



A location-inventory model for large three-level supply chains

Jean-Sébastien Tancrez^{a,b,*,1}, Jean-Charles Lange^{a,1}, Pierre Semal^{a,1}

^a Louvain School of Management, Université catholique de Louvain, Belgium

^b College of Management of Technology, Ecole Polytechnique Fédérale de Lausanne, Switzerland

ARTICLE INFO

Article history:

Received 31 October 2010

Received in revised form 24 August 2011

Accepted 25 September 2011

Keywords:

Location-inventory model

Facility location

Inventory management

Three-level supply chain

Reverse logistics

ABSTRACT

We study the location-inventory problem in three-level supply networks. Our model integrates three decisions: the distribution centers location, flows allocation, and shipment sizes. We propose a nonlinear continuous formulation, including transportation, fixed, handling and holding costs, which decomposes into a closed-form equation and a linear program when the DC flows are fixed. We thus develop an iterative heuristic that estimates the DC flows a priori, solves the linear program, and then improves the DC flow estimations. Extensive numerical experiments show that the approach can design large supply networks both effectively and efficiently, and a case study is discussed.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

In the last decades, supply chain management has proved to be a primary lever for companies to lower their costs and improve their overall competitiveness. In particular, the strategic design of the supply network is of crucial importance. It deeply impacts the supply chain planning and eventually the performance of the company. Thus, the facility location problem and its variants have been the focus of much attention from the scientific community. However, the problem is less often approached from a supply chain management perspective (Melo et al., 2009). In particular, while inventory costs may have a significant effect on the cost balance and the positioning of facilities, inventory management considerations are often neglected. Furthermore, the problem remains difficult to solve for large supply chains in reality.

Our research was inspired by the real-life case of a leading European glass manufacturer, mainly producing glass panes for the automotive and construction industries. Its supply chain includes 10 factories and around 500 customers (which can be retailers) throughout Europe. The case presented by the company concerns its reverse logistics network, and more precisely the return flow of reusable items (empty trestles) from the points of consumption to the factories. Currently, empty trestles are directly shipped back to the factories in the same truck that delivered the glass panes to the customer, whereas empty trestles can be folded and are thus less voluminous. The glass producer wishes to assess the advisability of an alternative strategy for the return flow: accumulating empty trestles in regional depots, to return them to factories in trucks that are better utilized. Consequently, inventory management decisions, such as the shipment size, play a central role, and have to be integrated with the location-allocation decisions within a single framework. We are therefore confronted with a fairly classic and difficult problem: the network design and inventory management of a three-level supply chain. In fact, this problem in the reverse logistics context bears a strong resemblance to that in a forward network. Solution methodologies can be

* Corresponding author. Address: Université catholique de Louvain, Chaussée de Binche, 151, 7000 Mons, Belgium. Tel.: +32 65 323 358; fax: +32 65 323 363.

E-mail addresses: js.tancrez@uclouvain-mons.be (J.-S. Tancrez), jean-charles.lange@uclouvain.be (J.-C. Lange), pierre.semal@uclouvain.be (P. Semal).

¹ The contributions of the three authors are equivalent.

applied equally well in both contexts, as indicated by Melo et al. (2009). We thus present our methodology in a general fashion, without distinguishing between the forward and reverse cases. Both will subsequently be discussed and exemplified.

Inspired by this real-life case, we study the location of intermediary facilities in a three-level network where factories and customers have fixed locations, and fixed constant production and demand rates. We also want to consider the impact of inventory decisions, and of shipment size in particular. Our approach assists with the making of location-allocation decisions as well as inventory management decisions in an integrated fashion. In other words, the cost function covers transportation and facility fixed costs as well as inventory holding and handling costs, and thus underlines the important trade-off between these costs. The targeted decision level is strategic. We consider a single-period planning horizon and a single product. Furthermore, we allow for direct flows between factories and customers and consider capacitated vehicles. In order to be able to analyze large real-life problems, we develop a continuous optimization formulation (and avoid using integer variables). The latter is shown to decompose when the flows through the DCs are fixed. In this case, the inventory decisions can be computed from a closed-form equation and the location-allocation decisions follow from solving a linear program. Based on this, we then propose an iterative heuristic which, at each iteration, estimates the DC flows, solves a linear program, and then improves the DC flow estimations.

The remainder of the paper is structured as follows. In the next section, we review the related literature. In Section 3, we present the problem and our mathematical modeling. A heuristic method is then proposed to solve it in Section 4. In order to assess the efficiency of the solution procedure, the heuristic is then tested on many different configurations in Section 5. In Section 6, we illustrate the application of the methodology in reverse logistics, to the case of the glass producer, and we discuss the application to forward supply networks. Finally, we conclude in Section 7.

2. Literature review

Our research can be related to several literature streams. First, it pertains to the facility location literature. Facility location models aim at finding the optimal placement of facilities and their links with other layers of the network, so that customer demand is satisfied at minimum cost. Melo et al. (2009) provide a thorough literature review focusing on supply chain management. The interested reader is also referred to the reviews provided by Klose and Drexl (2005) and ReVelle and Eiselt (2005). Our model can be classified as a deterministic single-period model with a single product, applied to a three-level network where the location decisions concern the intermediate layer. Other features are: capacitated vehicles, multiple sourcing, discrete locations, direct shipments from factories to customers, and inventory decisions.

Along with facility location decisions, an important characteristic of our model is that it incorporates inventory control. We rely on the classic economic order quantity (EOQ) control policy, which supposes deterministic demand and continuous review (see the seminal paper by Harris (1915), or Nahmias (2009)). Several EOQ extensions, solving similar continuous lot-sizing problems, have been studied (see the reference book by Axsäter (1980) for example). More specifically, our inventory control policy integrates multiple destinations and multiple sourcing (i.e. a facility may be supplied by several sources). As such, it is close to the EOQ extension known as “Single Resource Multi-Item Inventory Systems” (SRMIS). This EOQ extension considers a single factory that is replenished from multiple sources. The way the replenishments are coordinated impacts the computation of the average inventory at the factory. This problem was shown to be NP-complete by Gallego et al. (1992). The literature has primarily focused on finding the best possible coordination (Page and Paul, 1976; Rosenblatt, 1981; Gallego et al., 1996). In our work, as the problem at hand clearly has a broader scope, we will rely on an approximation of the average inventory at the factories, referred as perfect coordination (see Section 3.3). This approximation was implicitly used by Anily (1991) and Gallego et al. (1992) when they derived lower bounds for the problem. Lange and Semal (2010) discuss perfect coordination in detail and apply it for the allocation and lot sizing problem in a two-level supply chain (no location decision).

The case that motivated this research deals with the reverse network of a glass producer. Reverse logistics is now well established as a significant source of economies, and as an important lever to decrease the environmental impact of the supply chain. Some reference papers are those by Fleischmann et al. (1997) and Akçali et al. (2009), which review quantitative models for network design, and by Brito et al. (2002) which surveys over 60 case studies. The glass producer case concerns the return flow of reusable items, i.e. empty trestles, which have to be returned to factories to be refilled. One of the first papers to study the problem of reusable items is that by Crainic et al. (1993). It proposes an analysis of the inventory management and shipment planning, but does not deal with network design. Two interesting case studies concern the design of the network for reusable items. Inspired by the case of a Dutch logistics service organization, Kroon and Vrijens (1995) propose a classic facility location model (transportation and fixed costs) to choose the depots in a network designed to return the containers back to their original sender. Quite similarly, Jayaraman et al. (2003) propose a model for the design of reverse distribution networks. Due to the complexity of such mixed integer programs, they focus on the introduction and analysis of an evolved heuristic to solve the proposed model. These two papers concentrate on pure network design (location of depots), while the present work also deals with inventory management.

The main feature of our work is that it integrates supply network design with inventory management. Despite their acknowledged importance, so-called location-inventory models have only appeared quite recently and are still fairly scarce.

Early attempts date from around 2000. In their work, [Nozick and Turnquist \(1998\)](#) analyze how inventory costs, of safety stocks in particular, can be included in a classic distribution center location problem with one factory. They show that these costs are approximately linear in the number of DCs and can thus be included in the fixed costs. Subsequently, [Nozick and Turnquist \(2001\)](#) apply this approach to the real case of an automotive manufacturer. [Erlebacher and Meller \(2000\)](#) take a broader view. They formulate a highly nonlinear integer location-inventory model including facility fixed, transportation and inventory (cycle and safety stocks) costs. They do not attempt to solve it but rather propose a heuristic based on a stylized model, iterating on the number of opened DCs. They use a continuous space approximation and present computational results for problems with up to 16 customer areas. In their work, [Freling et al. \(2003\)](#) study a dynamic multi-period problem in a two-level supply chain where the demand is strongly seasonal. They focus on allocation and inventory decisions, considering transportation and inventory costs, but do not consider location decisions. [Huang et al. \(2005\)](#) propose a column generation for allocating demand in a two-level network. Their model integrates allocation, production and inventory (EOQ model) decisions but does not consider location decisions. More recently, [Üster et al. \(2008\)](#) focus on the location of a single warehouse. Their problem integrates the location decision and inventory replenishment decisions (reorder periods), and involves ordering, transportation and inventory costs. The authors propose three iterative heuristics to solve the problem. They also illustrate and discuss the interaction between location and inventory decisions, and the substantial cost savings that can be realized by integrating them. In their paper, [Sourirajan et al. \(2007\)](#) focus on the trade-off between the replenishment lead time and safety stock in the distribution center location problem, but neglect the transportation cost. The same problem is solved using genetic algorithms in ([Sourirajan et al., 2009](#)).

