Transportation Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Multicommodity vs. Single-Commodity Routing

Claudia Archetti, Ann Melissa Campbell, M. Grazia Speranza

TRANSPORTATION SCIENCE
Vol. 50, No. 2, May 2016, pp. 461–472
ISSN 0041-1655 (print) | ISSN 1526-5447 (online)

http://dx.doi.org/10.1287/trsc.2014.0528

Multicommodity vs. Single-Commodity Routing

Claudia Archetti
Department of Economics and Management, University of Brescia, 25122 Brescia, Italy, archetti@eco.unibs.it

Ann Melissa Campbell
Department of Management Sciences, University of Iowa, Iowa City, Iowa 52242, ann-campbell@uiowa.edu

M. Grazia Speranza
Department of Economics and Management, University of Brescia, 25122 Brescia, Italy, speranza@eco.unibs.it

In this paper we study a vehicle routing problem where customers request multiple commodities. We study the impact on transportation cost from using vehicles dedicated to a single commodity compared with using flexible vehicles capable of carrying any set of commodities. With vehicles that carry multiple commodities, we consider when the delivery to a customer can be made by more than one vehicle. If multiple vehicles can be used, we examine when deliveries of individual commodities may be split and when they may not be split. The latter problem has not previously been studied, and we present a mathematical programming model for it. We use worst case and computational analysis to compare these different models.

Keywords: multiple commodities; split deliveries; vehicle routing
History: Received: October 2012; revision received: July 2013; accepted: November 2013. Published online in *Articles in Advance* June 3, 2014.

1. Introduction

In most of the literature on vehicle routing problems (VRPs), a single numeric value expresses the demand of a customer even though it may represent a collection of different products. It is usually understood that all vehicles can serve all customers and deliver all products. If there are reasons to use different vehicles for different sets of products, it is assumed that it has been previously decided which set of vehicles will deliver which set of products. Then routing problems are solved, each limited to the customers that demand a set of products and to the vehicles assigned to this service.

In this paper, we compare the transportation cost implications from using vehicles dedicated to a set of products with using flexible vehicles capable of delivering all products. We will refer to a set of products that is served by a dedicated set of vehicles as a single commodity. Thus, if there are several sets of products with a set of vehicles dedicated to each, we will say that there are as many commodities as the number of sets of dedicated vehicles. This paper compares the situation when each of the multiple commodities is handled separately from the others with the situation when all commodities are handled jointly. In the latter situation, we examine several specific ways to handle the multiple commodities jointly. For example, when the same customer requests multiple commodities, the complete demand of an individual commodity may be requested to be served with a single vehicle, or this may not be necessary.

The comparison between handling commodities separately and handling them jointly is of interest in many different applications. The first category concerns situations where different sets of products are handled separately for organizational convenience and simplicity. Even though vehicles can carry multiple products, an individual vehicle is often loaded with one or only a small set of them. Having a single product or a small set of products on a vehicle speeds pickup from the warehouse, simplifies the loading and unloading phases, and avoids the need for reshuffling the load during the distribution process. Our study can be viewed as an examination of what this convenience translates to in terms of transportation costs.

There are also situations where different kinds of vehicles are available, namely vehicles that are specific to a commodity and vehicles that are flexible and may be used for different commodities. These flexible vehicles are often more expensive than the commodity-specific vehicles, and a decision has to be made about whether it is beneficial to invest in the flexible vehicles. A decision can be made only through the evaluation of the reduction of the operational costs with flexible vehicles with respect to the operational costs with commodity-specific vehicles. An example of this category of problems can be observed in waste management, where vehicles are available that collect one kind of waste only, and more expensive vehicles are available that may collect two or more different kinds of waste at a time. These latter vehicles have the container divided in as many compartments as the number of different commodities they may collect at the same time. The size of the compartments may be fixed, which implies that each compartment has a given capacity, or they may be totally or partly flexible. In many cases, the capacity of a compartment can be chosen among a predefined set of options. A similar situation can be found in the distribution of groceries, where a flexible vehicle may have compartments for frozen, refrigerated, and dry goods.

A multicompartment routing problem, with fixed capacity of each compartment, has been recently studied by Derigs et al. (2011), where research in multicompartment routing problems is discussed and related references can be found. Many of the analytical results in this stream of literature relax the problem to assume that the size of the compartments can be freely adjusted. These results provide upper bounds on the benefits from using vehicles with limited options on the size of the compartments. Similarly, in this paper we assume complete flexibility on the amount of each commodity (up to the vehicle capacity). In this way, our study provides an upper bound on the value of such flexible vehicles in terms of transportation costs.

The third category of applications concerns the evaluation of the opportunity for carriers to collaborate. Each carrier is associated with a commodity and serves a set of customers with its fleet of vehicles. Carriers want to know if it is beneficial to invest in a collaboration project where they would share customers and vehicles. This requires comparing a situation where commodities are handled separately (that is, each carrier handles its own customers) with one where multiple commodities are served jointly (that is, carriers share customers and vehicles). Based on a case study, it is estimated that the cost reduction achieved through cooperation ranges between 5% and 15% and can be higher (Cruijssen and Salomon 2004). The carrier collaboration problem has received an increasing amount of attention in recent years, and several different aspects have been investigated. We mention here some of the recent contributions to this area. A broad introduction to various forms of cooperation in logistics can be found in Cruijssen, Cools, and Dullaert (2007). When carriers collaborate, they have to determine which customers to share or exchange, and they must also decide how to share the resulting profit. The profit-sharing problem, together with the request allocation problem in a pickup and delivery context, is considered in Krajewska et al. (2008). A significant amount of work has been done for the truckload carrier collaboration problem. In Özener, Ergun, and Savelsbergh (2011) a good overview of the research in the area is provided and a lane exchange mechanism proposed. Our study examines the impact on transportation costs from collaboration both from a worst case and computational perspective without a truckload assumption.

In this paper, we consider a routing problem where multiple commodities are distributed to a set of customers with capacitated vehicles. Each vehicle starts from a depot, visits a set of customers, and returns to the depot at the end of the tour. Any customer may request any set of commodities. We assume that a vehicle that carries multiple commodities is totally flexible, in that it can carry any amount of any commodity, provided the constraint on the vehicle capacity is satisfied. We examine different ways of handling the commodities, depending on whether a vehicle is dedicated to a commodity or is flexible and depending on whether the demand of a customer may be satisfied by one or several vehicles. If the demand needs to be satisfied by more than one vehicle, we examine when the delivery to a customer of an individual commodity can be split among vehicles. Each situation gives rise to a different optimization problem. Whereas most of these problems are known, one is new. We compare the problems from a worst case perspective in terms of transportation cost and show in particular that, although it is intuitive that it may be highly beneficial to use flexible vehicles, there are situations where it is beneficial to use dedicated vehicles. We complement the worst case analysis with a computational study and test different classes of instances generated from benchmark instances.

