Inventory Routing

Luca Bertazzi[1], Martin Savelsbergh[2], and Maria Grazia Speranza[1]

[1] Department of Quantitative Methods
University of Brescia
Italy
bertazzi@eco.unibs.it, speranza@eco.unibs.it

[2] H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205
mwps@isye.gatech.edu

Summary. In this chapter, we introduce inventory routing problems. Inventory routing problems are among the more important and more challenging extensions of vehicle routing problems, in which inventory control and routing decisions have to be made simultaneously. The objective is to determine distribution policies that minimize the total cost, i.e., the sum of inventory holding and transportation costs, while avoiding stock-outs and respecting storage capacity limitations. All inventory routing problems have some common characteristics, but they may also have a number of significantly different characteristics. As a result, a variety of solution approaches has been developed. We discuss the various characteristics of inventory routing problems in order to create an understanding of and instill an appreciation for the complexities of inventory routing problems.

Key words: Logistics; supply chain management; inventory routing survey.

1 Introduction

The class of inventory routing problems (IRPs) is large and the number of solution approaches that have been proposed for their solution is even larger. Inventory routing problems all share some basic characteristics. They all consider environments in which products are shipped from a supplier to one or more customers by means of, usually capacitated, vehicles. Costs are incurred for the distance traveled by the vehicles and those costs are included in the objective function. This characteristic explains the word *routing* in the name of the problem class. What makes this class of problems significantly different from most other classes of routing problems is the presence of an inventory component, which explains the word *inventory* in the name of the problem class. The inventory component arises because customers consume product

B. Golden et al. (eds.), *The Vehicle Routing Problem*,
doi: 10.1007/978-0-387-77778-8_3,

over time and have only limited storage capacity. The supplier has to manage product inventory at customers to ensure that customers do not experience a stock-out. The inventory component thus adds a time dimension to the traditional spatial dimension of routing problems. The presence of inventory complicates the routing decisions in two fundamental ways. First, the limited inventory holding (storage) capacity at the supplier and/or at the customers has to be taken into account when deciding on delivery quantities. Second, inventory holding costs, at the supplier or at the customers, may be incurred which has to be accounted for in the objective function.

Beyond these basic characteristics, there is a variety of other characteristics that may significantly change the structure of a particular inventory routing problem, such as

- the planning horizon can be finite or infinite;
- inventory holding costs may or may not be considered;
- inventory holding costs may be charged at the supplier only, at the supplier and the customers, or at the customers only;
- the production and consumption rates can be deterministic or stochastic;
- production and consumption take place at discrete time instants or take place continuously;
- production and consumption rates are constant over time or vary over time;
- the optimal delivery policy can be chosen from among all possible policies or has to be chosen from among a specific class of policies.

In this chapter, we consider only inventory routing problems involving the distribution of a single product over a finite planning horizon with deterministic and stationary production and consumption rates. This simple, yet surprisingly complex, setting allows us to introduce the reader to some of the issues arising in inventory routing. Section 2 focuses on the impact of inventory holding costs and storage capacities on delivery policies in a discrete time setting. Section 3 focuses on the impact of consumption rates and storage capacities on delivery policies in a continuous setting. In Section 4 we provide a brief overview of the literature on inventory routing problems.

We start by introducing notation for the elements that are common to the inventory routing problems we discuss in Sections 2 and 3. Inventory routing problems are defined on a graph $G = (V, E)$, where $V = \{S, 1, \ldots, n\}$ is the set of vertices and E is the set of edges. Vertex S represents the supplier and vertices $1, \ldots, n$ represent the customers. A travel time t_{ij} and a cost c_{ij} are associated with edge $(i, j) \in E$. The capacity of each vehicle is Q. If time is discrete, we denote by q_i the quantity of product consumed per unit of time by customer i. If time is continuous, we denote by u_i the usage or consumption rate. The initial inventory level at the supplier is denoted by I_S^0. The initial inventory at customer i is denoted as I_i^0. Initial inventories can either be given or be decision variables. Considering initial inventories as decision variables may significantly improve the quality of the solution.

We denote by I_i^t and I_S^t the inventory level at customer i and at the supplier at time t, respectively. The inventory holding capacity at customer i is C_i and at the supplier is C_S. The inventory holding cost at the supplier is h_S and at customer i is h_i. Whenever a cost is not considered in a particular variant, its value is assumed to be 0. The length of the planning horizon is denoted by H.

The decisions to be made are (1) when to deliver to each customer, (2) how much to deliver to each customer each time it is served, and (3) how to route the vehicles so as to minimize the total cost. The total cost always includes the transportation costs incurred by the vehicles and may or may not include the holding costs incurred at the supplier and the customers. A delivery policy has to ensure that the supplier and the customers do not experience any stock-outs, that storage capacities at the supplier and the customers are not exceeded, and that the vehicle capacity limit is respected.

The routing component of inventory routing problems by itself already makes the problems hard. In fact, the problems reduce to the traveling salesman problem when the planning horizon is one, the inventory costs are zero, the vehicle capacity is infinite, and all customers need to be served. Furthermore, even when only one customer is considered, some variants remain computationally hard (see Speranza and Ukovich [55]).

We will not provide a detailed survey of the vast number of papers on inventory routing problems that have been published in the literature. A number of surveys already exist (e.g., [21, 30, 34]) and we refer the reader to these surveys for an in-depth overview of this fertile area of research. Instead, we focus on a few variants and hope to generate an appreciation for the variety and richness of inventory routing problems and for the challenges and pitfalls encountered when trying to construct optimal or high quality solutions, and to stimulate a desire to delve deeper into these fascinating problems. To facilitate such investigations, we briefly summarize the literature on IRPs in Section 4.

2 Inventory Holding Capacity and Costs

In this section, we consider a discrete time setting. We take the example introduced by Bell et al. [9] and investigate several simple modifications, involving inventory holding capacity and inventory holding costs, to demonstrate the impact of the presence of inventory on the routing decisions. We first investigate the case in which only the transportation costs are included in the objective function, then the case in which inventory holding costs are added to the transportation costs. In both these cases the initial inventory levels are fixed. Then, we study the case in which inventory holding costs are charged, together with the transportation costs, and the initial inventory levels are decision variables.

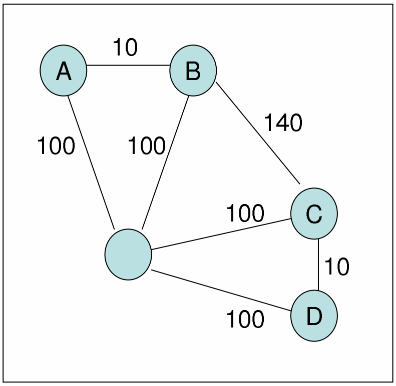

Fig. 1. Bell et al. example.

In the example of Bell et al. [9] there are four customers to be served by a single supplier. Time is discrete. Let $I = \{A, B, C, D\}$ be the index set of the customers. Figure 1 shows the available connections with their associated travel costs. The capacity C_i and quantity q_i of each customer $i \in I$ are given in Table 1.

The initial inventory level of each customer i is equal to its holding capacity, that is $I_i^0 = C_i$ for all $i \in I$. There is an unlimited number of vehicles with capacity $Q = 5000$. In the original description of the example, nothing is specified about the supplier. It is understood that there are no limitations on product availability, i.e., $C_S = \infty$. The objective is to find a periodic distribution policy that minimizes the total cost, that does not cause a stock out at any of the customers, and that does not exceed the storage capacity at the customers and the vehicle capacity. The periodicity of the policy implies that the inventory levels at the end of the period must be equal to the initial levels. The horizon is implicitly assumed to be infinity; the periodic policy is repeated over and over. From Figure 1 it is easy to see that a natural distribution policy is to combine the two pairs of close customers A and B and C and D, and serve each pair with a separate vehicle. The daily cost is 420. A better distribution policy repeats every two days. On the first day, A and B

Table 1. Capacity and consumption in the Bell et al. example.