The papers most resembling our work are the distribution center location models integrating transportation, fixed and inventory costs. In a forward network with a single factory, the location model with risk pooling (LMRP) proposed by [Shen et al. \(2003\)](#) aims at locating the distribution centers and assigning the customers to them. The candidate locations for the distribution centers are those of the customers. Facility fixed costs, variable transportation costs and ordering costs for the replenishments are considered. The inventory costs include working inventory (in an EOQ fashion) and safety stock (due to variable demand) allowing for risk pooling. This nonlinear integer model is solved to optimality (with identical variance-to-mean ratios) for instances with up to 150 customers via column generation on a set-covering formulation in ([Shen et al., 2003](#)), or more efficiently using Lagrangian relaxation in ([Daskin et al., 2002](#)). The LMRP has then been extended by [Shen \(2005\)](#) to consider for multiple products, by [Ozsen et al. \(2008\)](#) to include facility capacities, and by [Ozsen et al. \(2009\)](#) to study the impact of multisourcing. [Shu et al. \(2005\)](#) propose an elaborate column generation methodology exploiting special structures of the LMRP. Their approach is able to solve more general (concerning variance of demand) as well as larger problems (up to 500 retailers). Similarly, [You and Grossmann \(2008\)](#) propose algorithms to solve the LMRP without assuming identical variance-to-mean ratios, and present examples with up to 150 retailers. In parallel to this stream of research, [Miranda and Garrido \(2004\)](#) develop a similar model, bearing the same cost trade-offs. Their solution approach is based on the Lagrangian relaxation and the sub-gradient method and allows them to solve instances with 10 candidate warehouse locations and 20 customers. Likewise, [Teo and Shu \(2004\)](#) study the trade-off between inventory, transportation and fixed costs to locate warehouses and allocate the retailers in a supply network where one outside vendor is used to directly replenish the warehouses. Furthermore, they consider the joint replenishment of inventory at warehouses and retailers. They devise a set-partitioning integer program and solve moderate size problems (20 candidate warehouses and 100 retailers) using column generation. [Shu \(2010\)](#) adapts the set-covering greedy algorithm to propose a heuristic that solves this problem for large size instances (up to 5000 retailers) with errors within 3–4% on average. [Romeijn et al. \(2007\)](#) extend the model of [Teo and Shu \(2004\)](#) to take safety stocks and capacity constraints on the facilities into account. They provide computational results for problems with up to 20 warehouses and 70 retailers.

The previous models (except [Erlebacher and Meller, 2000](#)) assume a single supplier or plant, i.e. they study two-level distribution networks, while few papers consider multiple suppliers. [Vidyarthi et al. \(2007\)](#) propose a design model for a three-level production-inventory-distribution system with multiple products that includes safety stocks and decisions on the plant and DC locations. The resulting nonlinear MIP is linearized using piecewise-linear functions and solved by Lagrangian relaxation (5% accuracy, up to 50 retailers). Similarly, [Park et al. \(2010\)](#) propose a design model for a three-level network with safety stock, but they incorporate DC-to-supplier dependent lead times. Decisions on the suppliers/DC location, on the allocation, and on inventory are integrated. Their Lagrangian relaxation approach solves instances with up to 60 retailers. [Lin et al. \(2009\)](#) develop an extensive design model for a four-level distribution network with deterministic demand. The proposed hybrid evolutionary algorithm is applied on instances with up to 60 customers.

In this paper, we propose an inventory-location model for a three-level supply network in which factories and customers have fixed locations and fixed constant production and demand rates. Our contribution firstly differs in the extent of the networks our method can design. We propose a continuous model leading to a heuristic procedure that iteratively solves a series of easy linear programs to find an approximate solution. As a result, our approach allows the design of large realistic supply networks: up to 10 factories, 1000 candidate distribution centers and 1000 customers. The approaches proposed by [Shu et al. \(2005\)](#) and [Shu \(2010\)](#) are the only ones able to design networks of equivalent sizes. However, they are limited to one supplier, i.e. a two-level network, and do not include the features mentioned below. Indeed, besides the network size, our model incorporates four features rarely considered in the literature but that

are more general and common in practice. Transportation is capacitated (never considered in the literature). Multiple sourcing is allowed at each supply chain layer (only considered by Ozsen et al. (2009) for a two-level network). Direct shipments between factories and customers are allowed (only in Lin et al. (2009)). Inventory is modeled at the factories and at the customers (Lin et al., 2009; Vidyarthi et al., 2007; Teo and Shu, 2004). Furthermore, even if they are not considered in the base model, we argue that capacities of the distribution centers and safety stocks (as by Nozick and Turnquist (1998)) can quite easily be included in the model (see Section 3.5).

3. Problem and model

3.1. Location-inventory model

The problem we study concerns the location-allocation of intermediary facilities in a three-level supply network, including the impact of inventory management decisions. Our aim is to design the supply chain network at a strategic decision level. Our model is based on average constant rates (demand and production), has a continuous timeline, and assumes a single product context. The flows are balanced, i.e. the total production equals the total demand. These assumptions are rather classic at the strategic level. Such models are usually referred to as single period planning horizon models (Melo et al., 2009; Kroon and Vrijens, 1995). The model is applied to a forward as well as a reverse supply chain. In a forward chain, the sources of the distribution network are the factories whereas the destinations are the customers. The intermediary facilities are distribution centers. Such a network is illustrated in Fig. 1. In a reverse chain, the customers are the sources and the flows are directed towards the factories. Intermediary facilities are recollection depots. In the remainder of the paper, we will refer to an intermediary facility as a distribution center (DC).

The model integrates three decisions: DC location, flow allocation, and shipment sizes on each link. The decisions are integrated, i.e. all are made at the same time from one model. The location decisions aim at locating the distribution centers, which constitute the intermediary layer of the supply network. The DCs are selected from a finite set of candidate locations for potential DCs. It is thus a discrete location model. When a facility is indeed opened, fixed operational costs occur. These include every cost that is not proportional to the number of items handled in the DC – for example the rent or amortization of the building, fixed administrative/staff costs, and taxes. Note that, by default, we suppose that DCs are not capacitated (the extension to capacitated DCs is discussed in Section 3.5).

The allocation decisions enable the network links among the layers to be established, and the customers and DCs to be allocated in order to satisfy the demand. We allow for direct flows between factories and customers (as shown in Fig. 1), i.e. the customers may be allocated either to a DC or directly to a factory. This has scarcely been addressed in the literature, as revealed by Melo et al. (2009). When items flow through a distribution center, they induce an additional handling cost. In this allocation process, one objective is to minimize the distances traveled by the items, and thus the transportation cost.

Furthermore, we allow for multiple sourcing in every layer, i.e. in a forward network, a DC may be replenished by several factories and a customer may be served by several DCs and/or factories. This assumption is necessary since the factory production rates and the customer demand rates are fixed. In general, it cannot be assumed that the sum of an integer number of customer demands will equal the production of one factory. Some customers have to be served (directly or not) by more than one factory, so that the production and demand rates are respected. Moreover, this




Fig. 1. Three-level supply network with distribution centers (forward case).

assumption (plus others) allows a continuous model to be formulated. Ozsen et al. (2009) argue that, given modern information technology, multisourcing should be considered, as it can bring benefits (when transportation costs increase in particular) and does not necessarily increase the complexity of the network significantly (the number of multisourced retailers is less than the number of DCs).

The inventory decisions concern the shipment sizes, i.e. the number of items in one vehicle, on each link of the network. Vehicles are assumed to be capacitated. In practice, this assumption is most often true (and when it is not, one can just assume a very large capacity). The number of vehicles is assumed to be unlimited. Each shipment implies the payment of a transportation cost, independent of the actual load of the vehicle (in particular, this cost structure applies when a third party logistics provider is used). Consequently, using large shipment sizes increases the truckload and decreases the transportation cost per item. However, this will create a larger inventory of items in the network, and involve larger inventory holding costs (including cost of capital and storage). In short, frequently shipping the items allows inventories to be kept limited, while accumulating larger inventories in distribution centers leads to lower transportation costs. This illustrates the important trade-off between transport and storage costs, and justifies the integration of location-allocation decisions and inventory management decisions (shipment sizes).

The trade-offs included in this problem are multiple and complex. In a distribution network, a customer may be supplied either from a factory or from a distribution center. If the customer is supplied directly from a factory, a shipment size is determined which balances the transportation and holding costs in an EOQ fashion. If the customer is supplied via a DC, several benefits and costs have to be balanced. The benefit that can be generated comprises two elements. First, the rate at which items are accumulated at the DC is higher than the rate of accumulation at the individual customer. This leads to a larger shipment size between factory and DC than between factory and customer. Items passing through a distribution center will thus be carried in larger shipments for the larger part of the distance, and will incur lower transportation cost. Second, the distance between customer and DC is likely to be smaller than between customer and factory. The shipment size between the DC and the customer will thus be smaller. Thus, the average inventory at the customer will be reduced. However, additional costs are also implied if the customer is supplied from a DC. First, the total distance traveled by the items is longer, as items make a detour via a DC. Second, an additional inventory is created at each DC that is selected. Third, as the shipment size between a DC and a customer tends to be smaller, the transportation cost is increased over this (small) portion of the route. Fourth, fixed operational and handling costs are due for each item flowing through a distribution center. The facility fixed cost is shared between the customers allocated to this DC. In a sense, the choice to open a DC follows from the opportunity to find a sufficient number of supplied customers to share the fixed cost so that the aforementioned benefits overtake the costs.

Furthermore, our objective is to propose an approach allowing realistic facility location-inventory problems in large supply chains to be solved. For example, the glass producer's supply chain includes 10 factories and around 500 customers. For this, we need to carefully set up the mathematical model. The variables and parameters are first detailed in Section 3.2, followed by the cost function in Section 3.3. Finally, we formulate the continuous mathematical model in Section 3.4.