The paper is organized as follows. In §2, we introduce the different problems we consider and the notation we use. In §3, we analyze the savings that can be achieved by serving all commodities with the same vehicle with respect to using dedicated vehicles for each commodity and vice versa. In §4, we study the impact of splitting the demand of a customer over different vehicles. A mathematical formulation for the new problem is presented in §5. Section 6 presents the solution algorithms we adopted to solve the problems, and in §7 the computational study is presented. Finally, we draw some conclusions in §8.

2. Definitions and Notation

We consider a distribution network represented by a graph where the vertex set V is composed of vertex 0, which is the depot, and vertices $1, \ldots, n$, which are the customers. We will indicate by C the set of customers. The graph is undirected and we denote by c_{ij} the distance or cost of the edge that connects vertices i and j. We assume the triangle inequality holds. An unlimited fleet of vehicles of capacity Q is available. A vehicle may carry a single commodity or all commodities, depending on the problem we consider. Commodities $1, \ldots, m$ have to be distributed

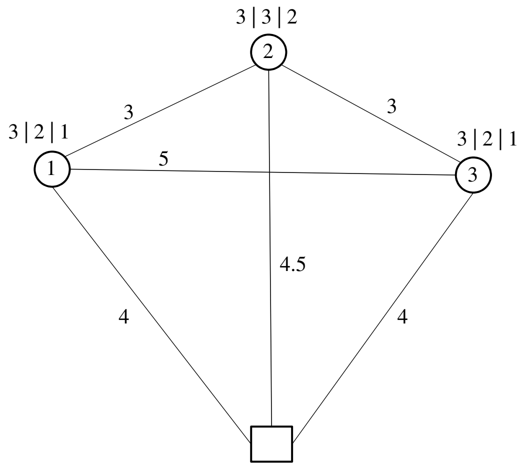

Figure 1 **Example Instance**

from the depot to the customers. We indicate by M the set of commodities. Let us denote by d_{ib} the demand of commodity b to be delivered to customer i. Let $t_b = \sum_i d_{ib}$ be the total demand of commodity b and let $d_i = \sum_b d_{ib}$ be the total demand of customer i. We assume that the total demand of each customer i is not greater than the vehicle capacity.

We use an example to highlight how the different restrictions change the optimal solution. The example instance is in Figure 1, where a square indicates the location of the depot and three customers are represented by circles. The number of commodities m is three. The number on each edge represents the cost to travel on the edge, and the numbers above each customer reflect the demand for each of the three commodities. For example, customer 1 demands three units of commodity one, two units of commodity two, and one unit of commodity three. Each vehicle has a capacity of 10 units.

In the *separate routing* problem (C-VRP), a specific set of vehicles is dedicated to each commodity, and any commodity is delivered to any customer by one visit of a vehicle only. In this case, a classical VRP has to be solved for each commodity (see Christofides, Mingozzi, and Toth 1979; Toth and Vigo 2002; Golden, Raghavan, and Wasil 2008). A customer will receive as many visits as the number of commodities requested. The solution to this problem is found in Figure 2(a). Three vehicles are needed, one for each commodity, and each vehicle visits all customers. Each trip costs 14, for a total cost of 42.

The *mixed routing* problem is one where any vehicle can deliver any set of commodities. No customer can be visited more than once; that is, if a customer requests one or more commodities, all the requested commodities are carried to the customer by one vehicle in one visit. This problem corresponds to a single classic VRP, and thus we will refer to this problem simply as VRP. The solution to this problem is in Figure 2(b). Three vehicles are needed, one for each customer. Each trip to customers 1 and 3 costs 8, and the trip to customer 2 costs 9, for a total cost of 25.

We also consider the problem where any vehicle can deliver any set of commodities and split deliveries of a commodity are allowed, and the same problem where split deliveries of a commodity are not allowed. The *split delivery mixed routing* is the problem where any vehicle can deliver any set of commodities and split deliveries are allowed; that is, the demand of a customer, requesting one or several commodities, can be served by one or several vehicles. The commodities can be split in any possible way. A customer can be visited several times if beneficial, even if he requests only one commodity. This problem corresponds to the split-delivery VRP (see Archetti and Speranza 2012 for a survey), and thus we will refer to this problem simply as SDVRP. The use of split deliveries has a big impact for our example problem, as shown in Figure 2(c). Now only two vehicles are needed to serve all of the demand. The first vehicle can use its capacity by serving all of customer 1's demand (six units) and four units of customer 2's demand. Customer 2 receives all of commodity one from the first vehicle, but the delivery of commodity two to customer 2 is split where only one unit is delivered by the first vehicle. This trip has a cost of 11.5. The second vehicle delivers the rest of customer 2's demand (four units) and all of customer 3's demand (six units). This trip also costs 11.5, for a total cost of 23 and use of only two vehicles.

Finally, we consider the problem where the vehicles are flexible and can deliver any set of commodities, and multiple visits of a customer are allowed only if the customer requests multiple commodities. When a commodity is delivered to a customer, the entire amount requested by the customer is carried. If the customer is visited more than once, the different vehicles will deliver different sets of commodities. We call this problem the *commodity-constrained split-delivery mixed routing* problem (C-SDVRP). Splitting the demand of a customer for different commodities on different vehicles is more natural and likely more acceptable to customers than splitting the delivery of an individual commodity. In terms of our example, the solution to this variant is found in Figure 2(d). The deliveries can still be accomplished by two vehicles, but the delivery of commodity two to customer 2 cannot be split among vehicles. Thus, vehicle one satisfies all of customer 1's demand (six units) and fills the rest of the truck with four units from customer 3. This can be accomplished by satisfying customer 3's complete demand for commodity one and three. This trip costs 13. Now vehicle two satisfies all of customer 2's demand (eight units) and the rest of customer 3's demand (two units), which represents its complete demand for commodity two.