Customer i	Capacity C_i	Consumption q_i
A	5000	1000
B	3000	3000
C	2000	2000
D	4000	1500

are served together with one vehicle and C and D are served together with another vehicle. A quantity 2000 (twice the daily consumption) is delivered to A, 3000 (the daily consumption) to B, 2000 (the daily consumption) to C and 3000 (twice the daily consumption) to D. Both vehicles fully utilize their available capacity. Since customers B and C have to be served daily, due to their storage capacity and daily consumption, on the second day customers B and C are served together in a tour, delivering a quantity of 3000 to B and of 2000 to C, while A and D are not visited. The average daily cost is 380. This solution has been proven to be optimal by Adelman [1] and Song and Savelsbergh [53].

Observe that the actual timing of deliveries during the day is not considered. Customer B consumes 3000 units of products per day and receives 3000 units of product per day. Implicitly, it is assumed that on any day delivery takes place after the inventory of 3000 units has been consumed. The inventory level is calculated at the end of the day, after consumption and delivery have taken place. A similar situation occurs at customer C.

Alternatively, we can assume that the initial inventories are zero, and that delivery takes place before consumption. The optimal delivery policy remains the same. For the remainder, this is the setting we consider, i.e., at each time instant delivery takes place before consumption and the inventory level is observed after that.

Next, we introduce various simple modifications to the example of Bell et al. and study the impact on the optimal solution. The length of the planning horizon H is four days. The initial inventory I_i^0 at each customer i is either zero or a decision variable. The inventory holding cost may or may not be accounted for in the objective function. In case it is, it can be either the inventory holding cost at the supplier, the inventory holding cost at the customers, or the sum of the inventory holding costs at the supplier and at the customers. The first papers that consider holding costs in the objective function are Blumenfeld et al. [19] for the case with continuous time and Speranza and Ukovich [54] for the case with discrete time. The following unit inventory costs are considered for both the supplier and the customers: 0, 0.01, and 0.1. If the inventory cost at the supplier is accounted for in the objective function, then the production rate at the supplier is assumed to be equal to the sum of the consumption rates at the customers, i.e., 7500, and the initial level of the inventory at the supplier I_S^0 is either zero or a decision variable of the problem. Deliveries can be performed at the discrete time instants $t \in \{1, \ldots, H\}$ by an unlimited fleet of vehicles. The sets of customers that can be served on a single route are: $R_1 = \{A\}$, $R_2 = \{B\}$, $R_3 = \{C\}$, $R_4 = \{D\}$, $R_5 = \{A, B\}$, $R_6 = \{B, C\}$, $R_7 = \{C, D\}$. The routes considered are the direct routes to a customer and the routes that serve adjacent customers. Let $k \in K = \{1, 2, \ldots, 7\}$ be an index of the routes and let r_{ik} be an indicator equal to 1 if customer i is served on route k and 0 otherwise. We assume, for the sake of simplicity, that each route can be performed at most once per day. The cost c_k of route k is the cost of the corresponding optimal traveling salesman problem.

We analyze different environments and show, for each of them, the optimal solution over the planning horizon of 4 days. The solutions are obtained by solving mixed integer linear programming models. These models are not representative of the models proposed to solve large, real-life instances of inventory routing problems. They are introduced here only to obtain optimal solutions for the examples used in the chapter. The optimal solutions are shown in figures. Each figure is organized in one or more rows. Each row shows four boxes, one for each day. Each box shows the routes traveled by the vehicles, the quantity delivered to each customer, and the level of the inventory at each customer at the end of the day (in the lower left corner; each small rectangle represents 500 units of product). If more than one row is shown, each corresponds to a different environment, labeled with a letter, starting from the top row.

Transportation cost only. We first analyze the case in which only transportation cost is considered in the objective function. This case corresponds to an environment in which the transportation cost represents the major cost component, for example due to the fact that the supplier and the consumers represent entities of one and the same company (e.g., a distribution center and retail outlets) and the differences in inventory holding costs at the supplier and the consumers are negligible. Inventory levels at the customers still have to be controlled to avoid a stock-out and to respect the holding capacity limits. Let $I_i^0 = 0$, $\forall i$, let x_{ik}^t be the quantity shipped to customer i at time t using route k, and let y_k^t be equal to 1 if route k is used at time t and 0 otherwise. The optimal solution is obtained by solving the following mixed integer linear programming model:

$$\min \frac{1}{H} \sum_{t \in T} \sum_{k \in K} c_k y_k^t \tag{1}$$

$$\sum_{i \in I} x_{ik}^t \leq Q y_k^t \qquad t \in T \quad k \in K \tag{2}$$

$$x_{ik}^t \leq Q r_{ik} \qquad t \in T \quad i \in I \quad k \in K \tag{3}$$

$$I_i^t = I_i^0 + \sum_{s=1}^{t} \sum_{k \in K} x_{ik}^s - t q_i \qquad i \in I \quad t \in T \tag{4}$$

Table 2. Optimal costs with different transportation capacity.

Q	C_i given	$C_i = +\infty$
5000	380.0	340.0
10000	377.5	190.0
15000	377.5	152.5
20000	377.5	105.0

$$I_i^t + q_i \leq C_i \qquad i \in I \quad t \in T \tag{5}$$

$$I_i^t \geq 0 \qquad i \in I \quad t \in T \tag{6}$$

$$x_{ik}^t \geq 0 \qquad t \in T \quad i \in I \quad k \in K \tag{7}$$

$$y_k^t \in \{0,1\} \qquad t \in T \quad k \in K. \tag{8}$$

The objective function (1) expresses the minimization of the average daily transportation cost. The constraints (2) guarantee that the total quantity delivered by each vehicle is not greater than its capacity. The constraints (3) guarantee that a delivery to customer i on route k only takes place if customer i is visited on route k. The constraints (4) define the level of the inventory of each customer i for each time instant t. The constraints can be also written as $I_i^t = I_i^{t-1} + \sum_{k\in K} x_{ik}^t - q_i$. The constraints (5) guarantee that the storage capacity of each customer is never exceeded. The constraints (6) guarantee that no stock–out occurs at any customer i during the planning horizon. Finally, the constraints (7)–(8) define the decision variables of the problem.

Figure 2 shows the optimal solutions obtained in two different environments. The first row corresponds to the environment in which the storage capacity C_i at each customer i has to be respected and cannot be exceeded, while the second row corresponds to the environment in which there are no storage capacities at customers, i.e., $C_i = +\infty$ for all $i \in I$. In other words, we show the effect of relaxing constraints (5). In the first case, the costs are 380, while in the second case the costs are 340. In both situations, the vehicle capacity is fully utilized in all tours.

Next, we examine what happens when the vehicle capacity is increased. The values of the optimal solutions are shown in Table 2 for $Q = 5000, 10000$, 15000, and 20000. In the case with holding capacities at customers, shown in Column 2, the cost decreases from 380 to 377.5. This happens for $Q \geq 7000$. In the case without holding capacities at the customers, shown in Column 3, the cost decreases from 340 to 105. This value is obtained for $Q \geq 16000$. As expected, the impact of the vehicle capacity on the optimal cost is substantially higher in the case without holding capacities at customers.

