

A Branch-and-Price Algorithm for the Vehicle Routing Problem with Roaming Delivery Locations

Gizem Ozbaygin^{1,2}, Oya Karasan², Martin Savelsbergh¹, and Hande Yaman²

¹H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332

²Department of Industrial Engineering, Bilkent University, Ankara 06800, Turkey

Abstract

We study the vehicle routing problem with roaming delivery locations in which the goal is to find a least-cost set of delivery routes for a fleet of capacitated vehicles and in which a customer order has to be delivered to the trunk of the customer's car during the time that the car is parked at one of the locations in the (known) customer's travel itinerary. We formulate the problem as a set-partitioning problem and develop a branch-and-price algorithm for its solution. The algorithm can also be used for solving a more general variant in which a hybrid delivery strategy is considered that allows a delivery to either a customer's home or to the trunk of the customer's car. We evaluate the effectiveness of the many algorithmic features incorporated in the algorithm in an extensive computational study and analyze the benefits of these innovative delivery strategies. The computational results show that employing the hybrid delivery strategy results in average cost savings of nearly 20% for the instances in our test set.

Keywords: Vehicle routing, trunk delivery, roaming delivery locations, branch-and-price, resource-constrained shortest path

1 Introduction

A recent survey revealed that in 2016, for the first time, online purchases have surpassed in-store purchases (Farber, 2016). The business-to-consumer retail segment continues to grow year-over-year with an ever-increasing push towards online shopping. As an example, Amazon, one of the e-commerce giants received 398 orders per second and 34.4 million orders in total during its Prime Day event on July 15th, 2015 (Garcia, 2015). A year later, on Prime Day 2016, the number of orders increased by more than 60%, making it the biggest sales day in the history of the company (Gustafson, 2016). Due to the sheer

volume, the huge number of delivery locations, and the aggressive service levels promised, last-mile logistics is a huge challenge for Amazon. Amazon’s first-quarter shipping costs in 2016 hit \$3.27 billion, a 42% increase from the same period in 2015 (Solomon, 2016).

Not surprisingly, Amazon is seeking and exploring innovative ideas to improve the efficiency of last-mile delivery operations. Among these innovative ideas is trunk delivery, where the orders of customers are delivered to the trunk of their cars, and introduced by Amazon in collaboration with Audi and DHL in certain areas (Popken 2015, Geuss 2015, Audi 2015). Volvo too is offering the required technology and has launched an in-car delivery service, initially limited to Stockholm, Sweden (Volvo 2015, Volvo 2016). Recently, Smart has partnered with DHL to start providing trunk delivery service to customers in Germany who order merchandise from Amazon in September 2016 (Behrmann and Weiss, 2016).

Motivated by the interest in trunk delivery, we study the variant of the vehicle routing problem (VRP) in which each customer has an itinerary specifying one or more locations with corresponding time windows where the customer’s order can be delivered to the trunk of his/her car (the car will be parked at these locations during the given time windows). This problem is known as the vehicle routing problem with roaming delivery locations (VRPRDL) and was introduced recently by Reyes et al. (2016), who developed various construction and improvement heuristics for the problem.

The VRPRDL combines two well-studied problems, namely, the VRP with time windows (VRPTW) and the generalized VRP (GVRP). The VRPTW is the problem of determining an optimal set of delivery routes serving the demand of a set of customers within their respective time windows. There is a vast body of literature on the VRPTW and its variants (see, for example, Savelsbergh 1985, Solomon 1987, Desrochers et al. 1988, Desrochers et al. 1992, Mitrović-Minić et al. 2004, Dabia et al. 2013, Agra et al. 2013, Schneider et al. 2014, Taş et al. 2014, Koç et al. 2015, Ghilas et al. 2016). The GVRP was introduced by Ghiani and Improta (2000) and it is another generalization of the VRP in which the set of delivery locations is partitioned into clusters and exactly one location from each cluster has to be visited in a solution. Despite its relatively recent introduction, the GVRP has already attracted the attention of many researchers, in part because it has many real-life applications (Baldacci et al. 2010, Bektas et al. 2011, Kovacs et al. 2014, Afsar et al. 2014, Quttineh et al. 2015, Biesinger et al. 2016, Louati et al. 2016).

The integration of the features of these two problems leads to the generalized vehicle routing problem with time windows (GVRPTW). To the best of our knowledge, the first study on the GVRPTW is due to Moccia et al. (2012), who present an incremental tabu search algorithm to solve the problem. The VRPRDL can be seen as a special case of the GVRPTW in which the sets of delivery locations for the customers form the clusters. However, the time windows exhibit a special structure, as the time windows of the locations in a cluster are non-overlapping. Our computational experiments reveal that it is this special structure that allows us to solve large instances. We note, however, that our solution approach does not explicitly exploit the time window structure and can, thus, also

be used to solve instances of the GVRPTW.

The main contribution of our paper is that it presents an effective branch-and-price algorithm for the VRPRDL. Branch-and-price (Barnhart et al., 1998) has established itself as an effective solution methodology for a variety of vehicle routing and scheduling problems, e.g., the VRPTW (Desrochers et al., 1992), the VRP with stochastic demands (Christiansen and Lysgaard, 2007), the VRP with pickup and deliveries (Dell’Amico et al., 2006), and the VRP with split deliveries (Salani and Vacca, 2011). As far as we know, this is the first study in which an exact solution approach for the VRPRDL, or, more generally, for the GVRPTW, is developed. An extensive computational study shows that the branch-and-price algorithm is capable of solving instances of up to 120 customers. We also demonstrate that the algorithm can be used to solve instances of the variant in which a customer order can be delivered either to the customer’s home or to the customer’s car, although only instances of up to 60 customers. Reyes et al. (2016) call this problem variant the VRP with home and roaming delivery locations (VRPHRDL). Finally, we use our algorithm to analyze the cost savings that can be achieved when making optimal use of the option to deliver to the trunk of a car.

The remainder of this paper is organized as follows. Section 2 presents a mixed-integer programming formulation as well as a set-partitioning model for the VRPRDL, describes the pricing problem and how to solve it. Section 3 discusses the techniques employed in our branch-and-price algorithm to increase its efficiency and gives some implementation details. Section 4 explains how the VRPHRDL can be solved with our algorithm. Section 5 provides the results of an extensive computational study, demonstrating the efficacy of the branch-and-price algorithm and analyzing the benefits of trunk delivery. Section 6 concludes with some final remarks.

2 Problem definition and formulations

The VRPRDL is formally defined as follows. Let $G = (N, A)$ with $N = \{0, 1, \dots, n\}$ be a complete directed graph in which node 0 corresponds to the depot and the remaining nodes correspond to the locations of interest. Each arc $(i, j) \in A$ has an associated travel time t_{ij} and cost w_{ij} both satisfying the triangle inequality. The set of customers that require a delivery during the planning period $[0, T]$ is represented by C . The delivery for a customer $c \in C$ is characterized by a demand quantity d_c and a geographic profile $N_c \subseteq N$ which specifies where and when a delivery can be made. By duplicating locations, we may assume $N_c \cap N_{c'} = \emptyset$ for different customers $c, c' \in C$. Note that we can express the set of nodes as $N = N_0 \cup \{i \in N_c \mid c \in C\}$, where $N_0 = \{0\}$. The locations $i \in N_c$ have non-overlapping time windows $[e_i, l_i]$ during which the delivery to customer $c \in C$ can take place and correspond to the customer’s vehicle itinerary during the planning horizon. We use $c(i)$ to denote the customer associated with location i and we let $c(0) = 0$. A fleet of m homogeneous vehicles, each with capacity Q , is available to make deliveries; vehicles

start and end their delivery routes at the depot. The goal is to find a set of delivery routes visiting each customer at one of the locations in the customer's itinerary, during the time that the customer is at that location, and such that the demand delivered on a route is no more than Q , the duration of a route does not exceed T , and the total cost is minimized.

The basic version of the VRPRDL is static and deterministic, i.e., it is assumed that all customer locations and the time spent at these locations is known with certainty for the entire planning horizon. We use $\delta^-(i)$ and $\delta^+(i)$ to denote the sets of incoming and outgoing arcs of node i , respectively. Also, we write $x(A') = \sum_{(i,j) \in A'} x_{ij}$ for $A' \subseteq A$. The basic version of the VRPRDL can be formulated as follows:

$$\begin{aligned} \min \quad & \sum_{(i,j) \in A} w_{ij} x_{ij} \\ & x(\delta^-(i)) - x(\delta^+(i)) = 0 && i \in N, && (1) \\ & \sum_{i \in N_c} x(\delta^-(i)) = 1 && c \in C, && (2) \\ & x(\delta^+(0)) \leq m, && && (3) \\ & s_j \geq s_i + t_{ij} x_{ij} + (e_j - l_i)(1 - x_{ij}) && (i, j) \in A, j \neq 0, && (4) \\ & s_i + t_{i0} x_{i0} \leq \min\{l_i + t_{i0}, T\} && (i, 0) \in A, && (5) \\ & e_i \leq s_i \leq l_i && i \in N, && (6) \\ & y_j \geq y_i + d_{c(j)} - Q(1 - x_{ij}) && (i, j) \in A, j \neq 0, && (7) \\ & d_{c(i)} \leq y_i \leq Q && i \in N, && (8) \\ & x_{ij} \in \{0, 1\} && (i, j) \in A, && \end{aligned}$$

where x_{ij} is a binary variable indicating whether arc (i, j) is used or not, s_i is the arrival time at location i and y_i is the total quantity delivered by a vehicle making a delivery at location i when it is leaving location i . Note that the values of the y_i variables are relevant only for those locations that are actually visited. The objective is to minimize the total cost of the delivery routes. Vehicle flow conservation at every location is captured by Constraints (1). Constraints (2) guarantee that each customer receives a delivery at exactly one of the locations in his/her itinerary. Constraints (3) limit the number of vehicle departures from the depot to at most m . Constraints (4) through (6) determine the arrival times at each of the locations while ensuring that all vehicles return to the depot by time T and the time windows at the locations are respected. At the same time, these constraints prevent subtours from occurring. Constraints (7) and (8) ensure that the required quantities are delivered to the customers (in combination with Constraints (2)) and that the capacity of the vehicles is respected. Note that this formulation is slightly different from the one presented in Reyes et al. (2016).