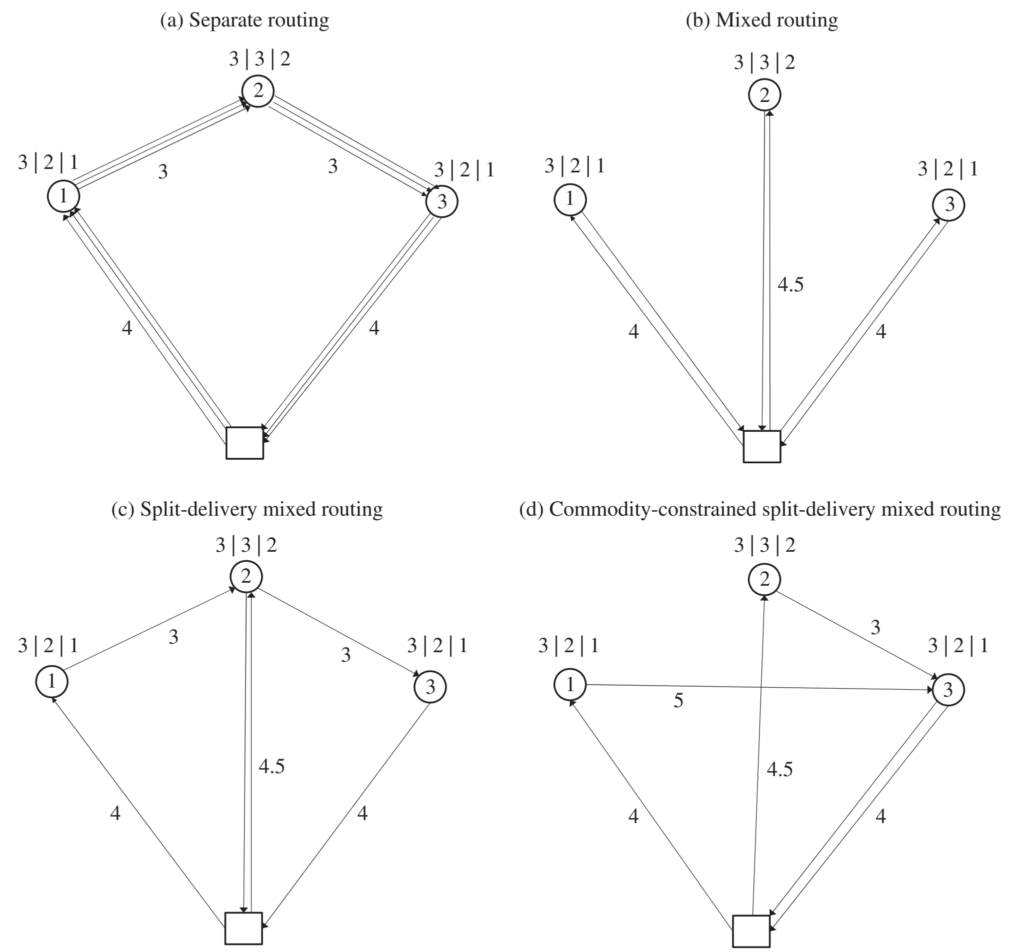

Figure 2 **Example Solutions**

Vehicle two's trip cost is 11.5. The total cost for the two trips is 24.5. The C-SDVRP is a new problem for which we will present an optimization model and a heuristic later in this paper.

In the following, we denote by $z^*(C\text{-}VRP)$, $z^*(VRP)$, $z^*(SDVRP)$, and $z^*(C\text{-}SDVRP)$ the cost optimum of the problems we study.

3. Separate vs. Mixed Routing

In this section, we study the savings that can be achieved by the VRP with respect to the C-VRP. We will evaluate the maximum possible savings through a worst case analysis. Although it is intuitive that the availability of flexible vehicles may be beneficial and may reduce the routing costs, it is interesting to understand the maximum amount of the savings. Moreover, we will show that there exist situations where the C-VRP may reduce the routing costs with respect to the VRP. This is a quite counterintuitive result. In the C-VRP and VRP, no split delivery is allowed; that is, either each commodity required by a customer is delivered by a single vehicle (in the C-VRP case) or each customer is visited by a single vehicle (in the VRP case). This is the simplest comparison to evaluate from the organizational point of view.

Let us start by considering a simple case. If there is only one customer, one vehicle is necessary and sufficient to serve all of the demand in the VRP, whereas in the C-VRP as many vehicles as the number of commodities are needed. For the sake of simplicity, let us denote by c the distance of the customer from the depot.

Theorem 1. *In the case of one customer and m commodities, $z^*(C\text{-}VRP) = z^*(VRP) + 2c(m-1)$.*

Proof. The claim follows from $z^*(C\text{-}VRP) = 2c\sum_b \lceil d_{1b}/Q \rceil = 2cm$ in the case where a customer receives a delivery for each commodity. Moreover, $z^*(VRP) = 2c(\lceil (\sum_b d_{1b})/Q \rceil) = 2c$ since we assumed $d_i \le Q$ for any customer i. Thus, $z^*(C\text{-}VRP) = 2c + 2c(m-1) = z^*(VRP) + 2c(m-1)$. $\square$

Corollary 1. *In the case of one customer and m commodities, $z^*(C\text{-}VRP)/z^*(VRP) = m$.*

Proof. As $z^*(C\text{-}VRP) = z^*(VRP) + 2c(m-1)$, then $z^*(C\text{-}VRP)/z^*(VRP) = 1 + 2c(m-1)/z^*(VRP) = 1 + 2c(m-1)/(2c) = m$. □

Now we build on these ideas as we consider the general case of n customers and analyze the maximum possible savings that can be achieved by the VRP with respect to the C-VRP.

Theorem 2. *In the case of n customers, with $n > 1$, and m commodities, $z^*(C\text{-}VRP)/z^*(VRP) \le m$ and the bound is tight.*

Proof. Consider an optimal solution to the VRP with optimum value $z^*(VRP)$. A feasible solution to the C-VRP can be obtained from this solution as follows. Take any vehicle used in the optimal solution to the VRP and use it to serve only the subset of customers served by that vehicle in the optimal solution to the VRP that have a demand of the first commodity, and serve the first commodity only to those customers. Obviously, the capacity constraint is satisfied as each vehicle serves less demand with respect to the mixed case. Also, because of the triangle inequality, the cost to serve the first commodity is not greater than $z^*(VRP)$. Having an unlimited fleet, we can repeat the procedure for the second commodity and for any commodity. Let $z(C\text{-}VRP)$ be the cost of this feasible solution to the C-VRP. Trivially, $z(C\text{-}VRP) \le mz^*(VRP)$. As $z^*(C\text{-}VRP) \le z(C\text{-}VRP)$, the claim follows.

To show that the bound is tight, consider the instance in Figure 3. The depot is located in the center of a circle of radius c. There are n customers spread out along the circle at a distance of $2c$ apart. There are m commodities and each customer requires each commodity with $d_{ib} = Q/m$. In the C-VRP, a single vehicle can handle each commodity by visiting all customers. The cost to serve each commodity is c (from the depot) $+\,2c(n-1)$ (visiting all customers) $+\,c$ (return to depot) $= 2cn$, and thus the optimal solution has value $z^*(C\text{-}VRP) = 2cnm$. In the VRP, a vehicle serves a single customer completely. The cost to serve each customer is the cost of a roundtrip from the depot ($2c$), and thus the cost of the optimal solution to the VRP is $2cn$. Thus, $z^*(C\text{-}VRP)/z^*(VRP) = m$. □

Thus, the C-VRP may cause an increase of the routing cost of up to m times the cost of the VRP. Surprisingly, however, there are situations where the C-VRP is more effective than the VRP.

Theorem 3. *$z^*(VRP)/z^*(C\text{-}VRP) \le 2$ and the bound is tight.*

Proof. Consider an instance of the problem with any number m of commodities. From Archetti, Savelsbergh, and Speranza (2006), we know that $z^*(VRP) \le 2z^*(SDVRP)$. On the other hand, we have $z^*(C\text{-}VRP) \ge z^*(SDVRP)$. Therefore, $z^*(VRP) \le 2z^*(C\text{-}VRP)$.