Transportation and inventory holding costs. When inventory holding costs at the customers are taken into account in the objective function, the total cost

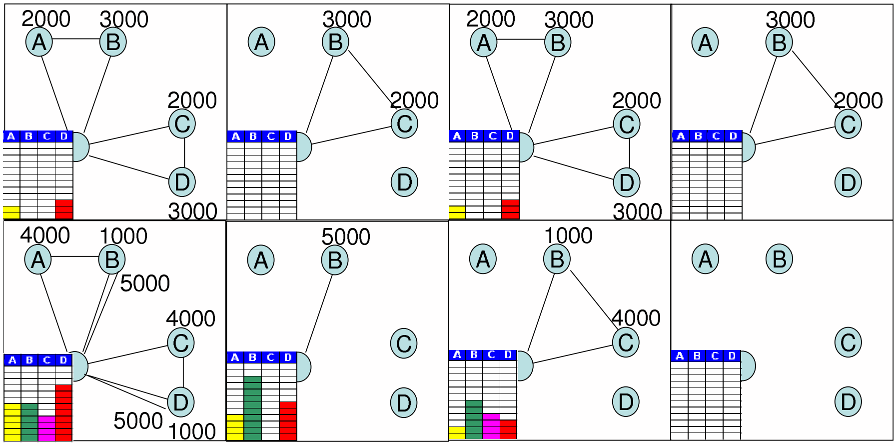

Fig. 2. Minimizing transportation costs only. a) With maximum inventory holding capacity at the customers: 380. b) Without maximum inventory holding capacity at the customers: 340.

of an optimal solution will obviously increase. An optimal solution can be obtained by solving the following mixed integer programming problem:

$$\min \frac{1}{H}\left(\sum_{t\in T}\sum_{k\in K} c_k y_k^t + \sum_{i\in I}\sum_{t\in T} h_i I_i^t\right)$$

subject to (2)–(8).

Figure 3 shows the solutions for two environments that differ only in terms of the inventory holding costs charged at the customers (storage capacity is limited at the customers). The first row shows the solution for the environment in which the inventory holding cost is small, $h_i = 0.01$ for all $i \in I$, while the second row shows the solution for the environment in which the inventory holding cost is large, $h_i = 0.1$ for all $i \in I$. It is interesting to observe that with higher inventory holding costs the optimal solution visits customers more frequently. Moreover, in this setting available vehicle capacity is no longer used fully.

To take inventory holding costs at the supplier into account, a term $\sum_{t\in T} h_S I_S^t$ needs to be incorporated in the objective function that becomes

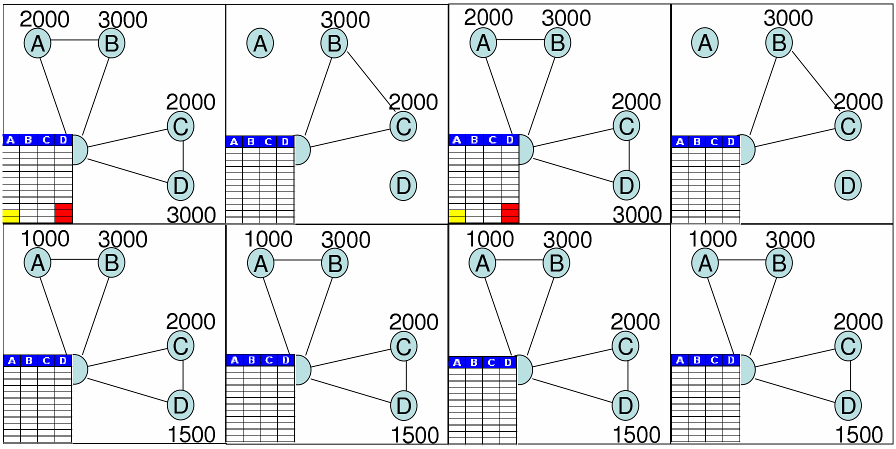

Fig. 3. Minimizing transportation cost and inventory cost at the customers. a) With $h_i = 0.01$: $380 + 12.5 = 392.5$. b) With $h_i = 0.1$: $420 + 0 = 420$.

$$\min \frac{1}{H} \sum_{t \in T} \sum_{k \in K} c_k y_k^t + \sum_{t \in T} h_S I_S^t.$$

The following constraints need to be added to the constraint set

$$I_S^t = I_S^0 + \sum_{i \in I} t q_i - \sum_{s=1}^{t} \sum_{i \in I} \sum_{k \in K} x_{ik}^s \qquad t \in T \tag{9}$$

$$I_S^t \geq 0 \qquad t \in T. \tag{10}$$

Constraints (9) state that the inventory level at the supplier at time t is obtained by the initial level increased by the production up to t and decreased by the quantity delivered to the customers up to t.

Figure 4 shows the solutions for environments in which, in addition to the transportation costs, inventory holding costs at customers only, at the supplier only, and at both customers and the supplier are taken into account, respectively (storage capacity is limited at the customers). In the figure, the inventory level of the supplier is also shown. While in the first environment the vehicles travel with full loads, in the second and the third environments this is no longer the case and, furthermore, each customer is served daily. Note that

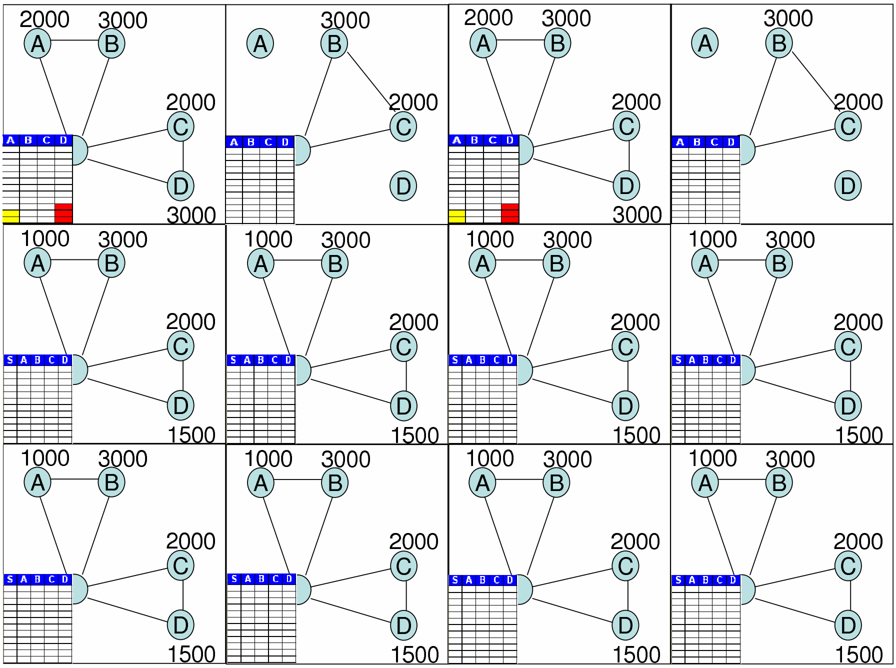

Fig. 4. a) Minimizing transportation cost and inventory holding cost at the customers, $h_i = 0.01$: $380 + 12.5 = 392.5$, b) Minimizing transportation cost and inventory holding cost at the supplier, $h_S = 0.01$: $420 + 0 = 420$, c) Minimizing transportation cost and inventory holding costs at the supplier and the customers, $h_i = 0.01$ and $h_S = 0.01$: $420 + 0 + 0 = 420$.

in the second and third environments the total inventory is always equal to zero. This is due to the fact that the total inventory is constant over time and that the initial inventory level is equal to zero both at the supplier and the customers. As a consequence, in an optimal policy we have $\sum_{k \in K} x_{ik}^t = q_i$, $t \in T$. In general, given that the total inventory is constant over time and equal to $\sum_{i \in I} I_i^0 + I_S^0$, if the initial inventory levels are given and $h_i = h_S$, $\forall i$, then the inventory holding costs can be ignored and only transportation costs need to be minimized. Therefore, the problem with inventory holding costs is interesting only when the inventory holding costs are different.

