Constant approximation algorithms for the one warehouse multiple retailers problem with backlog or lost-sales

J.-P. Gayon[a,*], G. Massonnet[b], C. Rapine[c], G. Stauffer[a]

[a]*Laboratoire G-SCOP, 46, avenue Flix Viallet, 38031 Grenoble Cedex 1, France*
[b]*INRIA Grenoble Rhône-Alpes, 655 Avenue de l'Europe, 38330 Montbonnot-Saint-Martin, France*
[c]*Université de Lorraine, laboratoire LGIPM, Ile du Saulcy, 57045 Metz Cedex 01, France*

Abstract

We consider the One Warehouse Multi-Retailer (OWMR) problem with deterministic time-varying demand in the case where shortages are allowed. Demand may be either backlogged or lost. We present a simple combinatorial algorithm to build an approximate solution from a decomposition of the system into single-echelon subproblems. We establish that the algorithm has a performance guarantee of 3 for the OWMR with backlog under mild assumptions on the cost structure. In addition, we improve this guarantee to 2 in the special case of the Joint-Replenishment Problem (JRP) with backlog. As a by-product of our approach, we show that our decomposition provides a new lower bound of the optimal cost. A similar technique also leads to a 2-approximation for the OWMR problem with lost-sales. In all cases, the complexity of the algorithm is linear in the number of retailers and quadratic in the number of time periods, which makes it a valuable tool for practical applications. To the best of our knowledge, these are the first constant approximations for the OWMR with shortages.

Keywords: Approximation algorithms, lot-sizing, inventory control,

*Corresponding author : jean-philippe.gayon@grenoble-inp.fr, phone : +33 4 76 57 47 46, fax : +33 4 76 57 46 95

Email addresses: jean-philippe.gayon@grenoble-inp.fr (J.-P. Gayon), guillaume.massonnet@inria.fr (G. Massonnet), christophe.rapine@univ-lorraine.fr (C. Rapine), gautier.stauffer@genoble-inp.fr (G. Stauffer)

distribution systems.

---

1. Introduction

We consider two multi-echelon inventory control problems: The *One Warehouse Multi-Retailer (OWMR)* problem and its special case the *Joint Replenishment Problem (JRP)*. Both problems share the same divergent network structure: A single warehouse replenishes N retailers facing customers demands, by ordering material from an external suppliers. Demands are fulfilled by units that are first ordered at the warehouse, then at the retailers. In particular, the divergent structure of the system requires that the retailers can only order material that is available (i.e. physically held) at the warehouse. The JRP is a special case of this problem in which the warehouse cannot hold any unit. In this work, we focus on a discrete and deterministic setting by assuming that the demands are known over a finite planning horizon of T periods. Costs are incurred when a location places an order to replenish its stock (*ordering costs*), when units are physically held in the system (*holding costs*) or when demand is not immediately satisfied (*penalty costs*). The goal is then to determine an ordering strategy in order to minimize the total cost incurred to fulfill the demands by moving the products through the network.

Most of the existing work on this problem focuses on the case where all demands must be satisfied on time. Even in this simpler setting the OWMR problem and the JRP are both known to be NP-hard (Arkin et al. (1989)). Those problems have attracted a lot of attention in the past and many heuristics have been developped. In particular Levi et al. (2008), Stauffer et al. (2011) and Bienkowski et al. (2014) introduced constant factor approximation algorithms for the former problem (1.8, 2, 1.791 respectively). We refer to these papers for a detailed review of the corresponding literature. Nonner and Souza (2009) have also proposed a 5/3-approximation algorithm for the JRP problem, for the special case where each demand has a strict deadline.

In this paper, we focus on extensions of these standard models in which shortages (either backlog or lost sales) are allowed. Such models have been extensively studied in the lot-sizing literature when a single location faces customers demands over a discrete, finite planning horizon. Zangwill (1966) was among the first to extend the pioneering work of Wagner and Whitin (1958) to incorporate the possibility to backlog units. Later, Federgruen and

Tzur (1993) and Aggarwal and Park (1993) used more advanced dynamic programming techniques to solve the problem and improved the time complexity. An interesting extension of the lot-sizing problem is the multi-item version, which considers that several items are distributed, each incurring its own holding and backlogging costs. Several papers (see Pochet and Wosley (1988, 1994) and Küçükyavuz and Pochet (2009)) have focused on the integer programming formulation of this problem in the case of single or multiple items. The special case of the JRP is a generalization of the multi-item problem in which additional ordering costs are incurred whenever a specific item replenishes its own stock. Levi et al. (2006) developed an algorithm based on primal-dual approach for this problem with backlogging and proved its cost is at most twice the optimal cost in the worst case.

The literature on lost sales models is more recent: Sandbothe and Thompson (1990) were the first to propose a forward algorithm to solve a lot-sizing problem with production capacity constraints and stockouts. Later, Aksen et al. (2003) introduced a dynamic programming approach to solve efficiently the uncapacitated version with time-varying costs. Liu et al. (2007) considered a model where inventory is bounded, present an optimal dynamic program for this model and test it on industrial problem successfully. Other papers deal with the multi-item versions of the basic lot-sizing problem. In particular, Absi and Kedad-Sidhoum (2008) extends the original work of Sandbothe and Thompson (1990) to a multi-item version and use a MIP approach to develop an effective method to find near optimal solutions.

Our contributions

This work extends the decomposition technique of Stauffer et al. (2011) to more general models allowing backorders or lost-sales and derive a constant approximation algorithm in each case. We present a new lower bound for each problem, based on a decomposition into simple single-echelon systems. We then recompose the optimal solutions to the subproblems into a feasible solution to the original problem. In the case of the OWMR problem with backlogging, we prove that the cost of the solution obtained is guaranteed to be at most three times the optimal cost under mild assumptions on the cost structure. In addition in the JRP case, some of these assumptions become irrelevant and we show that our algorithm can be modified to match the best known performance guarantee obtained by Levi et al. (2006) when backorders are allowed. Finally, we adapt our technique to the OWMR problem with lost-sales and build a 2-approximation for this problem. To the best of our knowledge, this paper presents the first constant approximation algorithms for the backlogging and the lost-sales version of the OWMR problem. The complexity of our algorithms is linear with the number of time periods and quadratic with the number of retailers.

The remainder of this paper is organized as follows. In §2, we formally introduce the assumptions and notations used throughout the following sections. In §3, we present the split and uncross technique and provide a 3-approximation algorithm for the backorder version of the OWMR problem. In §4, we show that in our cost structure, the lost-sales model is simply a special case of the backlogging problem. We are then able to improve the performance guarantee of our algorithm to two by simply modifying one step of our algorithm.

2. Assumptions and cost structure for the backlogging model

In this section, we focus on models where unmet demand is backordered. In this version of the OWMR problem, demands faced by the retailers are not necessarily satisfied on time but can instead be backlogged and served by an order placed later in time. We now discuss the assumptions and notations used troughout the remainder of the paper. In particular, we present a general cost structure that extends the cost structure introduced in Levi et al. (2008) and Stauffer et al. (2011).

We consider N retailers that face customers demands over a finite planning horizon, discretized into T periods. For each retailer $i = 1, \dots, N$ and each period $t = 1, \dots, T$, let d_t^i be the deterministic demand for retailer i in period t, to which we also refer as the *demand point* (i, t). Each retailer orders units from a central warehouse, which in turn orders from an external supplier. Recall that we assume that the decisions are centralized in order to minimize the total cost incurred by the system. As a consequence, orders placed by the retailers are always filled on time by units coming from the on-hand inventory of the warehouse. If the warehouse orders from its supplier in period t, it incurs a fixed ordering cost K_t^0, regardless of how many units are ordered. Similarly, retailer i pays a fixed ordering cost K_t^i for ordering from the warehouse in period t. Chan et al. (2000) have shown that if the ordering costs at the retailers vary over time, the OWMR problem is as hard to approximate as the set cover problem and thus admits no constant guarantee unless $P = NP$ (Feige (1998)). Therefore we assume in this paper that the ordering costs at each retailer $i > 0$ are stationary, i.e. $K_t^i = K^i$ for all periods $t = 1, \ldots, T$. The leadtimes are deterministic, thus we assume without loss of generality that the orders are delivered instantaneously from one location to another.

The most common assumption in the inventory literature is to consider that holding and backlogging units induce linear costs that are proportional to the inventory level in each location. However, it is worth noticing that the algorithms developped in the following sections only need to satisfy weaker properties to yield a constant performance guarantee for the OWMR problem with backorders. For clarity reasons, we start by introducing the assumptions on the cost structure in the linear case first. We then relate this basic parameters to the so-called metric carrying cost structure and generalize the assumptions to this more complex setting.

