

Branch-and-Price for the Active-Passive Vehicle-Routing Problem

Christian Tilk^a, Nicola Bianchessi^{a,b}, Michael Drexl^{*,a,c}, Stefan Irnich^a, Frank Meisel^d

^a*Chair of Logistics Management, Gutenberg School of Management and Economics, Johannes Gutenberg University, Jakob-Welder-Weg 9, D-55128 Mainz, Germany.*

^b*Department of Information Engineering, University of Brescia, Italy.*

^c*Fraunhofer Centre for Applied Research on Supply Chain Services SCS, Nuremberg, Germany.*

^d*Chair of Supply Chain Management, Faculty of Business, Economics and Social Sciences, University Kiel, Germany.*

Abstract

This paper presents a branch-and-price algorithm for the exact solution of the active-passive vehicle-routing problem (APVRP). The APVRP covers a range of logistics applications where pickup-and-delivery requests necessitate a joint operation of active vehicles (e.g., trucks) and passive vehicles (e.g., loading devices such as containers or swap bodies). The problem supports a flexible coupling and decoupling of active and passive vehicles at customer locations in order to achieve a high utilization of both resources. Accordingly, the operations of the vehicles have to be synchronized carefully in the planning. The contribution of the paper is twofold: Firstly, we present an exact branch-and-price algorithm for this class of routing problems with synchronization constraints. To our knowledge, this algorithm is the first such approach that considers explicitly the temporal interdependencies between active and passive vehicles. The algorithm is based on a non-trivial network representation that models the logical relationships between the different transport tasks necessary to fulfill a request as well as the synchronization of the movements of active and passive vehicles. Secondly, we contribute to the development of branch-and-price methods in general, in that we solve, for the first time, an ng-path relaxation of a pricing problem with linear vertex costs by means of a bidirectional labeling algorithm. Computational experiments show that the proposed algorithm delivers improved bounds and solutions for a number of APVRP benchmark instances. It is able to solve instances with up to 76 tasks, 4 active, and 8 passive vehicles to optimality within two hours of CPU time.

Key words: Vehicle-routing, Synchronization, Branch-and-price, Linear vertex costs

1. Introduction

Many applications in the area of transport logistics involve problems in which the execution of transport requests calls for a joint operation of different resources such as trucks, tractors, drivers, trailers, semi-trailers, swap bodies, containers, or accompanying staff. A typical example is found in the transportation of containerized goods, where not just a manned truck but also an empty container is required for executing a request. Further examples are found in the health care sector or in the security industry, where the transportation of patients and valuable items must be accompanied by medics and security guards respectively. In general, we can distinguish two classes of transport resources. The first class is constituted by means of transport that can move on their own from one location to another such as, for example, manned trucks. We refer to these resources as *active vehicles*. The second class consists of resources that cannot move autonomously but require an active vehicle for being repositioned. This class comprises trailers, semi-trailers, all kinds of loading devices, equipment, and accompanying staff. For simplicity, we refer to all these resources jointly as *passive vehicles*.

*Corresponding author.

Email address: drex1@uni-mainz.de (Michael Drexl)

In the literature on vehicle-routing problems (VRPs, see Irnich *et al.*, 2014, for an overview), the distinction of active and passive vehicles is usually ignored, and operations are planned for active vehicles only (cf. the recent survey by Lahyani *et al.*, 2015). This restricts the applicability of the developed models and algorithms to real-world problems in which either passive vehicles do not play a role or where active and passive vehicles are paired to fixed units. Although the latter eases the solution of the VRP, it may hinder an effective utilization of the resources. If, for example, a manned truck (active vehicle) and an empty container (passive vehicle) are considered a unit in operations planning, the truck and its driver have to wait at a customer location while the container is being stuffed. Therefore, in order to support a more flexible use of such resources, Meisel and Kopfer (2014) introduced what we denote here as the active-passive vehicle-routing problem (APVRP), in which an explicit distinction between active and passive vehicles is made. This distinction enables the modeling of complex transport operations in which an active vehicle carries a passive vehicle (e.g., an empty container) to some pickup location, drops it off there, and leaves this location for performing transports of other passive vehicles elsewhere. Later, when the container has been stuffed, the same or some other active vehicle returns to the customer, picks up the container, and carries it to the delivery location. The problem introduces multiple interdependencies between vehicles and raises the need to synchronize the operations and the movements of active and passive vehicles in time and space. According to the survey by Drexl (2012), such a combination of synchronization requirements is rarely addressed in the VRP literature. This fact is in marked contrast to the abovementioned practical relevance of the APVRP. It therefore seems appropriate to devote further studies to this generic and complex problem.

Our paper addresses this research gap and provides a twofold contribution: Firstly, we present an exact algorithm for the APVRP which, to the best of our knowledge, is the first branch-and-price approach that considers explicitly the temporal interdependencies between active and passive vehicles and the resulting synchronization requirements. Because of these synchronization requirements, the adaptation of the branch-and-price concept to the APVRP is not straightforward. Our algorithm is based on an extended set-partitioning formulation which, in turn, uses a non-trivial network representation that models the logical relationships between the different transport tasks necessary to fulfill a request as well as the synchronization of the movements of active and passive vehicles. Secondly, we contribute to the development of branch-and-price methods for routing with synchronization in general: We provide solutions to the pricing subproblem, which is an elementary shortest-path problem with time windows and with linear vertex costs, by solving, for the first time, its ng-path relaxation (Baldacci *et al.*, 2011) by means of a bidirectional labeling algorithm. We actually apply a refined ng-path relaxation, taking into account partial requests (henceforth called *tasks*) and precedences between these tasks instead of individual vertices or complete requests. The pricing problem structure and the use of the ng-path relaxation, moreover, require a sophisticated merge step in the labeling algorithm. Computational experiments show that the proposed algorithm delivers improved bounds and solutions for the APVRP benchmark suite of Meisel and Kopfer (2014), solving instances with up to 76 tasks, 4 active, and 8 passive vehicles to optimality within two hours of CPU time.

The paper is organized as follows. Related literature is reviewed in Section 2. In Section 3, we formally describe the APVRP. A corresponding extended set-partitioning formulation is provided in Section 4. The branch-and-price algorithm is presented in Section 5, and the method is computationally evaluated in Section 6. Finally, Section 7 summarizes the paper and discusses potential avenues for further research.

2. Literature

The manifold real-world logistics applications for vehicle-routing problems with synchronization requirements have motivated several studies. Most papers consider applications in which active

vehicles have to be synchronized. One example is the operations planning for cross-docks where different trucks deliver less-than-truckload shipments that are then merged to full-truckload shipments before being sent out to customers (see, e.g., Buijs *et al.*, 2014; Morais *et al.*, 2014). Another active research field is found in the management of home care operations. Here, service operations have to be synchronized if a person requires the help of two caregivers at the same time (see Bredström and Rönnqvist, 2008; Mankowska *et al.*, 2013; Labadie *et al.*, 2014; Afifi *et al.*, 2015). Further research addresses forestry applications where trucks have to be served by forest loaders (Hachemi *et al.*, 2013), intermodal transportation via ships, trains, and trucks for the supply of automotive factories (Mues and Pickl, 2005), scheduling of cooperating technician teams at customer locations (Dohn *et al.*, 2009), and ship routing where cargoes from different origins have to be delivered to one and the same client simultaneously (Andersson *et al.*, 2011). Eventually, there are also arc routing problems that include synchronization of active vehicles, e.g. when planning snow plowing operations (see Laporte, 2013) and road marking (see Salazar-Aguilar *et al.*, 2013).