3.2. Variables and parameters

The problem is now well established. Before presenting the mathematical model, we list the variables and parameters. The factories are given an index $f = 1 \dots n_F$ and the customers an index $c = 1 \dots n_C$. The candidate distribution center locations are given an index $d = 1 \dots n_D$. Below we introduce the decision variables considered, which are all continuous. There are the flow variables, which also reflect the network structure, and the shipment sizes. Note that a shipment size, i.e. the number of items per vehicle, is assumed to be continuous while in reality it is most often integral. The continuity assumption of shipment sizes is thus approximate. It is acceptable and has negligible impact on the results as long as the capacity of a vehicle is much larger than the volume of one item, i.e. when many items can be loaded into one vehicle.

λ_d^f is the flow between factory f and DC d , i.e. the average number of items sent per time period (e.g. one week) between f and d , in items/period.

λ_c^d is the flow between DC d and customer c , in items/period.

λ_c^f is the flow between factory f and customer c , in items/period.

q_d^f is the size of shipments sent between factory f and DC d (according to the EOQ policy), i.e. the load in vehicles traveling between f and d , in items/vehicle.

q_c^d is the shipment size between DC d and customer c , in items/vehicle.

q_c^f is the shipment size between factory f and customer c , in items/vehicle.

\mathcal{A}_d is the total flow passing through distribution center d , in items/period. We abbreviate this to DC flow in the remainder of the text (DC flows when referring to all DCs d). Note that not every candidate DC d is necessarily opened. Opened DCs correspond to candidate DCs for which $\mathcal{A}_d > 0$.

The parameters are listed below. They correspond to data, and thus depend on the application.

- A_f is the average production at factory f , in items/period.
- A_c is the average demand by customer c , in items/period. Note that the following balancing condition is assumed: $\sum_c A_c = \sum_f A_f$.
- O_d^f, O_c^d, O_c^f is the cost for using a vehicle in €, traveling between factory f and candidate DC d , between DC d and customer c and between factory f and customer c , respectively. It essentially depends on the distance between locations.
- H_f, H_d, H_c are the unit holding costs, in €/(item · period), at factory f , candidate DC d and customer c , respectively.
- F_d is the fixed operational cost per period (e.g. annual), in €, for locating a distribution center at candidate site d .
- M_d is the cost for handling one item flowing through DC d , in €/item.
- C_d^f, C_c^d, C_c^f is the capacity of a vehicle, in items/vehicle, between factory f and candidate DC d , between DC d and customer c and between factory f and customer c , respectively.

3.3. Cost function

In this subsection, we detail the costs (per period) at hand in the location-inventory problem we study. They will constitute the objective function of the optimization model. They include the transportation, holding, DC fixed and handling costs. The **transportation** cost between two facilities is given by the cost per vehicle (O_d^f, O_c^d, O_c^f) times the number of vehicles. The latter number equals the product flow ($\lambda_d^f, \lambda_c^d, \lambda_c^f$) divided by the number of items per vehicle (q_d^f, q_c^d, q_c^f). The transportation cost is thus nonlinear. The **handling** cost at one distribution center is simply given by the unit handling cost multiplied by the number of items passing through the DC, $M_d \cdot (\sum_c \lambda_c^d)$.

To compute the **fixed** operational cost of the distribution centers, we propose a formulation which avoids introducing binary variables. In other words, our formulation avoids the common expression, $\sum_d F_d \cdot y_d$, where y_d is a binary variable determining whether the DC d is opened. Instead, we only use the continuous flow variables introduced in the previous section. The sum $\sum_c \lambda_c^d$ is greater than zero if DC d is opened and otherwise equals zero. To mimic the binary variable, this sum can be divided by the flow through DC d , A_d (when $A_d > 0$). We thus get, for the total fixed cost, $\sum_{d, A_d > 0} F_d \cdot \frac{\sum_c \lambda_c^d}{A_d}$. This expression can be interpreted as the DC fixed operational cost shared among the customers served, according to their contribution in terms of flow through the distribution center. The fixed cost also turns out to be nonlinear. The reader is referred to Section 4 to see how this particular cost can be handled (A_d has to be estimated a priori).

The last cost to be defined is the total holding cost, which is crucial to reveal the impact of inventory decisions. It is given by the sum over all facilities of the unit holding cost times the average inventory level in this facility. To compute the inventory level, some modeling assumptions are required in order to keep the formulation tractable. The formulation has to be realistic and kept as simple as possible.

EOQ inventory. The demand rate is assumed to be constant and the inventory review system is assumed to be continuous. Accordingly, the classic EOQ inventory policy is applied. According to the EOQ model, shipments of items, of a given size, are regularly ordered to replenish the inventory when the latter is empty. They are shipped instantaneously (in other words the transportation time is negligible compared to the holding time). Even if based on simplifying assumptions, the EOQ policy has been shown to be quite robust, and valid at the strategic decision level (see Nahmias (2009) for example). It serves our purpose to study the trade-off between transportation and inventory at a strategic decision level.

Perfect coordination. We suppose that the partners in the supply chain are perfectly coordinated. By perfect coordination, we mean that shipments of items are only sent to a partner when the latter's inventory is empty. In other words, a partner (DC or customer) receives a new shipment only when it absolutely needs it, when its inventory is empty, but not earlier, and it thus minimizes the inventory it carries. The average inventory in the network is minimal under this assumption, corresponding to an ideal case. Perfect coordination is illustrated in Fig. 2: the customers supplied by the second DC receive a shipment when their inventory level reaches zero (same for the DCs, which receive it from the factory). Assuming perfect coordination in general is an approximation as specific proportionality between the production rate, consumption rates and shipment sizes is required for perfect coordination to be achievable. It allows the coordination issue to be simplified, which is not relevant at the strategic level. This assumption is reasonable as we focus on strategic decisions and as the distribution centers are centrally managed.

With these assumptions, the average inventory level at a location equals half of each shipment size built or consumed at this location weighted by the period of time it takes to build or consume it at this location. Next, we explain the computation of the inventory costs for the forward case, depicted in Fig. 2. The reverse case is very similar except that production and consumption are inverted. The reader can easily check that the formulas remain the same.




Fig. 2. Evolution of inventory level in each part of the supply chain, in the forward case (for one factory, the three distribution centers it serves, and the four customers the second DC serves).

Let us first consider the average inventory level at a factory. According to the previous assumptions, the factory builds an inventory of q_c^f (or q_d^f) items and then ships it to customer c (or DC d). Immediately after the shipment, the factory starts producing a shipment for the next customer or DC to be supplied, and so on. At the factory, the average inventory is thus the sum of half the shipment sizes weighted according to the period of time it takes to accumulate them, i.e.

$\frac{1}{\lambda_f} (\sum_c \lambda_c^f \frac{q_c^f}{2} + \sum_d \lambda_d^f \frac{q_d^f}{2})$. The same reasoning applies to the customers. They will receive a shipment from a factory or a DC at the precise moment when their inventory drops to zero. It follows that the average inventory at a customer is

$\frac{1}{\lambda_c} (\sum_f \lambda_f^c \frac{q_c^f}{2} + \sum_d \lambda_d^c \frac{q_c^d}{2})$. Regarding the inventories in distribution centers, there are flows coming from factories, and flows going to customers (see Fig. 2). In a “few to many” network, when there are more customers than factories (which is often the case), a distribution center will receive large shipments from the factories and then distribute the items to local customers (in the reverse case, items accumulate in the DC and are then transported in larger shipments to the factories). The inventory level is thus firstly determined by the sizes of the shipments coming from the factories, which are larger. Several shipments to customers progressively unload the inventory (the decrease is approximated as continuous, see Fig. 2). Consequently, the average inventory level in a DC is computed as $\frac{1}{\lambda_d} (\sum_f \lambda_f^d \frac{q_d^f}{2})$. Note that, similarly to the fixed operational cost expression, λ_d appears in the denominator of the inventory level formula. Thus, the particular expression that we choose for the fixed operational costs (without binary variable) does not complicate our mathematical model: the objective function would have been nonlinear anyway, because of the inventory cost.

3.4. Mathematical model

The problem we examine consists in minimizing the periodic transportation, inventory holding, DC fixed and handling costs detailed in the previous section, so that the flow conservation and capacity constraints are respected. We obtain the following nonlinear continuous mathematical model.

$$\min \sum_{d,f} O_d^f \cdot \frac{\lambda_d^f}{q_d^f} + \sum_{c,d} O_c^d \cdot \frac{\lambda_c^d}{q_c^d} + \sum_{c,f} O_c^f \cdot \frac{\lambda_c^f}{q_c^f} \tag{1}$$

$$+ \sum_f \frac{H_f}{2\lambda_f} \cdot \left(\sum_d \lambda_d^f q_d^f + \sum_c \lambda_c^f q_c^f \right) \tag{2}$$

$$+ \sum_{d, \lambda_d > 0} \frac{H_d}{2\lambda_d} \cdot \left(\sum_f \lambda_f^d q_d^f \right) \tag{3}$$

$$+ \sum_c \frac{H_c}{2\lambda_c} \cdot \left(\sum_f \lambda_f^c q_c^f + \sum_d \lambda_d^c q_c^d \right) \tag{4}$$

$$+ \sum_{d, \lambda_d > 0} \frac{F_d}{\lambda_d} \cdot \left(\sum_c \lambda_c^d \right) \tag{5}$$

$$+ \sum_d M_d \cdot \left(\sum_c \lambda_c^d \right) \tag{6}$$

$$\text{s.t. } \sum_d \lambda_d^f + \sum_c \lambda_c^f = A_f \quad \forall f = 1 \cdots n_F, \quad (7)$$

$$\sum_c \lambda_c^d = \sum_f \lambda_f^d = A_d \quad \forall d = 1 \cdots n_D, \quad (8)$$

$$\sum_d \lambda_d^c + \sum_f \lambda_f^c = A_c \quad \forall c = 1 \cdots n_C, \quad (9)$$

$$q_d^f \leq C_d^f \quad \forall f, d, \quad (10)$$

$$q_c^d \leq C_c^d \quad \forall d, c, \quad (11)$$

$$q_c^f \leq C_c^f \quad \forall f, c, \quad (12)$$

$$\lambda_d^f, \lambda_c^d, \lambda_c^f, q_d^f, q_c^d, q_c^f \geq 0 \quad \forall f, d, c. \quad (13)$$

The objective function aims to minimize the periodic transportation (1), inventory holding (2)–(4), fixed operational (5) and handling costs (6). The constraints (7)–(9) are respectively the flow conservation constraints at the factory, distribution center and customer levels. The capacity constraints (10)–(12) state that the shipment sizes cannot exceed the vehicle capacities. Finally, (13) gives the non-negativity constraints.