To show that the bound is tight consider the instance in Figure 4. The depot is located in the center of a circle of radius 1. There are k customers spread out along the circle at a distance ϵ apart. Furthermore, let there be k additional customers on a circle of radius $1+\epsilon$, perfectly aligned (along the radius) with the other k customers. The number of commodities is $m = 2$ and we assume that $Q \ge 2k$. Each customer in the inner circle has demand $Q/2-1$ for commodity 1 and demand 2 for commodity 2, whereas each customer in the outer circle has demand equal to $Q/2+1$ for commodity 1 and 0 for commodity 2. An optimal solution to the VRP requires visiting all customers with out-and-back tours, which results in a cost of $4k + 2k\epsilon$. In contrast, a feasible solution to the C-VRP visits two customers along the radius together, delivering $Q/2+1$ units of commodity 1 to the farthest customer and $Q/2-1$ units of commodity 1 to the closest customer. Commodity 2 is delivered to all customers on the inner circle with a single route. This results in a cost of $2k + 2k\epsilon + 2 + (k-1)\epsilon$. Therefore, the ratio $z^*(VRP)/z^*(C\text{-}VRP)$ is greater than or equal to $(4k + 2k\epsilon)/(2k + 2k\epsilon + 2 + (k-1)\epsilon)$. For k going to infinity (and ϵ going to 0) this ratio tends to 2. □

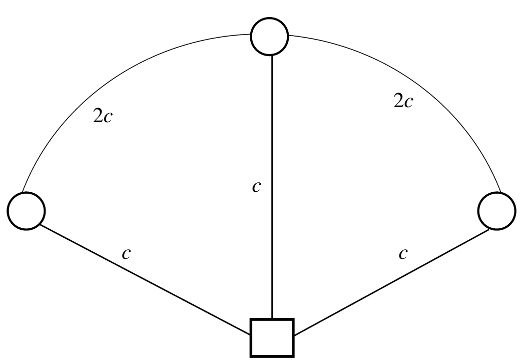

Figure 3 Tight Example for Theorem 2

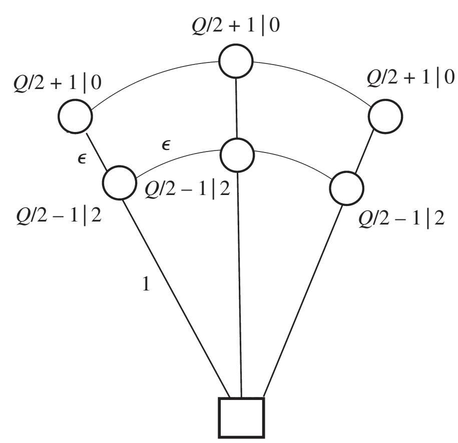

Figure 4 Tight Example for Theorem 3

Thus, although in general we expect the VRP to be more effective than the C-VRP, there is no strict dominance between them.

4. The Impact of Split Deliveries

If the commodities requested by a single customer can be split over different trips, it may create some savings with respect to $z^*(VRP)$ in terms of delivery cost. This may be beneficial even when the total delivery quantity to a customer is less than the truck capacity Q for the same reason that the split delivery routing problem can yield cheaper solutions than the standard VRP. In Archetti, Savelsbergh, and Speranza (2006), it is shown that $z^*(VRP) \le 2z^*(SDVRP)$ and that the bound is tight.

When only the different commodities can be split up but not the delivery quantity of a specific commodity to a customer, some of the gains of the SDVRP with respect to the VRP may be achieved, but not necessarily all. However, the following result shows that, at least from a worst-case perspective, the same gain can be achieved.

Theorem 4. *$z^*(VRP)/z^*(C\text{-}SDVRP) \le 2$ and the bound is tight.*

Proof. As

$$\frac{z^*(VRP)}{z^*(SDVRP)} \le 2 \quad \text{and} \quad z^*(C\text{-}SDVRP) \ge z^*(SDVRP),$$

it follows that

$$\frac{z^*(VRP)}{z^*(C\text{-}SDVRP)} \le \frac{z^*(VRP)}{z^*(SDVRP)} \le 2.$$

The tightness of the bound follows from the same instance used in the proof of Theorem 3. □

Theorem 5. *$z^*(C\text{-}SDVRP)/z^*(SDVRP) \le 2$ and the bound is tight.*

Proof. As

$$\frac{z^*(VRP)}{z^*(SDVRP)} \le 2 \quad \text{and} \quad z^*(C\text{-}SDVRP) \le z^*(VRP),$$

it follows that

$$\frac{z^*(C\text{-}SDVRP)}{z^*(SDVRP)} \le \frac{z^*(VRP)}{z^*(SDVRP)} \le 2.$$

The tightness of the bound follows from an instance adapted from Archetti, Savelsbergh, and Speranza (2006) to show that $z^*(VRP)/z^*(SDVRP) \le 2$. As illustrated in Figure 5, the depot is located in the center of a circle of radius 1. There are k customers spread out along the circle at a distance ϵ apart. Furthermore, let there be k additional customers on a circle of radius $1+\epsilon$, aligned along the radius with the other k customers. We assume that $Q \ge 2k$, and the number of

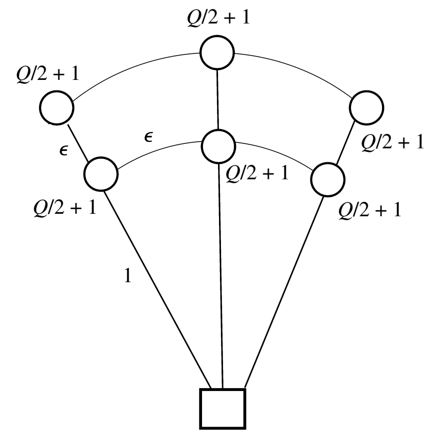

Figure 5 **Tight Example for Theorem 5**

commodities is k. Each customer demands a different commodity, each at a quantity of $Q/2+1$. An optimal solution to the C-SDVRP will have to visit all customers with out-and-back tours, which results in a cost of $4k+2k\epsilon$. Yet a feasible solution to the SDVRP visits two customers along the radius together, delivering $Q/2+1$ units to the farthest customer and $Q/2-1$ units to the closest customer. The remaining demand is delivered to all customers on the inner circle with a single route. This results in a cost of $2k+2k\epsilon+2+(k-1)\epsilon$. Therefore, the ratio $z^*(C\text{-}SDVRP)/z^*(SDVRP)$ is greater than or equal to $(4k+2k\epsilon)/(2k+2k\epsilon+2+(k-1)\epsilon)$. For k going to infinity (and ϵ going to 0) this ratio tends to 2. □

4.1. The k-Split Cycle Property

Consider a set $C=\{i_1, i_2, \ldots, i_k\}$ of customers and suppose that there exist k routes $r_1, \ldots, r_k$, $k \ge 2$, such that r_w contains customers i_w and i_{w+1}, $w=1,\ldots,k-1$, and r_k contains customers i_1 and i_k. Such a configuration is called a *k-split cycle*.