The mentioned applications all involve active vehicles that move autonomously to those locations where synchronized services are required. By contrast, in the APVRP, an active vehicle and a passive vehicle must both traverse route segments synchronously in order to perform a service. Typical applications are found in routing problems where trucks pull trailers or swap bodies (see Smilowitz, 2006; Cheung *et al.*, 2008; Drexl, 2013). Obviously, trailers and swap bodies are passive vehicles that are immobile without a truck. Usually, they play the role of optional capacity extensions of trucks rather than being mandatory for the execution of transport operations. An exceptional case is when an individual customer has such a large demand that it must be served jointly by a truck with a swap body (see Huber and Geiger, 2014). A further application of such a synchronization requirement is found in the drayage operations of container terminals, where empty and loaded containers are moved between customer locations and transshipment points. In this field of logistics, the containers constitute the passive vehicles that are mandatory for the transport of goods (see Cheung *et al.*, 2008; Xue *et al.*, 2014; Zhang *et al.*, 2010, 2013, 2014). Movement synchronization en route is also found in VRPs where driver crews can be assigned flexibly to trucks (see Hollis *et al.*, 2006; Drexl *et al.*, 2013). In such problems, the exchange of crews enables a better utilization of the trucks in compliance with work regulations for truck drivers. Further applications of a synchronization of vehicles and crews are considered in the papers by Kim *et al.* (2010), where technicians have to be carried to customer locations, and by Kergosien *et al.* (2011, 2013), where ambulances carry patients and accompanying physicians from one care unit to another.

Drexl (2007, 2014) studies the VRP with trailers and transshipments (VRPTT), a problem which also requires the synchronization of operations and movements of active and passive vehicles. Two branch-and-cut algorithms are presented, but only very small instances can be solved. For a deeper investigation of VRPs with synchronization of operations and vehicle movements, Meisel and Kopfer (2014) provide mixed-integer programming (MIP) formulations, a branch-and-cut algorithm, an adaptive large neighborhood search (ALNS) metaheuristic, and benchmark instances for the APVRP. To our knowledge, the branch-and-price algorithm by Smilowitz (2006) is the only other exact method for the APVRP so far. To facilitate understanding, the differences between the algorithm of Smilowitz (2006) and ours as well as the extensions and improvements our approach provides will be discussed in Section 5.2.1, when the difficulties arising from the synchronization requirements will have been thoroughly explained.

A few other papers provide exact methods (typically based on column generation techniques) for VRPs with synchronization requirements, but merely for problems that involve active vehicles only (see Mues and Pickl, 2005; Dohn *et al.*, 2009, 2011; Andersson *et al.*, 2011). However, the features of the VRPTT and the APVRP, which include simultaneous operations planning of active and passive vehicles together with the possibility to couple and decouple them flexibly on their routes, are not supported by any of these approaches.

3. Problem description and modeling

The problem considered in this paper can be formally described as follows. We are given a set of pickup-and-delivery requests R , a set A of classes of active vehicles and a set P of passive vehicles. For each class $a \in A$ of active vehicles, K_a denotes the number of vehicles in the class. (Aggregating identical active vehicles into classes while at the same time considering individual passive vehicles is convenient when setting up the network on which our problem formulation is based, ensures that each network arc is traversed by at most one active and/or at most one passive vehicle in any feasible solution, and makes branching easier.) Each request $r \in R$ consists of transporting a loading unit from a pickup location ℓ_r^+ to a delivery location ℓ_r^- . To fulfill a request r , an active vehicle must carry a passive vehicle to ℓ_r^+ for loading. Each passive vehicle can load only one request at a time, and each active vehicle can transport only one passive vehicle at a time. Hence, the loaded passive vehicle must then be transported directly to ℓ_r^- for unloading. Afterwards, the empty passive vehicle must be carried away from ℓ_r^- . From the point of view of an active vehicle, the fulfillment of a request comprises three transport tasks:

- (i) providing an empty passive vehicle at the pickup location of a request,
- (ii) direct transport of a loaded passive vehicle from the pickup to the delivery location, and
- (iii) carrying away the emptied passive vehicle from the delivery location.

One or two or three different active vehicles may perform these tasks for a request.

There are compatibility relationships between the requests and the active and the passive vehicles. P^r denotes the set of passive vehicles that can be used to perform request r . Likewise, R^p denotes the set of requests that can be performed with passive vehicle p . P^a indicates the set of passive vehicles that can be carried by an active vehicle from class a , and A^p is the set of classes of active vehicles compatible with passive vehicle p .

All active vehicles are initially based at the same start depot o and end their routes at the same end depot d . Each passive vehicle p has its own start and end positions: It is initially located at o_p and must be brought to d_p at the end, whether or not it is used to fulfill a request. Hence, there are two tasks associated with each passive vehicle p , namely, to pickup and to deliver the vehicle at o_p and d_p respectively.

Let s_r^+ indicate the time necessary to load an empty passive vehicle with request r and s_r^- indicate the time to unload r . Times for coupling and uncoupling of passive vehicles to or from active vehicles are assumed to be zero (but could easily be incorporated into the model). Picking up a loaded request r at its pickup location can be finished no earlier than at time e_r . Waiting is allowed. Unloading a request r at its delivery location must be finished no later than l_r . The overall planning horizon is $[0, t^{max}]$. Requests that cannot be fulfilled imply a penalty, but may be executed in a later period.

The objective is to minimize a weighted sum of the total distance traveled, the total completion time of the routes, and the number of unserved requests. The respective weights are $\alpha, \beta, \gamma \in \mathbb{R}_+$.

The APVRP as described above can be modeled as an optimization problem over a set of graphs $G^a = (V^a, E^a)$, one for each class a of active vehicles. Each set V^a of vertices contains o and d , i.e., the initial and the final location of all active vehicles, o_p and d_p , i.e., the initial and the final locations of all passive vehicles $p \in P^a$, and the set N^a containing the following four vertices for each passive vehicle $p \in P^a$ and each compatible request $r \in R^p$:

- v_{rp}^- delivery of an empty passive vehicle p into which request r is loaded;
- w_{rp}^+ pickup of request r loaded in passive vehicle p ;
- w_{rp}^- delivery of request r loaded in passive vehicle p ;
- v_{rp}^+ pickup of the empty passive vehicle p in which request r was transported.

Thus $N^a = \cup_{p \in P^a, r \in R^p} \{v_{rp}^-, w_{rp}^+, w_{rp}^-, v_{rp}^+\}$. For each request $r \in R$, we define $V_r^- = \{v_{rp}^- : p \in P^r\}$, $W_r^+ = \{w_{rp}^+ : p \in P^r\}$, $W_r^- = \{w_{rp}^- : p \in P^r\}$, and $V_r^+ = \{v_{rp}^+ : p \in P^r\}$. The vertices in

$N^R = \bigcup_{r \in R} (V_r^- \cup W_r^+ \cup W_r^- \cup V_r^+)$ are henceforth referred to as *request vertices*.