The VRPRDL can also be formulated as a set partitioning problem by applying Dantzig-Wolfe decomposition to the formulation above. Let R denote the set of all feasible delivery

routes (i.e., respecting capacity and time window constraints), let w_r be the cost of route $r \in R$, and let a_{ir} for every $i \in N$ and $r \in R$ indicate whether location i is visited on route r ($a_{ir} = 1$) or not ($a_{ir} = 0$). The set partitioning formulation of the VRPRDL is as follows:

$$\begin{aligned}
\min \quad & \sum_{r \in R} w_r z_r \\
& \sum_{r \in R} \sum_{i \in N_c} a_{ir} z_r = 1 && c \in C, \\
& \sum_{r \in R} z_r \leq m, \\
& z_r \in \mathbb{Z}_+ && r \in R,
\end{aligned}$$

where z_r is the number of times route r is used. This formulation has exponentially many variables as the number of routes is exponential in the size of N . Therefore, we use column generation to solve its LP relaxation, which will be referred to as the master problem from now on. Note that since the arc costs satisfy the triangle inequality, we can replace $\sum_{r \in R} \sum_{i \in N_c} a_{ir} z_r = 1$, $c \in C$ with $\sum_{r \in R} \sum_{i \in N_c} a_{ir} z_r \geq 1$, $c \in C$ and still obtain a solution in which each customer is visited exactly once. This restricts the associated dual variable to be nonnegative, which typically leads to faster convergence of the column generation procedure.

2.1 Pricing problem

Let $\bar{R} \subset R$ be such that there exists a feasible solution to the master problem when $z_r = 0$ for all $r \in R \setminus \bar{R}$. A formulation involving only routes in \bar{R} is called a restricted master problem (RMP). The dual of the RMP is:

$$\begin{aligned}
\max \quad & \sum_{c \in C} \lambda_c + m \lambda_0 \\
& \sum_{c \in C} \sum_{i \in N_c} a_{ir} \lambda_c + \lambda_0 \leq w_r && r \in \bar{R}, \\
& \lambda_0 \leq 0, \\
& \lambda_c \geq 0 && c \in C.
\end{aligned}$$

Denote the optimal solution to this dual problem by (λ_0^*, λ^*) . In the pricing problem, the goal is to identify a column with negative reduced cost with respect to the original master problem. In other words, we aim to find a column for which the associated dual constraint is violated. Such a column is a route r for which the following condition is satisfied:

$$w_r - \lambda_0^* - \sum_{c \in C} \sum_{i \in N_c} a_{ir} \lambda_c^* < 0.$$

Note that since $w_r = \sum_{(i,j) \in r} w_{ij}$ and $\lambda_0^* + \sum_{c \in C} \sum_{i \in N_c} a_{ir} \lambda_c^* = \sum_{i \in N} a_{ir} \lambda_{c(i)}^*$, we can rewrite the condition as:

$$\sum_{(i,j) \in r} w_{ij} - \sum_{i \in N} a_{ir} \lambda_{c(i)}^* < 0.$$

Thus, the pricing problem is an elementary shortest path problem with time window and capacity constraints (ESPPTWCC), where the cost of arc (i, j) is set to $w_{ij} - \lambda_{c(i)}^*$ for $i \in N_c$ and $j \in N \setminus N_c$, i.e., the goal is to find an elementary path of shortest length starting and ending at the depot and respecting capacity and time window constraints. The ESPPTWCC can be formulated as follows:

$$\min \sum_{(i,j) \in A} (w_{ij} - \lambda_{c(i)}^*) x_{ij} \tag{9}$$

$$x(\delta^-(0)) = 1, \tag{9}$$

$$x(\delta^+(0)) = 1, \tag{10}$$

$$x(\delta^-(i)) - x(\delta^+(i)) = 0 \quad i \in N \setminus \{0\}, \tag{11}$$

$$\sum_{i \in N_c} x(\delta^-(i)) \leq 1 \quad c \in C, \tag{12}$$

$$\sum_{c \in C} d_c \sum_{i \in N_c} x(\delta^+(i)) \leq Q, \tag{13}$$

$$s_j \geq s_i + t_{ij} x_{ij} + (e_j - l_i)(1 - x_{ij}) \quad (i, j) \in A, j \neq 0, \tag{14}$$

$$s_i + t_{i0} x_{i0} \leq \min\{l_i + t_{i0}, T\} \quad (i, 0) \in A, \tag{15}$$

$$e_i \leq s_i \leq l_i \quad i \in N, \tag{16}$$

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

Constraints (9)-(11), together with subtour elimination constraints (14), define a path starting from and ending at the depot. Constraints (12) force this path to be elementary, i.e., they guarantee that each customer is visited at most once. Constraints (13)-(16) enforce the capacity and time window restrictions.

2.2 Solving the pricing problem

Solving a mixed integer programming problem using an off-the-shelf solver at every pricing iteration is computationally expensive. Hence, we adopt the iterative label setting algorithm proposed by Boland et al. (2006) to solve the ESPPTWCC. Given a directed graph with arbitrary arc lengths and a specified pair of source and sink nodes s and t , ESPPTWCC is the problem of finding the shortest resource-feasible path from s to t . As implied by its name, we have two different resources in ESPPTWCC arising from time window and capacity restrictions. Hence, a resource-feasible path in our context is one that respects the availability of time and capacity resources. Moreover, this path should also

be elementary, i.e., free of cycles. A common way to ensure that an elementary resource-feasible path is obtained at the end of a label setting algorithm is to use node-visit resources, which were introduced for the first time by (Feillet et al., 2004). More precisely, a resource having a limit of one unit is introduced for each node and it is consumed upon a visit to the node. Note, however, that the definition of an elementary path is slightly different in the case of the VRPRDL, because multiple locations are associated with a single customer. If nodes i and j are both locations associated with customer $c \in C$, then a path containing both i and j should not be considered elementary because it visits customer c at least twice. Therefore, instead of using node-visit resources, we maintain a customer-visit resource for each customer, which is consumed when any one of his associated locations is included in the path.

Label setting is a dynamic programming approach which constructs resource-feasible paths originating at a source node s and ending at a sink node t . Starting with the trivial path containing only the source node s , the label setting procedure extends unprocessed partial paths along all feasible arcs to create new (partial) paths. A state is associated with each partial path capturing the cost and the resource consumptions along the path. The extension of a partial path along an arc is infeasible if the resulting path would not be resource-feasible or if it cannot be augmented to reach the sink node within the available resource limits.

The efficiency of a label setting algorithm depends on the dominance relation that is used to eliminate partial paths. Let P_{si} and P'_{si} be two distinct paths from the source node s to node i with their respective states given by the label vectors U and U' . Suppose that the first element of the label vector represents the cost and the remaining elements represent the resource consumptions (in the same order for both label vectors) along the corresponding path. Path P_{si} is said to dominate path P'_{si} if $U \neq U'$ and $U_k \leq U'_k$ for $k = 1, \dots, K + 1$, where K is the number of resources.

To be able to eliminate additional partial paths using the above dominance relation, Feillet et al. (2004) introduce the notion of *unreachable nodes*. When creating a new partial path, the idea is to consume not only the node-visit resources for the nodes in the partial path itself, but also the node-visit resources of nodes that are not in the path, but cannot be reached by any feasible extension of the current partial path. Because we use customer-visit resources (rather than node-visit resources), we consider the unreachability of customers and say that customer c is unreachable if every node $i \in N_c$ is unreachable. Thus, the following resource consumption rule is adopted in our algorithm when consuming customer-visit resources. The resource corresponding to customer c in a partial path is consumed either when a node $i \in N_c$ is in the path, or else when none of the nodes from N_c can be contained in any feasible extension of the current partial path.

As the label associated with a partial path is a vector representing the state of the path, i.e., its cost and resource consumptions, the size of the state-space increases with the number of resources. In order to limit the size of the state-space and to accelerate the label-setting algorithm, Boland et al. (2006) suggest to start the solution procedure with

a state-space relaxation in which node-visit resources are initially not present. When the label setting procedure ends, the multiplicity of each node (number of times each node is visited) in the optimal path is computed and node-visit resources are introduced for some or all of the nodes with multiplicity greater than one. The label setting procedure is iteratively executed, each time starting with the original resources and the set of all node-resources introduced during the previous iterations, until the optimal path returned at the end of an iteration is elementary. The advantage of this state-space augmenting approach is that the number of node-visit resources introduced to obtain a least-cost elementary path is usually much smaller than the number of nodes.

To implement the state-space augmenting approach of Boland et al. (2006), we define S to be the set of critical customers, that is, the customers for which a customer-visit resource is maintained. The multiplicity of customer c in a path is computed by summing the number of visits to every $i \in N_c$ in this path. We start with $S = \emptyset$ and update S at the end of each label setting iteration that did not return an elementary optimal path by adding the customer with the highest multiplicity. In case of ties, we add the customer with the smallest index.

A straightforward implementation of the state-space augmenting approach solves the ESPPTWCC optimally. Despite being more efficient than using an off-the-shelf software to solve the integer programming formulation of the pricing problem, employing a straightforward implementation of the state-space augmenting algorithm to solve the pricing problem may still be (too) time-consuming. However, to solve the master problem there is no need to identify a most negative reduced cost column at every pricing iteration; identifying any negative reduced cost column suffices. Finding a most negative reduced cost column is typically needed only towards the end of the column generation process, because few, if any, negative reduced cost columns exist at that time. Consequently, we invoke the algorithm described above in our implementation only if no negative reduced cost columns can be detected heuristically, and even in that case we terminate the algorithm as soon as a pre-specified number of negative reduced cost columns is constructed by the algorithm.