The linear cost structure

In the traditional linear cost structure, each unit physically held in location i in period t incurs a holding cost $h_t^i \geq 0$, while each backlogged unit for a specific retailer $i \geq 1$ induces a penalty cost of $b_t^i \geq 0$. Note that we consider that the penalty cost for backlogging a demand is entirely incurred at the retailer, where demands have to be served eventually. In other words, we restrict ourselves to policies in which the warehouse cannot backlog the orders of the retailers, i.e. retailers can only order units that are available at the warehouse. Another common assumption for the OWMR problem is that the set of retailers can be partitioned into two subsets I_J and I_W such that:

$$I_J = \left\{ i = 1, \ldots, N : h_t^i \leq h_t^0 \text{ for all } t \right\} \tag{1}$$

$$I_W = \left\{ i = 1, \ldots, N : h_t^i \geq h_t^0 \text{ and } b_t^i \geq h_t^0 \text{ for all } t \right\} \tag{2}$$

In the remainder of this paper, retailers in set I_J are also called J-retailers, while retailers from set I_W are called W-retailers. In simple words, it is cheaper to store units at J-retailers than at the warehouse, while it is cheaper to hold units at the warehouse than at W-retailers. Note that in addition we assume that for all $i \in I_W$ and $t = 1, \ldots, T$, it is also cheaper to hold unit at the warehouse rather than backlogging it at retailer i in period t. In other words, the backlogging cost of any W-retailer is greater than the holding cost at the warehouse. This assumption matches many practical situations, since backlogging units is often more expensive than holding them in the stock, which is in turn more expensive than holding them in the central warehouse in the case of W-retailers.

The (shelf-age) carrying cost structure

We now introduce a more general cost structure, called *carrying cost structure*, that encapsulate both holding and backlogging costs in a single notation. These parameters extend the shelf age dependent holding costs introduced in Levi et al. (2008) and Stauffer et al. (2011), which generalize the classical level-dependent holding costs usually considered in the literature (see Federgruen and Wang (2013) for further discussion on these cost structures). Namely, a unit ordered in period r at the warehouse and in period $s \geq r$ at retailer i (we denote by $\lceil r, s \rfloor$ such a pair of orders) to satisfy a unit of demand d_t^i incurs a per-unit carrying cost of ϕ_{rs}^{it}. That is, ϕ_{rs}^{it} encapsulates the cost for holding one unit at the warehouse from r to s, then either holding it at retailer i from period s to t (if $s \leq t$) to serve demand in period t, or backlogging one unit of demand (i,t) from period t to s (if $s > t$). Note that this cost structure can capture additional phenomena than the traditional linear holding/backlogging costs. For instance, it allows us to consider situations in which the holding cost incurred for a specific item at the central warehouse depends on which retailers it replenishes. It can also capture additional phenomena such as perishability, with a prohibitive holding cost if the period of storage in a location is longer than a certain threshold.

Clearly, linear cost parameters are a special case of carrying cost parameters: Indeed, for each demand point (i,t), each ordering period r of the warehouse and each ordering period $s \geq r$ of retailer i, one can define the total carrying cost incurred to serve demand (i,t) with the pair of orders $\lceil r, s \rfloor$ as follows:

$$\phi_{rs}^{it} = \begin{cases} \left(\sum_{u=r}^{s-1} h_u^0 + \sum_{v=s}^{t-1} h_v^i\right) d_t^i & \text{if } s \leq t \\ \left(\sum_{u=r}^{s-1} h_u^0 + \sum_{v=t}^{s-1} b_v^i\right) d_t^i & \text{otherwise} \end{cases} \tag{3}$$

As already mentioned, the carrying cost parameters do not need to satisfy equation (3) for us to derive constant approximations for the backlogging model we consider. In fact, we simply assume that parameters ϕ_{rs}^{it} satisfy the following five properties:

(P1) *Non-negativity.* The parameters ϕ_{rs}^{it} are nonnegative.

(P2) *Piecewise monotonicity with respect to s.* Every retailer i is in exactly one of the two following situations: Either ϕ_{rs}^{it} is non-increasing in $s \in [r,t]$ and non-decreasing in $s \in [t,T]$ for each demand point (i,t) and warehouse order r, or ϕ_{rs}^{it} is non-decreasing in $s \in [r,T]$ for each demand point (i,t) and warehouse order r. This property defines a partition of the set of retailers into two subsets: I_W and I_J, respectively.

(P3) *Monotonicity with respect to r.* For each retailer $i = 1, \dots, N$, each demand point (i,t) and retailer order in period $1 \leq s \leq T$, ϕ^{it}_{rs} is non-increasing in $r \in [1,s]$. Moreover if $i \in I_J$, we have: $\phi^{it}_{rr} \geq \phi^{it}_{r'r'}$ for $r \leq r' \leq t$ and $\phi^{it}_{rr} \leq \phi^{it}_{r'r'}$ for $t \leq r \leq r'$.

(P4) *Triangle inequality.* For each demand point (i,t) with $i \in I_W$, we have

$$\phi^{it}_{rs} \leq \begin{cases} \phi^{it}_{rt} + \phi^{it}_{ss} & \text{if } r \leq s \leq t \\ \phi^{it}_{rt} + \phi^{it}_{ts} & \text{if } r \leq t < s \\ \phi^{it}_{rr} + \phi^{it}_{ts} & \text{otherwise} \end{cases}$$

(P5) *Backlogging cost bounding.* If $i \in I_W$, for each demand point (i,t) and retailer i order in period $s > t$, we have $\phi^{it}_{ts} \leq 2\phi^{it}_{ss}$

In the remainder of the paper, we refer to the *metric* carrying cost structure when the parameters satisfy the five properties introduced above. It is straightforward to check that properties (P1), (P3) and (P4) capture the linear per-unit holding/penalty cost structure. Moreover if the set of retailers is partitioned into two subsets I_W and I_J as defined by (1) and (2), properties (P2) and (P5) are also satisfied and therefore the linear holding/penalty cost structure is a particular case of the metric carrying cost structure. Notice that property (P5) generalizes the inequality $b^i_t \geq h^0_t$ for all t and $i \in I_W$ to the case of carrying cost parameters. As already discussed, these assumptions cover many other practical cases.

In the remainder of this paper, we call a solution π to our problem a *policy*. Any policy π for the problem is represented by a $(N+1)$-uplet $\pi = (\pi_0, \pi_1, \dots, \pi_N)$ where each π_i is the set of pairs (ordering period, quantity) for location i. A policy π is *feasible* for the OWMR problem if the two following conditions are satisfied:

Condition 1. For all $i = 1, \dots, N$ policy π_i orders at least the sum of the demands d^i_t over the entire planning horizon.

Condition 2. Let r and r' be two consecutive ordering periods of the warehouse. Then the quantities ordered by π_0 in period r are sufficient to serve the orders placed by all the retailers in periods $r, \dots, r'-1$.

Given a policy π for the entire system, we denote $\mathscr{C}(\pi)$ the total cost incurred by π over the planning horizon. This cost can be split into two parts: The total ordering cost denoted $\mathcal{K}(\pi)$ and the total carrying (i.e. holding/backlogging) cost denoted $\Psi(\pi)$ and thus we have:

$$\mathscr{C}(\pi) = \mathcal{K}(\pi) + \Psi(\pi)$$

We conclude this section by pointing out three dominant properties on any feasible policy for the problem with the metric carrying cost structure. First, it is easy to prove that there exists an optimal policy $\pi^{OPT} = (\pi_0^{OPT}, \pi_1^{OPT}, \ldots, \pi_N^{OPT})$ such that for all $i = 1, \ldots, N$, π_i^{OPT} orders only when its inventory level is nonpositive. In addition, the metric carrying cost structure ensures that there exists an optimal policy in which each demand d_t^i is served from a unique pair of orders $\lceil r, s \rfloor$, which is a classical dominant property in most of the inventory models. This allows us to incorporate the amount d_t^i directly into the cost parameters by setting $\psi_{it}^{rs} \equiv \phi_{rs}^{it} d_t^i$, where ϕ_{rs}^{it} is a per-unit cost satisfying properties (P1)-(P5). Since $d_t^i \geq 0$ for all i, t, parameters ψ_{rs}^{it} also satisfy properties (P1)-(P5) and correspond to the total carrying cost incurred to serve the entire demand d_t^i from the pair of orders $\lceil r, s \rfloor$. Therefore in the remainder of the paper we use parameters ψ_{rs}^{it} to prove our approximation results. Finally, it is clear from property (P2) that there exists an optimal policy in which all J-retailers orders are synchronized with a warehouse order. In particular, we shall assume w.l.o.g. that the optimal policies considered in the following sections satisfy this property.

3. The split and uncross technique

Basically, the algorithm works in two steps: First decompose the original problem into several single-echelon subproblems, then solve each of them and recompose the resulting policies into a feasible policy for the OWMR problem. Each sub-problem is a relaxation of the OWMR problem where we remove the requirement that warehouse and retailer orders are synchronized but both just need to match the demand.