Arcs in E^a between each pair of vertex types are represented in Figure 1. In the figure, solid bold arcs (\rightarrow) can be traversed by an active vehicle while traveling together with a passive vehicle attached. Dotted arcs ($\cdots\rightarrow$) are service arcs and are traversed by the same passive vehicle p alone or by p together with an active vehicle, while solid and dashed arcs (\rightarrow and $--\rightarrow$) are traversed by the active vehicle alone. In particular, all dashed arcs ($--\rightarrow$) exist only if $p' \neq p \neq p''$. For the four request vertices in the shaded rectangle, all types of ingoing and outgoing arcs are depicted. It is easy to see that the graphs constructed in this way allow an active vehicle to perform any of the three transport tasks of any compatible request. For example, assume that active vehicle a_1 performs the first task of request r , and that active vehicle a_2 performs the second and the third task of r , i.e., vehicle a_2 transports some passive vehicle p loaded with r from w_{rp}^+ to w_{rp}^- and afterwards moves away the empty passive vehicle. In this case, a_2 reaches w_{rp}^+ without a passive vehicle, coming from either o or from a vertex $v_{r'p'}^+$, $w_{r'p'}^-$, or $d_{p'}$, with $r' \neq r$ and p' different from p . (If $p' = p$, $d_{p'}$ is the final location of passive vehicle p , and an arc from $d_{p'}$ to w_{rp}^+ for some $r \in R$ is impossible because of the uniqueness of the passive vehicles.) After that, a_2 visits vertices w_{rp}^- and v_{rp}^+ and leaves v_{rp}^+ heading towards either d_p or towards a vertex $v_{r''p''}^+$ with $r \neq r' \neq r'' \neq r$. The travel distances between any pair $(i, j) \in E^a$ of vertices are denoted by c_{ij} and are given by the distance between the associated locations. The same holds for the travel and service time t_{ij} for $(i, j) \in E^a$ with t_{ij} equal to s_r^+ and s_r^- for arcs (v_{rp}^-, w_{rp}^+) and (w_{rp}^-, v_{rp}^+) respectively. Following the specifications given by Meisel and Kopfer (2014), the time windows at the vertices are given in Table 1. The complete notation used throughout the paper is also summarized in the Appendix.

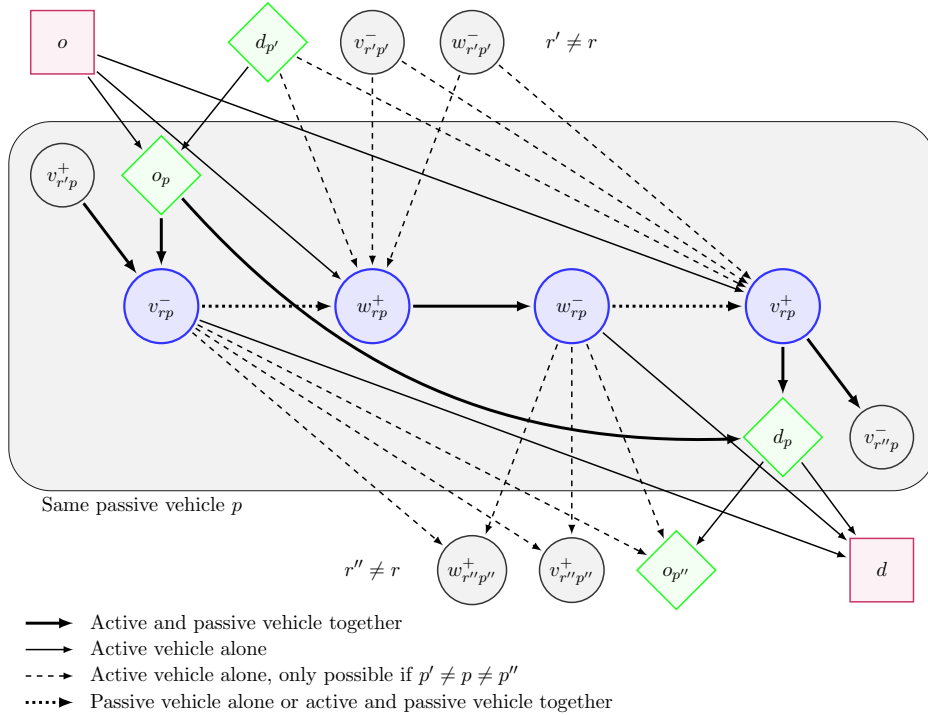


Figure 1: Arc set of graph $G^a = (V^a, E^a)$

The request vertices take up an idea described by Drexler (2007, Sect. 4.3.3). They essentially correspond to pairs of operations and passive vehicles. This ensures that the itineraries of the passive vehicles are implicit, i.e., need not be determined explicitly, but can be unequivocally reconstructed from given routes for the active vehicles. Hence, a feasible solution to the APVRP is a set of scheduled routes for the active vehicles that fulfills:

- Each route starts at vertex o and terminates at vertex d .

Vertex i	Earliest time e_i	Latest time l_i
v_{rp}^-	$t_{oo_p} + t_{o_p v_{rp}^-}$	$l_{w_{rp}^+} - (s_r^+ + t_{w_{rp}^+ w_{rp}^-} + s_r^-)$
w_{rp}^+	$\max\{e_{v_{rp}^-} + s_r^+, e_r\}$	$l_{w_{rp}^-} - (t_{w_{rp}^+ w_{rp}^-} + s_r^-)$
w_{rp}^-	$e_{w_{rp}^+} + t_{w_{rp}^+ w_{rp}^-}$	$\min\{l_{v_{rp}^+}, l_r\} - s_r^-$
v_{rp}^+	$e_{w_{rp}^-} + s_r^-$	$t^{max} - (t_{v_{rp}^+ d_p} + t_{d_p d})$
Other	0	t^{max}

Table 1: Vertex time windows

- All visited vertices are visited within their prescribed time windows.
- For each $r \in R$, at most one vertex per set is visited in sets V_r^- , W_r^+ , and V_r^+ . This ensures the synchronization of task fulfillment of requests within and between active vehicles. Note that it is not necessary to impose that at most one vertex is visited in set W_r^- , because vertices in this set can be reached only via an arc coming from a vertex in W_r^+ .
- For each $r \in R$, $p \in P^r$, vertex w_{rp}^+ is visited if and only if vertex v_{rp}^- has been visited, and vertex v_{rp}^+ is visited if and only if vertex w_{rp}^+ has been visited. This requirement is necessary to preserve consistency with respect to the passive vehicle used to satisfy a request in its different stages. Moreover, given that the final locations d_p of any passive vehicle $p \in P$ can only be reached from the corresponding initial location o_p or from vertices in $\bigcup_{r \in R} V_r^+$, this requirement preserves the flow of the unique passive vehicle used to satisfy a request through vertices v_{rp}^- , w_{rp}^+ , w_{rp}^- , and v_{rp}^+ .
- $T_{v_r^-} + s_r^+ \leq T_{w_r^+}$ and $T_{w_r^-} + s_r^- \leq T_{v_r^+}$ for each $r \in R$, where $T_{v_r^-}$, $T_{w_r^+}$, $T_{w_r^-}$, and $T_{v_r^+}$ are the service start times at the vertices possibly selected in sets V_r^- , W_r^+ , W_r^- , and V_r^+ respectively. This ensures temporal synchronization of tasks within and between active vehicles.
- Each passive vehicle is picked up at its initial location and placed at its final location.
- Each active vehicle performs at most one feasible route (the trivial route from o to d does not exist), so that at most K_a routes are performed for each class a of active vehicles.

The first two requirements are intra-route constraints, the subsequent four represent both intra-route and inter-route constraints, and the last one is an inter-route constraint.