Transportation and inventory holding costs, variable initial inventory levels. Next, we consider the case in which inventory holding costs are charged both at the supplier and at the customers and the initial inventory levels are not fixed, but are to be determined by the optimization. In this case it is necessary to include so–called *demand constraints* in the model, that is constraints that guarantee that the total quantity shipped to each customer is equal to the corresponding total consumption over the planning horizon:

$$\min \frac{1}{H}\left(\sum_{t\in T}\sum_{k\in K} c_k y_k^t + \sum_{i\in I}\sum_{t\in T} h_i I_i^t + \sum_{t\in T} h_S I_S^t\right)$$

subject to (2)–(8), (9)–(10) and the following constraints:

$$\sum_{t\in T}\sum_{k\in K} x_{ik}^t = q_i H \qquad i \in I \tag{11}$$

$$I_i^0 \geq 0 \qquad i \in I \tag{12}$$

$$I_S^0 \geq 0. \tag{13}$$

Figure 5 shows, in the first row, the optimal solution for the situation in which holding capacity is limited at the customers, and, in the second row, the optimal solution when there are no limits on the holding capacity at the customers (storage capacity is limited at the customers). We have seen that in the situation where the initial inventory is given and assumed to be 0, the optimal solution value is 420 (see the third row of Figure 4). The solution in the first row of Figure 5 has a value of 405. The savings are a result of the additional flexibility, i.e., the choice of initial inventory levels. The optimal initial inventory level is 0 for customers B and C and 1000 and 1500 for customers A and D, respectively. As a consequence, the total inventory is always equal to 2500 during the planning horizon. When we compare this solution to the solution shown in the second row of Figure 5, when there are no limits on the storage capacities at customers, we observe that the optimal initial inventory level is 0 for customers A, B and D and 2000 for customer C, the optimal initial level at the supplier is 2500. The total inventory is always 4500 during the planning horizon and the costs decrease even further to 390. The darker customers shown in Figure 5 indicate the customers that have an initial inventory in the optimal solution. The others have no initial inventory.

The models used in the above examples are built on the basis of a preselected set of routes. The models studied in the literature include the routes creation in the set of decisions to be taken. The only exact approach available for the solution of an inventory routing problem can be found in Archetti et al. [7]. In Bertazzi et al. [13] heuristics are used to study the impact of different objective functions on the solution of an inventory routing problem.

Figure 6 compares the optimal solution obtained in the case with equal inventory holding cost at the supplier and the retailers with respect to the case of unequal inventory holding costs. The first row shows the case with $h_S = h_i = 0$, $\forall i$, in which the sum of the transportation and inventory costs both at the supplier and the retailers is minimized, the initial inventory level at the supplier and the retailers is variable and with inventory holding capacity at the retailers. The second row shows the same case with the only exception of the inventory holding costs, which are $h_S = 0.1$ and $h_i = 0$, $\forall i$, respectively.

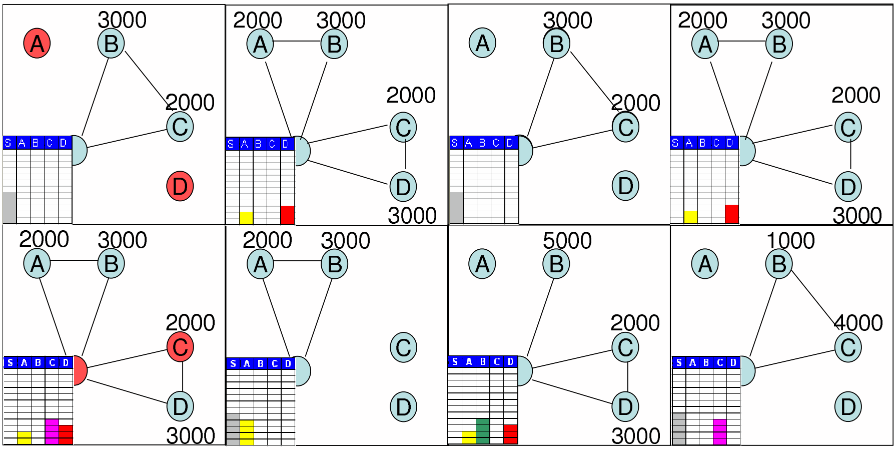

Fig. 5. Minimizing transportation cost and inventory holding costs at the supplier and customers; holding costs $h_i = 0.01$ and $h_S = 0.01$; variable initial inventory level. a) With inventory holding capacity at the customers: $380 + 12.5 + 12.5 = 405$ ($I_A^0 = 1000$, $I_D^0 = 1500$). b) Without inventory holding capacity at the customers: $345 + 32.5 + 12.5 = 390$ ($I_S^0 = 2500$, $I_C^0 = 2000$).

We note that in each time instant the quantity produced at the supplier is shipped to the retailers, even if not needed, in order to always have zero inventory levels at the supplier. Finally, the third row shows the case with $h_S = 0$ and $h_i = 0.1$, $\forall i$. We note that, in order to always have zero inventory levels at the retailers, the quantity delivered to each retailer in each day is equal to the quantity consumed by the retailer per time unit.

3 Continuous Production and Consumption

In the previous section, we focused on the impact that inventory holding costs had on delivery policies. We chose to work with a model in which deliveries take place at discrete time instants $t \in T = \{1, 2, 3, 4\}$, i.e., at the beginning of time periods $[t, t+1)$, and consumption takes place afterwards. The model assumes that (1) each vehicle performs exactly one delivery route during a time period, and that (2) the exact timing of the delivery during a time period is

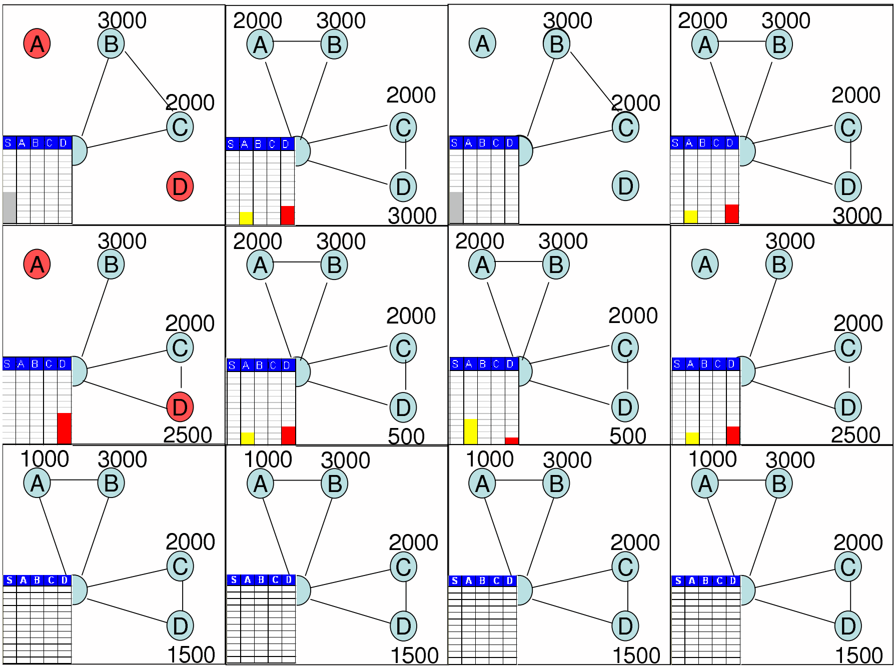

Fig. 6. Equal vs unequal inventory holding costs: a) $h_S = h_i = 0$: 380 + 12.5 + 12.5 = 405 ($I_A^0 = 1000$, $I_D^0 = 1500$). b) $h_S = 0.1$, $h_i = 0$: 415 + 0 + 0 = 415 ($I_A^0 = 1000$, $I_D^0 = 1500$). c) $h_S = 0$, $h_i = 0.1$: 420 + 0 + 0 = 420.

