

# The Vehicle Routing Problem with Occasional Drivers

Claudia Archetti<sup>\*1</sup>, Martin Savelsbergh<sup>2</sup>, and Grazia Speranza<sup>1</sup>

<sup>1</sup>*Department of Economics and Management, University of Brescia, Italy*

<sup>2</sup>*School of Industrial and Systems Engineering, Georgia Institute of Technology, USA*

September 16, 2015

1

## Abstract

We consider a setting in which a company not only has a fleet of capacitated vehicles and drivers available to make deliveries, but may also use the services of occasional drivers who are willing to make a single delivery using their own vehicle in return for a small compensation if the delivery location is not too far from their own destination. The company seeks to make all the deliveries at minimum total cost, i.e., the cost associated with its own vehicles and drivers plus the compensation paid to the occasional drivers. The option to use occasional drivers to make deliveries gives rise to a new and interesting variant of the classical capacitated vehicle routing problem. We design and implement a multi-start heuristic which produces solutions with small errors when compared with optimal solutions obtained by solving an integer programming formulation with a commercial solver. A comprehensive computational study provides valuable insight into the potential of using occasional drivers to reduce delivery costs, focusing primarily on the number and flexibility of occasional drivers and the compensation scheme employed.

Keywords: Vehicle Routing Problem, Crowdshipping, Occasional Drivers.

## 1 Introduction

The importance and competitive value of last-mile and same-day delivery has prompted many companies to seek creative and innovative solutions. One such innovative solution, considered among others by Walmart (see Barr and Wohl (2013)) and Amazon (see Bensinger (2015)), is “crowdshipping”, i.e., getting ordinary people, rather than delivery companies (e.g. UPS or FedEx) or company employed drivers, to drop-off packages en-route to their destination. This taps into the broader change taking place in our society, often referred to as the “sharing economy”, in which physical assets become services. Technology makes sharing assets cheaper and easier than ever – and therefore possible on a large scale. This is especially useful for expensive items that are not fully used by those who own them, such as automobiles. The collaborative consumption allows owners to make money from underused assets.

---

<sup>1\*</sup> Corresponding author. Department of Economics and Management, University of Brescia, 25122 Brescia, Italy. Telephone number: 00390302988587. E-mail address: claudia.archetti@unibs.it

The concept of collaborative consumption in transportation has been explored in several different directions recently. A discussion on sharing transportation resources between people and freight can be found in Trentini and Mahl  n   (2011), where concepts and solutions to improve urban mobility are investigated. The findings indicate great potential for sharing of transportation resources across traditionally unconnected modes (e.g., using a tram for freight transportation). In a recent paper Li et al. (2014) consider conceptual and mathematical models in which people and parcels are transported and served in an integrated way by a single taxi network. Several papers have analyzed problems arising in the general context of ridesharing, e.g., Agatz et al. (2011), Agatz et al. (2012), and Lee and Savelsbergh (2015), where the ridesharing refers to a mode of transportation in which individual travelers share a vehicle for a trip and split travel costs such as gas, toll, and parking fees with others that have similar itineraries and time schedules. The state of the art and future directions in ridesharing can be found in Furuhata et al. (2013).

The application of crowdshipping alluded to above gives rise to an interesting new variant of the vehicle routing problem, which we study in this paper: the Vehicle Routing Problem with Occasional Drivers (VRPOD). For a recent overview of more traditional problems encountered in the class of vehicle routing problems see Toth and Vigo (2014). In the VRPOD, a company not only has a fleet of capacitated vehicles and drivers available to make deliveries, but the company may also use the services of “occasional drivers” (ODs) who are willing to make a single delivery using their own vehicle in return for a small compensation if the delivery location is not too far from their own destination. The company seeks to satisfy customer demand (i.e., make all deliveries) at minimum total costs, i.e., the costs associated with the delivery routes of its own vehicles and drivers plus the compensation paid to the occasional drivers.

The VRPOD admittedly does not capture all aspects of the practical implementation of crowdshipping for last-mile delivery, but it does allow us to gain quantitative insights in its potential benefits. Specifically, it allows us to study the impact of the number of occasional drivers, the flexibility of the occasional drivers, and the employed compensation scheme on the benefits, i.e., cost savings, for a company. The main limitations of this variant of the vehicle routing problem are that it assumes an occasional driver can only make a single delivery and that it is a static problem. In reality, occasional drivers may be willing (and able) to make multiple deliveries and the availability of occasional drivers will vary over time. However, we believe that, despite these limitations, the insights obtained are valuable and informative.

In order to solve the VRPOD and generate the insight that we seek, we design and implement a multi-start heuristic, which we show produces solutions with small errors when compared with optimal solutions obtained by solving an integer programming formulation with a commercial solver.

Our computational study shows that employing occasional drivers can have significant benefits, and that choosing an appropriate compensation scheme is challenging, in fact more challenging than we initially anticipated. Obviously, the compensation scheme influences the number of available occasional drivers and their willingness to deviate from their intended travel route, but it also impacts which delivery locations will be assigned to occasional drivers (as opposed to the company’s drivers), which, in turn, affects the cost savings.

The remainder of the paper is organized as follows. In Section 2, we formally define the VRPOD, introduce and discuss possible compensation schemes. In Section 3, we present an integer programming formulation, which is used to benchmark our heuristic solution approach. In Section 4, we describe the heuristic we developed for the solution of instances of the VRPOD. In Section 5, we report and analyze the results of an extensive computational study. Finally, in Section 6, we

present some final remarks and discuss future research directions.

## 2 The Vehicle Routing Problem with Occasional Drivers

The basic variant of the VRPOD is defined as follows. Let  $G = (N, A)$  denote a complete directed graph with node set  $N$  and arc set  $A$ . The node set  $N$  is comprised of three sets of nodes: the location of the depot (node 0), the locations of the customers (set  $C$ ), and the locations of the destinations of the occasional drivers (set  $K$ ). For the sake of simplicity, we identify occasional drivers with their destinations. Each arc  $a \in A$  has a length  $d_a$  and a cost  $c_a$ . Each customer  $i \in C$  has a given demand  $q_i$ . Regular drivers that can make deliveries to customers using a vehicle of capacity  $Q$ . We assume that an unlimited number of regular drivers is available. This is consistent with practical applications where companies typically have a sufficient number of drivers to serve all customer requests. An occasional driver  $k \in K$  is willing to make a delivery at customer  $i$  when  $d_{0i} + d_{ik} \leq \zeta d_{0k}$  with  $\zeta \geq 1$ , i.e., when the extra distance traveled to reach the occasional driver's destination is less than or equal to  $(\zeta - 1)$  times the direct distance from the depot to the occasional driver's destination. An occasional driver can make at most one delivery to a customer. The demand of a customer has to be satisfied either on a route carried out by a regular driver (starting and ending at the depot) or on a trip carried out by an occasional driver, where a trip consists of two parts: traveling from the depot to the customer's location and traveling from the customer's location to the driver's destination. An occasional driver receives  $\rho c_{0i}$  as compensation for making a delivery to customer  $i$  with  $0 < \rho < 1$ . Regular driver routes need to satisfy a capacity constraint: the total demand of the customers served on a regular driver's route cannot exceed  $Q$ . It is implicitly assumed that an occasional driver  $k$  can accommodate the demand  $d_i$  for all customers the occasional driver is willing to serve. The cost of a regular driver route  $r$  is the sum of the costs of the arcs in the route, i.e.,  $\sum_{a \in r} c_a$ . The objective (from the company's perspective) is to minimize the total costs, that is the sum of the costs incurred by the regular drivers and the cost incurred for compensating occasional drivers. Note that if there are no occasional drivers, the problem becomes the standard capacitated vehicle routing problem (CVRP). Note too that for a given value of  $\zeta$ , it is easy to determine the set  $C_k$  of customers that an occasional driver  $k$  is willing to visit.