Note that, along a route, an active vehicle of class $a \in A$ can be associated with different passive vehicles $p \in P^a$. This is due to the structure of the graphs $G^a = (V^a, E^a)$. Moreover, a passive vehicle $p \in P$ can be associated with different routes traveled by different active vehicles.

4. An extended set-partitioning formulation

In order to solve the APVRP with branch-and-price, we use an extended set-partitioning formulation. This extensive formulation can be derived from a compact APVRP model, e.g., the one by Meisel and Kopfer (2014), using a Dantzig-Wolfe reformulation for integer programs (see Desaulniers *et al.*, 1998; Lübbecke and Desrosiers, 2005). For the sake of brevity, though, we omit the formal derivation.

For each class $a \in A$ of active vehicles, let Ω^a be the set of all feasibly scheduled routes. In contrast to many other branch-and-price algorithms for vehicle routing described in the literature, a column in the APVRP does not only represent a path, i.e., the sequence of visited vertices. A column in the APVRP provides a path (in the graph $G^a = (V^a, E^a)$) and a feasible schedule for this path. For simplicity, however, we will refer to *routes* in the following, but use the terms *path* and *schedule* to describe the routing and the scheduling components. The following attributes characterize a route q :

- X_{ij}^q number of times arc (i, j) is traversed by route q ;
- T_i^q service start time at a vertex $i \in N^R \cup \{d\}$;

b_i^q	number of times vertex i is visited on route q ;
c^q	total costs of route q .

The route costs are defined as $c^q = \alpha \sum_{(i,j) \in E^a} c_{ij} X_{ij}^q + \beta T_d^q$, i.e., they are a weighted sum of the length of the route and the arrival time at the destination. Note that service start times are well-defined because a feasible route is elementary.

The formulation uses the following types of variables: continuous variables λ^{aq} measuring the flow of active vehicles of class $a \in A$ along route $q \in \Omega^a$, binary variables x_{ij}^a indicating whether or not arc $(i, j) \in E^a$ is traversed by an active vehicle of class $a \in A$, and binary variables u_r indicating whether or not request $r \in R$ remains unfulfilled. Recall that the penalty for not performing a request is γ . Now, the extended set-partitioning formulation for the APVRP is as follows:

$$\min \sum_{a \in A} \sum_{q \in \Omega^a} c^q \lambda^{aq} + \gamma \sum_{r \in R} u_r \quad (1a)$$

$$\text{s.t.} \quad \sum_{a \in A} \sum_{q \in \Omega^a} \sum_{p \in P^r \cap P^a} b_{v_{rp}}^q \lambda^{aq} + u_r = 1 \quad r \in R \quad (1b)$$

$$\sum_{a \in A^p} \sum_{q \in \Omega^a} \left(b_{v_{rp}}^q - b_{w_{rp}}^q \right) \lambda^{aq} = 0 \quad r \in R, p \in P^r \quad (1c)$$

$$\sum_{a \in A^p} \sum_{q \in \Omega^a} \left(b_{w_{rp}}^q - b_{v_{rp}}^q \right) \lambda^{aq} = 0 \quad r \in R, p \in P^r \quad (1d)$$

$$\sum_{a \in A} \sum_{q \in \Omega^a} \sum_{p \in P^r \cap P^a} \left(T_{w_{rp}}^q - T_{v_{rp}}^q \right) \lambda^{aq} + s_r^+ u_r \geq s_r^+ \quad r \in R \quad (1e)$$

$$\sum_{a \in A} \sum_{q \in \Omega^a} \sum_{p \in P^r \cap P^a} \left(T_{v_{rp}}^q - T_{w_{rp}}^q \right) \lambda^{aq} + s_r^- u_r \geq s_r^- \quad r \in R \quad (1f)$$

$$\sum_{a \in A^p} \sum_{q \in \Omega^a} b_{o_p}^q \lambda^{aq} = 1 \quad p \in P \quad (1g)$$

$$\sum_{q \in \Omega^a} \lambda^{aq} \leq K_a \quad a \in A \quad (1h)$$

$$x_{ij}^a = \sum_{q \in \Omega^a} X_{ij}^q \lambda^{aq} \quad a \in A, (i, j) \in E^a \quad (1i)$$

$$\lambda^{aq} \geq 0 \quad a \in A, q \in \Omega^a \quad (1j)$$

$$x_{ij}^a \in \{0, 1\} \quad a \in A, (i, j) \in E^a \quad (1k)$$

$$u_r \in \{0, 1\} \quad r \in R \quad (1l)$$

(1a) is the objective function. (1b) are the set partitioning constraints, which ensure that each request is either performed exactly once or the penalty γ is paid for leaving the request unfulfilled. (1c) and (1d) preserve consistency with respect to the passive vehicle used to satisfy a request in its different stages. (1e) and (1f) are time synchronization constraints ensuring that the different vertices corresponding to a request are visited at correct points in time, taking into account their precedence relationships. (1g) and (1h) are quantity constraints for passive vehicles and for classes of active vehicles respectively. Constraints (1g) are valid only if it is not allowed to temporarily park a passive vehicle at its initial or final location. (1i) link the route and the arc variables. (1j)–(1l) indicate the domains of the variables.

Note that there are no binary restrictions on the λ^{aq} variables. The reason for this is that a solution may be fractional in terms of λ^{aq} variables for a class of active vehicles. The λ^{aq} variables may correspond to the same path with different schedules, i.e., to the same sequence of vertices

visited at different points in time, so that the arc variables are integral and no branching is necessary. This issue is discussed in detail by Desaulniers *et al.* (1998) and Jans (2010).

In the following, the linear relaxation of formulation (1) is denoted as the *master program*. It does not contain the coupling constraints (1i) and the integer constraints (1k)–(1l). For solving the master program, a column generation algorithm (Desaulniers *et al.*, 2005) is employed. Branching is required to finally ensure integer solutions of (1).

5. A branch-and-price algorithm

In this section, we present the branch-and-price algorithm we devised for solving the APVRP. First, in Subsection 5.1, we give a MIP model of the pricing problems for generating new columns. In Subsection 5.2, we describe a labeling algorithm for actually solving these pricing problems. Then, in Subsection 5.3, we describe our branching strategy.

5.1. Pricing problems

As is common for branch-and-price algorithms for vehicle-routing problems, the pricing problems for the APVRP are shortest-path problems with resource constraints (SPPRCs) on graphs with negative cost cycles (Irnich and Desaulniers, 2005). For each class $a \in A$ of active vehicles, there is one such pricing problem. Its goal is to find at least one route with negative reduced costs or to prove that no such route exists. The dual variables of those constraints in which the λ^{aq} variables occur in the restricted master program, the resulting linear vertex costs, and the resulting reduced costs of the arcs can be read from Table 2.