Dror and Trudeau (1989) prove that, if the distances satisfy the triangle inequalities, there always exists an optimal solution to the SDVRP with no k-split cycle. This means that there always exists an optimal solution where no two routes have more than one split delivery in common. This does not hold anymore for the C-SDVRP, as shown by the following example.

Example 1. Consider an instance with four customers and $m=2$ commodities, as in Figure 6. Customers 1 and 2 require seven units of commodity 1 and zero units of commodity 2. Customer 3 requires one unit of commodity 1 and one unit of commodity 2. Customer 4 requires two units of commodity 1 and two units of commodity 2. The capacity of the vehicles is $Q=10$. The distances between all customers and the depot are all equal to c, $c>1$, $c_{34}=\epsilon$, $c_{13}=c_{14}=c_{23}=c_{24}=1$, and $c_{12}=2$. The optimal solution is to make two routes as follows. The first route delivers seven units of commodity 1 to customer 1, one unit

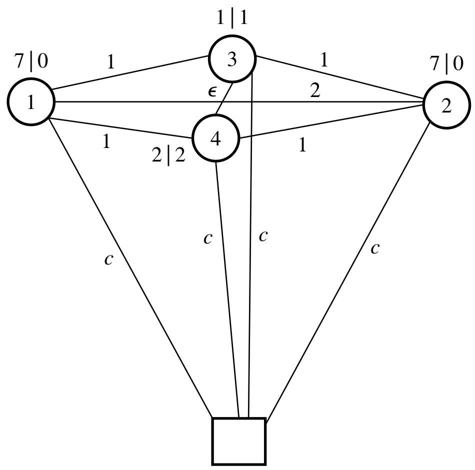

Figure 6 **Example Showing That the k-Split Property Does Not Hold for the C-SDVRP**

of commodity 1 to customer 3, and two units of commodity 1 to customer 4. The second route delivers seven units of commodity 1 to customer 2, one unit of commodity 2 to customer 3, and two units of commodity 2 to customer 4. The two routes share both customers 3 and 4. The cost of the optimal solution is $4c+2+2\epsilon$, and there is no other optimal solution that does not share the two customers.

5. A Mathematical Formulation for the C-SDVRP

We present in this section a mixed-integer linear programming model for the C-SDVRP. It will be used to find the optimal solutions presented in §7. Let K be the set of vehicles. We fix $|K|=n$ where n is a valid upper bound on the number of vehicles needed. Moreover, let $M_i=\{b\in M \mid d_{ib}>0\}$. We define the following decision variables:

- x_{ijk}: Binary variable equal to 1 if edge (i,j) is traversed by vehicle k, $i,j\in V$, $j<i$, $k\in K$
- z_{ik}: Binary variable equal to 1 if vertex i is visited by vehicle k, $i\in V$, $k\in K$
- y_{ikb}: Binary variable equal to 1 if commodity b of customer i is delivered by vehicle k, $i\in C$, $k\in K$, $b\in M_i$.

The mathematical formulation of the C-SDVRP is the following:

$$\min \sum_{i\in V}\sum_{j\in V,\, j<i}\sum_{k\in K} c_{ij}x_{ijk}, \tag{1}$$

$$\sum_{k\in K} y_{ikb}=1 \quad i\in C\ b\in M_i, \tag{2}$$

$$\sum_{j\in V,\, j<i} x_{ijk}=\sum_{j\in V,\, j>i} x_{jik}=2z_{ik} \quad i\in V\ k\in K, \tag{3}$$

$$y_{ikb}\le z_{ik} \quad i\in C\ k\in K\ b\in M_i, \tag{4}$$

$$\sum_{i\in C}\sum_{b\in M_i} d_{ib}y_{ikb}\le Qz_{0k} \quad k\in K, \tag{5}$$

$$\sum_{i\in S}\sum_{j\in S,\, j<i} x_{ijk}\le \sum_{i\in S} z_{ik}-z_{tk} \quad S\subseteq C\ t\in S\ k\in K, \tag{6}$$

$$x_{ijk}\in\{0,1\} \quad i,j\in V,\ j<i\ k\in K, \tag{7}$$

$$z_{ik}\in\{0,1\} \quad i\in V\ k\in K, \tag{8}$$

$$y_{ikb}\in\{0,1\} \quad i\in C\ k\in K\ b\in M_i. \tag{9}$$

The objective function (1) aims to minimize the total transportation cost. Constraint (2) requires that each commodity requested by each customer be delivered by a single vehicle. Constraint (3) provides degree constraints, and constraint (4) allows a vehicle to deliver a commodity to a customer only if it visits the customer. Capacity constraints are imposed by (5), whereas (6) establishes subtour elimination constraints. Constraints (7)–(9) are variable definitions.

6. Solution Algorithms

To create the computational study presented in the next section, we used different heuristic algorithms for the solution to the VRP and the SDVRP. In particular, for the solution to the SDVRP we adopted the optimization-based algorithm presented in Archetti, Savelsbergh, and Speranza (2008), which is one of the best performing heuristic solution procedures for the SDVRP. The reader is referred to Archetti, Savelsbergh, and Speranza (2008) for details on the algorithm. For the solution to the VRP, we used the open-access injection-ejection algorithm from the COIN-OR library (available at http://www.coin-or.org/projects/VRPH.xml). This algorithm is known to work well for the VRP.

To solve the C-SDVRP, we duplicated the node associated with each customer for the number of times that equals the number of commodities requested by the customer. We associated with each duplicated node the demand of the customer for the corresponding commodity. We then used the same VRP heuristic algorithm mentioned above.

To evaluate the quality of solutions given by the heuristic algorithms, we implemented exact solution approaches for all problems that are used to solve small instances. To solve the VRP and C-VRP to optimality, we implemented the classical three-index formulation for the VRP (for details, see Toth and Vigo 2002). Subtour elimination constraints are added when they are found to be violated using the max flow-min-cut algorithm presented in Padberg and Rinaldi (1991). Capacity cut constraints are also added after they are found to be violated through the separation algorithms presented in Ralphs et al. (2003). For the formulation of the capacity cuts, the reader is referred to Toth and Vigo (2002).

For the solution of the SDVRP, we adopted the branch-and-cut algorithm proposed in Archetti, Bianchessi, and Speranza (2014). Fractional capacity cuts are inserted to enhance the formulation.