From this model, the three decisions we target can be made: the allocation of the flows in the network (given by λ_d^f , λ_c^d and λ_c^f), the location of the selected distribution centers (for which $A_d = \sum_c \lambda_c^d > 0$) and the shipment sizes (q_d^f , q_c^d and q_c^f).

3.5. Remarks: capacitated DCs and safety stocks

Before showing how this model can be solved, we comment on two extensions that can be taken into account without changing the formulation: capacitated distribution centers and safety stocks. First, in our modeling, we consider capacities of the vehicles, but the distribution centers are uncapacitated. However, the capacities of the DCs could easily be appended. As we suppose perfect coordination, a DC is either empty when a full vehicle arrives (forward case, see Fig. 2), or emptied when a vehicle leaves (reverse case). Consequently, the maximum inventory at a distribution center equals the size of the largest shipment passing through it. If the capacity of a DC d is denoted C_d , we get the following extended capacity constraints (vehicles and DCs).

$$q_d^f \leq \min [C_d^f, C_d] \quad \forall f, d, \quad (14)$$

$$q_c^d \leq \min [C_c^d, C_d] \quad \forall d, c, \quad (15)$$

$$q_c^f \leq C_c^f \quad \forall f, c. \quad (16)$$

It is clear from these equations that considering the capacities of the distribution centers does not change the model. For simplicity, we only considered vehicle capacities in constraints (10)–(12), but capacities C_d^f and C_c^d can be modified to “include” distribution center capacities.

Second, concerning the inventory costs, the model only includes working inventory. However, the cost of safety stocks could be included. Nozick and Turnquist (1998) show that, for a given stockout probability, the safety stock held at each distribution center varies with the square root of the number of DCs, and the total safety stock held in the distribution centers can be accurately approximated as a linear function of the number of DCs. In other words, the cost of the safety stocks can be incorporated in the fixed operational cost of a DC. In our model, inventory costs related to safety stocks can thus be included by updating the fixed cost F_d in term (5).

4. Solution procedure

In this section, we propose a solution procedure for the problem presented in the previous section. The objective function (1)–(6) is non-convex and the problem is clearly difficult. Our goal is to propose a simple heuristic able to design large supply chains and reveal the benefits of good inventory management. For this purpose, we take advantage of the particular formulation that we devised. We first show that the problem simplifies when the variables A_d , the total flows through DCs, are fixed.

4.1. Simplification with DC flows fixed

To begin with, we show how the optimal shipment sizes q_d^{f*} , q_c^{d*} and q_c^{f*} can be computed when flow variables λ_d^f , λ_c^d and λ_c^f (and thus A_d) are fixed. Accordingly, let us assume that variables λ_d^f , λ_c^d and λ_c^f have been fixed. We ignore the capacity constraint initially but will include it later. With these assumptions, the optimal shipment sizes q_d^{f*} can easily be derived

as the problem becomes unconstrained. For this purpose, the objective function is simply derived with respect to q_d^f and equaled to zero.

$$-\frac{O_d^f \lambda_d^f}{(q_d^f)^2} + \frac{H_d \lambda_d^f}{2A_d} + \frac{H_f \lambda_d^f}{2A_f} = 0.$$

We note that the variables λ_d^f may be removed. It can easily be checked that the second derivative is positive for all $q_d^f \geq 0$, so that the cost function has a unique minimum as a function of q_d^f . Then, the capacity constraint has to be taken into account: the shipment size has to be kept lower than the capacity. We thus get the following formula for the optimal shipment sizes:

$$q_d^{f*} = \min \left[C_d^f, \sqrt{\frac{2 O_d^f}{\frac{H_f}{A_f} + \frac{H_d}{A_d}}} \right], \quad \forall f, d. \tag{17}$$

The optimal shipment size q_d^{f*} thus only depends on A_d , and can be readily deduced from it. Similarly, the optimal shipment sizes transported on the DC-customer links and factory-customer links can be computed. We get:

$$q_c^{d*} = \min \left[C_c^d, \sqrt{\frac{2 O_c^d}{\frac{H_c}{A_c}}} \right], \quad \forall d, c, \quad q_c^{f*} = \min \left[C_c^f, \sqrt{\frac{2 O_c^f}{\frac{H_f}{A_f} + \frac{H_c}{A_c}}} \right], \quad \forall f, c. \tag{18}$$

Note that these two expressions for the optimal shipment sizes are independent of the flow variables, and can thus be computed once a priori. This result shows that, with the proposed problem formulation, optimal inventory decisions are very easily made for DC-customer and factory-customers links. This constitutes a first interesting result regarding the inclusion of inventory management in the facility location problem. Furthermore, the shipment sizes q_d^{f*} transported on DC-factory links only depend on flows through DCs, A_d .

The problem (1)–(13) in λ_d^f , λ_c^d and λ_c^f is more difficult, but it becomes linear when the DC flows A_d are fixed. Indeed, if the variables A_d are fixed, we have shown that the shipment sizes q_d^{f*} , q_c^{d*} and q_c^{f*} can easily be computed, and it can be seen that the problem Eqs. (1)–(13) in the flow variables λ_d^f , λ_c^d and λ_c^f then becomes a linear program. In other words, only the total flows through DCs, A_d , prevent such a linear program. The variables A_d appear in only a few terms. They can be found in the objective function: in the DC holding costs (8) and in the DC fixed costs (5). They also appear in the flow constraint related to the distribution centers (11). Furthermore, A_d appears in the optimal shipment sizes q_d^{f*} , in Eq. (17).

In brief, we have shown that, when the DC flows A_d are fixed, the shipment sizes q_d^f , q_c^d and q_c^f can be computed from closed-form equations, and the flow variables λ_d^f , λ_c^d and λ_c^f can be found simply by solving a linear program.

4.2. Heuristic

Our solution procedure aims at taking advantage of the property described in the previous subsection, as linear programs are easy to solve and very efficient tools are easily available. For this, flows through DCs, A_d , have to be fixed whereas there are variables of the problem. Accordingly, we propose an iterative heuristic which estimates and fixes the flows through all DCs in each iteration, solves the resulting linear program, and then uses the solution to improve the A_d estimations for the next iteration. In the next paragraph, we present the main stages of the solution procedure in an integrated way. The heuristic is then explained in more detail: we describe how the initialization of the procedure starts, the updating modes of the DC flow estimations, the meaning of the cycles and their stopping criteria. The heuristic is formally presented in Algorithm 1. A_d^{est} denotes an estimation of the flow through DC d . In this section, we also provide default values for the algorithm parameters. These default values will be used in our computational experiments (see Section 5). They are the outcome of numerous tests to find the algorithm parameters that lead to a good balance between accuracy of results and computational time.

To start the heuristic, the DC flow estimations A_d^{est} are initialized to A_d^{init} . These estimations will then be improved in the next iterations. In each iteration, the shipment sizes $q_d^f(A_d^{est})$ are first computed using these estimations A_d^{est} and Eq. (17). The linear program (1)–(13) can then be solved with these shipment sizes $q_d^f(A_d^{est})$ and the estimations A_d^{est} as inputs. It gives a feasible solution in terms of the flow variables λ_d^f , λ_c^d and λ_c^f as the constraints in λ_d^f , λ_c^d and λ_c^f are satisfied. From the flow variables, the actual DC flows $A_d^{sol} = \sum_f \lambda_d^f$ (satisfying the flow constraint (11)) and the corresponding shipment sizes $q_d^f(A_d^{sol})$ (17) are easily deduced, to obtain a complete feasible solution of the original problem. Next, the DC flows of this feasible solution, A_d^{sol} , are used to find improved estimations A_d^{est} . The estimations A_d^{est} are updated for all DCs in each iteration. The process is then iterated, solving the linear program with these new estimations in the cost function. It is stopped when the best found solution has not been improved during numerous iterations.

Algorithm 1. Heuristic

$A_d^{est} = A_d^{init}, \forall d,$
OutNoImprov = 0, *PrevSol* and *BestSol* set very large, and preprocessing.
while *OutNoImprov* < *MaxOut* and *CompuTime* < *MaxTime* **do**
 InNoChange = 0, *OutNoImprov* = *OutNoImprov* + 1
 while *InNoChange* ≤ 3*MaxIn* and *CompuTime* < *MaxTime* **do**
 InNoChange = *InNoChange* + 1
 Solve LP (1)–(13) with $A_d = A_d^{est}, \forall d,$ and $q_d^f = \min \left[C_d^f, \sqrt{\frac{2 O_d^f}{H_f/A_f + H_d/A_d^{est}}} \right], \forall f, d.$
 Save the solution as *CurrSol*, and compute $A_d^{sol} = \sum_f \lambda_d^f, \forall d.$
 if *InNoChange* ≤ *MaxIn* **then**
 $A_d^{impro} = A_d^{sol} + \sum_{\{c \text{ s.t. } d=am(c), \sum_{d'} \lambda_c^{d'}=0\}} A_c, \forall d,$
 $A_d^{est} = A_d^{est} - meetFac \cdot (A_d^{est} - A_d^{impro}), \forall d.$
 else if *InNoChange* ≤ 2*MaxIn* **then**
 $A_d^{est} = A_d^{est} - meetFac \cdot (A_d^{est} - A_d^{sol}), \forall d.$
 else
 $A_d^{est} = A_d^{est} \cdot (1 + meetFac), \forall d \text{ s.t. } A_d^{sol} > 0.$
 end if
 if *CurrSol* is better than *PrevSol* **then**
 InNoChange = 0.
 end if
 if *CurrSol* is better than *BestSol* **then**
 BestSol = *CurrSol*, and *OutNoImprov* = 0.
 end if
 PrevSol = *CurrSol*.
 end while
 end while
 The best found solution is given by *BestSol*.