not important. Although this type of discrete time model is frequently used for inventory routing problems, it is too restrictive for many environments, for example environments in which customers use product continuously. In these environments the timing of deliveries is important as the available storage capacity, and thus the maximum delivery quantity, depends on the time of delivery. If product usage rates and storage capacities vary significantly among customers, then a mix of short and long routes may be inevitable, and the "single delivery route per vehicle per period" assumption may no longer be appropriate. In situations where there are customers with small storage capacities and high usage rates that require one or more deliveries per time period, the timing of such deliveries becomes important as the time of delivery determines the available storage capacity.

In this section, we will concentrate on environments with continuous product usage and, for convenience, we will ignore inventory holding costs.

We start by observing that when the number of vehicles is greater than the number of customers and the storage capacity at each customer is larger than the vehicle capacity, the problem is easy. Why? Because in that case, the optimal distribution strategy is to dispatch a vehicle to a customer in such a

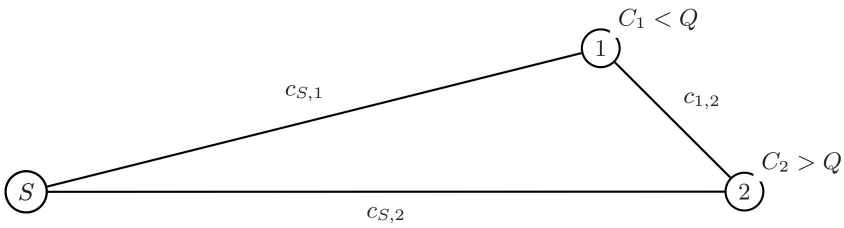

Fig. 7. IRP environment with a single supplier and two customers.

way that the vehicle arrives at the customer right at the time the customer is about to run out of product and to deliver an entire vehicle load. This minimizes the number of deliveries to a customer over the planning horizon and, at the same time, minimizes the distance traveled to do so.

As soon as the number of vehicles is smaller than the number of customers, the distribution strategy described above may no longer be optimal, because it may no longer be feasible; there may be a point in time when we need more vehicles than are available. As soon as the storage capacities at customers are smaller than the vehicle capacity a distribution strategy with only out-and-back trips to customers may be feasible, but is almost guaranteed to be non-optimal. We are not fully using the available vehicle capacity and are therefore likely missing opportunities to reduce transportation costs. The simplest example of such a situation involves two customers. Suppose two customers are located next to each other, each with a storage capacity that is half the size of the vehicle capacity. It is clear that it is advantageous to deliver to both customers each time one of them is visited, as no extra costs are incurred by doing so and the total number of visits is reduced.

The above discussion highlights the fact that storage capacities less than the vehicle capacity cause the inventory routing problem to become significantly more difficult. This complexity manifests itself in various ways. First, it may be advantageous to visit more than one customer on a single route. Deciding which customers to put together on route is non-trivial, among other reasons because evaluating the distance traveled on a route involves solving a traveling salesman problem. Second, because the available storage capacity at customers changes over time, because of product usage, deliveries need to be coordinated in time to be able to fully exploit available vehicle capacity.

Consider the distribution environment depicted in Figure 7, i.e., a single supplier and two customers. Since $C_1 < Q$, whenever a truck goes to customer 1 with a full load, at least $Q - C_1$ of product is left after the delivery at customer 1. That leftover product can be used to satisfy the need for product at customer 2. Is that needed? Is that cost effective? Suppose that the total product usage over the planning horizon is U_1 for customer 1 and U_2 for customer 2. To deliver U_1 to customer 1, at least $\frac{U_1}{C_1}$ deliveries have to be made. Therefore, at least $\frac{U_1}{C_1}(Q - C_1)$ leftover product is available for delivery at customer 2. If the leftover product is used to satisfy product need of customer 2, then $\frac{Q}{Q-C_1}$ trips with leftover product are necessary to deliver a quantity of Q to customer 2. Whenever leftover product is delivered to customer 2, an additional cost of $c_{12} + c_{S2} - c_{S1}$ is incurred. Therefore, the cost incurred to deliver Q to customer 2 with leftover product from customer 1 is $\frac{Q}{Q-C_1}(c_{12} + c_{S2} - c_{S1})$. The cost incurred to deliver Q to customer 2 directly from the supplier is $2c_{S2}$. Consequently, if $\frac{Q}{Q-C_1}(c_{12}+c_{S2}-c_{S1}) < 2c_{S2}$, it is better to use leftover product at customer 1 to satisfy the product need of customer 2. For the remainder, assume that this is the case, i.e., $\frac{Q}{Q-C_1}(c_{12}+c_{S2}-c_{S1}) < 2c_{S2}$. Two cases have to be considered: (1) the leftover product is sufficient to satisfy customer 2's needs, and (2) the leftover product is insufficient to satisfy customer 2's needs.

Case 1 : If $\frac{U_1}{C_1}(Q-C_1) \geq U_2$, then a cost of at least $\frac{U_2}{Q-C_1}(c_{S1}+c_{12}+c_{S2}) + \frac{U_1 - \frac{U_2}{Q-C_1}C_1}{C_1}2c_{S1}$ is incurred.

Case 2 : If $\frac{U_1}{C_1}(Q - C_1) < U_2$, then a cost of at least $\frac{U_1}{C_1}(c_{S1} + c_{12} + c_{S2}) + \frac{U_2 - \frac{U_1}{C_1}(Q-C_1)}{Q}2c_{S2}$ is incurred.

Note that the analysis above, although insightful, only considers the storage capacities at the customers; it ignores the timing of deliveries. It is not at all guaranteed that the suggested delivery scheme can be executed in practice as the quantity that can be delivered to a customer depends on the time of delivery. This reflects the intricate and complex relationship between customer usage rates and delivery route travel times. Thus, the analysis only provides a lower bound on the delivery costs incurred. A detailed discussion on computing lower bounds for inventory routing problems can be found in Song and Savelsbergh [53].

Two factors other than limited storage capacity at customers impact the complexity of the inventory routing problem: the ratio of the usage rate and storage capacity at a customer, i.e., $\frac{u_i}{C_i}$, which defines the maximum time between two consecutive visits to a customer, and, in situations with a finite planning horizon, the ratio of initial inventory and product use over the planning period, i.e., $\frac{I_i^0}{Hu_i}$, which represents the percentage of total product consumed during the planning horizon that can be served from initial inventory. Typically, the need to make frequent deliveries to customers increases the complexity, especially when these customers have limited storage capacities.