Finally, to obtain optimal solutions to the C-SDVRP, we implemented a branch-and-cut algorithm based on formulation (1)–(9), where subtour elimination constraints are added after they are found to be violated through the algorithm proposed in Padberg and Rinaldi (1991). To enhance the formulation, we added the following capacity cuts:

$$Q\sum_{k\in K}\sum_{(i,j)\in\delta(S),\, j<i} x_{ijk} \geq 2\sum_{i\in S}\sum_{b\in M_i} d_{ib} \quad S\subseteq C, \qquad (10)$$

where $\delta(S)=\{(i,j)\in E \mid (i\in S, j\notin S) \text{ or } (i\notin S, j\in S)\}$. The separation of the capacity cuts is performed through the algorithms presented in Ralphs et al. (2003).

7. A Computational Study

In this section, we present a computational study to compare the solutions of the C-VRP, VRP, C-SDVRP, and SDVRP over a variety of instances. Our aim is to understand the characteristics of the instances that make the different problems yield different relative results. If we can identify what types of instances make flexible vehicles or split commodities yield particular low-cost solutions, for example, it may be useful for managers in deciding when these types of delivery practices should be adopted.

We now present the set of instances we generated (§7.1), followed by the computational results (§7.2). All tests were run on a 64-bit Windows machine, with an Intel Xeon processor W3680, 3.33 GHz, and 12 GB of RAM.

7.1. Test Instances

We created a set of instances based on the R101 and C101 Solomon instances for the VRP. These data sets represent examples of when customers are distributed randomly and in a clustered manner. We generated three sets of instance sizes: small, midsize, and large. We first describe the settings we used to generate large instances. Small and medium instances are generated using a subset of the settings used to generate large instances.

7.1.1. Large Instances. For the large instances, the number of vertices (101) and the location of each vertex of the original instances (R101 and C101) is held constant, but we vary the remaining problem instance data. In particular, for both R101 and C101, we generated several instances according to the following characteristics:

- Number of commodities: Two and three.
- Probability that a customer requires a commodity: We considered two cases. Probability set to 100% means that each customer requires a delivery of each commodity. Probability set to 60% means that the probability that a customer demands a delivery of a particular commodity is 60%.
- Demand range: These data represent the interval in which we generate the value of the demand of each customer for each commodity. We considered two intervals: [1; 100] and [40; 60].
- Vehicle capacity: Denoted as $d_{\max}=\max_{i\in C}\sum_{j=1}^{m} d_{ij}$, the vehicle capacity is generated as $Q=\alpha d_{\max}$. We considered four values of α: 1.1, 1.5, 2, and 2.5.

For each of the 64 combinations of characteristics, we randomly generate five different instances, creating 320 large instances in total.

7.1.2. Midsize Instances. To generate midsize instances, we varied the number of customers n to see the influence on solution values. In particular, we tested $n=20$, 40, 60, and 80. We did not try every combination of characteristics, as for the large instances, but we fixed the number of commodities to three, the demand range to [1; 100], and the value of α to 1.5. For each combination of value of n, type of instance (R101 and C101), and value of the probability (60% and 100%), we generated five different instances by randomly choosing customers from the original instance, applying the probability to decide which commodities each customer demands (for the 60% probability), and generating the demands for each commodity in the interval [1; 100]. The total number of midsize instances is 80.

7.1.3. Small Instances. We also created small instances so that we could solve them to optimality and evaluate the performance of the heuristic algorithms. Small instances are defined here as having 15 customers, as preliminary experiments showed that instances of the C-SDVRP with more than 15 customers could not be solved to optimality using the branch-and-cut procedure described in §6. As mentioned, 64 combinations of characteristics were considered for the large instances, and five different instances were generated for each combination. Here, we take the first instance for each set of five random instances, and limit the data set to the first 15 customers. Thus, the total number of small instances evaluated is 64.

7.2. Computational Results

First, we present the results for the small instances to prove the reliability of the heuristic algorithms (§7.2.1). We then present the results on large (§7.2.2) and midsize (§7.2.3) instances, respectively.

7.2.1. Results for Small Instances. Small instances are solved by the branch-and-cut algorithms described in §6 with a maximum time limit of 30 minutes. We used CPLEX 12.5 as the solver. All 64 instances were

Table 1 Number of Small Instances Solved to Optimality for the C-SDVRP

	$m=2$	$m=3$
$\alpha=1.1$	1/8	0/8
$\alpha=1.5$	2/8	1/8
$\alpha=2$	5/8	4/8
$\alpha=2.5$	6/8	6/8

solved to optimality for the SDVRP, VRP, and C-VRP. The branch-and-cut algorithm for the C-SDVRP was able to solve 25 of the 64 instances. For the instances that are not solved, both the best feasible solution found by CPLEX and the lower bound are of very poor quality; thus, they cannot be used for further evaluation. In Table 1, we report the number of instances solved to optimality for each combination of values of m and α for the C-SDVRP. Table 1 clearly shows that instances with a larger value of α are easier to solve. This is because fewer vehicles are needed to serve all customers when α increases.

Table 2 compares the results for the 25 instances that have been solved to optimality for all problems. In particular, we report the average and the maximum gap of the C-VRP solution relative to the solution of the VRP, C-SDVRP, or SDVRP, respectively, in terms of cost and number of vehicles. We compare this with the C-VRP, as it is typically the problem with the largest cost and number of vehicles required. Cost gaps are calculated as follows:

$$100\frac{z(C\text{-}VRP)-z(P)}{z(C\text{-}VRP)},$$

where P corresponds to the VRP, C-SDVRP, or SDVRP, respectively. The calculation of the gaps in terms of number of vehicles is similar. Table 2 clearly shows the advantage of combining the distribution of the different commodities. Cost savings relative to the C-VRP average between 33% and 35% with a maximum savings over 56%. The solutions of the VRP, C-SDVRP, and SDVRP are not very different in terms of solution costs, but the C-SDVRP and SDVRP use fewer vehicles, on average, than the VRP solutions.

Table 3 looks at the performance of the heuristic algorithms. In particular, we report the number of optimal solutions found by the heuristic algorithms, as well as the average and the maximum error with respect to the optimal value. When reporting the number of optimal solutions for the heuristics, we give in parentheses the total number of instances solved to optimality by the branch-and-cut approaches. The results indicate that all of the heuristic algorithms provide high-quality solutions, and thus we can trust their results when using them to solve mid-size and large instances.

Table 2 Results on Small Instances Solved to Optimality

	Solution cost		
	VRP	C-SDVRP	SDVRP
Average % gap	33.53	34.17	34.53
Maximum % gap	56.27	56.82	56.85
	Number of routes		
	VRP	C-SDVRP	SDVRP
Average % gap	11.73	15.24	16.04
Maximum % gap	33.33	33.33	33.33

Table 3 Heuristic Solution Quality

	C-VRP	VRP	C-SDVRP	SDVRP
No. optimal	60 (64)	64 (64)	23 (25)	56 (64)
Average % error	0.02	0.00	0.05	0.13
Maximum % error	0.54	0.00	0.65	2.79

7.2.2. Results for Large Instances. In Tables 4–7, we report a summary of the results we obtained with large instances. Midsize and large instances have been solved heuristically. In each table, we report the percentage gap between the solution of the problem reported in the corresponding column and the solution of the C-VRP. We report both the average and the maximum gap. Instances are grouped based on the characteristic being examined, so each row represents results gathered across half of the instances (160), except for the last four rows, which refer to the values of α. Here, each row represents results gathered across 80 instances. Table 4 reports the percentage difference in solution cost, and Table 5 reports the percentage difference in the number of vehicles. The full results on large instances are available in the online appendix (available as supplemental material at http://dx.doi.org/10.1287/trsc.2014.0528) and at https://sites.google.com/site/orbrescia/instances.