3 A branch-and-price algorithm

We develop a branch-and-price algorithm to solve the VRPRDL, i.e., a branch-and-bound algorithm in which at each node of the search tree the LP relaxation is solved using column generation. The most time consuming component of a branch-and-price algorithm is typically the solution of the pricing problem. Therefore, the efficient detection of negative reduced cost columns is critical to the performance of any branch-and-price algorithm. In the following, we present the techniques we employ to speed up the solution of the pricing problem, we describe the adopted branching scheme, and we discuss how we deal with the tailing-off effect.

3.1 Heuristic pricing

It is not necessary to return a most-negative reduced cost column in each pricing iteration, it suffices to return (at least one) negative reduced cost column, if one exists. Even though the change in the value of the solution to the restricted master problem, when adding any negative reduced cost column rather than a most-negative reduced cost column, may be smaller, and more pricing iterations may have to be performed to reach an optimal solution, the reduction in solution time in each pricing step typically reduces the overall computation time.

A simple application of the above idea is based on the observation that usually many elementary negative reduced cost paths are found during the first few label setting iterations. Therefore, in our implementation of the state-space augmentation algorithm, we collect elementary negative reduced cost paths found in each iteration and will sometimes terminate the algorithm prematurely, i.e., before a most-negative reduced cost column has been found, and return all elementary negative reduced cost columns collected.

Another observation leads to a second useful idea: as the dual values are updated after each pricing iteration, some elementary paths with a non-negative reduced cost in one pricing iteration may have a negative reduced cost in a subsequent iteration. Therefore, at the end of a pricing iteration, efficient elementary paths with non-negative reduced costs found in the last label setting iteration are kept in a column pool. At the start of each pricing iteration, the columns in the pool are evaluated to see if they (now) have a negative reduced cost. To ensure that it does not become too costly to explore the column pool, we limit its size, i.e., we keep a maximum number of columns. At the end of a pricing iteration, if we add elementary non-negative reduced cost columns to the pool, we check its size, and, if the maximum pool size is exceeded, the oldest columns are removed until the desired pool size is reached.

Another strategy to detect negative reduced cost columns quickly is to make use of heuristics before invoking the exact pricing algorithm. To this end, we implement a truncated search version of the state-space augmenting approach. Instead of maintaining all non-dominated labels and their associated partial paths, we store only a small number of such labels at each node to speed up the search procedure. More specifically, we keep only a pre-specified number of efficient labels per node, and each time a new label is added to the list of efficient labels of a node, we discard the one with the largest cost if the number of labels exceeds the limit. In this way, fewer labels are treated at each label setting iteration which can facilitate faster detection of negative cost elementary paths. Again, excluding a part of the solution space may come at the expense of performing more iterations, i.e., there is a trade-off between the number of efficient labels maintained and the number of label setting iterations performed by the state-space augmenting approach.

Briefly, we try to identify negative reduced cost columns first by exploring the column pool, then by invoking the truncated-search version of the state-space augmentation algorithm, and finally by invoking the full-search version of the state-space augmentation

algorithm when the other approaches fail. To control the time spent in pricing iterations even further, we terminate any pricing iteration as soon as a predetermined number of elementary negative reduced cost columns has been found.

Finally, we keep track of the minimum cost γ of the elementary paths detected during the search, which gives an upper bound on the optimal value of the ESPPTWCC, and the cost η of the optimal path obtained at the end of each label setting iteration, which gives a lower bound on the optimal value of the ESPPTWCC. Once the gap between these bounds is closed, we can conclude that the pricing problem has been solved optimally.

As mentioned earlier, the state-space augmenting approach starts without any customer-visit resources. However, the customers added to the set S of critical customers during a pricing iteration are likely to be visited more than once in subsequent pricing iterations, and clearing S at the end of each pricing iteration may therefore result in an increase in the number of label setting iterations. Consequently, rather than initializing the algorithm with $S = \emptyset$ each time, we retain the set of critical customers throughout the column generation process at a node of the search tree. We set $S = \emptyset$ only at the start of processing a node in the search tree.

3.2 Branching

When the optimal solution to the master problem is fractional, we have to perform branching to find integer solutions. We branch on the arc variables, x_{ij} , in the original formulation. If an arc variable has a fractional value, we force the arc to be a part of the solution in one branch while prohibiting routes containing that arc in the other branch. Branching on variables in the original formulation has become a standard feature of many branch-and-price algorithms (Feillet, 2010).

Each variable of the master problem corresponds to the number of times a particular route is used in the optimal delivery plan. Therefore, we compute the value of an arc (i, j) in an optimal solution to the master problem by summing the optimal values of the routes that include the arc (i, j) . Once a branching decision is enforced, we have to ensure that the columns in the master problem and the ones that will be returned by the pricing subroutine are compatible with this decision. To guarantee compatibility at a child node, first we filter the columns coming from the restricted master problem associated with its parent node to exclude those that are in conflict with the branching decision and then update the pricing problem by forbidding the proper arc(s).

Ensuring compatibility is rather straightforward when the branching decision requires prohibiting a certain arc (i, j) , i.e., the columns containing this arc are removed from the restricted master problem and the arc (i, j) is ignored while solving the pricing problem. On the other hand, to ensure that (i, j) is included in the solution, we omit all columns (routes) containing any arc from the set $(\delta^+(i) \setminus (i, j)) \cup (\delta^-(j) \setminus (i, j)) \cup (\delta^-(l) : l \in (N_{c(i)} \setminus \{i\}) \cup (N_{c(j)} \setminus \{j\}))$ from the restricted master problem, and discard all the arcs in this set while solving the pricing problem (i.e., we do not extend partial paths through

these arcs). Note that as a result of column filtering, the master problem may become infeasible. Therefore, we always maintain a set of artificial columns with a high cost that guarantee the feasibility of the master problem. More specifically, we store the columns used to initialize the restricted master problem, and introduce these columns as “artificial” columns prior to starting column generation at a node of the search tree. By assigning the sum of the costs of these columns to each one of them individually, we ensure that these columns will not be in the optimal basis when column generation completes.

We have considered three strategies for selecting the arc to branch on. The first one is conventional branching, denoted by *CB*, where we select an arc with value closest to 0.5. The second one is to choose an arc that appears in the greatest number of routes in the solution to the master problem, denoted by *MF*. The last one is to choose an arc with a fractional value that occurs the earliest in time, denoted by *ED*. In particular, for an arc (i, j) , we determine the time at which a vehicle using this arc in its route departs from node i . There may be multiple routes containing (i, j) , in this case, we take the minimum of the departure times from node i over all such routes. In all of the arc selection strategies, we have the additional restrictions that the value of the chosen arc should be less than one and both of its endpoints should correspond to customer locations.

In *CB*, if there is more than one arc whose value is closest to 0.5, then we pick the first one encountered. In *MF*, we break ties by choosing either the arc encountered first during the search (*MF*) or the arc whose value is closest to 0.5 (*MFC*). Similarly, for *ED*, we either select the arc encountered first during the search (*ED*) or the arc whose value is closest to 0.5 (*EDC*).

3.3 Initial set of columns and feasible solutions

The column generation method starts by solving a restriction of the master problem. Therefore, we need to provide an initial set of columns that guarantees the feasibility of the master problem, i.e., a feasible starting solution. The quality of this solution can have a significant impact on the performance of the branch-and-price algorithm, especially for large problem instances. In our experiments, we use the feasible solution found by the heuristic of Reyes et al. (2016).

The value of any feasible solution provides an upper bound on the optimal objective function value and can thus be used to fathom nodes by bound. Of course, the higher the quality of the feasible solution, the more effective the fathoming becomes. Therefore, we also embed the following commonly used heuristic for producing a, hopefully high-quality, feasible solution. After solving the master problem at the root node, we solve the restricted master problem, i.e., including initial as well as generated columns, as an integer program (simply handing it to an off-the-shelf software package). This heuristic has proven to be quite successful in a number of different applications, and initial experimentation in our setting revealed that the resulting integer programs could be solved quickly, and, thus, did not impede the overall solution process.

3.4 Handling the tailing-off effect

Many pricing iterations may have to be performed with little or no improvement in the objective function value towards the end of the column generation process at a node in the search tree. This situation is quite common in branch-and-price algorithms and is known as the tailing-off effect. Below, we discuss two approaches for dealing with the tailing-off effect.

3.4.1 Using Lagrangian dual bounds for early pruning

The occurrence of tailing-off at a node is especially unfortunate if the node is fathomed by bound once the master problem has been solved, because in that case much time may have been “wasted” by “unnecessarily” solving pricing problems. This situation can, possibly, be prevented if an alternative lower bound can be computed that can be used to fathom the node before the column generation process at the node completes. Such a bound can be computed using concepts from Lagrangian relaxation. Note, however, that this may also mean that fewer columns are added to the column pool, which can have a negative impact on solution efficiency.

Dualizing the covering constraints in the master problem, we obtain the Lagrangian relaxation

$$\begin{aligned} \min \quad & \sum_{r \in R} w_r z_r + \sum_{c \in C} \lambda_c (1 - \sum_{r \in R} \sum_{i \in N_c} a_{ir} z_r) \\ \text{s.t.} \quad & \sum_{r \in R} z_r \leq m, \end{aligned} \tag{17}$$

$$z_r \geq 0, \quad r \in R, \tag{18}$$

which provides a lower bound on the optimal value of the master problem for any nonnegative vector of multipliers $(\lambda_1, \dots, \lambda_{|C|})$. Rearranging gives

$$\sum_{c \in C} \lambda_c + \left\{ \begin{array}{l} \min \sum_{r \in R} (w_r - \sum_{i \in N \setminus \{0\}} a_{ir} \lambda_{c(i)}) z_r \\ \text{s.t. (17), (18)} \end{array} \right\},$$

which, when taking the values of the optimal dual solution to RMP, gives

$$\sum_{c \in C} \lambda_c + \left\{ \begin{array}{l} \min \sum_{r \in R} (\bar{w}_r + \lambda_0) z_r \\ \text{s.t. (17), (18)} \end{array} \right\},$$

where \bar{w}_r is the reduced cost of route r and λ_0 is the value of the dual variable corresponding to (17). This quantity can be bounded from below by

$$\sum_{c \in C} \lambda_c + m(\bar{w}_{min} + \lambda_0) = \sum_{c \in C} \lambda_c + m\lambda_0 + m\bar{w}_{min} = z_{RMP}^* + m\bar{w}_{min},$$

where \bar{w}_{min} is the minimum reduced cost and z_{RMP}^* is the optimal value of RMP.