To start the algorithm, initial estimations of the DC flows are needed, $A_d^{est} = A_d^{init}, \forall d.$ Choosing a high estimation A_d^{est} tends to favor the opening of DC $d,$ as the holding and fixed costs are underestimated (see Eqs. (8) and (5)). On the other hand, a low A_d^{est} will penalize the opening of DC $d.$ In particular, when A_d^{est} tends to zero, the chance of DC d being opened also tends to zero. When starting the algorithm, we cannot presume where distribution centers should be opened. In a sense, each should have a chance. Accordingly, A_d^{init} should be sufficiently large for each DC. $A_d^{init} = 5A_c^{avg}$ may be used as a default value, where A_c^{avg} denotes the average demand of one customer. Note that the impact of A_d^{init} can be significant, and can be used to improve the algorithm efficiency (see below).

Once initialized, the heuristic aims at iteratively improving the estimations A_d^{est} of the DC flows, and getting closer to the optimal solution of the problem. In this sense, it can be compared to a local search (on $A_d,$ or a descent algorithm. In each iteration, the linear program (1)–(13) in λ_d^f, λ_c^d and λ_c^f is solved, using $A_d^{est} \forall d,$ and $q_d^f(A_d^{est})$ computed from (17) as inputs. The DC flows A_d^{sol} are computed from the feasible solution. The algorithm then progressively updates the DC flows A_d^{est} in three alternative ways.

- A_d^{sol} can be used as they potentially offer better estimations of the optimal flows through the distribution centers. However, with this solution some DCs are closed ($A_d^{sol} = 0$), whereas they could be opened in the optimal solution, and some DCs are opened whereas they should not be. It would thus be misleading to directly equal the new estimation A_d^{est} to the values A_d^{sol} computed. The latter nonetheless gives a good direction for the update of A_d^{est} . In a first update mode, A_d^{est} is modified in this direction, but only for a fraction of the way: $A_d^{est} = A_d^{est} - meetFac \cdot (A_d^{est} - A_d^{sol}).$
- DCs are worth opening only if a sufficient number of customers are supplied by them. Consequently, in order to promote the opening of DCs and their use by customers, the direction given by A_d^{sol} is improved to A_d^{impro} . The demand flow of every customer which is not linked to any DC is added to the flow A_d^{sol} of the closest DC opened in the current solution.

$$A_d^{impro} = A_d^{sol} + \sum_{\substack{c \text{ s.t. } d=am(c), \\ \sum_{\text{all DC } d'} \lambda_c^{d'}=0}} A_c, \quad \forall d, \quad (19)$$

where $am(c) = \arg \min_{\text{open DC } d'} (dist(c, d'))$ is the opened DC which is closest to customer c . Similarly to the previous mode, the estimation A_d^{est} is then simply updated by removing a fraction of the difference: $A_d^{est} = A_d^{est} - meetFac \cdot (A_d^{est} - A_d^{impro})$.

- A third way to update A_d^{est} is used, again to increase the attractiveness of the opened DCs, so that more customers get the chance to use them. The estimations A_d^{est} of the flows through open DCs are artificially increased: $A_d^{est} = A_d^{est} \cdot (1 + meetFac)$.

From our tests, we found that $meetFac = 0.5$ is a good default parameter. As shown in Algorithm 1, the three update modes are used sequentially in the inner “while” cycle (starting with the more evolved one using A_d^{impro} ; by default, $MaxIn = 3$). The sequence is started again each time the cost function is decreased compared to the previous iteration (this does not mean that the best found solution was improved). This cycle thus allows the cost function to deteriorate and resets the counters to zero when the cost function is improved again. In this manner, the heuristic can potentially move out of a local optimum (as a function of A_d). In addition to this, the inner cycle is repeated until it does not lead to any improvement of the best found solution $MaxOut$ times (by default, $MaxOut = 1$). In other words, even if the current solution was not improved for the last $3MaxIn$ iterations, and if this cycle previously led to improvement of the best found solution, we launch a new cycle hoping to improve the solution again (as the actualization mode will change). The algorithm is finally stopped when the best found solution does not improve for $MaxOut$ cycles. Moreover, a stopping criterion is set on the computational time to avoid any convergence issue (by default, $MaxTime = 12$ h). Note that this maximum computational time was never reached in our experiments. It must be remembered that this solution is of course a feasible solution: the DC flows A_d , the shipment sizes q_d^f and the cost function are computed from the variables λ_d^f , λ_c^d and λ_c^f found from the linear program with $A_d = A_d^{est}$.

The idea of the proposed heuristic is thus simple: as the proposed formulation simplifies to a linear program when the variables A_d are fixed, a good solution can be found by iteratively solving easy linear problems with an improving estimation of the “nonlinear” variables A_d . Note that the procedure could also be used as a way to improve an existing solution, with all or part of the DC locations and flows A_d given.

The heuristic can be seen as a local search, with several descent techniques, that allows the solution deterioration to move out of a local minimum. In this context, the initial estimation A_d^{init} may have a significant impact on the results of the heuristic. To improve the accuracy of the algorithm, it can thus be transformed into a multiple start local search. For this, Algorithm 1 is run with different initial estimations A_d^{init} , and the final solution is then the best solution found from the different runs. The user can choose the number of runs depending on the computational effort that can be afforded. Our computational experiments show that the multiple start approach significantly improves the efficiency of the heuristic (see Table 2).

5. Numerical results

In this section, our aim is to assess the usefulness of our approach. The problem presented in Section 3 is non-convex and very difficult to solve to optimality. The accuracy of the proposed heuristic can thus not be directly assessed. However, several arguments may be put forward in favor of the value of the heuristic.

First, the problem that most closely resembles ours and that can be solved to optimality is the problem in which the shipment sizes q_d^f are fixed at their capacity, $q_d^f = C_d^f$. This leads to a mixed integer program (MIP). Binary variables y_d are introduced, determining whether distribution center d is opened. The holding costs at DCs (8) become $\sum_{f,d} y_d H_d \frac{C_d^f}{2}$, and the DC fixed costs (5) become $\sum_{f,d} y_d F_d$. Solved to optimality, this MIP leads to an upper bound on the total cost as the shipment sizes are fixed at their capacity, but, of course, it requires more computational effort than the linear programs solved by our heuristic. The comparison between the results of the heuristic procedure and the MIP reveals the ability of our heuristic to reveal the benefits resulting from inventory decisions concerning shipments between factories and DCs. Note that the solved MIP includes inventory decisions concerning shipment sizes q_c^d and q_c^f . Indeed, we showed in Section 4.1 that the optimal values q_c^d and q_c^f can be computed a priori (Eq. (18)). The heuristic is implemented in MATLAB® (with glpk) on a 2.16 GHz PC, 2 GB RAM. The MIP is solved using CPLEX®.

To assess the ability to reveal the benefits resulting from inventory decisions, we carried out extensive numerical experiments, based on representative parameter values. The parameter values were chosen as being of the same magnitude as the parameters provided by the case study (see Section 6.1). In all our experiments, the duration of one period equaled one week. For simplicity's and clarity's sake, the candidate DC locations were assumed to coincide with the customer locations. This assumption is classic and balanced. The set of customer locations offers a suitable grid for the candidate depots, as the geographic spreading is realistic as well as proportional. Moreover, we started the algorithm with $A_d^{init} = A_{c|d} + 4A_c^{avg}$, where $A_{c|d}$ denotes the demand flow of customer c located in the location d , and A_c^{avg} denotes the average demand of one customer. In other words, to begin with, we supposed that the customer that shares the same location as well as four other average customers were linked to DC d .

We first focused on supply networks with 3 factories and 100 customers (these are usual values, see Section 2). For each instance, the factories and customers were randomly located in a square (uniformly), with width equaled 1000 km, 2500 km or 5000 km. The production at factories and demand at customers were randomly chosen so that the average weekly demand of one customer was one item. Note that what is referred to here as one item could represent a batch of products in reality (with no impact on the results as shipment sizes are assumed to be continuous). The productions were uniformly distributed (between zero and twice the mean). Concerning the demands, two distributions were tested: a

uniform one (between zero and twice the mean) and a distribution that is often observed in practice, i.e. with 80–20 repartition (20% of the customers make 80% of the demand, where the flow values are uniformly distributed in each set). For each of the six combinations, five instances were randomly generated, leading to a total of 30 problem instances.

Various cost parameters were also tested to show their influence. Transportation costs were set directly proportional to distances (which increase depending on problem instances). A 1.1 €/km factor was used. The DC fixed operational cost and the holding costs clearly play a central role. The following values were used for the fixed costs: 0, 25, 50, 100 or 200 €/week. The holding costs were first the same in any location and vary as follows: 5, 10, 20 or 40 €/(item · week). To illustrate the order of magnitude of these costs, if a 25% annual interest rate is supposed for the holding cost, it corresponds to a market value of 1000 € ($\approx 5 \cdot 52/25\%$) to 8000 € for one item. Two different handling costs were tested: 2 or 5 €/item. The capacity parameter is also important. To be significant and potentially constraining, it may vary depending on the other parameters. At the beginning of each test, a reference capacity was thus computed. For this, the location-inventory problem was first solved without capacity, and the maximum shipment size q_d^f was saved to serve as a benchmark. In the experiments, the capacity was then fixed to 100%, 80%, 60% or 40% of this benchmark, to be more and more constraining (same in any link). This basic test set thus led to $30 \cdot 5 \cdot 4 \cdot 2 \cdot 4 = 4800$ problems and applications of the heuristic.