3.1. Phase 1: Decomposition of the OWMR problem

The first step of our algorithm is to decompose a general OWMR problem with backlog into several single-echelon subproblems. In addition, we show that by splitting the carrying cost between the different subproblems, we can derive a new lower bound for the original problem.

($\widehat{S}_0$) The warehouse is regarded as a single-echelon, multi-item system with backlogging facing for each period t a demand d_t^i for item $i = 1, \ldots, N$. A fixed ordering cost K_r^0 is incurred for placing an order in period r. If a demand (i, t) is ordered in period $r \leq t$, it incurs a holding cost of $\frac{1}{3}\psi_{rt}^{it}$ if $i \in I_W$ and $\frac{1}{2}\psi_{rr}^{it}$ if $i \in I_J$. On the other hand if $r > t$, it incurs a backlogging cost of $\frac{1}{3}\psi_{rr}^{it}$ if $i \in I_W$ and $\frac{1}{2}\psi_{rr}^{it}$ if $i \in I_J$.

($\widehat{S}_i$) Retailer i is considered as a single-echelon system with backlog facing demand d_t^i with ordering cost K^i. The carrying cost incurred to order in period s to serve the demand in period t is equal to $\frac{2}{3}\psi_{ss}^{it}$ if $i \in I_W$ and $\frac{1}{2}\psi_{ss}^{it}$ if $i \in I_J$.

Problem ($\widehat{S}_0$) and problems ($\widehat{S}_i$) are all equivalent to single-echelon lot-sizing problems with backlog. However, note that the first one corresponds to the multi-item version of this problem. These problems have been widely studied in the literature and it is well-known that finding their optimal solution can be reduced to finding a shortest path in a graph $G_j = (V_j, E_j)$ representing problem ($\widehat{S}_j$) (see Figure 1). For all $j = 0, \ldots, N$, $V_j = \{0, \ldots, T+1\}$ is the set of periods, where 0 and $T+1$ are two artificial periods representing the beginning and the end of the planning horizon. Given two consecutive orders u and v and $t \in [u, v]$, properties (P2) and (P3) ensure that it is suboptimal to serve demand d_t^i with an order placed earlier than u or later than v. As a consequence for $i = 1, \ldots, N$ an arc $(u, v) \in E_i$ represents two consecutive orders in a solution and its length is equal to the minimal possible cost incurred for placing an order in period u and carrying the units necessary to serve demands $d_u^i, \ldots, d_{v-1}^i$ by ordering them either in period u

or in period v. Therefore its length $l^i_{u,v}$ is equal to:

$$l^i_{u,v} = \begin{cases} \sum_{t=1}^{v-1} \frac{2}{3}\psi^{it}_{vv} & \text{if } u=0 \text{ and } i \in I_W \\ \sum_{t=1}^{v-1} \frac{1}{2}\psi^{it}_{vv} & \text{if } u=0 \text{ and } i \in I_J \\ K^i + \sum_{t=u}^{v-1} \frac{2}{3}\min\left\{\psi^{it}_{uu}, \psi^{it}_{vv}\right\} & \text{if } u>0 \text{ and } i \in I_W \\ K^i + \sum_{t=u}^{v-1} \frac{1}{2}\min\left\{\psi^{it}_{uu}, \psi^{it}_{vv}\right\} & \text{otherwise} \end{cases}$$

The definition of an arc $(u,)$, $u \leq v$, is similar for the warehouse, except that in this case the system has to serve demands of all the items between two consecutive orders in periods u and v. Its length $l^0_{u,v}$ is then set to:

$$l^0_{u,v} = \begin{cases} \sum_{t=1}^{v-1} \left(\sum_{i \in I_W} \frac{1}{3}\psi^{it}_{vt} + \sum_{i \in I_J} \frac{1}{2}\psi^{it}_{vv} \right) & \text{if } u=0 \\ K^0_u + \sum_{t=u}^{v-1} \left(\sum_{i \in I_W} \frac{1}{3}\min\left\{\psi^{it}_{ut}, \psi^{it}_{vv}\right\} + \sum_{i \in I_J} \frac{1}{2}\min\left\{\psi^{it}_{uu}, \psi^{it}_{vv}\right\} \right) & \text{otherwise} \end{cases}$$

Note that these lengths are computed in time $O(T^2)$ for the single-item problem and $O(NT^2)$ for the multi-item case. Moreover, from the nonnegativity of parameters ψ^{it}_{rs}, they are nonnegative and therefore one can easily find a shortest path from node 0 to $T+1$ in graph G_j using classical shortest path algorithms. Since graph G_j is acyclic and has $O(T^2)$ arcs, a topological sorting or breadth-first search algorithm can compute the shortest path in time $O(T^2)$. Hence, finding an optimal policy for the single-echelon problems can be achieved in time complexity $O(NT^2)$ for $(\widehat{S}_0)$ and $O(T^2)$ for each $(\widehat{S}_i)$, $i = 1, \ldots, N$.

We denote $\hat{\pi}_j$ a feasible policy for system $(\widehat{S}_j)$. Note that if $j = 0$, $\hat{\pi}_0$ is a set of $(N+1)$-uplet (r, q_r), where $r \in \{1, \ldots, T\}$ and q_r is the N-uplet of quantities ordered for each item in period r. For all j, we denote $\widehat{\mathscr{C}}_j(\hat{\pi}_j)$ the total cost incurred by policy $\hat{\pi}_j$ in system $(\widehat{S}_j)$. In the same fashion as the one used for the global cost, it can be split between two components:

$$\widehat{\mathscr{C}}_j(\hat{\pi}_j) = \mathcal{K}_j(\hat{\pi}_j) + \widehat{\Psi}_j(\hat{\pi}_j)$$

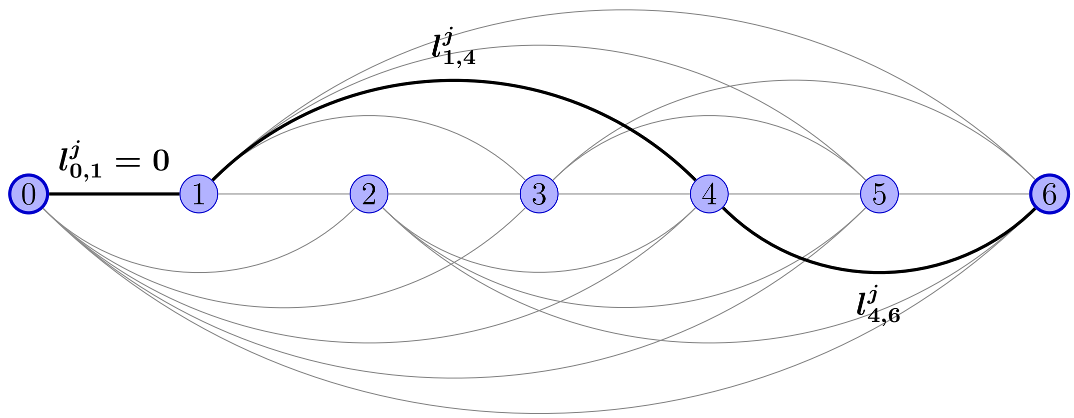

Figure 1: An example of graph modelizing problem $(\widehat{S}_j)$ with $T = 5$. The bold path corresponds to the solution in which the policy orders in periods 1 and 4, for a total cost of $l^j_{0,1} + l^j_{1,4} + l^j_{4,6}$.

where $\mathcal{K}_j(\hat{\pi}_j)$ (resp. $\widehat{\Psi}_j(\hat{\pi}_j)$) denotes the total ordering (resp. carrying) cost incurred by policy $\hat{\pi}_j$ in system $(\widehat{S}_j)$.

To conclude this section, we introduce a new lower bound for the optimal solution of the OWMR problem with backlog using the decomposition discussed above. Assume that we compute an optimal policy $\hat{\pi}^*_j$ to problem $(\widehat{S}_j)$ for all j. The next lemma uses the costs incurred by these policies in the single-echelon systems $(\widehat{S}_j)$ to define a lower bound on the cost of any policy for the original problem. In particular, the sum of these costs is a lower bound on its optimal cost.

Lemma 1. *Consider a OWMR problem with backlog and a metric carrying cost structure. Let $\mathscr{C}^*$ be the cost of an optimal policy for this problem and let $\hat{\pi}^*_0, \hat{\pi}^*_1, \ldots, \hat{\pi}^*_N$ be optimal single-echelon policies for problems $(\widehat{S}_0), (\widehat{S}_1), \ldots, (\widehat{S}_N)$ as defined in the decomposition. Then the following inequality holds:*

$$\mathscr{C}^* \geq \sum_{j=0}^{N} \widehat{\mathscr{C}}_j(\hat{\pi}^*_j)$$

Proof. The proof of this lemma is detailed in Appendix A. Basically, we start from an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \ldots, \pi_N^{\text{OPT}})$ and decompose it in $N+1$ single-echelon policies $\bar{\pi}_0, \bar{\pi}_1, \ldots, \bar{\pi}_N$, one for each system $(\widehat{S}_j)$. We then show that the sum of the costs incurred by policies $\bar{\pi}_j$ in systems $(\widehat{S}_j)$ is lower than the optimal cost $\mathscr{C}^*$ for the original problem. The proof then follows from the optimality of policies $\hat{\pi}_j^*$.