This lower bound can be computed easily at the end of a column generation iteration if the optimal value of the pricing problem is known. (Note that z_{RMP}^* always provides an upper bound on the optimal value of the master problem, and that when $\bar{w}_{min} = 0$ it provides a lower bound as well, in which case the master problem is solved optimally.) Observe that it is possible to obtain a lower bound on the optimal value of the master problem as soon as a complete label setting iteration is performed, because the state-space augmenting algorithm yields a lower bound on the optimal value of the pricing problem at the end of each label setting iteration. Therefore, whenever the full-search version of the state-space augmentation algorithm is used in a pricing iteration, we compute a lower bound for the master problem at the end of each label setting iteration and see if it can be used to prune the node.

3.4.2 Early branching

Another strategy to deal with the tailing-off effect is to prematurely terminate the column generation procedure and perform branching early. To accomplish this, one can stop generating columns and apply branching, for example, when the improvement in the objective function value is less than ϵ in the last Δ iterations. Note that the values of ϵ and Δ need to be set carefully, because branching too aggressively may grow the search tree significantly and may be counterproductive.

3.5 Implementation details

There are many algorithmic choices and parameter settings that impact the computational performance of the branch-and-price algorithm. In the following, we first summarize our implementation of a straightforward branch-and-price algorithm and then describe the algorithmic choices and parameter settings for the enhanced version that incorporates the ideas discussed above.

3.5.1 Straightforward branch-and-price

In order to evaluate the improvements achieved by our enhanced branch-and-price algorithm, we consider a straightforward approach, in which out-and-back routes from the depot to the first location of every customer are used to define the initial set of columns (these routes always correspond to a feasible solution because $m = |C|$ in VRPRDL instances), the state-space augmenting algorithm is initialized with $S = \emptyset$ at every pricing iteration, only a single column with the most negative reduced cost is added to the restricted master problem (i.e., the pricing problem is solved to optimality), and the conventional branching strategy (CB) is employed; there is no early pruning, no early branching, and the restricted master problem at the end of the column generation process at the root node is not used to seek a, possibly improved, feasible solution.

3.5.2 Enhanced branch-and-price

In our enhanced branch-and-price algorithm, the following parameter settings control the solution of the pricing problem: the column generation process terminates as soon as β elementary negative reduced cost columns are identified. The truncated-search version of the state-space augmentation algorithm restricts the number of non-dominated labels at each node to α . In the truncated-search version of the state-space augmentation algorithm, we also monitor the cost γ of the shortest elementary path detected so far, which yields an upper bound on the optimal ESPPTWCC value, and the cost η of the shortest path obtained at the end of each truncated label setting iteration, which provides a lower bound on the cost of the shortest elementary path that can be obtained by truncated-search, and terminate the truncated search when these bounds at the end of a label setting iteration are equal. The size of the column pool is 1000, i.e., at most 1000 columns are stored at any one time. Only when neither the exploration of the column pool nor the truncated search produces any elementary negative reduced cost columns do we invoke exact pricing. The set S of critical customers used in the state-space augmentation algorithm is retained throughout the column generation process at each node in the search tree, and it is cleared right before column generation starts at another node.

In some of our computational experiments and for some values of α and β , we also consider forcing the truncated search to stop upon completing a single label setting iteration even when the number of negative reduced cost columns detected is less than β and the gap between γ and η is not closed at the end of this label setting iteration. We use T to denote that we adopt this termination criterion, i.e., stop at the end of the first label setting iteration, instead of iterating until γ and η become equal, which is denoted by F .

Furthermore, the solution obtained by the heuristic of Reyes et al. (2016) is used to initialize the algorithm, the restricted master problem is solved as an integer program upon completing column generation at the root node in the hope of finding an improved feasible solution, early pruning is active, and the nodes in the branch-and-price tree are evaluated in a depth-first order (best-bound and depth-first search produce similar results on small and medium size instances, the latter performs better for large size instances). Early branching is not used as computational experiments revealed that the benefits are negligible.

4 Incorporating a home delivery option

Trunk delivery was introduced in the hope that it would create cost-saving opportunities compared to making home deliveries. It is easy to construct examples where this is indeed the case. However, it is also easy to construct examples where this is not the case and the cost actually increases. Thus, companies that are considering trunk delivery will likely deploy a hybrid model in which deliveries can either be made at the home location of the customer (during the entire planning horizon) or to trunk of the customer's car at one of the locations in the car's itinerary, if this creates cost savings.

Fortunately, our branch-and-price algorithm can be used for solving this hybrid model by simply making slight changes in the instance data. In particular, we replace the time windows associated with the home location of each customer, which, in the instances of VRPRDL, are the first and last location of the car’s itinerary by a single location with time window $[0, T]$, where T is the length of the planning horizon. Note that this means that the time windows of the locations visited by a customer are no longer non-overlapping. However, the branch-and-price algorithm has no components that exploit or rely on this non-overlapping property and thus it can be applied without any modification. Of course the performance may (and, as we will see, will) deteriorate.

5 Computational study

We conduct a computational study to (1) evaluate the performance of our branch-and-price algorithm, and to (2) assess the benefits of employing trunk delivery services.

5.1 Instances and preprocessing

In our computational experiments, we use slightly modified versions of the 40 VRPRDL instances introduced in Reyes et al. (2016); the travel times have been adjusted to satisfy the triangle inequality and, when necessary, the time windows at locations have been adjusted accordingly. In all the instances, the planning horizon is 12 hours, the vehicle capacity is 750, and the geographic profile of every customer consists of at least one and at most six locations (the first and last location always being the home location of the customer - in case there is only one location, it signals that the customer is at home all day).

Although there is no restriction on the fleet size in the original instances, we assume that m , the number of vehicles available for making the deliveries, is equal to the number of routes in the solution provided by the heuristic of Reyes et al. (2016) plus one when solving VRPRDL and VRPHRDL instances with our enhanced algorithm. We impose this restriction, because smaller values of m results in stronger lower bounds on the optimal value of the master problem, and preliminary experiments using small to medium sized instances revealed that the optimal costs remain unchanged even if we set $m = |C|$. Characteristics of the instances are provided in Table 1. For each instance, we present the number of customers, the average number of locations per customer, and, considering all customers, the minimum, average, and maximum distance from home location to the depot, and the minimum, average, and maximum fraction of time available for making a delivery. We observe that there are noticeable differences between instances, with some instances having, on average, only 2.47 locations per customer whereas others have, on average, 4.30 locations per customer, and some instances having, on average, only 49% of the planning horizon available to make deliveries whereas others have, on average, 77% of the planning horizon available.

Table 1: Characteristics of the instances

Instance	Number of customers	Average number of customer locations	Distance to depot (from home location)			Fraction of time available for delivery		
			min	avg	max	min	avg	max
1	15	4.07	3	67.07	166	0.05	0.58	1.00
2	15	3.73	28	83.73	148	0.31	0.65	1.00
3	15	3.40	7	72.33	167	0.08	0.66	1.00
4	15	3.27	3	70.80	159	0.08	0.59	1.00
5	15	3.40	13	101.07	174	0.14	0.68	1.00
6	20	3.25	2	76.45	152	0.16	0.67	1.00
7	20	3.35	15	76.25	174	0.08	0.64	1.00
8	20	3.95	6	87.10	161	0.04	0.56	1.00
9	20	3.75	3	73.35	171	0.01	0.49	1.00
10	20	3.10	0	80.55	169	0.14	0.66	1.00
11	30	3.40	2	85.43	176	0.03	0.62	1.00
12	30	3.73	20	78.03	174	0.01	0.53	1.00
13	30	3.90	3	71.27	171	0.04	0.54	1.00
14	30	3.53	5	73.07	171	0.12	0.61	1.00
15	30	4.10	3	81.70	176	0.04	0.54	1.00
16	30	3.93	2	72.77	160	0.03	0.59	1.00
17	30	4.30	2	85.23	165	0.08	0.50	1.00
18	30	3.50	17	89.70	154	0.01	0.58	1.00
19	30	3.23	12	82.47	179	0.08	0.63	1.00
20	30	2.47	12	75.10	172	0.05	0.77	1.00
21	60	3.73	4	78.88	173	0.04	0.60	1.00
22	60	3.53	1	80.07	170	0.01	0.58	1.00
23	60	3.90	2	89.03	179	0.04	0.55	1.00
24	60	3.80	8	93.12	175	0.03	0.63	1.00
25	60	3.88	3	79.60	165	0.06	0.59	1.00
26	60	3.73	1	89.77	178	0.01	0.56	1.00
27	60	3.45	8	90.27	176	0.03	0.63	1.00
28	60	3.63	1	75.07	180	0.04	0.56	1.00
29	60	3.75	2	91.95	179	0.02	0.59	1.00
30	60	3.95	4	93.97	177	0.01	0.59	1.00
31	120	3.83	0	76.23	178	0.01	0.60	1.00
32	120	3.67	0	83.98	174	0.04	0.59	1.00
33	120	3.92	0	69.69	179	0.02	0.57	1.00
34	120	3.78	2	79.45	176	0.01	0.59	1.00
35	120	3.75	1	77.35	174	0.07	0.56	1.00
36	120	3.51	0	89.57	178	0.05	0.62	1.00
37	120	3.88	3	83.14	177	0.02	0.55	1.00
38	120	3.51	0	86.11	177	0.03	0.62	1.00
39	120	3.56	0	89.30	178	0.03	0.59	1.00
40	120	3.84	0	79.73	171	0.07	0.55	1.00