Constraint	Dual variable	Range	Vertex i	Linear vertex costs \tilde{c}_i	Arc (i, j)	Reduced costs \tilde{c}_{ij}
(1b)	π_r	\mathbb{R}	v_{rp}^-	τ_r^+	(i, o_p)	$\alpha c_{io_p} - \mu_p^P$
(1c)	ϕ_{rp}^+	\mathbb{R}	w_{rp}^+	$-\tau_r^+$	(i, d_p)	αc_{id_p}
(1d)	ϕ_{rp}^-	\mathbb{R}	w_{rp}^-	τ_r^-	(i, v_{rp}^-)	$\alpha c_{iv_{rp}^-} - \pi_r - \phi_{rp}^+$
(1e)	τ_r^+	≥ 0	v_{rp}^+	$-\tau_r^-$	(i, w_{rp}^+)	$\alpha c_{iw_{rp}^+} + \phi_{rp}^+$
(1f)	τ_r^-	≥ 0	d	β	(i, w_{rp}^-)	$\alpha c_{iw_{rp}^-} - \phi_{rp}^-$
(1g)	μ_p^P	\mathbb{R}	otherwise	0	(i, v_{rp}^+)	$\alpha c_{iv_{rp}^+} + \phi_{rp}^-$
(1h)	μ_a^A	≤ 0			(i, d)	$\alpha c_{id} - \mu_a^A$

(a)
(b)
(c)

Table 2: (a) Dual variables and their ranges; (b) Linear vertex costs; (c) Reduced costs of arcs

We use binary variables x_{ij} indicating whether or not arc $(i, j) \in E^a$ is traversed and continuous T_i variables indicating the point in time when the service at vertex $i \in V^a$ begins. The symbols $\delta^+(i)$ and $\delta^-(i)$ denote the forward and backward star of vertex i respectively. The pricing problem for a class a of active vehicles can be specified as:

$$\min \sum_{(i,j) \in E^a} \tilde{c}_{ij} x_{ij} + \sum_{i \in N^a} \tilde{c}_i T_i + \beta T_d \quad (2a)$$

$$\text{s.t.} \quad \sum_{(o,j) \in \delta^+(o)} x_{oj} = 1 \quad (2b)$$

$$\sum_{(j,i) \in \delta^-(i)} x_{ji} - \sum_{(i,j) \in \delta^+(i)} x_{ij} = 0 \quad i \in V^a \setminus \{o, d\} \quad (2c)$$

$$\sum_{(i,d) \in \delta^-(d)} x_{id} = 1 \quad (2d)$$

$$T_i + t_{ij} \leq T_j + t^{max}(1 - x_{ij}) \quad (i, j) \in E^a \quad (2e)$$

$$e_i \leq T_i \leq l_i \quad i \in \{o, d\} \quad (2f)$$

$$e_i \sum_{(i,j) \in \delta^+(i)} x_{ij} \leq T_i \leq l_i \sum_{(i,j) \in \delta^+(i)} x_{ij} \quad i \in V^a \setminus \{o, d\} \quad (2g)$$

$$x_{ij} \in \{0, 1\} \quad (i, j) \in E^a \quad (2h)$$

(2a) is the objective of minimizing the reduced costs of the route. (2b)–(2d) are path-flow constraints and (2e)–(2g) ensure time-window feasibility. All arc variables are binary, as imposed by (2h).

5.2. Dynamic programming labeling algorithms

The usual solution approach for SPPRC pricing problems is a dynamic-programming based labeling algorithm. As has been shown by Salani (2005), solving such pricing problems by bidirectional dynamic programming, i.e., by propagating labels forward from the start depot vertex and backward from the end depot vertex, allows considerable speedups compared to a unidirectional procedure. We therefore adopt this approach and adapt it to the APVRP pricing problem.

In what follows, we first point out the specific characteristics resulting from the linear vertex costs in the APVRP pricing problem and briefly review similar problems and solution approaches encountered in the literature. We then discuss issues related to (non-)elementarity of SPPRC solutions and present our adaptation of the ng-path relaxation. Afterwards, we describe the forward label extension step and the dominance procedure we apply. Then, we elaborate on the backward label extension and the method for merging forward and backward labels, and, finally, we present some techniques for accelerating the pricing process.

5.2.1. Labeling algorithms for SPPRCs with linear vertex costs

Due to the linear vertex costs \tilde{c}_i in model (2), which are not present in pricing problems for standard VRPs, there is a tradeoff in the APVRP pricing problem: In a path in which the linear vertex costs are non-negative at all vertices, it is best to visit all vertices as early as possible. Similarly, in a path in which the linear vertex costs are non-positive at all vertices, it is best to visit all vertices as late as possible. However, paths in the APVRP may visit vertices with positive as well as vertices with negative vertex costs. Hence, determining a cost-optimal schedule even for a *given* path is a non-trivial optimization problem in itself. For a labeling algorithm, this means that the linear vertex costs and the resulting tradeoff between costs and time create an infinite number of possible states that do not dominate one another. Such a situation was first studied by Ioachim *et al.* (1998) on an acyclic time-space network in the context of aircraft fleet routing and scheduling. Their approach of storing cost functions (with time as the independent variable) as labels and propagating these functions instead of scalar values forms the basis of our method, which is described in detail in Section 5.2.3.

Other authors that consider shortest-path pricing problems with linear vertex costs in branch-and-price algorithms for routing problems are Christiansen and Nygreen (1998) (for a ship routing and scheduling problem with inventory constraints), Dohn *et al.* (2009, 2011) (for VRPs with precedences and temporal dependencies between visits to customers), Liberatore *et al.* (2011) (for a VRP with soft time windows) and Spliet and Gabor (2014) (for a VRP with time windows that have to be assigned before customer demand is known). The pricing problem of Christiansen and Nygreen (1998) is solved with the Ioachim *et al.* (1998) algorithm. Dohn *et al.* (2009, 2011) describe two alternative approaches for dealing with the linear vertex costs: They consider a so-called time-indexed formulation with an implicit representation of precedence constraints and the possibility of branching on time windows to handle temporal dependencies. Liberatore *et al.* (2011) and Spliet and Gabor (2014) also base their procedures on the Ioachim *et al.* (1998) approach.

The main differences between the problems and solution procedures presented by Ioachim *et al.* (1998), Liberatore *et al.* (2011), Spliet and Gabor (2014), and ours are as follows: Ioachim *et al.* (1998) consider an acyclic network, all others have networks with cycles. Due to the soft time windows, the Liberatore *et al.* (2011) problem is, on the one hand, simpler with respect to time windows, as the complete planning horizon is feasible. On the other hand, it is more difficult with respect to the cost (penalty) function *at each single vertex*: This penalty function is decreasing until the beginning of the desired time window, then constant until the end of the desired time window, and then increasing. By contrast, the cost functions in the papers by Ioachim *et al.* (1998), Spliet and Gabor (2014) as well as in our paper are either constant or only increasing or only decreasing at each vertex.

Both the Liberatore *et al.* (2011) and the Spliet and Gabor (2014) pricing problem feature a capacity constraint, which is not present in the Ioachim *et al.* (1998) problem or in the APVRP. Liberatore *et al.* (2011) handle the capacity constraint in the dominance procedure by allowing a label L to dominate another label L' only if the load of L is less than or equal to that of L' . Spliet and Gabor (2014) construct an auxiliary acyclic graph in which the capacity constraint is taken into account implicitly at the cost of a weaker dominance.

The branch-and-price algorithm by Smilowitz (2006) mentioned in the literature review section, although considering an APVRP use case, does not take into account the time synchronization aspect in an exact manner, contrary to our approach. Instead, the interdependencies between tasks are modeled away by assuming that a passive vehicle which is used to perform a task τ is available for performing other tasks at the earliest after the end of τ 's time window plus the service time for τ . In this way, the underlying network does not contain arcs between tasks with overlapping time windows, no linear vertex costs occur, and standard labeling approaches can be used for solving the pricing problem.