3.2. Phase 2: The uncrossing algorithm

In this section, we modify the uncrossing method and use the decomposition $(\widehat{S}_0), (\widehat{S}_1), \ldots, (\widehat{S}_N)$ to build a feasible solution for the original problem. In addition we prove that when the solution is built upon the optimal single-echelon policies for these independent systems, the policy obtained has a cost of at most three times the optimal cost.

Let $\hat{\pi}_0, \ldots, \hat{\pi}_N$ be feasible policies for problems $(\widehat{S}_0), \ldots, (\widehat{S}_N)$. Recall that in our model, the warehouse cannot backlog the orders of the retailers. Thus in order to build a feasible policy to the original problem from policies $\hat{\pi}_i$, we need to eliminate the problematic situations where policy $\hat{\pi}_i$ orders a demand d_t^i in period s while $\hat{\pi}_0$ orders the same demand in period $r > s$. We say that the pair of orders $\lceil r, s \rfloor$ is *crossing* (or not feasible) if $r > s$.

We now present the *uncrossing algorithm* for the OWMR problem with backlog. Let $R = \{r_1, \ldots, r_w\}$ denote the set of periods when the warehouse orders according to policy $\hat{\pi}_0$. For convenience we add to R an artificial period $r_{w+1} = T+1$ corresponding to the end of the planning horizon, with no ordering cost. Given a set R – induced by a policy π_0 – and a period $s = 1, \ldots, T$, let s^+ and s^- be the first period of R after s and the last period of R prior to s, respectively:

$$s^+ = \min \{r \in R : \ r \geq s\} \tag{4}$$

$$s^- = \max \{r \in R : \ r \leq s\} \tag{5}$$

In what follows, the set R used to define s^+ and s^- will be clear from the context.

For simplicity, we say that a retailer order in period s is crossing in period r if there exists a demand d_t^i served by $\hat{\pi}_i$ in period s and by $\hat{\pi}_0$ in period r, with $r > s$. As its name suggests, the following algorithm *uncrosses* such retailer orders by simply adding an order of retailer i in period s^+ in order to synchronize with the warehouse. Note that in the case of J-retailers, we also synchronize the orders when $r \leq s$ by placing an additional order in period s^-, which enables us to bound the holding cost incurred in the resulting policy.

The Uncrossing Algorithm

Input: A set of feasible policies $\hat{\pi}_i$ for each problem $(\hat{S}_i)$, $i = 0, \ldots, N$.

Output: A feasible policy π^u for the OWMR problem with backlog, defined as follows. Let r and s be the periods when $\hat{\pi}_0$ and $\hat{\pi}_i$ order to serve demand d_t^i, respectively. The final policy π^u then uses the pair of orders $\lceil r^{it}, s^{it} \rfloor$ to serve demand d_t^i, where

$$\lceil r^{it}, s^{it} \rfloor = \begin{cases} \lceil s^+, s^+ \rfloor & \text{if } r > s \\ \lceil s^-, s^- \rfloor & \text{if } i \in I_J \text{ and } r \leq s \\ \lceil s^-, s \rfloor & \text{if } i \in I_W \text{ and } r \leq s \end{cases} \tag{6}$$

Note that this operation is executed in time $O(NT)$, which leads to a final complexity (including the resolution of the single-echelon subproblems) of $O(NT^2)$.

We now focus on the total cost incurred by π^u, the final policy built by the uncrossing algorithm. The following lemma bounds the overcost incurred by the uncrossing algorithm compared to the costs of the independent single-echelon policies $\hat{\pi}_j$.

Lemma 2. *The uncrossing algorithm applied to single-echelon policies $\hat{\pi}_0, \ldots, \hat{\pi}_N$ produces an uncrossed feasible solution π^u for the OWMR problem. The total cost incurred by the resulting policy satisfies:*

$$\mathscr{C}(\pi^u) \leq 3 \sum_{i=0}^{N} \widehat{\Psi}_i(\hat{\pi}_i) + \mathcal{K}_0(\hat{\pi}_0) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i)$$

Moreover in the special case of the JRP (i.e. if $I_W = \emptyset$), we have:

$$\mathscr{C}(\pi^u) \leq 2 \sum_{i=0}^{N} \widehat{\Psi}_i(\hat{\pi}_i) + \mathcal{K}_0(\hat{\pi}_0) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i)$$

Proof. See Appendix B.

We now show that when the split and uncross algorithm builds a policy π^u upon optimal policies $\hat{\pi}_0^*, \hat{\pi}_1^*, \ldots, \hat{\pi}_N^*$ for single-echelon systems $(\widehat{S}_0), (\widehat{S}_1), \ldots, (\widehat{S}_N)$, $\mathscr{C}(\pi^u)$ is at most three times the optimal cost $\mathscr{C}^*$. More precisely, the 3-approximation algorithm works as follows:

Step 1 Compute optimal policies $\hat{\pi}_0^*, \hat{\pi}_1^*, \ldots, \hat{\pi}_N^*$ for problems $(\widehat{S}_0), (\widehat{S}_1), \ldots, (\widehat{S}_N)$ independently.

Step 2 Apply the uncrossing algorithm to policies $\hat{\pi}_0^*, \ldots, \hat{\pi}_N^*$.

Let π^{u*} be the resulting policy: From Lemma 1 and 2, we have:

$$
\begin{aligned}
\mathscr{C}(\pi^{u*}) &\leq 3 \sum_{i=0}^{N} \widehat{\Psi}_i(\hat{\pi}_i^*) + \mathcal{K}_0(\hat{\pi}_0^*) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i^*) \\
&\leq 3 \sum_{i=0}^{N} \widehat{\mathscr{C}_i}(\hat{\pi}_i^*) \\
&\leq 3\mathscr{C}^*
\end{aligned}
$$

We conclude that the following theorem holds:

Theorem 1. *The policy π^{u*} obtained after applying the split and uncross algorithm to policies $\hat{\pi}_0^*, \hat{\pi}_1^*, \ldots, \hat{\pi}_N^*$ has a performance guarantee of three for the OWMR problem with backlog. Its complexity is quadratic in the case of a metric carrying cost structure.*

Notice that in the special case of the JRP, property (P5) becomes irrelevant since $I_W = \emptyset$. In addition, the above inequality can be refined to reach a performance guarantee of two:

$$
\begin{aligned}
\mathscr{C}(\pi^{u*}) &\leq 2 \sum_{i=0}^{N} \widehat{\Psi}_i(\hat{\pi}_i^*) + \mathcal{K}_0(\hat{\pi}_0^*) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i^*) \\
&\leq 2 \sum_{i=0}^{N} \widehat{\mathscr{C}_i}(\hat{\pi}_i^*) \\
&\leq 2\mathscr{C}^*
\end{aligned}
$$

Theorem 2. *The policy π^{u*} obtained after applying the split and uncross algorithm to policies $\hat{\pi}_0^*, \hat{\pi}_1^*, \ldots, \hat{\pi}_N^*$ has a performance guarantee of two for the JRP with backlog.*

4. Lost sales

While in the backorder model we assume that customers are willing to wait for their demands to be fulfilled, another alternative is to consider a model in which customers are impatient and every unmet demand is simply lost. Most of the lost-sales literature focuses on cost structures in which it is dominant to serve demands on a first-come first-served basis. Although this is clearly the case for the linear cost setting, our model uses a general metric cost structure derived from the one introduced in §2, in which the latter property is not necessarily dominant.

In this section, we introduce a modified decomposition for the problem and exhibits a lower bound for the OWMR problem with lost-sales. We then show how to adapt the split and uncross technique to this problem and build a 2-approximation from the subproblems of the decomposition.

4.1. Cost structure

We use the general cost structure introduced in §2 to modelize the case of lost-sales problem and simply consider it as a special case of the backorder version, where parameters ϕ^{it}_{rs} are constant when demand d_{it} is not satisfied immediately. We have $\phi^{it}_{rs} = \phi^{it}$ for all $s > t$ and the following inequality holds for parameters ψ^{it}_{rs}:

$$\psi^{it}_{rs} = \psi^{it} = \phi^{it} d_{it} \quad \text{for all } s > t \tag{7}$$

Thus the parameter ψ^{it} corresponds to the total lost-sales cost associated to demand point (i, t). In other words, if the system serves demand d_{it} with the pair of orders $\lceil r, s \rfloor$ where $s \leq t$, it incurs a cost ψ^{it}_{rs} while if it does not serve this demand, the cost for loosing the sale is equal to ψ^{it}. An other way to see the lost-sales penalty cost is to consider that there exists an alternative stock from which we can order units to satisfy an unmet demand, for a total ordering cost of ψ^{it}. Note that in our model one can speculate on the penalty cost and deliberately choose to order from this alternative stock if the penalty cost is attractive compared to the cost of serving the demand using the inventory on hand.