In this basic variant of the VRPOD, the compensation paid to an occasional driver is independent of the destination of the occasional driver. This has practical advantages, since the company only needs to know the location of its customers, but may not be ideal from the perspective of an occasional driver, as it does not reflect the extra costs incurred by visiting the customer. An alternative is that occasional driver  $k$  receives  $\rho(c_{0i} + c_{ik} - c_{0k})$  with  $\rho \geq 1$  for making a delivery to customer  $i$ , i.e., the occasional driver is compensated for the extra mileage incurred. This, of course, is more difficult to implement in practice, because it requires knowledge of the occasional driver's destination. Should it be the location of the occasional driver's home? Should the occasional driver declare her/his destination at the time she/he declares her/his willingness to act as occasional driver? How can this be verified? Even though we are aware of the challenges associated with a compensation scheme that depends not only on the location of the customer, but also on the destination of the occasional driver, we believe it is interesting to see if such a compensation scheme offers any advantages (or disadvantages) in terms of cost savings for the company.

### 3 An integer programming formulation

To validate the performance of the multi-start heuristic presented in the next section on small to medium size instances, we solve these instances to optimality using a commercial integer programming solver and the formulation presented below.

Let  $x_{ij}$  be a binary variable indicating whether a regular vehicle traverses arc  $(i, j)$ . Let  $y_{ij}$  indicate the load a regular vehicle carries on the arc  $(i, j)$ . Let  $z_i$  be a binary variable indicating whether a customer is visited by a regular vehicle. Let  $w_{ik}$  be a binary variable indicating whether customer  $i$  is visited by occasional driver  $k$ . Finally, let  $\beta_{ik}$  indicate whether occasional driver  $k$  can serve customer  $i$ . For the ease of presentation, we let  $p_{ik}$  indicate the compensation paid to occasional driver  $k$  when delivering to customer  $i$ . Note that in the basic compensation scheme  $p_{ik} = p_i$  for all  $k \in K$ , i.e., the compensation for delivering to a customer does not depend on the destination of the occasional driver. The VRPOD can be formulated as follows

$$\min \sum_{(i,j) \in A} c_{ij} x_{ij} + \sum_{i \in C} \sum_{k \in K} p_{ik} w_{ik} \quad (1)$$

$$\sum_{j|(i,j) \in A} x_{ij} = \sum_{j|(j,i) \in A} x_{ji} = z_i \quad \forall i \in C \quad (1)$$

$$\sum_{j|(0,j) \in A} x_{0j} - \sum_{j|(j,0) \in A} x_{j0} = 0 \quad (2)$$

$$\sum_{j|(j,i) \in A} y_{ji} - \sum_{j|(i,j) \in A} y_{ij} = \begin{cases} d_i z_i & \forall i \in C \\ \sum_{i \in C} -d_i z_i & i = 0 \end{cases} \quad (3)$$

$$y_{ij} \leq Q x_{ij} \quad \forall (i, j) \in A \quad (4)$$

$$y_{i0} = 0 \quad \forall i \in C \quad (5)$$

$$w_{ik} \leq \beta_{ik} \quad \forall i \in C \quad \forall k \in K \quad (6)$$

$$\sum_{i \in C} w_{ik} \leq 1 \quad \forall k \in K \quad (7)$$

$$\sum_{k \in K} w_{ik} + z_i = 1 \quad \forall i \in C \quad (8)$$

$$x_{ij} \in \{0, 1\} \quad \forall (i, j) \in A$$

$$z_i \in \{0, 1\} \quad i \in C$$

$$w_{ik} \in \{0, 1\} \quad i \in C \quad k \in K$$

$$y_{ij} \geq 0 \quad \forall (i, j) \in A.$$

The objective function aims at minimizing the total cost. Constraints (1) and (2) are flow conservation constraints. Constraints (3) ensure demand is satisfied and that subtours are prevented. Constraints (4) ensure vehicle capacity is respected. Constraints (5) establish that the vehicles must return empty to the depot. Constraints (6) ensure that a customer is assigned to an occasional driver willing to serve that customer. Constraints (7) ensure that an occasional driver serves at most one customer. Constraints (8) ensure that each customer is served exactly once.

## 4 A multi-start heuristic for the VRPOD

We propose a multi-start heuristic for the solution of the VRPOD, which combines variable neighborhood search and tabu search. Integer programs are also solved to optimality to assign customers to occasional drivers, thus the solution approach may be seen as belonging to the class of matheuristics. A high-level overview of MATHOD, the name we have given to the approach, can be found in Algorithm 1.

---

### Algorithm 1: MATHOD

---

```

 $I \leftarrow \text{CONSTRUCTINITIALSOLUTIONS};$ 
forall the  $i \in I$  do
     $s \leftarrow i;$ 
     $k \leftarrow 1;$ 
    while  $k = k_{max}$  do
         $s \leftarrow \text{INTERNALTABUSEARCH}(s);$ 
         $s \leftarrow \text{JUMP}(k, s);$ 
        if  $s$  is better than  $s_{best}$  then
             $s_{best} \leftarrow \text{TWOOPT}(s);$ 
             $s \leftarrow s_{best};$ 
             $k \leftarrow 1;$ 
        else
             $k \leftarrow k + 1;$ 
        end
    end
end

```

---

Next, we describe the procedures that comprise MATHOD in more detail.