5.2.2. Elementarity and precedences

In an optimal solution to the set-partitioning formulation (1), no task will be performed more than once. This implies that all columns in an optimal solution correspond to elementary paths in the network. It is well known that the elementary shortest-path problem with resource constraints is NP-hard in the strong sense (Dror, 1994). Hence, in branch-and-price algorithms for VRPs, many authors solve as pricing problems non-elementary SPPRCs (see the survey by Desaulniers *et al.*, 2014). This can be done in pseudo-polynomial time (Irnich and Desaulniers, 2005). Although this yields weaker lower bounds, and although routes with cycles must be removed in the branching process, solving only a relaxed pricing problem often pays off with respect to overall computation time. One such approach that has been very successfully used for different types of VRPs is the ng-path relaxation introduced by Baldacci *et al.* (2011).

We adapt this approach to the APVRP as follows. Our ng-path relaxation is based on tasks instead of vertices, thus leading to a stronger relaxation on the network we use. For each vertex $i \in N^a \cup \{o_p, d_p : p \in P^a\}$, we use an ng-neighborhood \mathcal{N}_i containing the task associated to i (see Table 3) and the ν closest tasks associated to vertices j for which a cycle (j, \dots, i, \dots, j) would be feasible with respect to time windows and travel and service times. (We use different values of the parameter ν in our computational experiments.) In an ng-path in the APVRP, it is possible that a task τ is performed more than once if, between two visits to a vertex i associated with the task, at least one other vertex j is visited such that $\tau \notin \mathcal{N}_j$.

In addition, to foster elementarity of partial paths without weakening dominance, we use the following approach: Consider the second line of Table 3. As described in Section 3, we associate three tasks with each request, and two tasks with each passive vehicle, for pickup and delivery at the initial and final location. Obviously, these tasks must be performed in the chronological sequences $\tau_r^1, \tau_r^2, \tau_r^3$ and τ_p^o, τ_p^d (but not necessarily by the same active vehicle). Now, at each request vertex i , we consider two subsets of these tasks, $\mathcal{T}_i^{\text{test}}$ and $\mathcal{T}_i^{\text{set}}$, as indicated in the third

Vertex i	o_p	v_{rp}^-	w_{rp}^+	w_{rp}^-	v_{rp}^+	d_p
Associated task	τ_p^o	τ_r^1	τ_r^2	τ_r^2	τ_r^3	τ_p^d
Tasks to test $\mathcal{T}_i^{\text{test}}$	τ_p^o, τ_p^d	$\tau_r^1, \tau_r^2, \tau_r^3, \tau_p^d$	$\tau_r^2, \tau_r^3, \tau_p^d$	τ_r^3, τ_p^d	τ_r^3, τ_p^d	τ_p^d
Tasks to set $\mathcal{T}_i^{\text{set}}$	τ_p^o	τ_p^o, τ_r^1	τ_p^o, τ_r^1	$\tau_p^o, \tau_r^1, \tau_r^2$	$\tau_p^o, \tau_r^1, \tau_r^2, \tau_r^3$	τ_p^o, τ_p^d

Table 3: Chain of precedences for each pair $(p, r) \in P \times R^p$

and fourth line in the table. Before extending a label L at vertex i to vertex j , we test the tasks in the associated set $\mathcal{T}_j^{\text{test}}$. If any of these tasks has already been fulfilled along the partial path represented by L , then L is not extended to j , because such a path is not feasible for the APVRP. Similarly, upon extending a label L at vertex i to vertex j , we mark the tasks in $\mathcal{T}_j^{\text{set}}$ as fulfilled, because visiting them would lead to an infeasible path regarding task precedences. Note that it is unnecessary to include the tasks in the sets $\mathcal{T}^{\text{test}}$ and \mathcal{T}^{set} in the ng-neighborhood. This is because these tasks are irrelevant for dominance, as dominance checks occur only between labels resident at the same vertex, and as all tasks to test and all tasks to set are the same for all labels at the same vertex. With respect to the ng-path relaxation, on the one hand, the tasks to set allow making labels more comparable from the dominance point of view, and, on the other hand, the tasks to test prevent the construction of infeasible paths. This improvement can always be applied whenever a subset of tasks associated with some vertices of the graph must be executed according to a given precedence relation.

5.2.3. Forward label extension

As mentioned, our approach is based on the one by Ioachim *et al.* (1998), who store cost functions as labels instead of scalars. Such a cost function c provides the minimal costs $c(T)$ incurred by a path when the service at the last vertex of the path starts at time T . Ioachim *et al.* (1998) prove that this function is piecewise linear, convex, and contains at most as many linear pieces as there are vertices in the path. Moreover, they show that pieces with positive slope can be replaced by a single piece with slope zero. Hence, for the APVRP, a label comprises the following components:

- n number of time-slope pairs (in the following called pieces);
- $(t^p, s^p)_{p=1}^n$ the n pieces; $s^p < 0$ for $p = 1, \dots, n-1$, and $s^n \leq 0$
- c^1 (reduced) costs at start time of piece 1;
- t^{n+1} end time of last piece n ;
- S tasks that have already been performed or are unreachable along the route and which are taken into account for checking elementarity in the label extension step according to the ng-path relaxation;
- k number of intervals on which the label is dominated;
- $(I^d)_{d=1}^k$ the k intervals on which the label is dominated.

The following information can then be derived:

$$\begin{aligned}
c^* &= c^1 + \sum_{p=1}^n s^p(t^{p+1} - t^p); \text{ the optimal (reduced) costs;} \\
t^* &= t^n \text{ for } s^n = 0, \text{ and } t^{n+1} \text{ otherwise; the earliest time to obtain costs } c^*; \\
c(T) &= c^1 + \sum_{p=1}^{q-1} s^p(t^{p+1} - t^p) + s^q(T - t^q) \text{ for } t^q \leq T \leq t^{q+1}; \text{ the tradeoff curve;} \\
s^{n+1} &= 0 \text{ defined so for convenience.}
\end{aligned}$$

Figure 2 depicts two typical situations. In Figure 2(a), the slope of the function is strictly negative over its complete range, so that the minimal (reduced) costs are achieved when starting the service at the vertex as late as possible. In Figure 2(b), the slope of the function is zero on




Figure 2: Example of labels; (a) A label with $n = 3$ linear pieces and final slope $s^n < 0$; (b) A label with $n = 2$ linear pieces and final slope $s^n = 0$.

the positive-length interval from t^2 to t^3 , and any time value in this interval yields the minimal (reduced) costs c^* . Note that the figures imply that the paths whose cost functions are depicted in 2(a) and (b) contain at least three and two vertices respectively.

Now we describe the label extension step. (We present the extension step here because the paper by Ioachim *et al.* (1998) omits several details.) Let $L_i = ((t_i^p, s_i^p)_{p=1}^{n_i+1}, c_i^1, S_i, (I_i^d)_{d=1}^{k_i})$ be a label at vertex i . The extension of L_i along the arc (i, j) is feasible if $t_i^1 + t_{ij} \leq l_j$ and $S_i \cap \mathcal{T}_j^{\text{test}} = \emptyset$. In this case, a new label L_j at vertex j is created. The attributes of L_j are computed in the following way:

First, extending some of the existing pieces may be obsolete (cf. Ioachim *et al.*, 1998, p. 200), either because the extended pieces arrive too early or too late at j or because several resulting new slopes are zero. We compute the indices f and g of the first and last new piece to be kept:

$$f := \max\{p \in \{0, \dots, n_i + 1\} : t_i^p + t_{ij} \leq e_j\}$$

and

$$g := \min\{n_i, p \in \{0, \dots, n_i\} : s_i^p + \tilde{c}_j \geq 0 \text{ or } t_i^{p+1} + t_{ij} \geq l_j\}$$

It may also happen that a new piece must be created (cf. Ioachim *et al.*, 1998, p. 200). Thus, we define the new-piece indicator

$$\delta := \begin{cases} 1 & \text{if } (g = n_i, t_i^* = t_i^{n_i+1}, \text{ and } t_i^* + t_{ij} < l_j) \text{ or } (f = n_i + 1) \\ 0 & \text{otherwise} \end{cases},$$

so that the new label L_j comprises the

$$n_j := \max\{0, g - f + 1\} + \delta \quad (3)$$

new pieces.