Constraint (7) ensures that properties (P1)-(P4) are satisfied when $s > t$. On the other hand, property (P5) is not necessary in the lost-sales case and therefore we only assume that parameters ψ^{it}_{rs} satisfy properties (P1)-(P4) and equation (7) in the remainder of this section.

4.2. Decomposition in single-echelon systems

In the same fashion as the backorder model, we decompose the system into $N+1$ independent single-echelon problems with lost-sales. However in this version, the holding costs and lost-sales penalty costs are halved instead of divided by three. Similarly to the JRP case, increasing the fraction of the real costs used in the subproblems is crucial to obtain a performance guarantee of two.

The decomposition for the OWMR with lost-sales is defined as follows:

($\widetilde{S}_0$) The warehouse is regarded as a single-echelon, multi-item system with lost-sales, facing for each period t a demand d_{it} for item i. A fixed ordering cost K_r^0 is incurred for placing an order in period r. The holding cost incurred if demand (i,t) is ordered in period r is equal to $\frac{1}{2}\psi_{rr}^{it}$ if $i \in I_J$ and $\frac{1}{2}\psi_{rt}^{it}$ if $i \in I_W$. If the demand is lost, the system incurs a lost-sales cost equal to $\frac{1}{2}\psi^{it}$.

($\widetilde{S}_i$) Retailer i is considered as a single-echelon system with lost-sales, facing demand d_{it} with ordering cost K^i. The holding cost incured to order to serve demand (i,t) with an order in period s is equal to $\frac{1}{2}\psi_{ss}^{it}$, while the per-unit cost if the demand is lost is equal to $\frac{1}{2}\psi^{it}$.

Similarly to the previous sections, we define $\widetilde{\mathscr{C}}_i(\tilde{\pi}_i)$ (resp. $\widetilde{\Psi}_i(\tilde{\pi}_i)$) as the total (resp. carrying) cost incurred by $\tilde{\pi}_i$ in system $(\widetilde{S}_i)$ and we have for all $i = 0, \ldots, N$:

$$\widetilde{\mathscr{C}}_i(\tilde{\pi}_i) = \mathcal{K}_i(\tilde{\pi}_i) + \widetilde{\Psi}_i(\tilde{\pi}_i)$$

We first discuss how to solve these single-echelon problems. Similarly to the systems with backlog, this can be done using classical shortest path algorithms in the same type of graph (see Figure 1). We define a graph G_j corresponding to each single-echelon problem $(\widetilde{S}_j)$. The main difference comes from the computation of the length of each arc (u,v), $u < v$, that we detail below.

We first introduce the following virtual cost parameters:

$$\tilde{\omega}_r^{it} = \begin{cases} \min\{\psi_{rr}^{it}, \psi^{it}\} & \text{if } i \in I_J \\ \min\{\psi_{rt}^{it}, \psi^{it}\} & \text{otherwise} \end{cases} \tag{8}$$

$$\tilde{\rho}_s^{it} = \min\left\{\psi_{ss}^{it}, \psi^{it}\right\} \tag{9}$$

These are the virtual minimum cost between satisfying or loosing demand d_{it} when the last order is placed in period r or s. One can precompute these parameters in time $O(T^2)$ for one retailer and $O(NT^2)$ for the warehouse. We then define the length of each edge (u, v) in graphs G_i and G_0 as follows:

$$\forall i = 1, \dots, N, \quad \tilde{l}^i_{u,v} = \begin{cases} \displaystyle\sum_{t=1}^{v-1} \frac{1}{2}\psi^{it} & \text{if } u = 0 \\ \displaystyle K^i + \sum_{t=u}^{v-1} \frac{1}{2}\tilde{\rho}^{it}_u & \text{otherwise} \end{cases}$$

$$\tilde{l}^0_{u,v} = \begin{cases} \displaystyle\sum_{i=1}^{N}\sum_{t=1}^{v-1} \frac{1}{2}\psi^{it} & \text{if } u = 0 \\ \displaystyle K^0_u + \sum_{i=1}^{N}\sum_{t=u}^{v-1} \frac{1}{2}\tilde{\omega}^{it}_u & \text{otherwise} \end{cases}$$

These lengths are computed in time $O(NT^2)$. As for the backorder version, the time complexity to compute the $N+1$ optimal policies to problems $(\widetilde{S}_j)$ is equal to $O(NT^2)$.

We now focus on the lower bound resulting from this decomposition. The following lemma derives from Lemma 1:

Lemma 3. *Consider a OWMR problem with lost-sales. Let $\mathscr{C}^*$ be the cost of an optimal policy for this problem and let $\tilde{\pi}^*_0, \tilde{\pi}^*_1, \dots, \tilde{\pi}^*_N$ be optimal single-echelon policies for problems $(\widetilde{S}_0), (\widetilde{S}_1), \dots, (\widetilde{S}_N)$ as defined in the decomposition. Then the following inequality holds:*

$$\mathscr{C}^* \geq \sum_{i=0}^{N} \widetilde{\mathscr{C}}_i(\tilde{\pi}^*_i)$$

Proof The proof is similar to the one of Lemma 1 (see Appendix C).

4.3. The algorithm

To conclude this section, we show how to adapt the split and uncross technique to the lost-sales case and build a 2-approximation algorithm. Let $\tilde{\pi}_0, \dots, \tilde{\pi}_N$ be feasible policies for problems $(\widetilde{S}_0), \dots, (\widetilde{S}_N)$ and let $R = \{r_1, \dots, r_w, r_{w+1}\}$ the set of periods when policy $\tilde{\pi}_0$ places order, where again $r_{w+1} = T + 1$ is an artificial period corresponding to the end of the planning horizon. The algorithm is very similar to the one presented in §3.2, with the following simple modification. Let $i = 1, \ldots, N$ and $t = 1, \ldots, T$: If neither $\tilde{\pi}_0$ nor $\tilde{\pi}_i$ satisfies d_t^i, then the demand is lost in the final policy π^u. Otherwise the ordering periods for demand d_t^i at the warehouse and at the retailer are defined according to equation (6).

It is straightforward that the lower bound from Lemma 2 holds in this special case of the problem. However, since we only halve the holding and penalty costs in the subproblems instead of dividing them by three, we can improve upon this result to get the following lower bound:

Lemma 4. *The uncrossing algorithm applied to single-echelon policies $\tilde{\pi}_0, \ldots, \tilde{\pi}_N$ produces an uncrossed feasible solution $\tilde{\pi}^u$ for the OWMR problem with lost-sales. The total cost incurred by the resulting policy satisfies:*

$$\mathscr{C}(\tilde{\pi}^u) \leq \mathcal{K}_0(\tilde{\pi}_0) + 2\widetilde{\Psi}_0(\tilde{\pi}_0) + 2\sum_{i=1}^{N} \widetilde{\mathscr{C}}_i(\tilde{\pi}_i)$$

Proof See Appendix D.

Finally, assume the solution built by the split and uncross algorithm when applied to the optimal single-echelon policis $\tilde{\pi}_0^*, \tilde{\pi}_1^*, \ldots, \tilde{\pi}_N^*$ for problems $(\widetilde{S}_0), (\widetilde{S}_1), \ldots, (\widetilde{S}_N)$. Let $\tilde{\pi}^{u*}$ be the resulting policy: From Lemma 3 and 4, we have:

$$\begin{aligned}
\mathscr{C}(\tilde{\pi}^{u*}) &\leq \mathcal{K}_0(\tilde{\pi}_0^*) + 2\widetilde{\Psi}_0(\tilde{\pi}_0^*) + 2\sum_{i=1}^{N} \widetilde{\mathscr{C}}_i(\tilde{\pi}_i^*) \\
&\leq 2\sum_{i=0}^{N} \widetilde{\mathscr{C}}_i(\tilde{\pi}_i^*) \\
&\leq 2\mathscr{C}^*
\end{aligned}$$

and the following theorem states the approximation ratio of the technique:

Theorem 3. *The policy $\tilde{\pi}^{u*}$ obtained after applying the split and uncross algorithm to policies $\tilde{\pi}_0^*, \tilde{\pi}_1^*, \ldots, \tilde{\pi}_N^*$ has a performance guarantee of two for the OWMR problem with lost-sales. Its complexity is linear in the case of linear cost and quadratic in the case of metric carrying cost parameters.*

5. Conclusion and perspectives

Although the literature on discrete deterministic divergent inventory systems is already substantial, almost all the existing works assume that demands must be satisfied on time. Our work propose the first constant approximation for the OWMR problem with backlog or lost-sales. We derive new, intuitive lower bounds for the problems considered and prove that our algorithms have a constant performance guarantee under a general cost structure that can capture many practical situations. In addition we show that the performance guarantee can be improved in special cases of our model, namely the JRP and a lost-sales model.