CONSTRUCTINITIALSOLUTIONS. MATHOD embeds a multi-start approach with five different initial solutions. Any initial solution is characterized by a set  $S \subseteq C$  of customers to be served by regular drivers. The routes for the regular drivers are constructed using a sequential greedy insertion algorithm, which inserts customers into the active route in non-decreasing order of their distance to the depot, starting a new route when the vehicle capacity constraint would be violated, and, after all customers in  $S$  have been inserted, improves each route using a 2-exchange neighborhood. The five different sets  $S$  are generated as follows:

1. Simply take  $S = C$ , i.e., all customers are served by regular drivers.
2. Solve the following integer program to determine the set of customers to be served by occasional drivers, where, as before,  $w_{ik}$  is a binary variable indicating whether customer  $i$  is visited by occasional driver  $k$  and  $\beta_{ik}$  indicates whether occasional driver  $k$  can serve customer

$i$ :

$$\begin{aligned} \max \sum_{k \in K} \sum_{i \in C} (2c_{0i} - p_{ik})w_{ik} \\ w_{ik} \leq \beta_{ik} \end{aligned} \quad \forall i \in C \quad \forall k \in K \quad (9)$$

$$\sum_{i \in C} w_{ik} \leq 1 \quad \forall k \in K \quad (10)$$

$$\sum_{k \in K} w_{ik} \leq 1 \quad \forall c \in C \quad (11)$$

$$w_{ik} \in \{0, 1\} \quad i \in C \quad k \in K. \quad (12)$$

That is, we seek to serve as many far-away customers with occasional drivers. Let  $S'$  be the set of customers served by an occasional driver, i.e., for which  $\sum_{k \in K} w_{ik} = 1$  in the optimal solution. Then  $S = C \setminus S'$ .

3. Solve the integer program

$$\max \sum_{i \in C} \sum_{k \in K} d_i w_{ik}$$

subject to constraints (9)-(12). That is, maximize the demand delivered by the occasional drivers. Again, if  $S'$  is the set of customers served by an occasional driver in the optimal solution, then  $S = C \setminus S'$ .

4. Solve the integer program

$$\max \sum_{i \in C} \sum_{k \in K} w_{ik}$$

subject to constraints (9)-(12). Let  $\theta$  be the value of an optimal solution. Note that  $\theta$  corresponds to the maximum number of occasional drivers that can be engaged to serve customers. Next, solve the following integer program:

$$\min \sum_{i \in C} \sum_{k \in K} p_{ik} w_{ik}$$

subject to constraints (9)-(12) and

$$\sum_{k \in K} w_{ik} = \theta,$$

i.e., find the minimum cost solution that uses the maximum possible number of occasional drivers. As before, if  $S'$  is the set of customers served by an occasional driver in the optimal solution, then  $S = C \setminus S'$ .

5. Solve a relaxation of the integer program presented in Section 3 obtained by ignoring the integrality constraint on variables  $\mathbf{x}$ . The set  $S$  contains those customers  $i$  for which  $z_i = 1$ .

INTERNALTABUSEARCH( $s$ ). For a given solution  $s$ , let  $C'$  denote the set of customers served by occasional drivers and let  $r_i$ , for  $i \in C \setminus C'$ , denote the route of the regular driver serving customer  $i$ . For a customer  $i$  and a route  $r \neq r_i$ , we denote by  $r + i$  the route obtained by inserting  $i$  into  $r$

(using the cheapest insertion criterion to determine where to insert  $i$ ). Similarly, given a route  $r_i$  of a regular driver and a customer  $i$  served on  $r_i$ , we denote  $r_i - i$  the route obtained from  $r_i$  by deleting  $i$  and joining its predecessor with its successor. The tabu search performed on  $s$  uses four feasible moves:

- *1-move*: A customer  $i \in C \setminus C'$  is moved from its route  $r_i$  to a route  $r \neq r_i$ , where  $r$  may be the empty route. Hence,  $r_i$  and  $r$  are replaced by  $r_i - i$  and  $r + i$ , respectively.
- *swap-move*: Customers  $i, j \in C \setminus C'$ , with  $r_i \neq r_j$ , are exchanged. Hence,  $r_i$  and  $r_j$  are replaced by  $(r_i - i) + j$  and  $(r_j - j) + i$ , respectively.
- *in-move*: A customer  $i \in C'$  is inserted into a route  $r$ . Hence,  $r$  is replaced by  $r + i$ .
- *out-move*: A customer  $i \in C \setminus C'$  is assigned to an occasional driver (if possible, i.e., there exists an occasional driver  $k$ , which does not serve a customer in  $s$  and for which  $\beta_{ik} = 1$ ). Hence,  $r_i$  is replaced by  $r_i - i$ .

All non-tabu moves are evaluated and the best one is chosen.

A temporary tabu status forbids customers to be inserted in routes from which they have been recently removed. Also, when a customer  $i \in C'$  ( $C \setminus C'$ ) is moved to  $C \setminus C'$  ( $C'$ ), then it is temporarily tabu to move it back to  $C'$  ( $C \setminus C'$ ). Each time a new best solution is found, each route is improved with a local search procedure using a 2-exchange neighborhood. The tabu search terminates after  $n_{max}$  iterations without improvement. At each iteration all moves are evaluated and the best non-tabu move is chosen. A tabu move is chosen only if it improves the best solution found so far.

JUMP( $k, s$ ). For a given solution  $s$ , a customer  $i \in C'$  is randomly selected and inserted in a route  $r$  (using the cheapest insertion criterion to determine where to insert  $i$ ). If no route can accommodate customer  $i$ , then a new route is created. The procedure is repeated  $\min\{k, |C'|\}$  times. Let  $C''$  be the set of customers which are still served by an occasional driver plus those customers served by a regular driver in a route which visits a single customer, but excluding those customers that have just been removed from  $C'$ . Furthermore, let  $K'$  be the set of occasional drivers who are not used in the current solution. Then, we solve the following integer program:

$$\begin{aligned}
& \max \sum_{k \in K'} \sum_{i \in C''} \gamma_i w_{ik} \\
& w_{ik} \leq \beta_{ik} && \forall i \in C'' \quad \forall k \in K' \\
& \sum_{i \in C''} w_{ik} \leq 1 && \forall k \in K' \\
& \sum_{k \in K'} w_{ik} \leq 1 && \forall i \in C'' \\
& w_{ik} \in \{0, 1\} && i \in C'' \quad k \in K',
\end{aligned}$$

where  $\gamma_i$  is randomly drawn from a uniform distribution over  $[0, 10]$ . That is, we seek to generate a new assignment of customers to occasional drivers and thus a solution that considerably differs from the one obtained at the end of the previous internal tabu search phase.

TWOOPT( $s$ ). This procedure aims at improving each regular driver route using a 2-exchange neighborhood.

## 5 A computational study

As mentioned in the introduction, the aim of our research is to gain initial quantitative insights in the potential benefits of crowdshipping for last-mile delivery. More specifically, we aim to understand the impact of the number of occasional drivers, the flexibility of the occasional drivers, and the employed compensation scheme on the potential benefits, i.e., cost savings. To be able to do so, we conducted a series of computational experiments using randomly generated instances, because the use of occasional drivers to make deliveries is still in its infancy and at a conceptual stage, and no real-life data exist.

