions to difficult instances.

Future research should address a deeper structural analysis of the functions studied in the paper, relaxing some of the assumptions made in this paper. For example, extending the BB algorithm to the case $MET < MTTR + OS$ would enlarge the range of applicability of the proposed technique. To this aim, the convexity of function $f_3(S)$ has to be investigated.

Future research should also address the development of faster exact methods and effective metaheuristics for the solution of large and difficult instances, as well as on the application of the ideas proposed in this paper to manage the maintenance of different critical infrastructures, such as medical equipments in hospitals or communication or energy distribution networks and so on.

References

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

[2] Axsater,S.. Modelling emergency lateral transshipments in inventory systems. *Management Science* 6, (1990) 13291338.

[3] Cesaro, A., Pacciarelli, D.. Performance assessment for single echelon airport spare part management. *Computers And Industrial Engineering*, in press (2011) doi: 10.1016/j.cie.2011.03.005.

[4] Cesaro, A., Pacciarelli, D.. Evaluation of peaked lateral transshipment in inventory system subject to a service constraint. *MSOM 2007 conference* (June 18–19,2007).

[5] Chiou, C.C.. Transshipment problems in supply chain systems: review and extensions. *Supply Chains, Theory and Application, Vedran Kordic* (2008).

[6] Fisher, M.L. The Lagrangian relaxation method for solving integer programming problems. *Management Science* 27, (1981).

[7] Grahovac, J., A. Chakravarty. Sharing and Lateral Transshipment of Inventory in a Supply Chain with Expensive Low-Demand Items. *Management Science* 47, (2001) 579–594

[8] Huiskonen, J.. Maintenance spare parts logistics: special characteristics and strategic choices. *Int. J. Production Economics* 71, (2001) 125–133.

[9] Kemeny, J., J. L. Snell. Finite Markov Chains. *D. Van Nostrand*, (1960).

[10] Kennedy, W.J., J. Wayne Patterson, and L. D. Fredendall. An overview of recent literature on spare parts inventories. *Int. J. Production Economics* 76, (2002) 201–215.

[11] Kranenburg, A.A., Van Houtum, G.J.. A multi-item spare parts inventory model with customer differentiation. *Working Paper WP 110, Research School Beta, Technische Universiteit Eindhoven* (2004).

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

[13] Kukreja, A., C.P. Schmidt, D.M. Miller. Stocking decisions for low-usage items in a multilocation inventory system *Management Science* 47, (2001) 1371–1383.

[14] Kutanoglu, E., M. Mahajan. An inventory sharing and allocation method for a multi location service parts logistics network with time based service levels. *European Journal of Operational Research* 194, (2009) 728-742.

[15] Paterson, C., G. Kiesmuller, R. Teunter, K. Glazebrook. Inventory models with later transshipments: a review. *European Journal of Operational Research* 210, (2011) 125-136.

[16] Sherbrooke, C.C.. Optimal inventory modeling of systems: multi-echelon techniques. *Wiley* (2004).

[17] Taragas, G., D. Vlachos. Effectiveness of stock transshipment under various demand distributions and nonnegligible transshipment times. *Production And operations Management* 11, (2002) 183–198.

[18] Wong, H.,D. Cattrysse,D. Van Houdeusden, Inventory pooling of repairable spare parts with non-zero lateral transshipment time and delayed lateral transshipments. *European Journal of Operational Research* 165 (2005) 207–218.

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

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