Although these results fill a gap in the literature on deterministic inventory theory, there remains some room for improvement. One obvious way for future research is to try to improve the performance guarantee of three in the backlogging case. One could also want to generalize the assumptions on the cost structure by eliminating the backlogging cost bounding constraint or considering nonlinear carrying costs.

Appendix A. Proof of Lemma 1

Consider an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \ldots, \pi_N^{\text{OPT}})$ for the OWMR system with backlog. For all $j = 0, \ldots, N$, let $\bar{\pi}_j$ be the single-echelon policy that places the same orders as π_j^{OPT}. That is, if π^{OPT} orders demand (i,t) in period r at the warehouse and s at retailer i, $\bar{\pi}_0$ and $\bar{\pi}_i$ orders demand (i,t) in period r and s, respectively.

Note that by definition of $\bar{\pi}_j$ for all j, the ordering costs remain untouched: $\mathcal{K}(\pi^{\text{OPT}}) = \sum_{j=0}^{N} \mathcal{K}_j(\bar{\pi}_j)$. Therefore we simply have to prove that $\sum_{j=0}^{N} \Psi_j(\bar{\pi}_j) \leq \Psi(\pi^{\text{OPT}})$ and the proof will follow from the optimality of policies $\hat{\pi}_0^*, \hat{\pi}_1^*, \ldots, \hat{\pi}_N^*$. Consider a demand point (i,t) and let r and s be the periods when π^{OPT} orders the demand at the warehouse and at retailer i, respectively. The carrying cost incurred by π^{OPT} for this demand is therefore equal to ψ_{rs}^{it}. On the other hand, the total cost incurred by policies $\bar{\pi}_0$ and $\bar{\pi}_i$ to carry this demand in systems $(\widehat{S}_0)$ (resp. $(\widehat{S}_i)$) is equal to $\frac{1}{3}\omega_{rr}^{it}$ (resp. $\frac{2}{3}\psi_{ss}^{it}$) if $i \in I_W$ and $\frac{1}{2}\omega_{rr}^{it}$ (resp. $\frac{1}{2}\psi_{ss}^{it}$) if $i \in I_J$. Hence we want to show that ψ_{rs}^{it} is bounded as follows:

$$\psi_{rs}^{it} \geq \begin{cases} \dfrac{1}{3}\psi_{rr}^{it} + \dfrac{2}{3}\psi_{ss}^{it} \text{ if } i \in I_W \\ \dfrac{1}{2}\psi_{rr}^{it} + \dfrac{1}{2}\psi_{ss}^{it} \text{ otherwise.} \end{cases} \tag{A.1}$$

which is precisely the sum of the costs incurred by policies $\bar{\pi}_0$ and $\bar{\pi}_j$ to carry demand d_t^i. We now prove that inequality (A.1) holds depending on the nature of retailer i.

If $i \in I_J$, we have $r = s$ since π^{OPT} is optimal and therefore

$$\begin{aligned} \psi_{rs}^{it} &= \psi_{rr}^{it} = \psi_{ss}^{it} && \text{(A.2)} \\ &= \frac{1}{2}\psi_{rr}^{it} + \frac{1}{2}\psi_{ss}^{it} && \text{(A.3)} \end{aligned}$$

If $i \in I_W$, we distinguish between two cases, listed below:

- $s \leq t$ or $r > t$: Properties (P2) and (P3) ensure that

$$\begin{aligned} \psi_{rs}^{it} &\geq \max\left\{\psi_{rt}^{it}, \psi_{ss}^{it}\right\} \\ &\geq \frac{1}{3}\psi_{rr}^{it} + \frac{2}{3}\psi_{ss}^{it} && \text{(A.4)} \end{aligned}$$

- $r \leq t < s$: We have from properties (P2) and (P3)

$$\begin{aligned} \psi_{rs}^{it} &\geq \max\left\{\psi_{rt}^{it}, \psi_{ts}^{it}\right\} \\ &\geq \frac{1}{3}\psi_{rr}^{it} + \frac{2}{3}\psi_{ts}^{it} \\ &\geq \frac{1}{3}\psi_{rr}^{it} + \frac{2}{3}\psi_{ss}^{it} && \text{(A.5)} \end{aligned}$$

Hence in all cases, inequality (A.1) is satisfied and we have $\Psi(\pi^{\text{OPT}}) \geq \sum_{j=0}^{N} \Psi_j(\bar{\pi}_j)$. As a consequence the following inequality holds:

$$\mathscr{C}^* = \mathscr{C}(\pi^{\text{OPT}}) \geq \sum_{j=0}^{N} \mathscr{C}_j(\bar{\pi}_j) \tag{A.6}$$

Inequality (A.6) and the optimality of policies $\hat{\pi}_j^*$ for problems $(\widehat{S}_j)$ then conclude the proof:

$$\begin{aligned} \sum_{j=0}^{N} \mathscr{C}_j(\hat{\pi}_j^*) &\leq \sum_{j=0}^{N} \mathscr{C}(\bar{\pi}_j) \\ &\leq \mathscr{C}^* \end{aligned}$$

Appendix B. Proof of Lemma 2

First, note that for all $i = 1, \dots, N$ and $t = 1, \dots, T$, the pair of orders $\lceil r^{it}, s^{it} \rfloor$ chosen by the algorithm to serve d_t^i satisfies $r^{it} \leq s^{it}$ and therefore policy π^u is uncrossed and feasible.

We now prove that the final cost incurred by π^u is bounded. By construction, each ordering period of $\hat{\pi}_i$ is replaced by at most two ordering periods at retailer i in π^u. More precisely, the set of ordering periods of retailer i in π^u is included in $\{s^-, s^+ | s \text{ an ordering period of } \hat{\pi}_i\}$ if $i \in I_J$, and in $\{s, s^+ | s \text{ an ordering period of } \hat{\pi}_i\}$ if $i \in I_W$. Thus the number of ordering periods at each retailer is at most doubled by the uncrossing algorithm. Moreover the uncrossing algorithm leaves the orders of $\hat{\pi}_0$ untouched, hence using the assumption of stationary ordering costs K_i at the retailers, we have that the ordering costs incurred by π^u are at most twice the sum of the ordering costs incurred by policies $\hat{\pi}_i$.

Therefore if $I_W \neq \emptyset$ it only remains to prove that the following inequalities hold:

$$\Psi(\pi^u) \leq 3 \sum_{j=0}^{N} \widehat{\Psi}_j(\hat{\pi}_j) \tag{B.1}$$

Let $i = 1 \dots, N$ and $t = 1, \dots, T$ and focus on demand point (i, t). Let r be the period when $\hat{\pi}_0$ orders d_t^i and s the period when $\hat{\pi}_i$ orders d_t^i in their single-echelon solution. We now bound the final carrying cost incurred in π^u to serve d_t^i. Recall that in π^u, d_t^i is served by the the pair $\lceil r^{it}, s^{it} \rfloor$, as defined in equation (6). Hence we want to show that for a retailer i and periods t, r, s, the following inequality is satisfied:

$$\psi^{it}_{r^{it}, s^{it}} \leq \begin{cases} 3 \cdot \frac{1}{3}\psi^{it}_{rt} + 3 \cdot \frac{2}{3}\psi^{it}_{ss} = \psi^{it}_{rt} + 2\psi^{it}_{ss} & \text{if } i \in I_W \text{ and } r \leq t \quad \text{(B.2)} \\ 3 \cdot \frac{1}{3}\psi^{it}_{rr} + 3 \cdot \frac{2}{3}\psi^{it}_{ss} = \psi^{it}_{rr} + 2\psi^{it}_{ss} & \text{if } i \in I_W \text{ and } r > t \quad \text{(B.3)} \\ 2 \cdot \frac{1}{2}\psi^{it}_{rr} + 2 \cdot \frac{1}{2}\psi^{it}_{ss} = \psi^{it}_{rr} + \psi^{it}_{ss} & \text{if } i \in I_J \quad \text{(B.4)} \end{cases}$$

If $i \in I_J$, the order of retailer i that serves d_t^i in policy π^u is synchronized with a warehouse order either in period s^- or s^+. Note that by definition we necessarily have $r \leq s^-$ if $r \leq s$ and $r \geq s^+$ if $r > s$. Hence if $r \leq s$, we have from property (P4)

$$\begin{aligned}\psi^{it}_{s^- s^-} &\leq \psi^{it}_{rr} \\ &\leq \psi^{it}_{rr} + \psi^{it}_{ss}\end{aligned}$$

On the other hand if $r > s$, π^u uses the pair $\lceil s^+, s^+ \rfloor$ to serve d^i_t and from property (P2) and (P3) the following inequality holds

$$\begin{aligned}\psi^{it}_{s^+ s^+} &\leq \begin{cases} \psi^{it}_{rr} & \text{if } r > t \\ \psi^{it}_{ss} & \text{otherwise} \end{cases} \\ &\leq \max\left\{\psi^{it}_{rr}, \psi^{it}_{ss}\right\} \\ &\leq \psi^{it}_{rr} + \psi^{it}_{ss}\end{aligned}$$

Hence if $i \in I_J$ inequality (B.4) is satisfied for all t, r, s.