Second, non-obsolete pieces (t, s) are extended using the function $f_{ij}(t, s) := (\max\{e_j, t + t_{ij}\}, \min\{0, s + \tilde{c}_j\})$ so that the new pieces are

$$(t_j^p, s_j^p) := f_{ij}(t_i^{f+p-1}, s_i^{f+p-1}) \quad \text{for all } p = 1, \dots, n_j \quad (4a)$$

and the new end time of the pieces associated with slope 0 is

$$(t_j^{n_j+1}, s_j^{n_j+1}) := (l_j, 0). \quad (4b)$$

Third, with the help of the tradeoff curve $c_i(T)$ of L_i and the already computed attributes, the costs at the start time t_j^1 of the new pieces can be expressed as

$$c_j^1 := c_i(\min\{t_i^*, t_j^1 - t_{ij}\}) + \tilde{c}_{ij} + \tilde{c}_j t_j^1. \quad (4c)$$

Note that the formula for the cost update given by Ioachim *et al.* (1998, p. 202) is not always correct.

Fourth, the remaining attributes are

$$S_j := (S_i \cap \mathcal{N}_j) \cup \mathcal{T}_j^{\text{set}} \quad \text{and} \quad k_j := 0. \quad (4d)$$

The latter definition means that at the time of its creation, the new label is undominated. This completes the description of the new label $L_j = ((t_j^p, s_j^p)_{p=1}^{n_j+1}, c_j^1, S_j, (I_j^d)_{d=1}^{k_j})$.

From a label L_d at the end depot d , the corresponding o - d -path is reconstructed, as usual, by iterating backward through the predecessor labels. Let $(d = i + 1, i, \dots, 0 = o)$ be the backtracked path from a label L_d . The optimal schedule (T_{i+1}, \dots, T_0) is then given by the following recursion:

$$T_{i+1} = t_d^*, \quad (5a)$$

$$T_i = \min\{t_i^*, T_{i+1} - t_{i,i+1}\}. \quad (5b)$$

We now proceed to the discussion of the dominance procedure and details of its implementation.

5.2.4. Dominance between labels

When checking a label L for dominance, it is compared one by one with all other undominated labels residing at the same vertex. During each such comparison between L and another label L' , first, the sets S_L and $S_{L'}$ are compared. If $S_{L'} \subseteq S_L$, a pointwise dominance is performed with respect to each point on the tradeoff curve of L , and the intervals on which L is dominated, i.e., where $c_{L'}(T) \leq c_L(T)$ holds, are tentatively stored. We compute these domination intervals efficiently by determining the intersection points of $c_L(T)$ and $c_{L'}(T)$. Ioachim *et al.* (1998) have shown that even if the intersection points may be non-integer real numbers, it suffices to describe the intervals with integer bounds. Also Liberatore *et al.* (2011) exploit domination intervals, however, they delete the dominated intervals from the tradeoff curves so that they can become discontinuous.

After having compared L with all other labels in this way, these intervals are used to update $(I^d)_{d=1}^k$, the intervals on which the tradeoff curve of L is dominated, and k , the number of intervals on which L is dominated. If the complete tradeoff curve of a label is dominated, the label itself is dominated and can be discarded. Note that this means that the decision on whether or not a label is obsolete will regularly be based on comparisons with more than one other label. Put differently, in most cases, only several other labels together will make one label obsolete. This is in contrast to dominance procedures for classical VRPs, where a pairwise dominance is applicable and one label dominates another one or not.

To avoid that two labels L and L' dominate each other on an interval I^d when $S_L = S_{L'}$ and $c_L(T) = c_{L'}(T)$ for all $T \in I^d$, we use a tie-breaking rule, similar as in standard VRPs where it must be avoided that two identical labels eliminate each other.

Refinements of the dominance procedure that speed up the pricing process are the elimination of dominated pieces at the beginning or at the end of the range of the tradeoff curve and the replacement of consecutive dominated pieces by one aggregated piece.

5.2.5. Bidirectional labeling

We start by briefly describing the backward label extension before we detail the bidirectional labeling approach.




Figure 3: Example of a backward label with $n = 3$ pieces




Figure 4: Example of a merge step, forward and backward label with optimal times T_i and T_j for their concatenation

Given the forward label extension process, the backward label extension process is rather simple: It is sufficient to invert all time windows and to invert the linear vertex costs. To be precise, a time window $[e_r, l_r]$ for a request r is replaced by $[t^{max} - l_r, t^{max} - e_r]$, \tilde{c} is replaced by $-\tilde{c}$, and then the same algorithm as in the forward labeling is applied. Note that it suffices to compute $t^{max} - T$ to recalculate the real time for a given point in time T on the tradeoff curve of the backward label. For simplicity, when a point in time on the tradeoff curve of the backward label is mentioned in the following, we refer to the real time. Hence, the pieces are numbered from 1 to n with decreasing, non-negative slopes, and t^1 is the latest feasible time for a backward label. Figure 3 summarizes the notation for backward labels.

In bidirectional labeling algorithms, forward labels are not necessarily propagated until the end depot, and backward labels are not necessarily propagated until the start depot. Instead, labels are propagated only up to a so-called half-way point, thus limiting the overall number of created labels. Suitable forward and backward labels must then be merged to obtain complete o - d -paths. As described by Salani (2005, Sect. 4.6.4), this is done using a half-way point test to avoid creating the same path from different pairs of forward and backward labels. Setting $t^{max}/2$ as half-way point, we propagate forward labels at a vertex i only if $t_i^1 \leq t^{max}/2$, and backward labels at a vertex j only if $t_j^1 > t^{max}/2$.

We then merge on vertices, i.e., we consider forward and backward labels at the same vertex i . A forward label at a vertex i qualifies for merging if $i = d$ or its earliest service start time is $t_{fw}^1 > t^{max}/2$. This condition prevents the creation of identical paths from different pairs of forward and backward labels. We check all backward labels at vertex i for whether or not they can be merged with the chosen forward label. The procedure becomes simpler if instead of the backward

label itself, its predecessor label is considered. Let this predecessor be resident at vertex j . Note that this is a kind of merging along the resulting arc (i, j) . However, with the merge on vertices we benefit from the fact that both the forward and the backward label are Pareto-optimal. This is in contrast to the situation when merging is directly performed over arcs, where neither the forward extension of the forward label along (i, j) nor the backward extension of the backward label along (i, j) is guaranteed to be Pareto-optimal, leading to considerably more pairs of labels to be checked.