In the remainder of this section, we discuss how we generated instances, how we evaluated the performance of the multi-start heuristic, and what we learned about the potential benefits of crowdshipping for last-mile delivery. All experiments were conducted on an Intel(R) Xeon(R) CPU E5-1650 v2 3.50 GHz and 16 GB of RAM. Cplex 12.6 was used as exact solver. The algorithms were coded in Visual C++.

### 5.1 Instance generation

Our starting point for generating instances are the well-known and widely available Solomon instances of the CVRP with time windows. Solomon’s 56 instances are generated from 6 basic instances by imposing time windows of different widths. Since we do not consider time windows, we consider only the six basic instances: two clustered instances (C101 and C201), two instances where customers are randomly distributed (R101 and R201), and two instances where customers are partially clustered and partially randomly distributed (RC101 and RC201). For a given CVRP instance, with customer locations  $(x_i, y_i)$ , we generate destinations for the occasional drivers uniformly randomly in the square with lower left hand corner  $(\min_i x_i, \min_i y_i)$  and upper right hand corner  $(\max_i x_i, \max_i y_i)$ .

### 5.2 Performance evaluation of the multi-start heuristic

To assess the performance of the multi-start heuristic, we compare the solution it produces to an optimal solution found by solving the formulation presented in Section 3 (using CPLEX). The internal tabu search is terminated after 400 iterations without an improvement, i.e.,  $n_{max} = 400$ . The value of  $k_{max}$  is set to  $|C|/3$ . In Table 1, we report, for six basic instances and ten variations, the number and flexibility of the occasional drivers ( $|K|$  and  $\zeta$ ), the total cost, the fraction of the total cost representing compensation to occasional drivers, the number of regular driver routes, and the number of deliveries made by occasional drivers. The occasional drivers are compensated for the distance traveled to a customer’s location at a rate of  $\rho = 0.2$ . We observe that the multi-start heuristic performs very well. In all but one instance, the cost gap is less than 1%, and for instances with 25 occasional drivers, the average cost gap is only 0.05%. We also observe that, as expected, the number of occasional drivers used (and, thus, the fraction of the total cost representing compensation to occasional drivers) increases as the number of available occasional drivers increases.

Table 1: Performance of the multi-start heuristic.

		multi-start heuristic				IP				% cost gap
		total cost	% OD cost	#routes	#OD used	total cost	% OD cost	#routes	#OD used	
	$\zeta = 1.1$	166.7	10	2	5	166.7	10	2	5	0.01
		172.3	6	1	3	171.4	6	1	3	0.53
		312.7	7	2	6	312.7	7	2	6	0.00
		282.4	9	1	7	282.4	9	1	7	0.00
		295.0	0	3	0	295.0	0	3	0	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
	$\zeta = 1.2$	165.7	11	2	5	165.7	11	2	5	0.00
		172.3	6	1	3	171.4	6	1	3	0.53
		274.6	12	1	8	274.6	12	1	8	0.00
		274.6	12	1	8	274.6	12	1	8	0.00
		284.5	12	2	5	284.5	12	2	5	0.00
		227.0	0	1	0	226.1	0	1	0	0.40
$ K  = 13$	$\zeta = 1.3$	165.4	11	2	7	165.4	11	2	7	0.00
		170.4	10	1	4	169.5	10	1	4	0.54
		268.1	16	1	9	268.1	16	1	9	0.00
		268.1	16	1	9	268.1	16	1	9	0.00
		284.5	12	2	5	284.5	12	2	5	0.00
		227.0	0	1	0	226.1	0	1	0	0.40
	$\zeta = 1.4$	165.4	11	2	7	165.4	11	2	7	0.00
		169.2	12	1	6	168.3	12	1	6	0.54
		268.1	16	1	9	266.4	17	1	10	0.63
		268.1	16	1	9	266.4	17	1	10	0.63
		284.5	12	2	5	284.5	12	2	5	0.00
		227.0	0	1	0	226.1	0	1	0	0.40
	$\zeta = 1.5$	146.7	28	1	11	146.1	28	1	11	0.46
		169.2	12	1	6	166.6	22	1	9	1.56
		258.0	17	1	10	256.5	17	1	9	0.56
		258.0	17	1	10	256.5	17	1	9	0.56
		281.7	12	2	5	281.7	12	2	5	0.00
		227.0	0	1	0	226.1	0	1	0	0.40
<b>average</b>			10		5.4		10		5.5	0.27
	$\zeta = 1.1$	165.4	11	2	7	165.4	11	2	7	0.00
		170.4	9	1	5	170.2	9	1	5	0.12
		276.7	12	1	9	276.7	12	1	9	0.00
		276.6	11	1	8	274.5	10	1	8	0.77
		280.3	12	2	5	280.3	12	2	5	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
	$\zeta = 1.2$	165.4	11	2	7	165.4	11	2	7	0.00
		170.2	9	1	5	170.2	9	1	5	0.00
		262.9	13	1	9	262.9	13	1	9	0.00
		264.5	16	1	10	262.9	13	1	9	0.62
		280.3	12	2	5	280.3	12	2	5	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
$ K  = 25$	$\zeta = 1.3$	165.4	11	2	7	165.4	11	2	7	0.00
		165.5	20	1	9	165.5	20	1	9	0.00
		257.1	18	1	11	257.1	18	1	11	0.00
		257.1	18	1	11	257.1	18	1	11	0.00
		280.3	12	2	5	280.3	12	2	5	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
	$\zeta = 1.4$	143.5	30	1	12	143.5	30	1	12	0.00
		163.0	26	1	10	163.0	26	1	10	0.00
		241.8	28	1	14	241.8	28	1	14	0.00
		241.8	28	1	14	241.8	28	1	14	0.00
		280.3	12	2	5	280.3	12	2	5	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
	$\zeta = 1.5$	143.5	30	1	12	143.5	30	1	12	0.00
		163.0	26	1	10	163.0	26	1	10	0.00
		229.6	31	1	15	229.6	31	1	15	0.00
		229.6	31	1	15	229.6	31	1	15	0.00
		280.3	12	2	5	280.3	12	2	5	0.00
		226.1	0	1	0	226.1	0	1	0	0.00
<b>average</b>			15		7.5		15		7.47	0.05

### 5.3 Potential benefits of employing occasional drivers

The potential benefits of employing occasional drivers to make deliveries depend on three factors:

- how many occasional drivers there are relative to the number of customers that need to be served, i.e.,  $\frac{|K|}{|C|}$ ;
- how much flexibility an occasional driver has, i.e., the value of  $\zeta$ ; and
- how much an occasional driver is paid for making a delivery, i.e., the compensation scheme and the value of  $\rho$ .