We compute the Euclidean distance between each pair of locations based on their coordinates and then round this distance to the nearest integer. In order to ensure that arc costs satisfy the triangle inequality, we assign the shortest distance between each pair of nodes to the cost of the arc connecting these two nodes. Furthermore, we reduce the size of the network by eliminating nodes and arcs that cannot be a part of any feasible solution. Our preprocessing steps are as follows:

1. Eliminate node i and all arcs incident to this node if $t_{0i} > l_i$ or $e_i + t_{i0} > T$; and
2. Eliminate arc (i, j) if $\max\{t_{0i}, e_i\} + t_{ij} > l_j$ or $\max\{t_{0i}, e_i\} + t_{ij} + t_{j0} > T$

Essentially, we eliminate a node i if an out-and-back route from the depot to the node, i.e., route $0 - i - 0$, leads to a time window violation. Similarly, we eliminate an arc (i, j) if route $0 - i - j - 0$ is not time-feasible. This simple preprocessing is very effective, it reduces the number of nodes and arcs substantially: the minimum, average, and maximum reduction in the number of arcs across the VRPRDL instances are 72.5%, 87.5% and 93.09%, respectively. The reduction is smaller for the VRPHRDL instances as the home locations now have wide time windows, and thus fewer nodes and arcs are eliminated.

The algorithm is implemented in Java using the branch-and-price framework of Java OR Library, and CPLEX 12.6.3 is employed for solving the restricted master problems through Concert Technology. All experiments are performed on a 64-bit machine with Intel Xeon E5-2650 v3 processor at 2.30 GHz. The time limit is set to two hours for instances with up to 60 customers and to six hours for instances with 120 customers. In the computational results tables, we will indicate that a time limit was reached with TL.

5.2 Evaluating the performance of the branch-and-price algorithm

To be able to evaluate the benefits of the various techniques introduced in the branch-and-price algorithm, we start by solving the instances with the straightforward implementation described in Section 3.5.1. The results can be found in Table 2. For each instance, we report the name of the instance, the cost of the initial solution, the cost of the best solution, the integrality gap, the solution time (in seconds), the total number of pricing iterations performed during the execution of the algorithm, and the number of nodes evaluated during the search, respectively. When the algorithm terminates within the time limit, the best solution is optimal, otherwise it may or may not be optimal. The integrality gap is computed as the ratio $(z_{IP}^* - z_{LP}^*)/z_{LP}^*$, where z_{LP}^* is the optimal value of the master problem at the root node of the search tree and z_{IP}^* is the cost of the best solution found during the execution of the algorithm. When the master problem at the root node cannot be solved within the time limit, it is not possible to compute an integrality gap, which is indicated with a dash. If an instance is solved to optimality (i.e., the algorithm terminates within the time limit), the integrality gap provides some additional insight into the relative difficulty of the instance. If, on the other hand, an instance cannot be solved within the time limit, then the integrality gap provides a quality guarantee of the best solution found.

Table 2: Results with the straightforward BAP

Instance	Initial solution	Best solution	Integrality gap (%)	Solution time (s)	Iterations	Nodes
1	2074	901	0	0.60	39	1
2	2316	1286	0	0.32	29	1
3	2236	991	0	0.48	31	1
4	1982	1062	0	0.24	30	1
5	3327	1833	0	0.05	26	1
6	3330	1294	1.73	2.18	250	33
7	3208	1156	1.34	31.98	983	85
8	3172	1455	0	0.16	39	1
9	2840	1260	1.82	3.24	298	32
10	3270	1684	0	0.52	33	1
11	4934	1922	0.29	12.94	175	13
12	4620	2326	2.48	72.47	6889	1017
13	4872	1749	0	20.85	104	1
14	4085	1273	0	29.56	100	1
15	4660	1696	0	24.17	100	1
16	4775	1940	0	58.92	88	1
17	4504	1968	0.10	4.46	125	4
18	5394	1827	0	2.50	88	1
19	5286	2083	2.46	126.17	4974	503
20	4238	1822	0	81.48	108	1
21	10380	3762	0	660.09	309	4
22	9327	2830	0	224.92	312	2
23	10349	4448	0	218.88	171	2
24	10551	3383	0	344.39	280	2
25	9797	9797	-	TL	1	0
26	10828	4538	0	68.55	150	1
27	10647	2868	0	281.85	283	1
28	9459	9459	-	TL	1	0
29	10999	4085	3.70	TL	138541	18415
30	11851	4112	0	53.14	201	1
31	18172	18172	-	TL	1	0
32	19558	19558	-	TL	1	0
33	18063	18063	-	TL	4	0
34	19941	19941	-	TL	6	0
35	19244	19244	-	TL	2	0
36	21803	21803	-	TL	4	0
37	20056	20056	-	TL	2	0
38	20078	20078	-	TL	3	0
39	20190	5903	1.55	TL	1984	77
40	20077	20077	-	TL	2	0

We observe that the straightforward implementation is able to solve all instances with 15, 20, and 30 customers, but fails to find an optimal solution or prove its optimality for three of the instances with 60 customers. In fact, the first pricing problem cannot be solved within two hours for Instances 25 and 28. None of the instances with 120 customers can be solved within six hours. In fact, the algorithm can only solve the master problem at the root node for Instance 39.

Next, we solve the instances with up to 60 customers using the enhanced implementation introduced in Section 3.5.2 with different combinations of control parameters for the solution of the pricing problem and with different branching schemes. More specifically, six configurations for solving the pricing problem were evaluated: $(5, 5, F)$, $(5, 10, F)$, $(10, 10, F)$, $(15, 10, F)$, $(10, 10, T)$ and $(15, 10, T)$, where the first parameter specifies the number of labels kept in the truncated-search (α), the second parameter specifies the number of negative reduced cost columns required before terminating the solution of the pricing problem early (β), and the third parameter specifies whether or not the state-space augmenting algorithm is terminated after a single label setting iteration (T or F). Furthermore, all branching schemes were evaluated: CB , MF , MFC , ED , and EDC . Thus, in total, the 30 instances were solved 30 times.

We found that 18 of the instances can be solved optimally without branching by each of the tested configurations. The solution to the master problem at the root node is integer in 17 cases, and solving the restricted master problem at the end of the column generation process at the root node as an integer program produced a solution with cost equal to the optimal objective value of the master problem in the other case (Instance 21). Furthermore, we found that the branching schemes ED and EDC performed poorly compared to the other branching schemes. Therefore, in Table 3, we present the results for six configurations and three branching schemes for 12 instances. We report the solution time (in seconds), the number of pricing iterations, and the number of nodes evaluated.

To analyze the results and to choose a default branching scheme and default control parameters for the solution of the pricing problem, we will focus on the solution times. First, we observe that both the MF and the MFC branching schemes outperform the CB branching scheme. The choice between the different configurations of control parameters for the solution of the pricing problem is not obvious, but schemes $MF(5, 10, F)$, $MF(15, 10, T)$, and $MFC(15, 10, F)$ appear to be most “robust”, in the sense that they lead to the minimum total solution time across the instances (with up to 60 customers).

We also examined how often a pricing iteration completes after (1) exploring the column pool, (2) applying truncated-search, and (3) applying full-search. For all branching schemes and all parameter configurations, most pricing iterations complete after truncated-search, which implies that truncated-search returns at least one elementary negative reduced cost column. Although rare, pricing iterations sometimes complete after exploring the column pool, especially when many pricing iterations have to be performed, which is not surprising since the column pool fills up gradually and the more columns, the better the chance of having negative reduced cost columns in the column pool in subsequent iterations.