We have mentioned a few times already in this section that the maximum quantity that can be delivered to a customer depends on the time of delivery. Consequently, the selection of actual delivery times, between the earliest and latest delivery times, will affect the total volume deliverable on a trip. Because

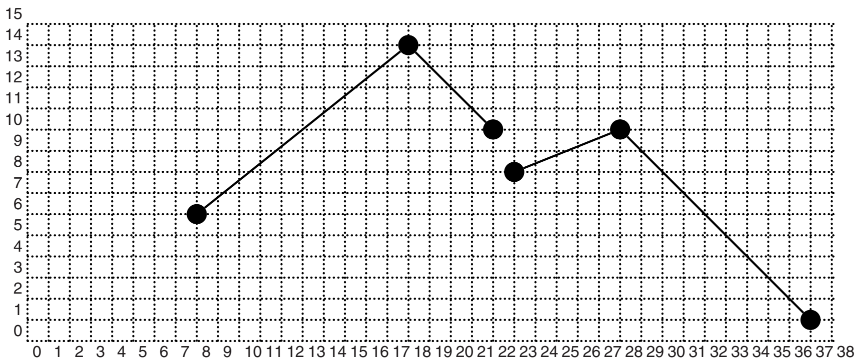

Fig. 8. Maximum delivery quantities at two customers as a function of delivery times.

the customer consumes product over time, the later a vehicle arrives, the more inventory holding capacity is available and the more product can be delivered. On the other hand, delivery is typically not instantaneous, but depends on the size of the delivery. Furthermore, when a vehicle arrives later at a customer, less time may be available for making a delivery due to delivery time restrictions of customers to be visited later on the trip. These two dueling effects make determining an optimal delivery schedule for a given sequence of customer visits on a trip not as easy as it may seem at first glance.

As an example consider the situation depicted in Figure 8. It shows the maximum delivery quantity as a function of delivery time for two consecutive customers on a route; the first part of the graph, between times 8 and 22, relates to the first customer and the second part of the graph, between times 23 and 37, relates to the second customer. The earliest delivery time for the first customer is 8 and the latest delivery time is 22. The earliest delivery time for the second customer is 23 and the latest delivery time is 37. The slope of the line between times 8 and 18 is 0.8, the usage rate of the first customer, and the slope of the line between times 23 and 28 is 0.4, the usage rate of the second customer. The slope of the lines between times 18 and 22 and 28 and 37 is -1, corresponding to the rate at which product is discharged at the customers. The travel time between the first and second customer is 5. Initially, the amount of product that can be delivered to a customer increases, because product is consumed and the available storage capacity increases. The

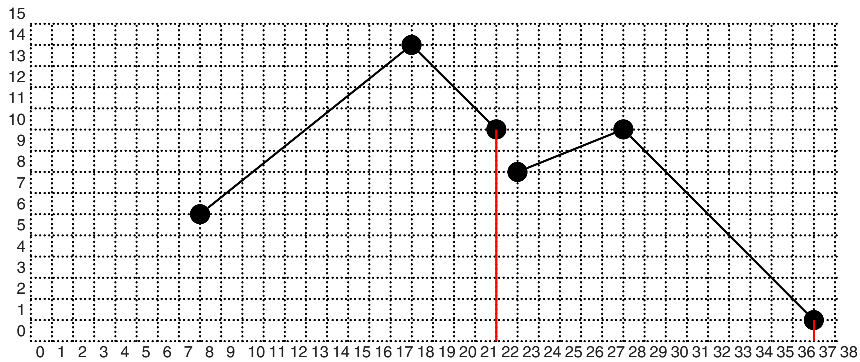

Fig. 9. Deliver as late as possible.

amount of product that can be delivered reaches a peak when the discharge time plus the travel time will be just enough to reach the next customer in time to make a delivery there. After the peak, the amount of product that can be delivered to a customer decreases with the discharge rate. The peak for the first customer is at time 18, since 18 plus 14 (the discharge time) plus 5 (the travel time to the second customer) is 37 (the latest delivery time at the second customer). We want to select the delivery times that result in the maximum total delivery quantity at both customers. Below we consider a few options. First, we consider delivering to both customers as late as possible. This situation is shown in Figure 9. The total quantity delivered is 11 (10 at the first customer and 1 at the second customer). Second, we consider delivering the maximum possible to the customer with the highest usage rate. This situation is shown in Figure 10. The total quantity delivered is 15. Finally, in Figure 11 we show the optimal delivery times, i.e., deliver 10 units at the first customer at time 13 and deliver 10 units at the second customer at time 28 (13 + 10 + 5) for a total quantity of 20.

A detailed discussion of how to maximize the delivery volume on a route in the context of inventory routing problems can be found in Campbell and Savelsbergh [24]. This simple example illustrates that few decisions are easy in inventory routing problems.

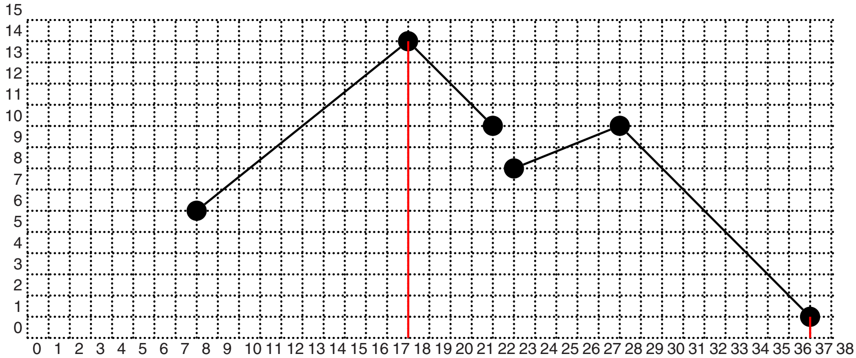

Fig. 10. Deliver maximum at the customer with highest usage rate.

4 The Literature

In this section, we briefly summarize the literature on IRPs. There are several surveys of inventory routing problems, and we point the user to these for more detailed information [21, 30, 34]. There is such a variety of problems and solution approaches that even structuring a literature review is challenging. We have adopted the following scheme. We start with the pioneering papers in the eighties that generated the initial interest in inventory routing problems. Next, we group papers according to two basic characteristics: whether or not inventory costs are considered in the objective, and whether or not product usage is deterministic or stochastic. Finally, we mention a few papers that are related, but do not clearly belong to one of the four categories that we created.

As mentioned, the first papers on inventory routing problems appeared in the nineteen eighties. These papers, in most cases, discuss and are inspired by applications in which both inventory and distribution have to be considered. For example, Bell et al. [9], Blumenfeld et al. [19], Burns et al. [20], Chien et al. [29], Dror and Ball [32], Dror et al. [33], Federgruen and Zipkin [35], Fisher et al. [36], Golden et al. [40], and Hall [41].

The four categories described next contain a more varied class of papers, as the papers discuss applications, solution approaches, and worst-case and asymptotic analysis.

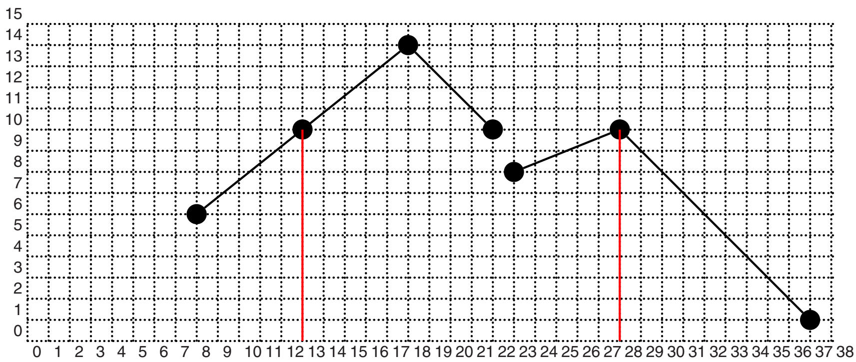

Fig. 11. Optimal delivery times.