To obtain quantitative insights into the impact of these three factors, we solve each instance in the test set with every possible combination of the following parameter values:  $\frac{|K|}{|C|} = 0.5$  and  $1.0$ ,  $\zeta = 1.1, 1.2, 1.3, 1.4$ , and  $1.5$ , and  $\rho = 0.05, 0.1$ , and  $0.2$  for the default compensation scheme (occasional driver-independent) and  $\rho = 1.2, 1.4$ , and  $1.6$  for the alternative compensation scheme (occasional driver-dependent). Thus, the test set consists of 180 instances.

We assess the potential benefits of employing occasional drivers by examining and comparing the total cost, the fraction of the total cost contributed by regular drivers, and the number of occasional drivers employed for the different solutions.

In our analysis and comparisons, we also include the solution obtained when no occasional drivers are available (which provides an upper bound on total cost) and the solution obtained when occasional drivers are free, i.e., do not have to be compensated, and have the largest flexibility  $\zeta = 1.5$  (which provides a lower bound on total cost). The first solution is obtained by solving a classical CVRP. For this, we use the open-access injection-ejection algorithm from the COIN-OR library available at the following URL: <http://www.coin-or.org/projects/VRPH.xml>. The second solution is obtained using MATHOD.

#### 5.3.1 Analysis of a specific instance

We start by examining the results for one specific instance in more detail. The instance is instance C201 with 100 customers and 100 occasional drivers. The customer locations, the locations of the destinations of the occasional drivers, and the location of the depot can be seen in Figure 1.

An area of interest on the right side of the instance is indicated by a red ellipse. The reason this area is of interest is that it is relatively far away from the depot and that there are more occasional drivers going to the area than there are customers. As such, the area represents an opportunity for cost-savings for the company. By exploiting occasional drivers for deliveries in that area, it may not be necessary for regular drivers to visit the area. This is quite different, for example, from the area at the top right where few, if any, occasional drivers go, but where there are a number of customers.

Before analyzing solutions involving occasional drivers, we show, in Figure 2, the VRP solution, i.e., the solution that only considers regular drivers for customer deliveries, which consists of three routes.

Next, we take a look at the solution obtained when occasional drivers are fairly flexible, i.e.,  $\zeta = 1.5$ , and the default compensation scheme (Compensation Scheme I) with parameter  $\rho = 0.2$  is used, i.e., the scheme that compensates occasional drivers for the distance from the depot to the customer at the highest compensation rate we consider in the study. Recall that the intuition behind this compensation scheme is that we want occasional drivers to make deliveries to far-away

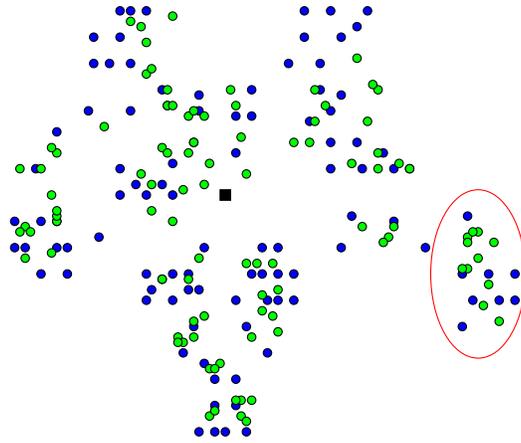


Figure 1: Customers and destinations of ODs.

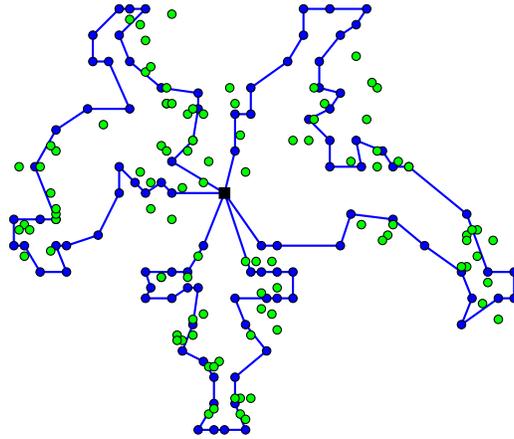


Figure 2: VRP solution.

customers and therefore we compensate them based on the distance from the depot to the customer. The results can be found in Figure 3, where we show the routes of the regular drivers, the customers served by occasional drivers (shown in blue), and the destinations of occasional drivers that are *not* used (in green).

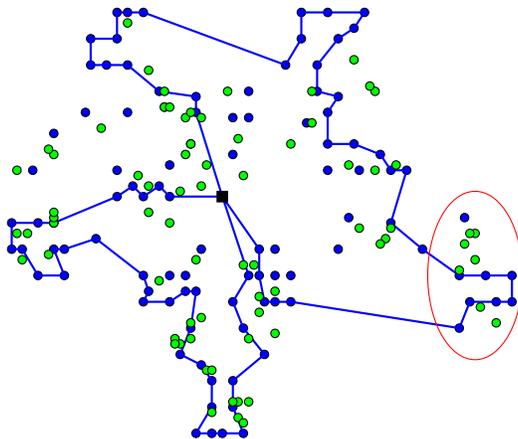


Figure 3: Routes, customers served by ODs, and destinations of ODs not used for Compensation Scheme I with  $\zeta = 1.5$  and  $\rho = 0.2$ .

Unfortunately, even though the occasional drivers are fairly flexible, the compensation scheme does not have the desired effect. One of the regular drivers, for example, visits the area indicated by the red ellipse contrary to what we want to achieve. The reason is that once a regular driver reaches an area, it is relatively cheap to visit one more customer. The extra mileage to visit another customer in the same area is small, which means that it is not worth paying an occasional driver to make the delivery instead (it would lead to a cost increase rather than result in a cost saving).

The alternative compensation scheme (Compensation Scheme II), which is based on the extra mileage for an occasional driver, might be more appropriate. The results (for the same occasional driver flexibility) for two compensation rates are shown in Figure 4.

We see that at the highest compensation rate, the same effect occurs. It is still cheaper to send a regular driver to the area indicated by the red ellipse. However, at the lowest compensation rate, we do see the desired effect. All the customers in the area indicated by the red ellipse are served by occasional drivers.

The three solutions shown in Figures 3 and 4 clearly demonstrate the sensitivity of the solution to the compensation employed and the parameters used and highlights the challenge associated with defining an appropriate compensation scheme.

Finally, in Figure 5, we show the solution that results when the use of occasional drivers is free (but they still have limited flexibility given by  $\zeta = 1.5$ ), i.e., the penultimate situation for the company.

As expected, the customers at the top right are still served by a regular driver, because there are simply no occasional drivers that are willing to go that far out of their way to make a delivery there.