Table 3: Results with CB , MF , and MFC arc selection rules

	(5, 5, F)			(5, 10, F)			(10, 10, F)			(15, 10, F)			(10, 10, T)			(15, 10, T)		
	Time	Iters.	Nodes	Time	Iters.	Nodes	Time	Iters.	Nodes	Time	Iters.	Nodes	Time	Iters.	Nodes	Time	Iters.	Nodes
6	0.97	172	39	0.74	84	21	1.23	126	39	1.42	156	45	0.93	131	29	1.22	187	43
7	1.37	220	21	1.75	174	21	2.01	169	25	2.67	195	27	2.26	169	15	3.74	252	33
9	0.47	107	20	0.48	98	21	0.45	71	17	0.48	73	17	0.85	132	25	1.00	147	25
11	0.82	109	9	0.67	65	7	0.65	63	9	0.79	64	9	0.73	79	9	1.29	98	7
12	221.84	36554	10335	121.89	19783	6119	69.23	10931	3533	178.34	27649	9005	17.85	2578	707	57.82	9824	3015
17	0.55	76	3	0.98	97	7	0.58	58	7	0.52	47	3	0.70	65	7	0.51	63	3
19	5.74	534	79	6.10	516	99	7.66	537	97	8.35	496	97	9.58	588	95	10.64	704	121
21	4.40	107	1	5.20	82	3	8.98	92	3	5.45	55	1	8.50	88	1	5.56	84	1
23	2.77	110	2	2.36	68	2	3.46	77	3	5.72	101	7	6.81	116	3	6.12	116	3
25	794.84	11898	773	433.27	4937	530	529.03	5195	571	192.00	1423	116	969.33	9390	835	741.13	6643	593
28	104.24	1098	87	30.96	226	16	86.42	470	49	66.03	364	25	103.38	539	48	143.32	794	77
29	2,031.47	42760	9103	1,286.81	25733	6070	619.81	7624	1869	4,003.29	59256	15556	6,250.11	128035	27197	1,069.00	20528	4300
Select the arc for branching using the MF rule																		
6	1.08	172	39	0.74	84	21	1.36	126	39	1.46	156	45	0.90	131	29	1.19	187	43
7	3.11	449	69	2.38	273	37	1.95	171	25	3.78	327	65	3.51	254	28	3.87	323	27
9	0.47	102	18	0.39	78	17	0.53	88	19	0.53	88	18	0.63	108	20	0.75	121	20
11	0.82	109	9	0.67	65	7	0.65	63	9	0.80	64	9	0.74	79	9	1.08	98	11
12	8.08	1158	267	10.63	1486	375	13.81	1787	543	18.82	2463	759	3.24	397	95	15.18	2107	499
17	0.56	76	3	1.01	97	7	0.64	58	7	0.51	47	3	0.70	65	7	0.53	63	3
19	10.92	1111	199	8.85	839	169	13.48	1007	211	10.03	784	171	9.96	833	155	14.90	1109	205
21	4.66	107	1	4.60	77	3	9.12	92	3	5.54	55	1	8.347	88	1	5.74	84	1
23	2.86	110	2	2.64	71	3	3.49	77	3	4.374	80	4	6.68	116	3	6.10	116	3
25	1,127.09	16716	1321	464.36	5355	577	1,207.38	12406	1693	739.68	7080	811	1,336.75	14418	1506	89.88	299	7
28	79.72	1047	77	27.84	199	10	97.98	566	67	65.78	364	25	91.30	484	44	131.17	783	71
29	994.52	26220	6785	40.21	749	151	33.27	592	131	200.60	3619	997	191.94	3887	928	57.03	804	133
Select the arc for branching using the MFC rule																		
6	0.99	172	39	0.71	84	21	1.13	126	39	1.32	156	45	0.90	131	29	1.21	187	43
7	2.33	346	43	1.72	206	23	2.10	178	25	2.15	185	23	4.26	274	33	3.13	225	23
9	0.45	102	18	0.58	102	23	0.64	88	21	0.67	90	21	0.81	132	25	1.03	147	25
11	0.80	109	9	0.73	65	7	0.71	63	9	0.942	64	9	0.74	79	9	1.08	98	11
12	12.54	1733	451	12.15	1627	399	8.77	992	287	15.94	1902	561	6.69	820	207	18.90	2667	677
17	0.66	76	3	0.99	97	7	0.72	58	7	0.52	47	3	0.704	65	7	0.50	63	3
19	9.37	929	163	8.30	819	159	12.24	886	185	8.74	679	147	10.60	904	159	13.94	1069	199
21	4.58	107	1	5.26	80	3	9.04	92	3	5.49	55	1	8.52	88	1	5.56	84	1
23	2.89	110	2	2.92	80	3	3.56	77	3	3.41	76	3	6.60	116	3	6.15	116	3
25	314.06	4321	318	248.12	3180	345	181.03	1669	201	105.35	834	87	226.06	2340	235	90.28	279	5
28	62.40	691	41	27.64	199	10	76.18	419	45	65.72	364	25	93.83	492	46	110.35	510	43
29	162.38	3914	831	350.66	6436	1658	277.86	4252	1006	158.38	2757	703	231.84	4155	825	362.94	6210	1305

We use the three most promising settings, i.e., $MF(5, 10, F)$, $MF(15, 10, T)$, and $MFC(15, 10, F)$ to solve the 120 customer instances. The results can be found in Table 4. For each instance, the first four columns correspond to the name of the instance, the cost of the initial solution, the cost of the solution obtained by solving the restricted master problem at the root node as an integer program, and the cost of the best solution found, respectively. The remaining columns provide the integrality gap, the solution time, the total number of pricing iterations, the average time per pricing iteration, and finally, the number of nodes evaluated. The best solution for each instance is highlighted in bold.

We observe that the algorithm finds an optimal solution to four instances with the setting $MF(5, 10, F)$, whereas the number of instances solved optimally is only two with $MF(15, 10, T)$ and three with $MFC(15, 10, F)$. Moreover, the setting $MF(5, 10, F)$ produces the best solution for nine instances, whereas only four best solutions are produced by the other two schemes. Finally, the setting $MF(5, 10, F)$ solves the two instances that are solved to optimality with all settings much faster.

We also observe that solving the restricted master problem at the root node as an integer program is quite effective. With the setting $MF(5, 10, F)$, the solution value improves, on average, by almost 5%, and for Instance 39 the improvement exceeds 10%. The integrality gap values reveal that high-quality solutions are produced; with the setting $MF(5, 10, F)$, the average integrality gap is only 1.12%. Finally, we observe, as expected, that when the number of labels kept at a node during truncated-search is larger, the time per pricing iteration increases (compare the time per iterations for $MF(5, 10, F)$ and $MFC(15, 10, F)$).

To evaluate the benefits of the various techniques introduced in the branch-and-price algorithm to improve its performance, we next compare the straightforward implementation with the enhanced implementation with the setting $MF(5, 10, F)$. The results can be found in Table 5. A dash in the column “Root IP solution” means that the master problem has an integer optimal solution.

We observe that the benefits of implementing the techniques described earlier are significant. Instances that can be solved to optimality are solved orders of magnitude faster, and for the instances that cannot be solved to optimality, solutions of much better quality are obtained. Specifically, for the instances that can be solved optimally by both approaches, the reduction in average solution time is 97.3%. Furthermore, when the straightforward implementation fails to find an optimal solution within the time limit, it is usually unable to solve the root LP; this never happens with the enhanced implementation. Finally, we note that the average integrality gap, when using the enhanced implementation, is less than 1.8%, which demonstrates that high-quality solutions can be obtained for VRPRDL instances with up to 120 customers.

5.3 Assessing the benefits of employing trunk delivery services

To be able to comprehensively assess the benefits of trunk delivery services, we need to be able to solve instances of the VRPHRDL as well. Therefore, we start by investigating how

Table 4: Results obtained with the selected settings on the large instances

Instance	Initial solution	Root IP solution	Best solution (5, 10, F)	Integrality gap (%) with MF branching rule	Solution time (s) with MF branching rule	Iterations in total	Time per iteration (s)	Nodes
31	5193	4943	4943	0.02	5,349.30	3747	1.42	149
32	5696	5360	5276	2.63	TL	29135	0.74	3252
33	5165	5100	5100	3.02	TL	31216	0.69	1620
34	5497	5232	5225	0.19	1,673.29	1166	1.43	33
35	5701	5554	5545	1.69	TL	37717	0.57	4180
36	7098	-	6506	0	122.79	308	0.40	1
37	4975	4841	4836	0.22	935.58	1051	0.88	39
38	5752	5618	5618	1.07	TL	50287	0.43	4511
39	6564	5873	5857	0.76	TL	66286	0.32	8437
40	5274	5017	5017	1.59	TL	34302	0.63	3968
(15, 10, T) with MF branching rule								
31	5193	4943	4943	0.02	10,685.09	4195	2.54	173
32	5696	5312	5296	3.02	TL	9234	2.34	836
33	5165	5105	5105	3.12	TL	52356	0.41	4876
34	5497	5227	5225	0.19	TL	8266	2.61	550
35	5701	5551	5551	1.80	TL	33910	0.63	3309
36	7098	6506	6506	0	993.29	358	2.77	1
37	4975	4851	4851	0.53	TL	36341	0.59	2882
38	5752	5615	5615	1.02	TL	36571	0.59	2692
39	6564	5939	5857	0.76	TL	41676	0.51	4676
40	5274	5060	5060	2.46	TL	15292	1.41	1233
(15, 10, F) with MFC branching rule								
31	5193	4943	4943	0.02	11,814.27	5206	2.26	255
32	5696	5390	5303	3.16	TL	21301	1.01	2326
33	5165	5112	5103	3.08	TL	24882	0.86	1971
34	5497	5233	5227	0.23	TL	8530	2.52	898
35	5701	5558	5558	1.93	TL	33635	0.64	3982
36	7098	-	6506	0	299.59	364	0.82	1
37	4975	4839	4836	0.22	2,856.11	2351	1.21	138
38	5752	5623	5611	0.95	TL	35926	0.59	4567
39	6564	5925	5907	1.62	TL	64529	0.33	13093
40	5274	5022	5021	1.67	TL	22311	0.96	2397

Table 5: Straightforward branch-and-price vs. the default branch-and-price with parameter configuration $(5, 10, F)$ and the MF branching rule

Instance	Straightforward		$MF(5, 10, F)$				
	Best solution	Solution time (s)	Initial solution	Root IP solution	Best solution	Integrality gap (%)	Solution time (s)
1	901	0.60	957	-	901	0	0.22
2	1286	0.32	1292	-	1286	0	0.05
3	991	0.48	1004	-	991	0	0.07
4	1062	0.24	1069	-	1062	0	0.04
5	1833	0.05	1833	-	1833	0	0.02
6	1294	2.18	1300	1300	1294	1.73	0.74
7	1156	31.98	1159	1159	1156	1.34	2.38
8	1455	0.16	1555	-	1455	0	0.06
9	1260	3.24	1272	1265	1260	1.82	0.39
10	1684	0.52	1733	-	1684	0	0.03
11	1922	12.94	1932	1922	1922	0.29	0.67
12	2326	72.47	2328	2328	2326	2.48	10.63
13	1749	20.85	1758	-	1749	0	0.52
14	1273	29.56	1281	-	1273	0	0.52
15	1696	24.17	1698	-	1696	0	0.43
16	1940	58.92	1943	-	1940	0	0.59
17	1968	4.46	1969	1968	1968	0.10	1.01
18	1827	2.50	1831	-	1827	0	0.18
19	2083	126.17	2121	2109	2083	2.46	8.85
20	1822	81.48	1889	-	1822	0	0.57
21	3762	660.09	3776	3762	3762	0	4.60
22	2830	224.92	2880	-	2830	0	7.96
23	4448	218.88	4455	4449	4448	0	2.64
24	3383	344.39	3481	-	3383	0	8.17
25	9797	TL	3376	3169	3162	0.82	464.36
26	4538	68.55	4591	4538	4538	0	1.63
27	2868	281.85	2880	-	2868	0	8.39
28	9459	TL	4222	4177	4174	0.08	27.84
29	4085	TL	4019	3976	3968	0.73	40.21
30	4112	53.14	4128	-	4112	0	1.47
31	18172	TL	5193	4943	4943	0.02	5349.34
32	19558	TL	5696	5360	5276	2.63	TL
33	18063	TL	5165	5100	5100	3.02	TL
34	19941	TL	5497	5232	5225	0.19	1673.29
35	19244	TL	5701	5554	5545	1.69	TL
36	21803	TL	7098	-	6506	0	122.79
37	20056	TL	4975	4841	4836	0.22	935.58
38	20078	TL	5752	5618	5618	1.07	TL
39	5903	TL	6564	5873	5857	0.76	TL
40	20077	TL	5274	5017	5017	1.59	TL

well our branch-and-price algorithm performs on VRPHRD instances. We experimented with the three settings that proved most effective for the VRPRDL instances and found that although no setting clearly dominated, most best-known solutions are found with the setting $MFC(15, 10, F)$. Therefore, we report the results obtained with this setting in Table 6, where, for completeness sake, we include also the value of the best solution found by any of the three settings (highlighting in bold the values of the solutions that have a smaller cost than the value of the best solution obtained with the setting $MFC(15, 10, F)$).