In this basic test set, the parameters were identical for every location and every link. This is clearly not always the case in practice. We thus added new experiments by slightly modifying the basic test set. In the first case, the holding costs at customers were five times higher than at factories and DCs. In practice, customers may dislike stocking items (at the producer's request) and may charge a higher rate for this. In the second case, the capacity of vehicles traveling to customers' locations (from factories or DCs) was smaller than that of vehicles traveling between factories and distribution centers. In practice, this may happen if larger vehicles are used to go to DCs (boat vs. trucks for example, see also the case in Section 6.1). In the tests, the capacities in $d - c$ and $f - c$ links were divided by 4, compared to the capacities in $f - d$ links, which were limited to 100% and 60% of the benchmark capacity (and thus 25% and 15% for the other links).

In total, we thus carried out 12000 experiments. The results are summarized in Table 1, showing the impact of the main parameters: DC fixed operational cost, holding cost and capacity. The average (over all the other parameter combinations, problem instances and handling costs) differences between the heuristic and the MIP results are given. Table 1 shows that our heuristic outperforms the MIP in around half of the parameter combinations. Unsurprisingly, the benefit procured by our heuristic increases when the holding cost increases and the capacity is less constraining. These situations correspond to the cases where the potential benefits of wise inventory management are larger. When the holding cost is low or when the vehicle capacity is small, it is advisable to load the vehicles close to capacity. In this case, estimating the shipment sizes, as opposed to equaling them to capacity, cannot result in much benefit. The heuristic does not find better solutions. In the other cases, when substantial benefits are possible, the heuristic is able to reveal them. To illustrate this, Table 1 shows the average loading rate of a vehicle (number of items compared to the capacity). It can be seen that when it is sufficiently smaller than 100% (say less than 85%), the heuristic improves the network design. Some exceptions can be found when the DC fixed cost is high, due to low ratios (flow through DCs/total flow) (also shown in Table 1), i.e. few items pass through DCs so that computing q_d^f does not result in much benefit. Looking at the impact of the DC fixed operational cost in more detail, we observe a similar effect. When the fixed cost increases, fewer DCs can be opened and only the most efficient are left, i.e. those which offer efficient transportation and thus those with large shipment sizes (see Section 3.1). Accordingly, a high fixed cost goes with large shipment size, as seen in Table 1, and also with less potential benefit resulting from inventory decisions. However, besides these arguments, we acknowledge that the heuristic seems to perform less well for high fixed costs. To overcome this drawback, part of the cost of a distribution center can be included in the handling costs.

The test sets with higher holding cost at customers or smaller vehicle capacity for shipments to customers tend to offer higher potential for benefits resulting from inventory management and DC opening. Thus, the heuristic is able to reveal benefits from inventory decisions q_d^f and often find better network designs than the MIP. In line with the previous arguments, it can also be observed in Table 1 that the loading rates are lower and the flow through DCs larger than in the basic test set. As to the other parameters, we observed from our experiments that distances and handling costs have little impact on the quality of the results obtained using the heuristic. The results tend to be slightly better for problem instances with an 80–20 repartition of customer demand.

A second set of experiments was run to assess the impact of the initial estimation A_d^{init} and show the improvement resulting from the multiple start local search. Our heuristic (Algorithm 1) is started with 20 different initial vectors of DC flows. For each candidate location d , A_d^{init} is randomly generated according to a uniform distribution with mean $fac \cdot A_c^{avg}$ (between 0 and $2 \cdot fac \cdot A_c^{avg}$), where fac is 1 or 5 (i.e. the default A_d^{init} we previously used thus corresponds to the mean in this case). The tested problem configurations were chosen from those previously used. The DC fixed operational cost can be 0 or 50 €/week. The handling cost is fixed to 2 €/item. The capacity is equal to 100% or 50% of the maximum shipment size observed in the noncapacitated case. The holding cost is 10 €/(item · week) and can be the same at any location (basic set) or five times larger at customers. For each parameter combination, 10 problem instances are randomly generated in a 2500 km-width square, half with a uniform customer flow repartition, half with an 80–20 repartition. This makes a total of 80 problem instances. For each instance, the heuristic ran with 20 different starting points. The minimum, the mean and the maximum solution are saved.

Table 2 gives the average results (among the network instances) for the various parameter combinations. It shows that the efficiency of the heuristic may vary significantly, but the average results of the 20 heuristic runs are similar to those observed with the default starting point heuristic. The default starting point is thus good, but it is difficult to isolate a starting

Table 1

Average gaps $(100(Res_{MIP} - Res_{Heur})/Res_{Heur})$ in percentage between heuristic and MIP (with $q_d^f = C_d^f$), when fixed cost F , holding cost H and capacity (percentage of benchmark) are varied. Results are given for the basic test set (identical values), the second case (higher holding cost at customers), and the third case (lower capacity of vehicles serving customers). Below the gaps, the average vehicle loading rate and the ratio (flow passing through DCs/total flow) are also given, both in percentages.

F	H	Basic test set				$H_c = 5H_d = 5H_f$				$C_c = C_d^f/4$		
		100%	80%	60%	40%	100%	80%	60%	40%	100%	60%	
0	5	2.6	1.7	-0.4	-1.3	24.1	18.4	12.4	6.9	3.2	0.0	
		66 51	78 54	87 50	95 47	48 100	59 100	72 100	86 100	64 63	81 76	
		3.3	1.6	-0.1	-1.3	24.6	18.8	12.7	7.0	3.7	0.0	
		64 56	75 58	84 58	94 52	48 100	59 100	72 100	86 100	61 67	79 77	
	10	3.7	2.3	0.5	-0.9	25.0	19.1	12.8	7.1	4.1	0.4	
		62 65	74 62	85 65	95 57	48 100	59 100	72 100	86 100	58 73	79 80	
		4.1	2.5	0.7	-0.7	25.2	19.2	12.9	7.2	4.4	0.6	
		59 68	72 67	83 70	93 64	48 100	59 100	72 100	86 100	59 72	76 82	
	25	5	0.1	-1.0	-1.5	-2.4	7.7	4.1	1.4	-0.7	1.0	-0.6
			75 44	85 44	94 41	98 38	50 98	59 99	74 98	88 99	74 51	91 63
			0.9	-0.1	-1.2	-2.1	10.4	6.8	3.2	0.4	1.9	-0.3
			69 56	83 53	92 55	97 52	50 99	62 99	74 99	88 99	71 57	89 68
10		2.0	0.7	-0.8	-1.9	13.5	9.3	5.1	1.6	2.5	0.0	
		69 56	81 58	90 57	96 56	49 99	61 99	73 100	87 100	67 65	86 72	
		3.0	1.2	-0.5	-1.4	17.3	12.3	7.3	3.1	3.4	0.4	
		67 62	76 65	88 61	96 62	48 100	58 100	71 100	85 100	64 66	85 73	
50		5	-1.5	-2.2	-2.7	-3.3	4.7	1.6	-0.9	-2.9	-0.5	-1.0
			77 33	88 32	87 29	86 18	55 88	64 89	76 91	89 92	83 40	94 59
			-0.2	-1.0	-1.5	-2.4	5.8	2.8	0.1	-2.0	0.8	-0.9
			75 43	88 42	96 41	99 33	50 96	59 97	73 96	89 97	78 47	92 63
	10	0.7	-0.4	-1.3	-2.7	8.1	4.6	1.6	-0.8	1.4	-0.5	
		71 53	84 51	94 48	98 47	49 98	60 98	73 98	87 99	71 56	90 66	
		1.8	0.5	-0.8	-1.9	10.8	7.0	3.3	0.5	2.0	-0.2	
		71 56	82 56	92 57	97 56	50 99	62 99	74 99	88 100	67 63	88 71	
	100	5	-3.0	-3.5	-3.7	-2.7	2.4	0.2	-1.4	-2.7	-1.6	-1.4
			53 15	54 14	56 13	44 10	59 80	72 79	84 80	95 82	86 35	95 51
			-2.4	-3.1	-3.7	-3.6	4.0	1.3	-1.4	-2.9	-1.6	-1.4
			73 28	80 26	76 26	66 16	54 86	69 85	79 88	93 87	84 36	95 56
10		-1.0	-2.1	-2.2	-3.3	5.0	2.0	-1.1	-2.8	-0.3	-0.9	
		78 38	87 37	92 32	94 27	53 90	65 90	77 91	88 94	81 45	93 63	
		0.2	-0.8	-1.4	-2.7	6.6	3.1	0.3	-2.1	1.0	-0.9	
		75 51	86 51	96 43	96 41	51 95	59 96	74 96	86 98	74 54	90 66	
200		5	-2.1	-1.9	-1.7	-0.8	0.4	-0.9	-2.1	-2.3	-2.4	-2.7
			28 9	27 8	22 5	3 0	70 63	82 64	92 63	99 63	82 29	97 42
			-3.1	-3.4	-3.1	-2.0	1.4	-0.5	-1.7	-2.8	-1.9	-1.8
			47 15	50 15	46 12	39 9	66 73	77 73	90 73	98 74	86 33	99 47
	10	-2.6	-3.0	-3.8	-3.2	2.8	0.3	-1.7	-2.6	-1.7	-1.5	
		62 21	68 20	61 16	53 12	60 81	73 80	85 81	96 82	83 36	97 52	
		-2.2	-3.1	-3.7	-3.7	4.3	1.3	-1.8	-3.4	-1.7	-1.3	
		79 35	88 33	91 34	79 20	55 87	67 87	78 87	90 89	81 39	95 56	

point that is a priori better than others. However, starting with large DC flows ($fac = 5$) tends to give better results. The results presented in Table 2 also show that the multiple start approach is a valuable way of improving our algorithm. It can be seen in Table 2 that, when the best result (max) found from the various runs is selected, the efficiency is significantly

Table 2

Differences in percentages observed between our heuristic and the MIP. First, the default heuristic is applied. Second, we give the mean, minimum and maximum differences when 20 random initial vectors A_d^{init} are used.