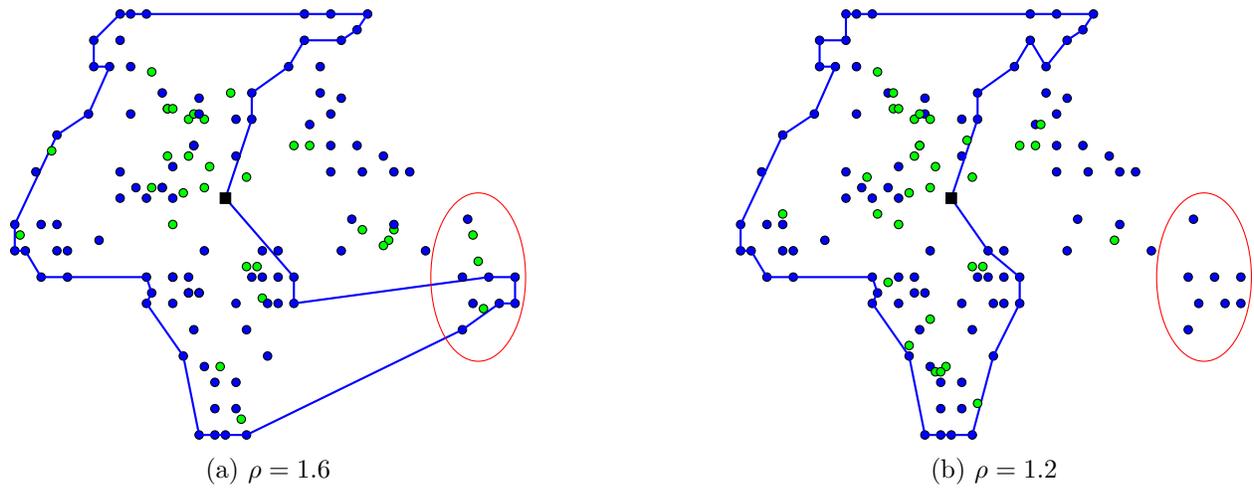


Figure 4: Routes, customers served by ODs, and destinations of ODs not used for Compensation Scheme II with  $\zeta = 1.5$  and two different values of  $\rho$ .

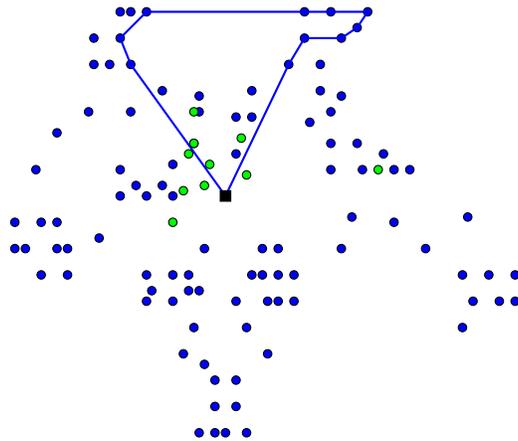


Figure 5: Routes, customers served by ODs, and destinations of ODs not used when there is no charge for using ODs.

### 5.3.2 Analysis of all instances

Here, we provide summary results for all instances in our test set. Tables 2 and 3 give the results when Compensation Scheme I is employed. Table 2 shows the percentage reduction in total cost and regular drivers used with respect to the VRP solution where no occasional driver is available, whereas Table 3 shows the percentage of available occasional drivers used and the compensation paid to occasional drivers as a percentage of the total cost of the solution. Each table is organized as follows. The top part of the table presents results for instances with 50 ODs, while the bottom part of the table presents results for instances with 100 ODs. For each part, rows correspond to different values of  $\zeta$  and columns correspond to different values of  $\rho$ . Note that the values reported for each  $(\zeta, \rho)$  pair are averages over six instances. The last row and the last column present averages over all instances with the same value of  $\rho$  and  $\zeta$ , respectively. Tables 4 and 5 give the same results for Compensation Scheme II.

Table 2: Percentage reduction of total cost and of number of routes with respect to VRP: Compensation Scheme I.

$ K  = 50$								
	$\rho = 0.05$		$\rho = 0.1$		$\rho = 0.2$		<b>average</b>	
	% cost reduction	% route reduction						
$\zeta = 1.1$	28.08	46.30	21.28	46.30	10.65	46.30	20.00	46.30
$\zeta = 1.2$	33.08	50.05	26.27	50.05	13.75	48.38	24.37	49.49
$\zeta = 1.3$	34.16	49.81	26.91	49.81	14.56	50.05	25.21	49.89
$\zeta = 1.4$	36.38	49.81	28.60	51.90	14.80	41.90	26.59	47.87
$\zeta = 1.5$	38.55	49.81	30.45	49.81	15.52	51.90	28.17	50.51
<b>average</b>	34.05	49.16	26.70	49.57	13.85	47.70	24.87	48.81

$ K  = 100$								
	$\rho = 0.05$		$\rho = 0.1$		$\rho = 0.2$		<b>average</b>	
	% cost reduction	% route reduction						
$\zeta = 1.1$	40.11	62.87	30.19	61.20	13.90	53.80	28.07	59.29
$\zeta = 1.2$	46.11	66.81	34.71	66.81	16.78	57.31	32.54	63.64
$\zeta = 1.3$	58.14	68.24	43.55	68.24	18.99	60.83	40.23	65.77
$\zeta = 1.4$	60.47	72.18	45.37	68.24	20.08	58.98	41.97	66.47
$\zeta = 1.5$	65.56	72.18	48.55	70.32	20.85	64.77	44.98	69.09
<b>average</b>	54.08	68.45	40.47	66.96	18.12	59.14	37.56	64.85

We first comment on the results obtained when Compensation Scheme I is employed, i.e., the results in Tables 2 and 3. As expected, the cost reduction increases as the ODs become more flexible, i.e., when  $\zeta$  increases, and the cost reduction decreases when ODs are compensated at a higher rate, i.e., when  $\rho$  increases. Also as expected, when the number of available ODs is larger, the reduction in total cost is greater, especially when their compensation rate is low ( $\rho = 0.05$ ). With 50 flexible ODs ( $\zeta = 1.5$ ) compensated at a low rate, the cost reduction is 38.55%; with 100 flexible ODs compensated at a low rate, the cost reduction is 65.56%.

When ODs are cheap ( $\rho = 0.05$ ) and there are 50 ODs ( $|K| = 50$ ), almost all ODs are used. In fact, the percentage of available ODs used does not tell the complete story, because it is not always possible to use all available ODs. For example, if the flexibility of ODs is small ( $\zeta = 1.1$ ), it may happen that two ODs can only make a delivery at the same customer. The average over the six instances with 50 ODs and  $\zeta = 1.1$  of the maximum percentage of available ODs that can be used is 88%. This average goes up to more than 97% for the six instances with 50 ODs and  $\zeta = 1.5$ . This phenomenon is even more pronounced when there are 100 ODs. The average over the six instances

Table 3: Percentage of ODs used and of ODs cost: Compensation Scheme I.