For the VRPHRD instances with up to 60 customers, 24 of the 30 instances can be solved to optimality within two hours and optimal solutions to most of these instances are found at the root node. Note that, as in previous experiments, the initial solutions are obtained by the heuristic of Reyes et al. (2016).

For the VRPHRD instance with 120 customers, the algorithm does not finish solving the root node LP within six hours, except for Instance 39, demonstrating that VRPHRD instances are harder to solve than VRPRDL instances. Even when the time limit is reached before solving the root node LP, a large number of columns has been generated, which enables us to find an improved solution by solving the restricted master problem as an integer program when the time limit is reached. We observe that for these difficult instances, the branch-and-price algorithm finds solutions that, on average, improve the initial solution by a little over 4%. As mentioned at the start of this section, the primary reason that VRPHRD instances are more difficult than VRPRDL instances is that time windows at the home locations are much wider, which reduces the effectiveness of the preprocessing, and, as a consequence, leads to denser graphs and more arc variables.

Having alternative locations to deliver a customer order, e.g., to the trunk of the customer’s car when it is parked somewhere other than at home, provides additional flexibility to a delivery company. However, it is easy to construct examples where allowing deliveries only to the trunk of a customer’s car may not lead to cost savings, but may, in fact, lead to an increase in cost, compared to being able to only deliver at the customer’s home. This is the reason that companies are more likely to deploy a hybrid model in which deliveries can either be made to the customer’s home or to the customer’s car.

To assess the benefits of trunk delivery, we compute, for the same instance, a VRP solution, in which deliveries can only be made at the customer’s home location, a VRPRDL solution, in which deliveries can only be made to the trunk of the customer’s car, and a VRPHRD solution, in which deliveries can be made either to the customer’s home or to the trunk of the customer’s car. The results can be found in Table 7, where the VRP solution is also computed with our branch-and-price algorithm. We report the cost, the fraction of deliveries made at a customer’s home, and the number of delivery routes. We use a star to indicate that the best solution is not necessarily optimal, i.e., the algorithm reached the time limit.

We observe that having the flexibility to deliver either to a customer’s home or to the trunk of the customer’s car has huge advantages. The average percentage of cost-savings is about 19% (with a minimum cost savings of 4.85% and a maximum cost savings of 36.34%).

Table 6: Results for the VRPHRDL instances with parameter configuration $(15, 10, F)$ and the *MFC* branching rule

Instance	Initial solution	Root IP solution	Best solution	Best known solution	Integrality gap (%)	Solution time (s)	Iterations	Nodes
1	777	-	773	773	0	0.67	20	1
2	1111	-	1065	1065	0	0.08	9	1
3	992	-	989	989	0	0.20	16	1
4	917	-	915	915	0	0.16	15	1
5	1710	-	1710	1710	0	0.05	12	1
6	1099	1099	1099	1099	2.04	5.29	283	57
7	1020	1000	996	996	0.67	10.20	215	29
8	1458	-	1346	1346	0	0.40	27	1
9	998	-	997	997	0	0.62	19	1
10	1167	-	1166	1166	0	0.17	26	1
11	1608	-	1587	1587	0	6.03	62	1
12	1836	-	1809	1809	0	6.53	49	1
13	1563	-	1563	1563	0	3.94	39	1
14	1058	-	1058	1058	0	3.54	47	1
15	1363	1354	1347	1347	2.06	121.36	868	107
16	1575	1566	1566	1521	5.24	TL	187147	51035
17	1446	-	1445	1445	0	2.76	59	1
18	1680	1628	1628	1628	0.97	27.87	260	35
19	1468	-	1461	1461	0	1.20	34	1
20	1730	-	1715	1715	0	2.87	47	1
21	2860	-	2580	2580	0	425.16	431	1
22	2214	2214	2214	2214	5.01	TL	11746	1088
23	3395	3371	3364	3364	0.16	288.36	530	37
24	2589	2574	2569	2569	4.47	TL	10328	892
25	2430	2383	2383	2383	4.59	TL	10671	1257
26	2882	2849	2846	2846	0.76	TL	16605	1961
27	2525	-	2523	2523	0	34.43	105	1
28	2832	2760	2760	2760	0.04	3,875.62	532	3
29	2917	2902	2902	2901	6.53	TL	16330	2352
30	2714	-	2695	2695	0	57.43	116	1
31	3994	3941	3941	3941	-	TL	2581	0
32	4016	3967	3967	3967	-	TL	1242	0
33	3833	3640	3640	3640	-	TL	2003	0
34	4070	3873	3873	3873	-	TL	2185	0
35	3294	3238	3238	3238	-	TL	2663	0
36	4611	4527	4527	4527	-	TL	1165	0
37	3586	3397	3397	3390	-	TL	1872	0
38	4132	3986	3986	3978	-	TL	1554	0
39	4691	4318	4318	4318	2.29	TL	1766	10
40	3981	3649	3649	3649	-	TL	2528	0

Table 7: Comparison of the VRP, the VRPRDL, and the VRPHRDL solutions

Instance	VRP			VRPRDL			VRPHRDL			Savings wrt best VRP solution (%)
	Routes	Best solution	Home/total	Routes	Best solution	Home/total	Routes	Best solution		
1	3	864	0.47	4	901	0.60	3	773	10.53	
2	4	1187	0.47	5	1286	0.73	4	1065	10.28	
3	5	1307	0.60	4	991	0.67	3	989	24.33	
4	3	975	0.53	5	1062	0.87	3	915	6.15	
5	7	2171	0.53	6	1833	0.67	6	1710	21.23	
6	4	1247	0.60	5	1294	0.75	4	1099	11.87	
7	4	1156	0.55	4	1156	0.85	3	996	13.84	
8	5	1637	0.60	6	1455	0.70	5	1346	17.78	
9	4	1216	0.35	5	1260	0.60	4	997	18.01	
10	6	1444	0.60	8	1684	0.85	4	1166	19.25	
11	8	2493	0.47	7	1922	0.67	5	1587	36.34	
12	6	2003	0.50	8	2326	0.67	6	1809	9.69	
13	6	1774	0.60	6	1749	0.73	6	1563	11.89	
14	5	1649	0.50	6	1273	0.80	4	1058	35.84	
15	6	1938	0.47	6	1696	0.60	5	1347	30.50	
16	6	1892	0.47	7	1940	0.80	5	1521*	19.61	
17	7	2203	0.33	8	1968	0.57	5	1445	34.41	
18	6	1837	0.43	7	1827	0.73	5	1628	11.38	
19	6	1890	0.70	7	2083	0.83	5	1461	22.70	
20	6	1823	0.83	6	1822	0.83	6	1715	5.92	
21	9	3049	0.47	13	3762	0.82	8	2580	15.38	
22	8	2542*	0.58	10	2830	0.78	7	2214*	12.90	
23	12	3954	0.48	16	4448	0.87	10	3364	14.92	
24	9	2926	0.43	11	3383	0.80	8	2569*	12.20	
25	8	2649*	0.47	11	3162	0.77	8	2383*	10.04	
26	11	3665	0.48	16	4538	0.75	9	2846*	22.35	
27	11	3444*	0.60	10	2868	0.75	8	2523	26.74	
28	10	3425	0.55	14	4174	0.85	8	2760	19.42	
29	12	4022*	0.50	14	3968	0.75	9	2901*	27.87	
30	12	3930	0.38	14	4112	0.77	8	2695	31.42	
31	16	4926	0.55	18	4943	0.78	14	3941*	20.00	
32	15	4601*	0.48	19	5276*	0.78	14	3967*	13.78	
33	14	4306*	0.55	18	5100*	0.75	13	3640*	15.47	
34	14	4596*	0.53	17	5225	0.82	13	3873*	15.73	
35	12	3855*	0.48	20	5545*	0.88	10	3238*	16.01	
36	18	5758*	0.52	22	6506	0.83	15	4527*	21.38	
37	14	4825*	0.50	17	4836	0.73	11	3390*	29.74	
38	15	4408*	0.54	21	5618*	0.80	15	3978*	9.75	
39	22	6040*	0.45	23	5857*	0.84	16	4318*	28.51	
40	12	3835*	0.45	19	5017*	0.83	13	3649*	4.85	
Average	9.03	2836.55	0.51	11.08	3067.40	0.76	7.70	2288.65	18.50	

This is, in a large part, because fewer delivery routes are required; on average 9.03 for the VRP solutions and on average 7.70 for the VRPHRDL solutions. Interestingly, the average number of delivery routes in the VRPRDL solutions is 11.08. Even though there are more delivery locations to choose from, the time during which deliveries can be made is less, because deliveries cannot be made while the car is driving. Also interesting to note is that the fraction of deliveries made at a customer’s home location in VRPHRDL solutions is quite high, on average 0.76.