	$F = 0$				$F = 50$			
	Basic set		$H_c = 5H_{df}$		Basic set		$H_c = 5H_{df}$	
	100%	50%	100%	50%	100%	50%	100%	50%
Heuristic	3.6	-0.4	24.0	9.0	0.5	-2.1	4.7	-1.7
Multistart (fac)								
Avg. (5)	3.9	-0.6	23.7	7.0	-0.7	-2.5	6.7	-1.8
min max	1.2 5.4	-2.3 0.7	16.3 31.6	1.6 11.7	-3.7 1.6	-4.5 -1.1	2.5 10.4	-4.9 0.3
Avg. (1)	3.2	-1.9	21.2	1.8	-0.6	-3.1	7.4	-2.2
min max	-0.2 5.0	-4.0 0.3	11.6 35.1	-6.3 13.7	-3.9 1.2	-5.2 -1.3	3.2 11.4	-7.1 0.1




Fig. 3. Average computational time (on a logarithmic scale) over four problem instances with three factories (2500 km, half uniform, half 80–20 customer flow repartition), with $F_f = F_d = F_c = 0$, $M_d = 2$, $H_f = H_d = H_c = 10$, and the capacity 80% of the noncapacitated benchmark.

improved. In several configurations, it even allows benefits to be revealed from inventory management, which was not the case with the one-start heuristic.

A first objective of our approach is to integrate inventory management decisions and location/allocation decisions within a single framework. We have shown that the heuristic is able to reveal benefits resulting from inventory decisions. Furthermore, an important objective of the methodology is to be able to design large supply networks, of realistic size. This is why we proposed a continuous formulation (with no integer variable), and a formulation that resembles a linear program. Fig. 3 illustrates the increase in the computational time when the size of the supply network increases. Unsurprisingly, the computational time increases faster for the mixed integer program than for our heuristic based on linear program solving. Less time is required to solve the LP and the number of iterations does not increase dramatically (89, 103, 113, 122, and 131 with $n_c = 100, 200, 300, 500, 1000$). For networks comprising more than 300 customers, i.e. when computing time becomes important, our heuristic is faster. This represents an important advantage of our approach: it can handle large supply chains. For example, a network with 10 factories, 1000 candidate DC locations and 1000 customers can be designed in around four hours, making location, allocation and inventory decisions, which is reasonable for a strategic problem (and can be compared to problems involving one factory and 100 customers handled in the literature, see Section 2). Also note that the heuristic could be stopped before the end, based on a maximum computational time criterion, as the main improvements of the objective function are of course made at the beginning of the algorithm. In a sense, our heuristic can thus also be seen as an alternative to mixed integer programs for solving location problems. It enables large supply networks that are beyond the scope of other methods to be designed with a low error level (see Table 1).

6. Application

In this section, we discuss the application of our methodology to a particular problem. First, we show how our approach can be applied to reverse supply chains and we analyze the real case proposed by a glass manufacturer. Then, we discuss the application to forward supply chains, i.e. three-level distribution systems.

6.1. A reverse logistics problem

Our methodology can clearly be applied in the case of a reverse supply chain, where the flow is generated in the customer layer and goes upstream to the factories, possibly via the distribution centers. For instance, the flow variable λ_{fd}^f , hitherto denoting the flow between factory f and depot d , represents the flow from d to f in a reverse network. Traditionally, forward supply chains have a “few-to-many” shape, i.e. from few factories to many customers. Recovery networks thus generally have a “many-to-few” configuration. Flows are generated in small quantities from many sources. In this context, of course, the flow conservation constraints (7)–(9) remain valid, and the inventory levels can be computed in the same way (in Fig. 2, the orientation of the sawteeth is simply reversed).

Furthermore, in the design of a recovery network for reusable items, an important decision concerns the fleet size, in our case the number of trestles required for the flow of products to move freely. Our model allows the number of empty trestles to be computed as a side benefit, giving an estimation of the additional number of trestles needed in the fleet. Indeed, the number of reusable items can easily be estimated as follows:

$$I = \sum_{f,d} \frac{\lambda_d^f q_d^f}{A_f} \frac{1}{2} + \sum_{f,c} \frac{\lambda_c^f q_c^f}{A_f} \frac{1}{2} + \sum_{d,f} \frac{\lambda_d^f q_d^f}{A_d} \frac{1}{2} + \sum_{c,f} \frac{\lambda_c^f q_c^f}{A_c} \frac{1}{2} + \sum_{c,d} \frac{\lambda_c^d q_c^d}{A_c} \frac{1}{2}.$$

We now illustrate the application of our methodology to the case of the glass producer. The company mainly manufactures glass panes for the automotive and construction industries. Presently, the glass panes are carried on trestles by especially designed trucks to customers all over Europe. The transportation is outsourced to several transport companies. The latter are also responsible for returning the empty trestles. However, when empty, the trestles can be folded and therefore require much less room in the truck. Two loaded trestles are shipped in a forward flow truck, while 20 trestles can be shipped when folded. Accordingly, the transport company charges the return trip at only a fraction of the fully loaded trip cost. However, due to the inconvenience caused, this fraction is clearly higher than one tenth, the fraction of space allocated to empty trestles in the truck. Generally, the transport company would try to make an additional profit from the return trip by loading the free space, but it cannot carry a standard full truckload. In this context, the glass producer wishes to assess the advisability of an alternative strategy for the return flow: accumulating empty trestles in regional depots, in order to return up to 20 of them in a truck to the factories. We write “up to” deliberately as the trucks are not necessarily fully loaded (depending on the best balance of costs). The fact that more (folded) reusable items can be shipped by one truck in the return flow than in the forward flow is naturally an important motivation for the introduction of depots.

The supply network of the glass producer comprises 10 factories and 479 customers (which may be retailers), spread throughout Europe. The data includes the location coordinates and dates from 2008. The total flow, among all factories, equals 458 trucks (loaded with glass) per week, i.e. 22,883 per year, corresponding to 915 trestles sent per week. Factory productions A_f range from 10 to 176 trestles per week, while customer demands A_c range from 2 reusable items per year to 88 items per week. As in the numerical experiments, the potential locations for installing the depots coincide with the locations of the customers.

The various cost parameters were assessed in collaboration with the company’s supply chain team. The transportation cost matrices are deduced from the distance matrices. The maximum distance between a customer and a factory (respectively between two customers) is 3252 (respectively 3857) kilometers. For transport from depots to factories, standard trucks can be used to ship up to 20 folded trestles ($C_d^f = 20$), and a 1.1 €/km factor is applied. When trestles are loaded with glass panes however, special trucks have to be used, shipping two trestles. The same trucks are then used to ship the empty trestles back from the customers, either to the distribution centers or the factories, without stocking at customer locations ($H_c = 0$). The size of shipments leaving customers are thus fixed at $q_c^f = q_c^d = 2$. The special trucks cost 1.35 €/km. Moreover, for direct shipments from customers to factories, when the distance is sufficiently long, the special trucks can be loaded with another charge, as they are only 10% full. In this case, a supplementary factor of 0.3 is thus used, when the distance is longer than 150 km (i.e. 30% of the cost, compared with 10% of the space). Furthermore, the unit holding cost (H_d and H_f) of an empty container is estimated to be 10 €/week. The fixed operational cost of a depot equals 100 €/week and the handling cost is 5 €/unit.

The methodology proposed in this paper has been used to help decide whether depots should be included in the recovery supply network of the glass producer. It also provides guidelines as to how the network should be designed, how many depots should be opened and how loaded the trucks should be. The cost of the best network found by our heuristic is 46,667 €/week. The largest part of the costs is due to the transportation (32,584 €/week), particularly from depots to factories (21,669 €/week). The holding cost accounts for 5518 €/week.

The reverse supply network comprises 51 depots (out of 479 candidate locations). On average, 4.3 customers are linked to a depot. Around 76% of the total flow, coming from 46% (219/479) of the customers, passes through the depots. The empty trestles that have to be shipped back from the other 262 customers (54%), corresponding to 24% of the total flow, go directly to the factories. It is clearly apparent that customers with higher demand (average flow of 3.2 items/week) tend to go through depots, while customers with a lower flow (0.9 items/week on average) ship the empty trestles directly to factories. This can be explained as follows: the profit resulting from aggregating inventory is higher for the former, and the detour and the depot fixed cost are thus justified.

Of course, the distance factor also plays an important role. When a customer directly ships the items to a factory, the distance is not great: 180 km on average. When a customer is linked to a distribution center, the average distance from a DC to a factory (respectively from a customer to a depot) is 544 km (respectively 44 km, for non-zero distances). For such long trips, it is worth aggregating items in depots and loading the trucks with nearly 20 folded trestles. Indeed, from depots to factories, the number of items in a truck is 19.3 on average and equals 20 (full load) for 88% of the trucks. Also note that, concerning the fleet sizing, the optimal number of empty trestles in the network is estimated to be 1031.

In Fig. 4, we illustrate the reverse supply network, including the depots, as found by our heuristic and as presented to the glass producer. However, for confidentiality reasons, the network depicted is not exactly the network found in the glass producer case. The positions and flow values are randomly modified but the total flow and number of facilities are the same. Fig. 4 shows the solution found by our heuristic to this slightly modified problem. It can be seen that the depots are spread all over Europe.

The proposed solution can be compared to the actual recovery network, without depots. The optimal network without depots can be found by solving a classic transportation problem, leading to a total cost of 89,031 €/week. The estimated profit, 42,318 €/week, i.e. more than 2 millions € per year, is thus considerable. The cost of recovering the empty trestles



Fig. 4. Illustrative map of the reverse network. Stars represent factories, diamonds depots, and dots customers.

is reduced by 48%. This clearly motivates the glass manufacturer to make the effort to incorporate depots in the recovery network.