$ K  = 50$								
	$\rho = 0.05$		$\rho = 0.1$		$\rho = 0.2$		<b>average</b>	
	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost
$\zeta = 1.1$	81.33	8.64	78.00	14.89	71.00	22.60	76.78	15.37
$\zeta = 1.2$	88.33	10.70	86.33	18.09	75.67	25.05	83.44	17.95
$\zeta = 1.3$	92.33	11.79	89.00	19.03	75.67	25.73	85.67	18.85
$\zeta = 1.4$	91.33	12.60	91.67	20.52	70.00	24.83	84.33	19.32
$\zeta = 1.5$	96.00	14.20	90.33	22.64	82.00	30.41	89.44	22.42
<b>average</b>	89.87	11.59	87.07	19.03	74.87	25.72	83.93	18.78

$ K  = 100$								
	$\rho = 0.05$		$\rho = 0.1$		$\rho = 0.2$		<b>average</b>	
	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost
$\zeta = 1.1$	61.50	17.55	56.67	27.03	44.67	32.44	54.28	25.67
$\zeta = 1.2$	68.67	23.37	64.00	34.06	50.00	38.60	60.89	32.01
$\zeta = 1.3$	78.33	37.29	72.83	49.72	57.00	48.76	69.39	45.26
$\zeta = 1.4$	82.50	41.52	75.17	53.50	55.17	49.26	70.94	48.09
$\zeta = 1.5$	88.67	52.03	80.17	61.53	62.33	56.87	77.06	56.81
<b>average</b>	75.93	34.35	69.77	45.17	53.83	45.19	66.51	41.57

with 100 ODs and  $\zeta = 1.1$  of the maximum percentage of available ODs that can be used is just over 70%, but goes up to just over 94% for the six instances with 100 ODs and  $\zeta = 1.5$ . We also see that when ODs are cheap ( $\rho = 0.05$ ) the reduction in total cost can be substantial, especially if the number of available ODs is large (100) and they are flexible ( $\zeta = 1.5$ ). When ODs are expensive ( $\rho = 0.2$ ), a large number of available and flexible ODs is necessary to be able to achieve a high total cost reduction, because fewer ODs will be used and they have to be chosen carefully.

We now comment the results obtained when Compensation Scheme II is employed, i.e., the results in Tables 4 and 5. Interestingly, the total cost reduction and the number of available ODs used do not vary much when the flexibility ( $\zeta$ ) or the rate of compensation ( $\rho$ ) changes. The cost reduction is 24.70% with a low flexibility ( $\zeta = 1.1$ ) and a high compensation rate ( $\rho = 1.6$ ) and is 28.69% with a high flexibility ( $\zeta = 1.5$ ) and a low compensation rate ( $\rho = 1.2$ ). Also, we see that the number of available ODs used no longer increases when the flexibility of ODs ( $\zeta$ ) increases. In fact, because the maximum number of available ODs that can be used does increase with the flexibility, this means that effectively a smaller fraction of the available ODs is used when the flexibility increases. However, these ODs are chosen more carefully, which is one of the reasons that we see that the compensation paid to ODs, as a percentage of the total cost, increases when the flexibility of ODs increases, especially when there is a large number of available ODs. (Another factor that contributes to the increase is that the total cost itself decreases.)

When comparing the two compensation schemes, we see that the Compensation Scheme II has the advantage that it is much less sensitive to the flexibility of the ODs and the to the rate at which they are compensated.

In Tables 6 and 7, the same results are reported in a different way, for Compensation Scheme I and II, respectively. We show the percentage reduction in the total cost and the number of regular drivers used with respect to the VRP solution, as well as the percentage of available ODs used and the compensation paid to the ODs as a percentage of total cost, when averaged over groups of instances. The groups are determined by the base instance, the number of available ODs, the flexibility of the ODs, and the rate of compensation of the ODs. The last row gives the average over all instances. Looking at the averages over all instances for the two compensation schemes,

Table 4: Percentage reduction of total cost and of number of routes with respect to VRP: Compensation Scheme II.

$ K  = 50$								
	$\rho = 1.2$		$\rho = 1.4$		$\rho = 1.6$		<b>average</b>	
	% cost reduction	% route reduction						
$\zeta = 1.1$	25.55	46.30	25.08	46.30	24.70	46.30	25.11	46.30
$\zeta = 1.2$	27.44	47.96	26.79	50.05	26.10	46.30	26.78	48.10
$\zeta = 1.3$	28.19	47.96	27.12	47.96	26.50	47.96	27.27	47.96
$\zeta = 1.4$	28.19	50.05	27.58	50.05	26.66	47.96	27.48	49.35
$\zeta = 1.5$	28.69	47.96	27.58	50.05	26.61	47.96	27.63	48.66
<b>average</b>	27.61	48.05	26.83	48.88	26.12	47.30	26.85	48.07
$ K  = 100$								
	$\rho = 1.2$		$\rho = 1.4$		$\rho = 1.6$		<b>average</b>	
	% cost reduction	% route reduction						
$\zeta = 1.1$	38.98	62.87	38.19	62.87	37.44	62.87	38.20	62.87
$\zeta = 1.2$	40.70	64.72	39.52	64.72	38.40	62.87	39.54	64.10
$\zeta = 1.3$	43.27	66.39	41.12	66.39	39.40	64.54	41.26	65.77
$\zeta = 1.4$	43.56	66.39	42.20	66.39	40.26	64.54	42.01	65.77
$\zeta = 1.5$	43.98	66.39	41.71	66.39	40.22	64.54	41.97	65.77
<b>average</b>	42.10	65.35	40.55	65.35	39.14	63.87	40.60	64.86

Table 5: Percentage of ODs used and of ODs cost: Compensation Scheme II.