The difference in VRP, VRPRDL, and VRPHRDL solutions can be seen quite well in Figure 1, where we show the optimal VRP, VRPRDL, and VRPHRDL solutions for Instance 28. The empty rectangles represent the home locations while the filled circles represent the other potential delivery locations. The optimal VRP solution has 10 delivery routes and a cost of 3425, whereas the number of delivery routes in the optimal VRPRDL solution is 14 with cost 4174 (55% of the deliveries are made at home locations), and the number of delivery routes in the optimal VRPHRDL solutions is only 8 with cost 2760 (85% of the deliveries are made at home locations).

6 Concluding remarks

Trunk delivery is one of the innovative ideas being explored to reduce delivery cost in the business-to-consumer retail market sector. Our computational study shows that a hybrid model in which a delivery can be made either at a customer’s home or to the trunk of the customer’s car has huge potential. On the instances in our test suite, the cost savings were in the order of 20%.

We have assumed deterministic travel times and complete knowledge of the itinerary of a customer (more specifically the itinerary of the customer’s car). These are strong assumptions, but they are, in our opinion, not completely unreasonable. There are well-funded and highly-regarded start-up companies, e.g., Roadie (www.roadie.com), with a business proposition that is entirely based on the assumption that in the future, there will be reliable and predictable information on the travel patterns of people (based on 24-hour monitoring of the location of their smart phone and predictive analytics). Starting to explore new business models assuming this information is available is important, and we are among the first to do so.

However, we recognize that an important next step is to investigate operational strategies and related optimization models that can dynamically handle deviations from the planned customer itineraries. This is precisely what we are currently pursuing.

Acknowledgements

We would like to thank Luis Damian Reyes Rodriguez for generating new test instances and providing associated heuristic solutions. The research of the first author is supported by




Figure 1: The optimal VRP, VRPRDL and VRPHRDL solutions for Instance 28 from top to bottom, respectively

the Scientific and Technological Research Council of Turkey, grant number BIDEB-2214-A.

References

- H Murat Afsar, Christian Prins, and Andréa Cynthia Santos. Exact and heuristic algorithms for solving the generalized vehicle routing problem with flexible fleet size. *International Transactions in Operational Research*, 21(1):153–175, 2014.
- Agostinho Agra, Marielle Christiansen, Rosa Figueiredo, Lars Magnus Hvattum, Michael Poss, and Cristina Requejo. The robust vehicle routing problem with time windows. *Computers & Operations Research*, 40(3):856–866, 2013.
- Audi. Audi, dhl and amazon deliver convenience. <https://www.audiusa.com/newsroom/news/press-releases/2015/04/audi-dhl-and-amazon-deliver-convenience>, 2015. Accessed: 2016-07-24.
- Roberto Baldacci, Enrico Bartolini, and Gilbert Laporte. Some applications of the generalized vehicle routing problem. *Journal of the Operational Research Society*, 61(7):1072–1077, 2010.
- Cynthia Barnhart, Ellis L Johnson, George L Nemhauser, Martin WP Savelsbergh, and Pamela H Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46(3):316–329, 1998.
- Elisabeth Behrmann and Richard Weiss. Soon your smart car will also be an amazon locker. <http://www.bloomberg.com/news/articles/2016-07-25/smart-city-cars-to-become-delivery-stations-in-dhl-german-test>, 2016. Accessed: 2016-08-07.
- Tolga Bektas, Günes Erdogan, and Stefan Røpke. Formulations and branch-and-cut algorithms for the generalized vehicle routing problem. *Transportation Science*, 45(3):299–316, 2011.
- Benjamin Biesinger, Bin Hu, and Günther Raidl. An integer l-shaped method for the generalized vehicle routing problem with stochastic demands. *Electronic Notes in Discrete Mathematics*, 52:245–252, 2016.
- Natashia Boland, John Dethridge, and Irina Dumitrescu. Accelerated label setting algorithms for the elementary resource constrained shortest path problem. *Operations Research Letters*, 34(1):58–68, 2006.
- Christian H Christiansen and Jens Lygaard. A branch-and-price algorithm for the capacitated vehicle routing problem with stochastic demands. *Operations Research Letters*, 35(6):773–781, 2007.

- Said Dabia, Stefan Ropke, Tom Van Woensel, and Ton De Kok. Branch and price for the time-dependent vehicle routing problem with time windows. *Transportation Science*, 47(3):380–396, 2013.
- Mauro Dell’Amico, Giovanni Righini, and Matteo Salani. A branch-and-price approach to the vehicle routing problem with simultaneous distribution and collection. *Transportation Science*, 40(2):235–247, 2006.
- Martin Desrochers, Jan K Lenstra, Martin WP Savelsbergh, and François Soumis. Vehicle routing with time windows: optimization and approximation. *Vehicle routing: Methods and studies*, 16:65–84, 1988.
- Martin Desrochers, Jacques Desrosiers, and Marius Solomon. A new optimization algorithm for the vehicle routing problem with time windows. *Operations Research*, 40(2):342–354, 1992.
- Madeline Farber. Consumers are now doing most of their shopping online. <http://fortune.com/2016/06/08/online-shopping-increases/>, 2016. Accessed: 2016-07-24.
- Dominique Feillet. A tutorial on column generation and branch-and-price for vehicle routing problems. *4OR*, 8(4):407–424, 2010.
- Dominique Feillet, Pierre Dejax, Michel Gendreau, and Cyrille Gueguen. An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks*, 44(3):216–229, 2004.
- Ahiza Garcia. Amazon prime day shattered global sales records. <http://money.cnn.com/2015/07/15/news/amazon-walmart-sales/>, 2015. Accessed: 2015-09-22.
- Megan Geuss. Amazon, audi and dhl want to turn a car trunk into a delivery locker. <http://arstechnica.com/business/2015/04/amazon-audi-and-dhl-want-to-turn-a-car-trunk-into-a-delivery-locker>, 2015. Accessed: 2016-08-07.
- Gianpaolo Ghiani and Gennaro Improta. An efficient transformation of the generalized vehicle routing problem. *European Journal of Operational Research*, 122(1):11–17, 2000.
- Veaceslav Ghilas, Emrah Demir, and Tom Van Woensel. A scenario-based planning for the pickup and delivery problem with time windows, scheduled lines and stochastic demands. *Transportation Research Part B: Methodological*, 91:34–51, 2016.
- Krystina Gustafson. Amazon just had its biggest sales day ever. <http://www.cnbc.com/2016/07/13/amazon-prime-day-is-biggest-day-for-online-retailer-ever.html>, 2016. Accessed: 2016-07-24.

- Çağrı Koç, Tolga Bektaş, Ola Jabali, and Gilbert Laporte. A hybrid evolutionary algorithm for heterogeneous fleet vehicle routing problems with time windows. *Computers & Operations Research*, 64:11–27, 2015.
- Attila A Kovacs, Bruce L Golden, Richard F Hartl, and Sophie N Parragh. The generalized consistent vehicle routing problem. *Transportation Science*, 49(4):796–816, 2014.
- Amal Louati et al. Modeling municipal solid waste collection: A generalized vehicle routing model with multiple transfer stations, gather sites and inhomogeneous vehicles in time windows. *Waste Management*, 52:34–49, 2016.
- Snežana Mitrović-Minić, Ramesh Krishnamurti, and Gilbert Laporte. Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B: Methodological*, 38(8):669–685, 2004.
- Luigi Moccia, Jean-François Cordeau, and Gilbert Laporte. An incremental tabu search heuristic for the generalized vehicle routing problem with time windows. *Journal of the Operational Research Society*, 63(2):232–244, 2012.
- Ben Popken. Amazon tests delivery to your car trunk. <http://www.nbcnews.com/business/autos/amazon-testing-delivery-your-car-trunk-n346886>, 2015. Accessed: 2016-08-07.
- Nils-Hassan Quttineh, Torbjörn Larsson, Jorne Van den Bergh, and Jeroen Beliën. A time-indexed generalized vehicle routing model and stabilized column generation for military aircraft mission planning. In *Optimization, Control, and Applications in the Information Age*, pages 299–314. Springer, 2015.
- Damián Reyes, MWP Savelsbergh, and Alejandro Toriello. Vehicle routing with roaming delivery locations. *Optimization Online*, pages 01–5281, 2016.
- Matteo Salani and Ilaria Vacca. Branch and price for the vehicle routing problem with discrete split deliveries and time windows. *European Journal of Operational Research*, 213(3):470–477, 2011.
- M.W.P. Savelsbergh. Local search in routing problems with time windows. *Annals of Operations research*, 4(1):285–305, 1985.
- Michael Schneider, Andreas Stenger, and Dominik Goeke. The electric vehicle-routing problem with time windows and recharging stations. *Transportation Science*, 48(4):500–520, 2014.
- M. M. Solomon. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2):254–265, 1987.

Mark B. Solomon. Amazon's first-quarter shipping costs hit \$3.27 billion, 42-percent jump from 2015. <http://www.dcvelocity.com/articles/20160428-amazons-first-quarter-shipping-costs-hit-327-billion-42-percent-jump-from-2015>, 2016. Accessed: 2016-08-08.

D Taş, M Gendreau, N Dellaert, Tom Van Woensel, and AG De Kok. Vehicle routing with soft time windows and stochastic travel times: A column generation and branch-and-price solution approach. *European Journal of Operational Research*, 236(3):789–799, 2014.

Volvo. Volvo in-car delivery. <https://incardelivery.volvocars.com>, 2015.

Volvo. Volvo's solution for the package theft epidemic: Your car's trunk. <http://fortune.com/2016/05/10/volvo-urb-it-delivery>, 2016.