First, we will examine the results in Table 4. The most obvious result is the dramatic cost increase, both in

Table 4 Solution Cost: Percentage Gap with Respect to the C-VRP

	Average			Maximum		
	VRP	C-SDVRP	SDVRP	VRP	C-SDVRP	SDVRP
C101	14.67	19.22	19.83	32.25	33.29	33.36
R101	15.54	19.60	19.95	38.11	38.75	38.75
2 commodities	11.10	14.66	15.58	25.30	26.46	30.26
3 commodities	19.10	24.16	24.20	38.11	38.75	38.75
Probability 0.6	16.24	17.56	18.10	35.61	35.61	35.61
Probability 1.0	13.96	21.26	21.68	38.11	38.75	38.75
Demand range [1; 100]	20.79	22.38	22.44	38.11	38.75	38.75
Demand range [40; 60]	9.42	16.44	17.34	27.76	28.40	30.26
$\alpha=1.1$	11.76	15.78	17.03	25.43	29.57	30.26
$\alpha=1.5$	9.47	18.88	19.25	32.25	33.29	33.36
$\alpha=2.0$	18.07	19.95	20.15	32.63	33.67	33.67
$\alpha=2.5$	21.11	23.04	23.13	38.11	38.75	38.75

Table 5 Number of Routes: Percentage Gap with Respect to the C-VRP

	Average			Maximum		
	VRP	C-SDVRP	SDVRP	VRP	C-SDVRP	SDVRP
C101	−5.03	1.87	3.15	11.76	11.76	19.39
R101	−5.28	1.87	2.79	11.76	11.76	14.93
2 commodities	−4.11	1.54	3.62	7.69	8.00	19.39
3 commodities	−6.20	2.21	2.32	11.76	11.76	11.76
Probability 0.6	1.00	2.92	3.99	11.76	11.76	17.46
Probability 1.0	−11.31	0.82	1.95	4.35	6.67	19.39
Demand range [1; 100]	−0.07	2.22	2.30	11.76	11.76	11.76
Demand range [40; 60]	−10.24	1.52	3.64	8.00	8.00	19.39
$\alpha = 1.1$	−4.99	1.26	3.93	3.13	5.97	19.39
$\alpha = 1.5$	−13.44	1.15	1.97	4.17	8.00	9.30
$\alpha = 2.0$	−1.49	1.90	2.48	5.56	6.67	6.67
$\alpha = 2.5$	−0.69	3.18	3.50	11.76	11.76	11.76

terms of average and maximum values, associated with the C-VRP versus all of the remaining problems. In fact, we always get the relative performance we expect with the average percentage gap for SDVRP > C-SDVRP > VRP across all of the instance characteristics. This shows computationally the value of combining commodities on trucks and allowing deliveries to be split. We see slightly higher gaps with random data sets, with the most noticeable difference for the VRP. This indicates that it may be slightly more important to combine commodities when customers are distributed randomly. We can observe significant increases in gaps for all problems with the increase in the number of commodities from two to three. This is not surprising, as more commodities yield more opportunities for improvements through combined deliveries. The results with the different probabilities are interesting. The VRP with probability 1 has a lower gap than with probability 0.6, since many customers now receive total delivery quantities that do not combine well with other customers. With probability 0.6, customers receive different numbers of commodities and thus different size deliveries, allowing for nice combinations that do not exist with probability 1. These same issues do not occur when deliveries are allowed to be split. When considering the range of delivery size of each commodity, we see that a wider range of delivery sizes allows for better combinations than more restrictive, fairly large delivery sizes. Finally, we see interesting results with increasing sizes of α. For the C-SDVRP and SDVRP, we see increasing gaps as the vehicle capacity grows. This is not surprising, as the splitting of deliveries allows the larger-capacity vehicles to be used well. It is interesting that this pattern is not maintained for the VRP. This actually makes sense, as with $\alpha = 1.5$, few vehicles, especially in the probability 1 case, can make deliveries to two customers since deliveries may not be split.

Next, Table 5 reveals the percentage difference in the number of vehicles used relative to the C-VRP. Here the most interesting result is related to the negative average gaps of the VRP, meaning that, on average, the VRP uses a higher number of routes than the C-VRP in these situations. Note that, by combining this information with that in Table 4, we can conclude that the VRP often has routes with fewer customers than the others, including the C-VRP. Here we see that the SDVRP yields more substantial savings than the C-SDVRP in terms of the vehicles.

7.2.3. Results for Midsize Instances. In Tables 6 and 7 we report the results from the midsize instances in terms of solution cost and number of vehicles, respectively.

The two tables are organized in the following way. The first column reports the kind of instance (R101 versus C101 and probability 60% versus probability 100%). The next 12 columns are grouped in three sets of four columns each, where the first set refers to the VRP, the second to the C-SDVRP, and the last to the SDVRP. In all sets, each column refers to a given value of n. Each cell in the top (bottom) part of the table

Table 6 Results on Midsize Instances: Solution Cost

	VRP				C-SDVRP				SDVRP			
	$n=20$	$n=40$	$n=60$	$n=80$	$n=20$	$n=40$	$n=60$	$n=80$	$n=20$	$n=40$	$n=60$	$n=80$
					Average % gap with respect to the C-VRP							
C101—Probability 60%	34.18	32.45	25.65	22.68	36.46	33.91	27.30	23.43	37.30	34.42	27.56	23.61
C101—Probability 100%	38.16	25.90	19.26	16.85	40.17	29.18	22.68	19.44	41.55	29.72	23.01	19.46
R101—Probability 60%	30.59	30.03	25.55	23.56	33.43	30.54	26.36	24.40	33.71	30.84	26.82	24.55
R101—Probability 100%	29.60	28.36	25.50	24.31	33.70	31.36	27.64	26.44	34.50	31.79	27.85	26.58
					Maximum % gap with respect to the C-VRP							
C101—Probability 60%	40.11	35.86	28.22	25.40	44.34	36.23	29.41	25.76	44.34	36.98	30.71	25.94
C101—Probability 100%	39.70	30.04	24.00	18.91	42.63	32.53	26.22	20.87	43.58	32.53	26.22	20.87
R101—Probability 60%	34.34	32.03	28.72	25.21	36.46	33.12	29.62	25.81	36.46	33.12	29.83	26.24
R101—Probability 100%	36.21	30.69	28.97	27.37	39.10	33.83	30.23	28.86	39.76	33.83	30.26	29.28