In the case where i is a W-retailer, we again derive a lower bound for the carrying cost depending on how r, s and t are ordered. If $r \leq s$ and $r \leq t$, the final cost incurred in π^u to serve d^i_t is $\psi^{it}_{s^- s}$. Notice that in this case, we have $r \leq s^-$. Therefore we can bound the carrying cost incurred to serve demand d^i_t using properties (P3), (P4) and (P5) in their turn:

$$\begin{aligned}\psi^{it}_{s^- s} &\leq \begin{cases} \psi^{it}_{s^- t} + \psi^{it}_{ss} & \text{if } r \leq s \leq t \\ \psi^{it}_{s^- t} + \psi^{it}_{ts} & \text{if } r \leq t < s \end{cases} \\ &\leq \begin{cases} \psi^{it}_{rt} + \psi^{it}_{ss} & \text{if } r \leq s \leq t \\ \psi^{it}_{rt} + \psi^{it}_{ts} & \text{if } r \leq t < s \end{cases} \\ &\leq \psi^{it}_{rt} + 2\psi^{it}_{ss}\end{aligned}$$

On the other hand if $t < r \leq s$, properties (P3) and (P5) ensure that

$$\begin{aligned}\psi^{it}_{s^- s} &\leq \psi^{it}_{ts} \\ &\leq 2\psi^{it}_{ss} \\ &\leq \psi^{it}_{rr} + 2\psi^{it}_{ss}\end{aligned}$$

Finally if $r > s$, we use properties (P2) and (P3) to bound the final carrying cost:

$$\begin{aligned}\psi^{it}_{s^+ s^+} &\leq \begin{cases} \psi^{it}_{rr} & \text{if } t < s^+ \\ \psi^{it}_{ss} & \text{otherwise} \end{cases} \\ &\leq \max\left\{\psi^{it}_{rr}, \psi^{it}_{ss}\right\} \\ &\leq \psi^{it}_{rr} + \psi^{it}_{ss}\end{aligned}$$

We conclude that if $i \in I_W$ inequalities (B.2) and (B.3) are satisfied.

Therefore if π^u uses the pair of orders $\lceil r^{it}, s^{it} \rfloor$, inequalities (B.2), (B.3) and (B.4) are always satisfied. As a consequence inequality (D.1) holds and the proof follows:

$$\mathscr{C}(\pi^u) \leq 3 \sum_{j=0}^{N} \widehat{\Psi}_j(\hat{\pi}_j) + \mathcal{K}_0(\hat{\pi}_0) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i)$$

To conclude the proof, notice that when $I_W = \emptyset$, inequality (B.4) implies that

$$\Psi(\pi^u) \leq 2 \sum_{j=0}^{N} \widehat{\Psi}_j(\hat{\pi}_j)$$

and thus one can modify the upper bound in the special case of the JRP:

$$\mathscr{C}(\pi^u) \leq 2 \sum_{j=0}^{N} \widehat{\Psi}_j(\hat{\pi}_j) + \mathcal{K}_0(\hat{\pi}_0) + 2 \sum_{i=1}^{N} \mathcal{K}_i(\hat{\pi}_i)$$

Appendix C. Proof of Lemma 3

We prove this inequality in a similar manner as we did for Lemma 1. Consider an optimal policy $\pi^{\text{OPT}} = (\pi_0^{\text{OPT}}, \pi_1^{\text{OPT}}, \ldots, \pi_N^{\text{OPT}})$ for the OWMR system with lost-sales. For all $j = 0, \ldots, N$, we define $\bar{\pi}_j$ as the policy that orders in the same periods as π_j^{OPT} and fills or looses demands exactly as the latter does. Obviously, we again have $\mathcal{K}(\pi^{\text{OPT}}) = \sum_{j=0}^{N} \mathcal{K}_j(\bar{\pi}_j)$. We now prove that $\sum_{j=0}^{N} \Psi_j(\bar{\pi}_j) \leq \Psi(\pi^{\text{OPT}})$. Consider a retailer i and a demand point (i, t). If d_t^i is lost in π^{OPT}, then by construction it is lost in $\bar{\pi}_0$ and $\bar{\pi}_i$ and we directly have:

$$\psi^{it} = \frac{1}{2}\psi^{it} + \frac{1}{2}\psi^{it} \tag{C.1}$$

Now assume the demand is filled and let r and s be the periods when π^{OPT} orders d_t^i at the warehouse and at the retailer, respectively. The holding cost incurred by π^{OPT} for d_t^i is equal to ψ_{rs}^{it}. If i is a J-retailer, we have $r = s$ and therefore

$$\psi_{rs}^{it} = \psi_{rr}^{it} = \psi_{ss}^{it} = \frac{1}{2}\psi_{rr}^{it} + \frac{1}{2}\psi_{ss}^{it} \tag{C.2}$$

On the other hand if i is a W-retailer, we can bound this cost using properties (P2) and (P3):

$$\begin{aligned} \psi_{rs}^{it} &\geq \max\left\{\psi_{rt}^{it}, \psi_{ss}^{it}\right\} \\ &\geq \frac{1}{2}\psi_{rt}^{it} + \frac{1}{2}\psi_{ss}^{it} \end{aligned} \tag{C.3}$$

In both cases, the sum of the holding costs incurred by $\bar{\pi}_0$ and $\bar{\pi}_i$ to serve demand (i,t) in systems $(\widetilde{S}_0)$ and $(\widetilde{S}_i)$ is lower than or equal to the cost incurred in π^{OPT} to serve the same demand. Therefore from equations (C.1), (C.2) and (C.3) we conclude $\Psi(\pi^{\text{OPT}}) \geq \sum_{j=0}^{N} \Psi_j(\bar{\pi}_j)$. The optimality of policies $\tilde{\pi}_j^*$ for problems $(\widetilde{S}_j)$ then yields:

$$\begin{aligned} \sum_{j=0}^{N} \mathscr{C}_j(\tilde{\pi}_j^*) &\leq \sum_{j=0}^{N} \mathscr{C}(\bar{\pi}_j) \\ &= \sum_{j=0}^{N} \mathcal{K}_j(\bar{\pi}_j) + \Psi_j(\bar{\pi}_j) \\ &\leq \mathcal{K}(\pi^{\text{OPT}}) + \Psi(\pi^{\text{OPT}}) \\ &\leq \mathscr{C}^* \end{aligned}$$

Appendix D. Proof of Lemma 4

To prove the feasibility of the resulting policy, we refer the reader to the proof of Lemma 2. Similarly, it is straightforward to see that $\mathcal{K}(\tilde{\pi}^u) \leq \mathcal{K}_0(\tilde{\pi}_0) + 2\sum_{i=1}^{N} \mathcal{K}_i(\tilde{\pi}_i)$, since for every order in $\tilde{\pi}_i$ the algorithm places at most two orders in the final policy $\tilde{\pi}^u$. Therefore we only have to prove that

$$\Psi(\tilde{\pi}^u) \leq 2\sum_{j=0}^{N} \Psi_j(\tilde{\pi}_j) \tag{D.1}$$

Let $i = 1\ldots,N$ and $t = 1,\ldots,T$ and focus on demand point (i,t). If the d_t^i is lost in $\tilde{\pi}_0$ or $\tilde{\pi}_i$, it incurs a cost of at least $\frac{1}{2}\psi^{it}$ in these policies. Since by construction it is also lost in $\tilde{\pi}^u$, it incurs a cost of exactly ψ^{it} in the final solution. Now assume the demand is filled both in $\tilde{\pi}_0$ and $\tilde{\pi}_i$ and let r and s be the respective ordering periods of these policies for this demand. We want to prove that the final holding cost incurred in $\tilde{\pi}^u$ to serve d_t^i is at most

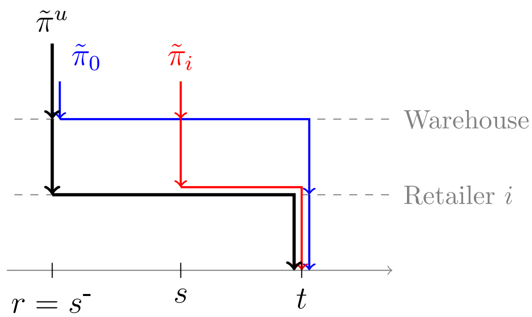

(a) Example with $i \in I_J$ and $r \leq s$ ($t \geq s$)