We first consider papers covering inventory routing problems with deterministic product usage in which inventory holding costs are considered in the objective function. This category contains papers concerned with supplying just a single retailer as well as papers concerned with supplying multiple retailers. Starting with the paper by Speranza and Ukovich [54], in which the single retailer case with given discrete shipping frequencies was introduced and modeled, several papers have appeared studying computational complexity, analysis of shipping policies, and heuristic and exact solution approaches, e.g., Bertazzi et al. [11, 18], Bertazzi and Speranza [15, 16] and Speranza and Ukovich [55]. The case of multiple retailers has been studied by Archetti et al. [7], Bertazzi [10], Bertazzi et al. [13, 14, 17], Cousineau-Ouimet [31] and Rabah and Mahmassani [49]. Bertazzi [10], Gallego and Simchi-Levi [37, 38], and Hall [42] study the performance of direct shipping policies. Anily [6], Anily and Federgruen [3, 4, 5], Chan et al. [26, 27] and Chan and Simchi-Levi [28] analyze the asymptotic performance of certain policies. Other papers in this category include Herer and Roundy [43] and Viswanathan and Mathur [57].

Next, we consider papers covering inventory routing problems with stochastic product usage in which inventory holding costs are considered in the objective function. This category contains the work by of Kleywegt et al. [45, 46] and the work of Minkoff [48] on Markov Decision Process models and solution approaches.

Next, we consider papers covering inventory routing problems with deterministic product usage in which no inventory holding costs are considered in the objective function. This category contains a variety of papers. Savelsbergh and Song [51] compute lower bounds on the optimal distribution costs. Campbell and Savelsbergh [23, 24, 25] as well as Gaur and Fisher [39] study time-indexed formulations. Jaillet et al. [44] and Trudeau and Dror [56] analyze d-day policies and use this analysis to develop rolling-horizon approaches.

Next, we consider papers covering inventory routing problems with stochastic product usage in which no inventory holding costs are considered in the objective function. This category contains the work by Adelman [1, 2] on price-directed approaches and the work of Berman and Larson [12] on stochastic dynamic programming.

Finally, we list papers that cannot clearly be associated with one of the above categories. Campbell and Hardin [22] study periodic delivery policies. Savelsbergh and Song [52] and Song and Savelsbergh [53] consider inventory routing problems with continuous moves, i.e., where vehicles do not return to a designated depot, but following multi-day tours visiting many supply and demand points along the way. Bard et al. [8] consider satellite facilities for temporary storage of product. Webb and Larson [58] study the strategic inventory routing problem where the goal is to determine the fleet size necessary to deliver a set of customers from a single depot. Reiman et al. [50] use queueing control problems to model and analyze stochastic inventory routing problems. Lau, Liu, Ono [47] combine ideas from local search and network flows to solve inventory routing problems.

5 Conclusions

We have illustrated the trade-off between transportation and inventory holding costs, and the impact of inventory holding capacity and time-dependent delivery quantities on distribution strategies. Our goal was to familiarize our readers with the type of issues and challenges encountered when studying or solving inventory routing problems. We hope these examples demonstrate the richness of the class of inventory routing problems and the many opportunities they offer for exciting high impact research. The summary of the literature is a showcase of most of the work that has already been done in this area and the interested reader should take the time to study these in more detail. It will be a rewarding experience.

Acknowledgement

The authors wish to acknowledge the valuable contributions of the anonymous referees which helped to improve the presentation of the material in this chapter.

References

1. Adelman, D. (2003), Price-directed Replenishment of Subsets: Methodology and its Application to Inventory Routing, *Manufacturing & Service Operations Management* 5, 348-371.
2. Adelman, D. (2004), A Price-directed Approach to Stochastic Inventory/Routing, *Operations Research* 52, 499-514.
3. Anily, S., Federgruen, A. (1990), One Warehouse Multiple Retailer Systems with Vehicle Routing Costs, *Management Science* 36, 92-114.
4. Anily, S., Federgruen, A. (1991), Rejoinder to 'One Warehouse Multiple Retailer Systems with Vehicle Routing Costs', *Management Science* 37, 1497-1499.
5. Anily, S., Federgruen, A. (1993), Two–Echelon Distribution Systems with Vehicle Routing Costs and Central Inventories, *Operations Research* 41, 37-47.
6. Anily, S. (1994), The General Multi–Retailer EOQ Problem with Vehicle Routing Costs, *European Journal of Operational Research* 79, 451-473.
7. Archetti, C., Bertazzi, L., Laporte, G., Speranza, M.G. (2007), A Branch-and-Cut Algorithm for a Vendor Managed Inventory Routing Problem, *Transportation Science*, to appear.
8. Bard, J., Huang, L., Jaillet, P., Dror, M. (1998), A Decomposition Approach to the Inventory Routing Problem with Satellite Facilities, *Transportation Science* 32, 189-203.
9. Bell, W., Dalberto, L., Fisher, M., Greenfield, A., Jaikumar, R., Kedia, P., Mack, R., Prutzman, P. (1983), Improving the Distribution of Industrial Gases with an On-line Computerized Routing and Scheduling Optimizer, *Interfaces* 13, 4-23.
10. Bertazzi, L. (2006), Analysis of Direct Shipping Policies in an Inventory Routing Problem with Discrete Shipping Times, Tecnical Report n. 265, Department of Quantitative Methods, University of Brescia (under second revision in *Management Science*).
11. Bertazzi, L., Chan, L.M.A., Speranza, M.G. (2007), Analysis of Practical Policies for the Single Link Distribution System, *Naval Research Logistics*, to appear.
12. Berman, O., Larson, R.C. (2001), Deliveries in an Inventory/Routing Problem Using Stochastic Dynamic Programming, *Transportation Science* 35, 192-213.
13. Bertazzi, L., Paletta, G., Speranza, M.G. (2002), Deterministic Order-up-to Level Policies in an Inventory Routing Problem, *Transportation Science* 36, 119-132.
14. Bertazzi, L., Paletta, G., Speranza, M.G. (2005), Minimizing the Total Cost in an Integrated Vendor–Managed Inventory System, *Journal of Heuristics* 11, 393-419.
15. Bertazzi, L., Speranza, M.G. (2002), Continuous and Discrete Shipping Strategies for the Single Link Problem, *Transportation Science* 36, 314-325.
16. Bertazzi, L., Speranza, M.G. (2005), Worst-case Analysis of the Full Load Policy in the Single Link Shipping Problem, *International Journal of Production Economics* 93-94C, 217-224.
17. Bertazzi, L., Speranza, M.G., Ukovich, W. (1997), Minimization of Logistic Costs with Given Frequencies, *Transportation Research B* 31, 327-340.
18. Bertazzi, L., Speranza, M.G., Ukovich, W. (2000), Exact and Heuristic Solutions for a Shipment Problem with Given Frequencies, *Management Science* 46, 973-988.