Table 7 Results on Midsize Instances: Number of Routes

	VRP				C-SDVRP				SDVRP			
	$n=20$	$n=40$	$n=60$	$n=80$	$n=20$	$n=40$	$n=60$	$n=80$	$n=20$	$n=40$	$n=60$	$n=80$
					Average % gap with respect to the C-VRP							
C101—Probability 60%	7.22	3.08	4.67	3.38	10.08	6.15	5.85	5.20	12.58	6.15	6.96	6.20
C101—Probability 100%	7.27	−3.29	2.37	−0.54	11.09	4.34	6.19	3.65	13.09	3.52	6.19	3.65
R101—Probability 60%	4.72	3.08	1.05	2.82	10.08	6.15	3.50	3.69	12.58	6.15	6.91	3.69
R101—Probability 100%	5.64	−2.22	0.72	−0.48	9.45	4.46	5.36	2.52	11.45	5.46	6.13	3.09
					Maximum % gap with respect to the C-VRP							
C101—Probability 60%	12.50	15.38	6.67	8.00	14.29	15.38	6.67	8.70	14.29	15.38	11.11	8.70
C101—Probability 100%	18.18	0.00	4.17	3.23	18.18	5.88	8.33	6.45	18.18	5.88	8.33	6.45
R101—Probability 60%	12.50	15.38	5.26	5.00	14.29	15.38	6.67	8.70	14.29	15.38	11.11	8.70
R101—Probability 100%	18.18	0.00	3.85	3.23	18.18	5.88	7.69	6.45	18.18	10.00	11.54	6.45

reports the average (maximum) gap with respect to the solution given by the C-VRP.

The results on the solution cost show a fairly similar behavior among the VRP, C-SDVRP, and SDVRP. What we are interested in here is how the value of n changes the results. We can observe across both averages and maximums that the gaps with respect to the C-VRP decrease as the dimension increases. Even though all of the values of n show what appear to be significant improvements by combining commodities on a vehicle, it appears particularly important with low values of n. This makes sense as lower n values yield fewer good routing combinations without the option of split deliveries.

In terms of number of vehicles, again the results for the VRP are different than for the C-SDVRP and SDVRP. The gaps are lower, and in some cases the VRP uses a higher number of routes than the C-VRP.

8. Conclusions

In this paper, we studied different strategies of distributing a set of commodities to a number of customers. In particular, we examined the impact of using flexible vehicles, which can transport multiple commodities at one time, versus vehicles dedicated to a single commodity. We also examined the impact of making split deliveries to customers. We showed that these decisions can have a big impact in terms of costs and number of vehicles needed, both from a worst case analysis point and experimental results from a wide computational study.

Computationally we saw that:

- In practice, we often get the relative performance we expect with costs from SDVRP < C-SDVRP < VRP < C-VRP. This shows the value of combining commodities on trucks and allowing deliveries to be split.
- It may be slightly more important to combine commodities when customers are distributed randomly.
- Larger numbers of commodities yield more opportunities for improvements through combined deliveries. Thus, we can infer that collaboration may yield more benefits, as more parties participate in the collaboration.
- If all customers receive all commodities, the VRP has much less room for improvement relative to the C-VRP than when customers receive a varied number of commodities.
- If all customers receive deliveries of each commodity of a similar size, it allows for much less room for improvement relative to the C-VRP than when the range of delivery size of each commodity is more varied. This may indicate that flexible vehicles may be more beneficial with flexible, rather than standardized, compartments.
- As the vehicle capacity grows, we see increasing gaps for the C-SDVRP and SDVRP. This is because splitting deliveries allows the larger-capacity vehicles to be used well.
- As the vehicle capacity grows, the VRP does not always yield increasing improvements relative to the C-VRP. This indicates that vehicle capacity needs to be carefully considered when evaluating whether to deliver commodities separately or together.
- The SDVRP yields more substantial savings than the C-SDVRP in terms of the number of vehicles required. If vehicles are expensive, as they often are, this is an important reason to consider splitting the delivery of individual commodities.
- The cost gaps with respect to the C-VRP decrease as the number of customers increases. Thus, it may be more important to consider combined deliveries, collaborations, or flexible vehicles with smaller customer sets than larger ones to create better routes.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/trsc.2014.0528.

Acknowledgments

The authors wish to acknowledge the comments and suggestions of two anonymous reviewers who helped us improve a former version of this paper.

References

Archetti C, Speranza MG (2012) Vehicle routing problems with split deliveries. *Internat. Trans. Oper. Res.* 19:3–22.

Archetti C, Bianchessi N, Speranza MG (2014) Branch-and-cut algorithms for the split delivery vehicle routing problem. *Eur. J. Oper. Res.* 238:685–698.

Archetti C, Savelsbergh MWP, Speranza MG (2006) Worst-case analysis for split delivery vehicle routing problems. *Transportation Sci.* 40:226–234.

Archetti C, Savelsbergh MWP, Speranza MG (2008) An optimization-based heuristic for the split delivery vehicle routing problem. *Transportation Sci.* 42:22–31.

Christofides N, Mingozzi A, Toth P (1979) *The Vehicle Routing Problem* (John Wiley & Sons, Chichester, UK), 315–338.

Cruijssen F, Salomon M (2004) Empirical study: Order sharing between transportation companies may result in cost reductions between 5 to 15 percent. Technical report, CentER Discussion Paper 2004-80, Faculty of Economics and Business Administration, Tilburg University, Tilburg, Netherlands.

Cruijssen F, Cools M, Dullaert W (2007) Horizontal cooperation in logistics: Opportunities and impediments. *Transportation Res. Part E* 43:129–142.

Derigs U, Gottlieb J, Kalkoff J, Piesche M, Rothlauf F, Vogel U (2011) Vehicle routing with compartments: Applications, modelling and heuristics. *OR Spectrum* 33:885–914.

Dror M, Trudeau P (1989) Savings by split delivery routing. *Transportation Sci.* 23:141–145.

Golden B, Raghavan R, Wasil E (2008) *The Vehicle Routing Problem. Latest Advances and New Challenges*, Oper. Res./Comput. Sci. Interfaces Ser., Vol. 43 (Springer-Verlag, New York).

Krajewska MA, Kopfer H, Laporte G, Ropke S, Zaccour G (2008) Horizontal cooperation among freight carriers: Request allocation and profit sharing. *J. Oper. Res. Soc.* 59:1483–1491.

Özener OÖ, Ergun Ö, Savelsbergh M (2011) Lane-exchange mechanisms for truckload carrier collaboration. *Transportation Sci.* 45:1–17.

Padberg M, Rinaldi G (1991) A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems. *SIAM Rev.* 33:60–100.

Ralphs TK, Kopman L, Pulleyblank WR, Trotter LE (2003) On the capacitated vehicle routing problem. *Math. Programmming* 94:343–359.

Toth P, Vigo D (2002) *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Its Applications (SIAM, Philadelphia).