6.2. Forward supply chain

Our approach is of course also able to design forward supply networks. It can locate the distribution centers, choose a good clustering, and assess shipment frequencies. In distribution networks, the factories are the source of the flow. The items are shipped from the factories to the customers, potentially through the distribution centers. Most often, a distribution network thus has a “few-to-many” configuration.

Our methodology supposes a continuous mono-item production, so that items are produced and shipped directly to the customers when a vehicle can be filled with the appropriate quantity (q_d^f or q_c^f). This is acceptable for products like glass panes, fresh milk or televisions for example (one manufacturing line, one product). However, when lines manufacture several products, the continuous mono-item production assumption can be questioned. In this case, and particularly when set-up costs/times are high, items are produced in large batches, much larger than vehicle sizes. The products are thus stored at factories in relatively large quantities before being transported to the DCs and customers. In this sense, inventory in factories is not created for the efficient shipment of items through the network, but rather for production reasons. In this case, holding costs should thus not be taken into account in the design of the distribution network, or, more precisely, they should be reinterpreted. Our approach can easily be adapted. The holding costs at factories are fixed to zero, $H_f = 0$, and the other holding costs, H_d and H_c , are reinterpreted as the supplementary cost of storing at DCs and customers rather than in factories (space cost, different countries, etc.). Note that in the limit case where $H_f = H_d = H_c = 0$, the proposed mathematical model (1)–(13) of course reduces to a classic facility location problem. Our solution methodology remains applicable and particularly to large networks.

7. Conclusion

In this paper, we present a new approach to design supply networks while integrating inventory decisions. Assumptions regarding inventory control (EOQ logic and perfect coordination) are made to integrate tactical/operational information at an appropriate level of detail in this strategic problem. They allow us to formulate a model that is both realistic and tractable. An iterative heuristic is proposed to solve the model and is shown to be both effective and efficient using extensive numerical experiments. Finally, its applicability to real-life problems is illustrated by means of a case study.

The contribution of the paper to the literature lies mainly in the size and features of the supply networks that can be designed using our approach. First, a new location-inventory model is proposed and includes features that are generally applicable and common in practice, but rarely considered in the literature: capacitated transportation, multiple sourcing, direct shipments and inventory in every layer. Second, an iterative heuristic is devised to solve our continuous non-convex formulation. This heuristic solves one linear program at each iteration and is thus able to design large networks (1000 customers). To our knowledge, there is no other approach in the literature able to design three-level supply networks of such sizes. Furthermore, we show that the sizes of shipments between factories and customers, and between DCs and customers, can be computed a priori, independently of the location-allocation decisions.

The heuristic we propose is intended to be simple and to directly take advantage of the specificities of our problem formulation. However, in the future, more evolved solution procedures could be devised, based on non-convex optimization techniques. Moreover, our approach could quite easily be extended to take intra-layer flows into account. Future research could also aim to analyze vehicle tours, another way to locally consolidate shipments.

Acknowledgments

This research was funded by the Walloon Region of Belgium in the context of the TransLogisTIC project. The authors are grateful to the supply chain and distribution department team of AGC Flatglass Europe, and in particular to O. Herinck, E. Lejeune and D. Feron for their active collaboration in providing the case. The authors are also grateful to the anonymous referees for their helpful comments.

References

- AkÇali, E., Çetinkaya, S., Üster, H., 2009. Network design for reverse and closed-loop supply chains: an annotated bibliography of models and solution approaches. *Networks* 53, 231–248.
- Anily, S., 1991. Multi-item replenishment and storage problems (MIRSP): heuristics and bounds. *Operations Research* 39, 233–239.
- Axsäter, S., 1980. Economic lot sizes and vehicle scheduling. *European Journal of Operational Research* 4, 395–398.
- Brito, de, M., Dekker, R., Flapper, S., 2002. Reverse Logistics: A Review of Case Studies. Tech. Rep. Econometric Institute EI 2002–21, Erasmus University Rotterdam.
- Crainic, T., Gendreau, M., Dejax, P., 1993. Dynamic and stochastic models for the allocation of empty containers. *Operations Research* 41 (1), 102–126.
- Daskin, M., Coullard, C., Shen, Z.-J., 2002. An inventory-location model: formulation, solution algorithm and computational results. *Annals of Operations Research* 110, 83–106.
- Erlebacher, S., Meller, R., 2000. The interaction of location and inventory in designing distribution systems. *IIE Transactions* 32, 155–166.
- Fleischmann, M., Bloemhof-Ruwaard, J., Dekker, R., van der Laan, E., van Nunen, J., Wassenhove, L.V., 1997. Quantitative models for reverse logistics: a review. *European Journal of Operational Research* 103 (1), 1–17.
- Freling, R., Romeijn, H., Romero Morales, D., Wagelmans, A., 2003. A branch-and-price algorithm for the multiperiod single-sourcing problem. *Operations Research* 51, 922–939.
- Gallego, G., Shaw, D., Simchi-Levi, D., 1992. The complexity of the staggering problem and other classical inventory problems. *Operation Research Letters* 12, 47–52.
- Gallego, G., Queranne, M., Simchi-Levi, D., 1996. Single resource multi-item inventory systems. *Operations Research* 44, 580–595.
- Harris, F., 1915. What quantity to make at once. *The Library of Factory Management* 5, 47–52.
- Huang, W., Romeijn, H., Geunes, J., 2005. The continuous-time single-sourcing problem with capacity expansion opportunities. *Naval Research Logistics* 52, 193–211.
- Jayaraman, V., Patterson, R., Rolland, E., 2003. The design of reverse distribution networks: models and solution procedures. *European Journal of Operational Research* 150 (1), 128–149.
- Klose, A., Drexl, A., 2005. Facility location models for distribution system design. *European Journal of Operational Research* 162, 4–29.
- Kroon, L., Vrijens, G., 1995. Returnable containers: an example of reverse logistics. *International Journal of Physical Distribution and Logistics Management* 25 (2), 56–68.
- Lange, J.-C., Semal, P., 2010. Design of a Network of Reusable Logistic Containers. CORE Discussion Paper 2010/9, Université catholique de Louvain.
- Lin, L., Gen, M., Wang, X., 2009. Integrated multistage logistics network design by using hybrid evolutionary algorithm. *Computers and Industrial Engineering* 56, 854–873.
- Melo, M., Nickel, S., Saldanha-da-Gama, F., 2009. Facility location and supply chain management – a review. *European Journal of Operational Research* 196 (2), 401–412.
- Miranda, P., Garrido, R., 2004. Incorporating inventory control decisions into a strategic distribution network design model with stochastic demand. *Transportation Research Part E* 40, 183–207.
- Nahmias, S., 2009. *Production and Operations Analysis*, sixth international ed. McGraw-Hill/Irwin, New York.
- Nozick, L., Turnquist, M., 1998. Integrating inventory impacts into a fixed-charge model for locating distribution centers. *Transportation Research Part E* 34 (3), 173–186.
- Nozick, L., Turnquist, M., 2001. Inventory, transportation, service quality and the location of distribution centers. *European Journal of Operational Research* 129 (2), 362–371.
- Ozsen, L., Coullard, C., Daskin, M., 2008. Capacitated facility location model with risk pooling. *Naval Research Logistics* 55, 295–312.
- Ozsen, L., Daskin, M., Coullard, C., 2009. Facility location modeling and inventory management with multisourcing. *Transportation Science* 43, 455–472.
- Page, E., Paul, R., 1976. Multi-product inventory situation with one restriction. *Journal of the Operational Research Society* 27, 815–834.
- Park, S., Lee, T., Sung, C., 2010. A three-level supply chain network design model with risk-pooling and lead times. *Transportation Research Part E* 46, 563–581.
- ReVelle, C., Eiselt, H., 2005. Location analysis: a synthesis and survey. *European Journal of Operational Research* 165, 1–19.
- Romeijn, H., Shu, J., Teo, C.-P., 2007. Designing two-echelon supply networks. *European Journal of Operational Research* 178, 449–462.
- Rosenblatt, M., 1981. Multi-item inventory system with budgetary constraint: a comparison between the lagrangian and the fixed cycle approach. *International Journal of Production Research* 19, 331–339.
- Shen, Z.-J.M., 2005. A multi-commodity supply chain design problem. *IIE Transactions* 37, 753–762.
- Shen, Z.-J.M., Coullard, C., Daskin, M., 2003. A joint location-inventory model. *Transportation Science* 37, 40–55.
- Shu, J., 2010. An efficient greedy heuristic for warehouse-retailer network design optimization. *Transportation Science* 44, 183–192.
- Shu, J., Teo, C., Shen, Z., 2005. Stochastic transportation-inventory network design problem. *Operations Research* 53, 48–60.

- Sourirajan, K., Ozsen, L., Uzsoy, R., 2007. A single-product network design model with lead time and safety stock considerations. *IIE Transactions* 39, 411–424.
- Sourirajan, K., Ozsen, L., Uzsoy, R., 2009. A genetic algorithm for a single product network design model with lead time and safety stock considerations. *European Journal of Operational Research* 197, 599–608.
- Teo, C.-P., Shu, J., 2004. Warehouse-retailer network design problem. *Operations Research* 52, 396–408.
- Üster, H., Keskin, B., Çetinkaya, S., 2008. Integrated warehouse location and inventory decisions in a three-tier distribution system. *IIE Transactions* 40, 718–732.
- Vidyarthi, N., elebi, E., Elhedhli, S., Jewkes, E., 2007. Integrated production-inventory-distribution system design with risk pooling: model formulation and heuristic solution. *Transportation Science* 41, 392–408.
- You, F., Grossmann, I., 2008. Mixed-integer nonlinear programming models and algorithms for large-scale supply chain design with stochastic inventory management. *Industrial and Engineering Chemistry Research* 47, 7802–7817.