$ K  = 50$									
	$\rho = 1.2$		$\rho = 1.4$		$\rho = 1.6$		<b>average</b>		
	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD cost
$\zeta = 1.1$	68.42	9.90	68.92	9.75	64.67	11.03	67.58	10.23	12.31
$\zeta = 1.2$	72.33	11.98	71.50	12.24	59.33	11.71	69.40	11.97	12.31
$\zeta = 1.3$	73.25	12.49	73.25	12.52	65.00	13.66	71.60	12.89	12.31
$\zeta = 1.4$	75.50	13.36	75.00	13.74	64.50	13.01	73.10	13.37	12.31
$\zeta = 1.5$	73.83	12.15	74.92	13.72	64.50	13.38	72.40	13.09	12.31
<b>average</b>	72.67	11.98	72.72	12.39	63.70	12.56	70.75	12.31	12.31
$ K  = 100$									
	$\rho = 1.2$		$\rho = 1.4$		$\rho = 1.6$		<b>average</b>		
	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD used	% OD cost	% OD cost
$\zeta = 1.1$	60.83	18.18	60.50	18.89	60.00	19.29	60.44	18.79	28.77
$\zeta = 1.2$	63.00	22.83	61.33	22.55	59.33	21.39	61.22	22.25	28.77
$\zeta = 1.3$	66.17	32.85	65.17	30.12	65.00	36.87	65.44	33.28	28.77
$\zeta = 1.4$	66.33	34.26	66.33	35.78	64.50	32.79	65.72	34.28	28.77
$\zeta = 1.5$	66.67	36.84	66.83	36.00	64.50	32.84	66.00	35.23	28.77
<b>average</b>	64.60	28.99	64.03	28.67	62.67	28.64	63.77	28.77	28.77

we see that, as expected, the main difference is in the compensation paid to ODs as a percentage of total cost.

Table 6: Summary: Compensation Scheme I.

	% cost reduction w.r.t. VRP	% route reduction w.r.t. VRP	% OD used	% OD cost w.r.t. total cost
C101	37.40	70.67	87.33	40.61
C201	18.61	44.44	63.63	24.57
R101	39.55	67.92	82.93	33.52
R201	31.53	50.00	75.77	30.71
RC101	34.76	59.63	74.57	26.83
RC201	25.43	48.33	67.10	24.81
$ K  = 50$	24.87	48.81	83.93	18.78
$ K  = 100$	37.56	64.85	66.51	41.57
$\zeta = 1.1$	24.03	52.79	65.53	20.52
$\zeta = 1.2$	28.45	56.57	72.17	24.98
$\zeta = 1.3$	32.72	57.83	77.53	32.05
$\zeta = 1.4$	34.28	57.17	77.64	33.71
$\zeta = 1.5$	36.58	59.80	83.25	39.61
$\rho = 0.05$	44.06	58.81	82.90	22.97
$\rho = 0.1$	33.59	58.27	78.42	32.10
$\rho = 0.2$	15.99	53.42	64.35	35.46
<b>average</b>	31.21	56.83	75.22	30.17

Table 7: Summary: Compensation Scheme II.

	% cost reduction w.r.t. VRP	% route reduction w.r.t. VRP	% OD used	% OD cost w.r.t. total cost
C101	43.85	71.67	85.50	35.36
C201	20.49	50.00	66.42	17.04
R101	40.79	64.17	74.92	21.27
R201	33.70	50.00	71.32	20.58
RC101	33.47	52.96	64.20	14.70
RC201	30.05	50.00	61.56	14.28
$ K  = 50$	26.85	48.07	80.73	12.31
$ K  = 100$	40.60	64.86	63.77	28.77
$\zeta = 1.1$	31.66	54.58	67.58	14.51
$\zeta = 1.2$	33.16	56.10	69.40	17.11
$\zeta = 1.3$	34.27	56.87	71.60	23.09
$\zeta = 1.4$	34.74	57.56	73.10	23.82
$\zeta = 1.5$	34.80	57.21	72.40	24.16
$\rho = 0.05$	34.86	56.70	72.67	20.48
$\rho = 0.1$	33.69	57.12	72.72	20.53
$\rho = 0.2$	32.63	55.58	63.70	20.60
<b>average</b>	33.72	56.47	70.75	20.54

## 6 Final remarks

The goal of our investigation has been to gain an initial understanding of the potential benefits and the implementation challenges associated with crowdshipping. The results of our computational study are both encouraging and dispiriting. We have found that substantial cost savings can be realized when there is a large number of people with a generous amount of flexibility available to make deliveries. That, of course, depends to a large extent on the compensation offered. Designing an appropriate and cost-effective compensation scheme is one of a major implementation challenges

associated with crowdshipping. We have experimented with two natural compensation schemes and found that the performance of both of them was passable, but sensitive to the choice of parameters. Compensation schemes based on the “cost-to-serve” of a customer may be most appropriate. Of course determining the cost-to-serve of a customer is notoriously difficult in routing problems, and will be even more complicated than usual in this setting as it needs to account for the fact that some of the customers may be served by crowdshippers. Research into more sophisticated compensation schemes is both interesting and necessary. Another avenue for further research, maybe even more important and more interesting, is the study of variants of the VRPOD, in which aspects of the highly dynamic nature of the setting are captured. In reality, occasional drivers become available over time and their services will only be available for a short period of time after they become available. Of course, orders from customers become available over time too, and, these orders imply new *deliveries*. Almost all the literature on dynamic vehicle routing focuses on orders to be picked up, orders to be picked up and delivered, or to a service performed by the driver. The underlying structure of the dynamic routing problem changes significantly when new orders have to be delivered, since there are few, if any, opportunities to accommodate additional deliveries after a delivery vehicle has left the depot, because the vehicle would have to return to the depot to pick up the additional deliveries.

## References

- N. Agatz, A. Erera, M. Savelsbergh, and X. Wang. Dynamic ride-sharing: a simulation study in metro atlanta. *Procedia - Social and Behavioral Sciences*, 17:532 – 550, 2011.
- N. Agatz, A. Erera, M. Savelsbergh, and X. Wang. Optimization for dynamic ride-sharing: A review. *European Journal of Operational Research*, 223:295 – 303, 2012.
- A. Barr and J. Wohl. Exclusive: Wal-mart may get customers to deliver packages to online buyers. *REUTERS – Business Week*, 2013.
- G. Bensinger. Amazon’s next delivery drone: You. *Wall Street Journal*, 2015.
- M. Furuhata, M. Dessouky, F. Ordóñez, M.-E. Brunet, X. Wang, and S. Koenig. Ridesharing: The state-of-the-art and future directions. *Transportation Research Part B: Methodological*, 57:28–46, 2013.
- A. Lee and M. Savelsbergh. Dynamic ridesharing: Is there a role for dedicated drivers? *Transportation Research Part B: Methodological*, 2015. doi: <http://dx.doi.org/10.1016/j.trb.2015.02.013>.
- B. Li, D. Krushinsky, H. Reijers, and T. Van Woensel. The share-a-ride problem: People and parcels sharing taxis. *European Journal of Operational Research*, 238(1):31–40, 2014.
- P. Toth and D. Vigo. *Vehicle Routing: Problems, Methods, and Applications, Second Edition, MOS-SIAM Series on Optimization 18*. SIAM, Philadelphia, 2014.
- A. Trentini and N. Mahléne. Toward a shared urban transport system ensuring passengers & goods cohabitation. *TeMA. Journal of Land Use, Mobility and Environment*, 3, 2011.