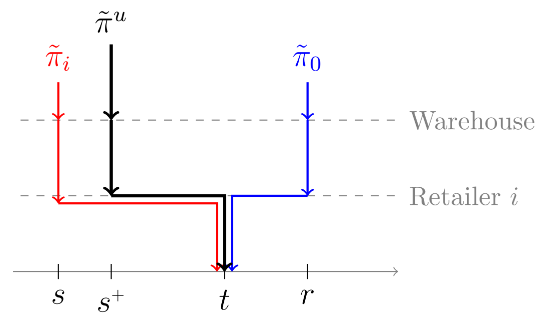

(b) Example with $i \in I_J$ and $r > s$ ($s \leq t \leq r$)

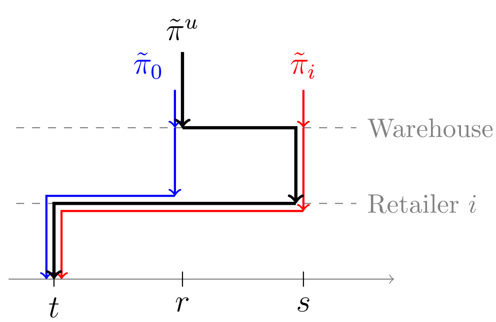

(c) Example with $i \in I_W$ and $r \leq s$ ($t \leq r$)

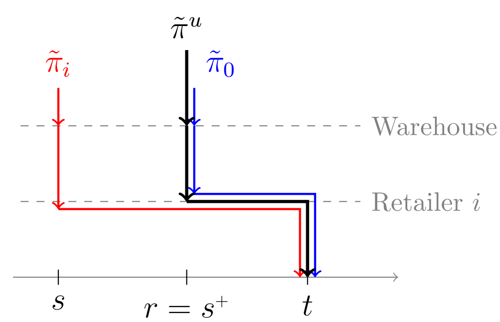

(d) Example with $i \in I_W$ and $r > s$ ($t \geq r$)

Figure D.2: Examples of paths followed by the units after applying the uncrossing policy

the sum of the holding cost it incurs in policies $\tilde{\pi}_0$ and $\tilde{\pi}_i$. This final cost depends on the path followed by the units in the system after applying the uncrossing policy. Specifically, we want to compare this cost with the cost of the paths the same units follow in the single echelon policies $\tilde{\pi}_0$ and $\tilde{\pi}_i$. We distinguish several cases depending on the type of retailer and on the respective position of r and s (see also Figure D.2 for a specific example of units path in each case):

- If $i \in I_J$ and $r \leq s$, $\tilde{\pi}^u$ orders d_t^i in period s^- at the warehouse and at the retailer for a final cost of $\psi_{s^- s^-}^{it}$. The holding cost incurred by policy $\tilde{\pi}_0$ (resp. $\tilde{\pi}_i$) to serve this demand is equal to ψ_{rr}^{it} (resp. ψ_{ss}^{it}).

Since $s^- \geq r$, we have from property (P3):

$$\psi^{it}_{s^-s^-} \leq \psi^{it}_{rr} \leq \psi^{it}_{rr} + \psi^{it}_{ss} \tag{D.2}$$

- If $i \in I_J$ and $r > s$, $\tilde{\pi}^u$ orders the demand in period s^+ at the warehouse and at the retailer for a final cost of $\psi^{it}_{s^+s^+}$. The holding cost incurred by policy $\tilde{\pi}_0$ (resp. $\tilde{\pi}_i$) to serve this demand is equal to ψ^{it}_{rr} (resp. ψ^{it}_{ss}). Since $s^+ \geq s$, we have from property (P3):

$$\psi^{it}_{s^+s^+} \leq \psi^{it}_{ss} \leq \psi^{it}_{rr} + \psi^{it}_{ss} \tag{D.3}$$

- If $i \in I_W$ and $r \leq s$, $\tilde{\pi}^u$ orders the demand in period s^- at the warehouse and in period s at the retailer for a final cost of $\psi^{it}_{s^-s}$. Recall that by definition, we have $r \leq s^-$. The holding cost incurred by policy $\tilde{\pi}_0$ (resp. $\tilde{\pi}_i$) to serve this demand is equal to ψ^{it}_{rt} (resp. ψ^{it}_{ss}). Thus we have from property (P3) and (P4):

$$\psi^{it}_{s^-s} \leq \psi^{it}_{rs} \leq \psi^{it}_{rt} + \psi^{it}_{ss} \tag{D.4}$$

- If $i \in I_W$ and $s < r$, $\tilde{\pi}^u$ orders the demand in period s^+ at the warehouse and at the retailer for a final cost of $\psi^{it}_{s^+s^+}$. The holding cost incurred by policy $\tilde{\pi}_0$ (resp. $\tilde{\pi}_i$) to serve this demand is equal to ψ^{it}_{rt} (resp. ψ^{it}_{ss}). Since $s^+ > s$ we have from property (P3):

$$\psi^{it}_{s^+s^+} \leq \psi^{it}_{ss} \leq \psi^{it}_{rt} + \psi^{it}_{ss} \tag{D.5}$$

Note that in inequalities (D.2)-(D.5), the right-hand side corresponds exactly to twice the sum of the holding cost incurred in policy $\tilde{\pi}_0$ and $\tilde{\pi}_i$ to serve demand d^i_t. As a consequence in all cases the holding or penalty cost incurred to serve or loose this demand in $\tilde{\pi}^u$ is at most twice the sum of the costs incurred for the same demand in policies $\tilde{\pi}_0$ and $\tilde{\pi}_i$. We conclude that inequality (D.1) holds and the proof follows.

Absi, N., Kedad-Sidhoum, S., 2008. The multi-item capacitated lot-size problem with setup times and shortage costs. Eur. J. Oper. Res. 185, 1351–1374.

Aggarwal, A., Park, J. K., 1993. Improved algorithm for economic lot-size problems. Oper. Res. 41 (3), 549–571.

Aksen, D., Altinkemer, K., Chand, S., 2003. The single-item lot-sizing problem with immediate lost sales. Eur. J. Oper. Res. 147, 558–566.

Arkin, E., Joneja, D., Roundy, R., 1989. Computational complexity of uncapacitated multi-echelon production planning problems. Oper. Res. Letters 8, 61–66.

Bienkowski, M., Byrka, J., Chrobak, M., Jeż, Ł., Nogneng, D., Sgall, J., 2014. Better approximation bounds for the joint replenishment problem. Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, 42–54.

Chan, L. M. A., Muriel, A., Shen, Z.-J. M., Simchi-Levi, D., Teo, C.-P., 2000. Effective zero-inventory-ordering policies for the single-warehouse multiretailer problem with piecewise linear cost structures. Management Sci. 48 (11), 1446–1460.

Federgruen, A., Tzur, M., 1993. The dynamic lot-sizing model with backlogging: A simple o(n log n) algorithm and minimal forecast horizon procedure. Naval Research Logistics 40 (4), 459–478.

Federgruen, A., Wang, M., 2013. Inventory models with shelf age and delay dependent inventory costs. In: MSOM Conference 2013, Fontainebleau, France.

Feige, U., 1998. A treshold of $\log(n)$ for approximating set-cover. J. of the ACM 45, 634–652.

Küçükyavuz, S., Pochet, Y., 2009. Uncapacitated lot sizing with backlogging: The convex hull. Math. Programming 118 (1), 151–175.

Levi, R., Roundy, R., Shmoys, D., 2006. Primal-dual algorithms for deterministic inventory problems. Math. Oper. Res. 31 (2), 267–284.

Levi, R., Roundy, R., Shmoys, D., Sviridenko, M., 2008. A constant approximation algorithm for the one-warehouse multiretailer problem. Management Sci. 54 (8), 763–776.

Liu, X., Chu, F., Chu, C., Wang, C., 2007. Lot-sizing with bounded inventory and lost-sales. International J. of Production Res. 45 (24), 5881–5894.

Nonner, T., Souza, A., 2009. A 5/3-approximation algorithm for joint replenishment with deadlines. Proc. of the 3rd Annual International Conference on Combinatorial Optimization and Applications, Huangshan, China, 24–35.

Pochet, Y., Wosley, L., 1988. Lot-size models with backlogging: Strong reformulations and cutting planes. Math. Programming 40, 317–335.

Pochet, Y., Wosley, L., 1994. Polyhedra for lot-sizing withwagnerwhitin costs. Math. Programming 67, 297–323.

Sandbothe, R. A., Thompson, J. L., 1990. A forward algorithm for the capacitated lot size model with stockouts. Oper. Res. 38 (3), 474–486.

Stauffer, G., Massonnet, G., Rapine, C., Gayon, J.-P., 2011. A simple and fast 2-approximation algorithm for the one-warehouse multi-retailers problem. Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, 67–79.

Wagner, H. M., Whitin, T., 1958. Dynamic version of the economic lot size model. Management Sci. 5, 89–96.

Zangwill, W., 1966. A deterministic multi-period production scheduling model with backlogging. Management Sci. 13 (1), 105–119.