19. Blumenfeld, D.E., Burns, L.D. Diltz, J.D., Daganzo, C.F. (1985), Analyzing Trade-offs between Transportation, Inventory and Production Costs on Freight Networks, *Transportation Research B* 19, 361-380.
20. Burns, L.D., Hall, R.W., Blumenfeld, D.E., Daganzo, C.F. (1985), Distribution Strategies that Minimize Transportation and Inventory Cost, *Operations Research* 33, 469-490.
21. Campbell, A.M., Clarke, L., Kleywegt, A., Savelsbergh, M.W.P. (1998), The Inventory Routing Problem, in: *Fleet Management and Logistics*, T.G. Crainic and G. Laporte (eds), 95-113, Kluwer, Boston.
22. Campbell, A.M., Hardin, J. (2005), Vehicle Minimization for Periodic Deliveries, *European Journal of Operational Research* 165, 668-684.
23. Campbell, A.M., Savelsbergh, M.W.P. (2004), A Decomposition Approach for the Inventory–Routing Problem, *Transportation Science* 38, 488-502.
24. Campbell, A.M., Savelsbergh, M.W.P. (2004), Delivery Volume Optimization, *Transportation Science* 38, 210–223.
25. Campbell, A.M., Savelsbergh, M.W.P. (2004), Efficiently Handling Practical Complexities in Insertion Heuristics, *Transportation Science* 38, 369-378.
26. Chan, L.M.A., Federgruen, A., Simchi-Levi, D. (1998), Probabilistic Analyses and Practical Algorithms for Inventory–Routing Models, *Operations Research* 46, 96-106.
27. Chan, L.M.A., Muriel, A., Shen, Z.J.M., Simchi-Levi, D., Teo, C.P. (2002), Effective Zero–Inventory–Ordering Policies for the Single–Warehouse Multiretailer Problem with Piecewise Linear Cost Structures, *Management Science* 48, 1446-1460.
28. Chan, L.M.A., Simchi-Levi, D. (1998), Probabilistic Analyses and Algorithms for Three-Level Distribution Systems, *Management Science* 44, 1562-1576.
29. Chien, T., Balakrishnan, A., Wong, R. (1989), An Integrated Inventory Allocation and Vehicle Routing Problem, *Transportation Science* 23, 67-76.
30. Cordeau, J.-F., Laporte, G., Savelsbergh, M.W.P., Vigo, D. (2007), Short-Haul Routing, in *Handbooks in Operations Research and Management Science: Transportation* G. Laporte and C. Barnhart (eds.). Forthcoming.
31. Cousineau-Ouimet, K. (2002), A Tabu Search Heuristic for the Inventory Routing Problem, in *Proceedings of 37th Annual ORSNZ Conference*.
32. Dror, M., Ball, M. (1987), Inventory/Routing: Reduction from an Annual to a Short-Period Problem, *Naval Research Logistics Quarterly* 34, 891-905.
33. Dror, M., Ball, M., Golden, B. (1985), A Computational Comparison of Algorithms for the Inventory Routing Problem, *Annals of Operations Research* 4, 3-23.
34. Federgruen, A., Simchi–Levi, D. (1995), Analysis of Vehicle Routing and Inventory–Routing Problems, in: *Handbooks in Operations Research and Management Science* (Ball, M.O., Magnanti, T.L., Monma, C.L. and Nemhauser, G.L., eds.), Vol. 8, 297-373, North–Holland.
35. Federgruen, A., Zipkin, P. (1984), A Combined Vehicle Routing and Inventory Allocation Problem, *Operations Research* 32, 1019-1032.
36. Fisher, M., Greenfield, A., Jaikumar, R., Kedia, P. (1982), Real-time Scheduling of a Bulk Delivery Fleet: Practical Application of Lagrangean Relaxation. Technical report, The Wharton School, University of Pennsylvania, Department of Decision Sciences.

37. Gallego, G., Simchi–Levi, D. (1990), On the Effectiveness of Direct Shipping Strategy for the One-Warehouse Multi-Retailer R-Systems, *Management Science* 36, 240-243.
38. Gallego, G., Simchi–Levi, D. (1994), Rejoinder to "A Note on Bounds for Direct Shipping Costs", *Management Science* 40, 1393.
39. Gaur, V., Fisher, M. (2004), A Periodic Inventory Routing Problem at a Supermarket Chain, *Operations Research* 52, 813-822.
40. Golden, B., Assad, A., Dahl, R. (1984), Analysis of a Large Scale Vehicle Routing Problem with an Inventory Component, *Large Scale Systems* 7, 181-190.
41. Hall, R.W. (1985), Determining Vehicle Dispatch Frequency when Shipping Frequency Differs among Suppliers, *Transportation Research B* 19, 421-431.
42. Hall, R.W. (1992), A Note on Bounds for Direct Shipping Costs, *Management Science* 38, 1212-1214.
43. Herer, Y., Roundy, R. (1997), Heuristic for One–Warehouse Multiretailer Distribution Problem with Performance Bounds, *Operations Research* 45, 102-115.
44. Jaillet, P., Bard, J., Huang, L., Dror, M. (2002), Delivery Costs Approximations for Inventory Routing Problems in Rolling Horizon Framework, *Transportation Science* 36, 292-300.
45. Kleywegt, A.J., Nori, V.S., Savelsbergh, M.W.P (2002), The Stochastic Inventory Routing Problem with Direct Deliveries, *Transportation Science* 36, 94-118.
46. Kleywegt, A.J., Nori, V.S., Savelsbergh, M.W.P (2004), Dynamic Programming Approximations for a Stochastic Inventory Routing Problem, *Transportation Science* 38, 42–70.
47. Lau, H.C., Liu, Q., Ono, H. (2002), Integrating Local Search and Network Flow to Solve the Inventory Routing Problem, *American Association for Artificial Intelligence* 2, 9-14.
48. Minkoff, A. (1993), A Markov Decision Model and Decomposition Heuristic for Dynamic Vehicle Dispatching, *Operations Research* 41, 77-90.
49. Rabah, M.Y., Mahmassani, H.S. (2002), Impact of Electronic Commerce on Logistics Operations: A Focus on Vendor Managed Inventory (VMI) Strategies, Technical Report N. SWUTC/02/167227–1, Center for Transportation Research, University of Texas at Austin.
50. Reiman, M.I., Rubio, R., Wein, L.M. (1999), Heavy Traffic Analysis of the Dynamic Stochastic Inventory-Routing Problem, *Transportation Science* 33, 361-380.
51. Savelsbergh, M.W.P., Song, J.-H. (2006), Inventory Routing with Continuous Moves, *Computers and Operations Research* 34, 1744-1763.
52. Savelsbergh, M.W.P., Song, J.-H. (2007), An Optimization Algorithm for Inventory Routing with Continuous Moves, *Computers and Operations Research*, to appear.
53. Song, J.-H., Savelsbergh, M.W.P. (2006), Performance Measurements for Inventory Routing, *Transportation Science*, to appear.
54. Speranza, M.G., Ukovich, W. (1994), Minimizing Transportation and Inventory Costs for Several Products on a Single Link, *Operations Research* 42, 879-894.
55. Speranza, M.G., Ukovich, W. (1996), An Algorithm for Optimal Shipments with Given Frequencies, *Naval Research Logistics* 43, 655-671.
56. Trudeau, P., Dror, M. (1992), Stochastic Inventory Routing: Route Design with Stockouts and Route Failures, *Transportation Science* 26, 171-184.

57. Viswanathan, S., Mathur, K. (1997), Integrating Routing and Inventory Decisions in One Warehouse Multiretailer Multiproduct Distribution System, *Management Science* 43, 294-312.
58. Webb, I.R., Larson, R.C. (1995), Period and Phase of Customer Replenishment: A New Approach to the Strategic Inventory/Routing Problem, *European Journal of Operational Research* 85, 132-148.