We consider forward labels at a vertex i for a merge with predecessor backward labels at a vertex j if the following conditions are fulfilled to obtain a feasible path:

$$S_{fw} \cap S_{bw} = \emptyset \quad \text{and} \quad t_{fw}^1 + t_{ij} \leq t_{bw}^1 \quad (6)$$

The first condition is necessary to avoid paths violating the partial elementarity required by the ng-path relaxation, the second is needed to ensure overall temporal feasibility. Now assume that these conditions are fulfilled for a forward label at a vertex i and a backward label at a vertex j , and consider Figure 4. On the left hand side of Figure 4, an exemplary forward cost function at i is depicted. The predecessor backward label with an exemplary backward linear cost function at j is depicted on the right hand side. We have to compute the optimal service start times T_i and T_j at the vertices i and j , respectively, so that all other service start times at the vertices in the path can be derived from these values, see Eqs. (5). There are two cases:

(i) $t_{fw}^* + t_{ij} \leq t_{bw}^*$.
This is unproblematic: $T_i = t_{fw}^*$ and $T_j = t_{bw}^*$ are the optimal service start times for vertices i and j .

(ii) $\Delta = t_{fw}^* + t_{ij} - t_{bw}^* > 0$.
In this case, T_i must be set to an earlier point in time than t_{fw}^* , and/or T_j must be set to a later point in time than t_{bw}^* . To ensure optimality, we compute the optimal times T_i and T_j by distributing the total required time shift Δ between the forward and the backward label. To distribute Δ , we initialize $T_i = t_{fw}^*$ and $T_j = t_{bw}^*$ and update them in the following manner. We consider the slopes of the cost function pieces. If, for example, $t_{fw}^* = t_{fw}^{n+1}$, $t_{bw}^* = t_{bw}^n$ and $|s_{fw}^n| \leq |s_{bw}^{n-1}|$, as it is the case in Figure 4, setting T_i to an earlier time than t_{fw}^* will increase the costs of the merged path by at most as much as pushing T_j to a later point in time than t_{bw}^* . Hence, we shift T_i backward by $\partial_1 := \min\{T_i - t_{fw}^n, \Delta\}$ units, i.e., $T_i := T_i - \partial_1$. The remaining time shift is $\Delta := \Delta - \partial_1$. Now, if $\Delta = 0$, we are done. Otherwise, we compare the slope of the preceding forward piece, here s_{fw}^{n-1} , with s_{bw}^{n-1} and select the piece with the smaller absolute slope. If, as depicted in the figure, $|s_{bw}^{n-1}| < |s_{fw}^{n-1}|$, we shift T_j forward by $\partial_2 := \min\{t_{bw}^{n-1} - T_j, \Delta\}$ units, i.e., $T_j := T_j + \partial_2$. Again, the remaining time shift is $\Delta := \Delta - \partial_2$. We iterate until $\Delta = 0$, and this is guaranteed to happen because of condition (6). The overall algorithm is depicted in Algorithm 1. Note that if the next forward piece to consider does not exist, i.e., if the merge point T_i is already equal to the start time t_{fw}^1 of the forward tradeoff curve, we take the backward one, and vice versa. This can be achieved by defining $s_{fw}^0 := -\infty$ and $s_{bw}^0 := +\infty$.

In the merge procedure just described, the cost functions of backward labels are relevant only in the interval $[t^{max}/2, t^{max}]$ from the half-way point to the end of the planning horizon. Therefore, it is sufficient to construct the tradeoff curve of a backward label only up to $t^{max}/2$. This saves some computation time when constructing the linear cost function while maintaining optimality of the procedure.

The reduced cost of a merged path can be obtained by $c_m^* := c_{fw}(T_i) + c_{bw}(T_j) + \tilde{c}_{ij}$. After the merge, we perform a final dominance test as described in Subsection 5.2.4 for all routes resulting from merged labels. Finally, we add all undominated negative reduced-cost routes to the master problem.

Algorithm 1: Computation of optimal service start times T_i and T_j when merging

```

1  $\Delta := t_{fw}^* + t_{ij} - t_{bw}^*$ ,  $T_i := t_{fw}^*$ ,  $T_j := t_{bw}^*$ 
2 if ( $s_{fw}^n < 0$ ) then  $p_{fw} := n$  else  $p_{fw} := n - 1$ 
3 if ( $s_{bw}^n > 0$ ) then  $p_{bw} := n$  else  $p_{bw} := n - 1$ 
4 while ( $\Delta > 0$ ) do
5   if ( $|s_{fw}^{p_{fw}}| < |s_{bw}^{p_{bw}}|$ ) then
6      $\partial := \min\{T_i - t_{fw}^{p_{fw}}, \Delta\}$ 
7      $T_i := T_i - \partial$ ,  $p_{fw} := p_{fw} - 1$ 
8   else
9      $\partial := \min\{t_{bw}^{p_{bw}} - T_j, \Delta\}$ 
10     $T_j := T_j + \partial$ ,  $p_{bw} := p_{bw} - 1$ 
11   $\Delta := \Delta - \partial$ 

```

Result: Optimal service start times T_i and T_j

5.2.6. Acceleration techniques

To accelerate the pricing process, we use a heuristic pricing procedure, the so-called *limited discrepancy search* (LDS) introduced by Feillet *et al.* (2007), and a heuristic dominance.

LDS works as follows. In each pricing iteration, the outgoing arcs of each vertex are partitioned into “good” and “bad” arcs. The arcs with the lowest current modified arc costs and all arcs incident to the start or the end depot are considered good arcs, the others are regarded as bad arcs. The labels have an additional attribute storing the number of bad arcs used in the associated path. The value of this attribute is incremented by one each time a label is extended along a bad arc. Only those labels are considered feasible for which the number of bad arcs used is below a specified upper bound, the discrepancy limit. This limit is set to a fixed value during the complete solution process.

As for heuristic dominance, we use a pairwise comparison of two labels L and L' , and if $c_L^* < c_{L'}^*$ and $t_L^1 < t_{L'}^1$, we discard L' . The limited discrepancy search as well as the heuristic dominance are applied in each column generation iteration. Only if they fail to produce a negative reduced cost column, the exact pricing and dominance algorithms are applied.

5.3. Branching strategy

Let $\tilde{\lambda}^{aq}$, \tilde{u}_r , and $\tilde{x}_{ij}^a = \sum_{q \in \Omega^a} X_{ij}^q \tilde{\lambda}^{aq}$ be the values of the corresponding decision variables λ^{aq} , u_r , and x_{ij}^a . We apply the following five-stage hierarchical branching scheme: First, if the overall number of unserved requests, i.e., $u_\Sigma = \sum_{r \in R} \tilde{u}_r$, is fractional, we create the two branches $\sum_{r \in R} u_r \leq \lfloor u_\Sigma \rfloor$ and $\sum_{r \in R} u_r \geq \lceil u_\Sigma \rceil$. If the objective (1a) prioritizes request fulfillment, i.e., $\gamma \gg \alpha, \beta$, the latter branch is obsolete. Second, we branch on the individual u_r variables, where one request r^* with \tilde{u}_{r^*} closest to $1/2$ is selected, and the two branches $u_{r^*} = 0$ and $u_{r^*} = 1$ are created. Third, we branch on the overall number of active vehicles in use: If $a_\Sigma = \sum_{a \in A} \sum_{q \in \Omega^a} \tilde{\lambda}^{aq}$ is fractional, the resulting branches are given by $\sum_{a \in A} \sum_{q \in \Omega^a} \lambda^{aq} \leq \lfloor a_\Sigma \rfloor$ and $\sum_{a \in A} \sum_{q \in \Omega^a} \lambda^{aq} \geq \lceil a_\Sigma \rceil$. All these branching rules are put into effect by adding a constraint to the master program (1). In addition, when the second branching rule is applied and u_{r^*} is set to 1, all the subgraphs induced by vertices in $V_{r^*}^- \cup W_{r^*}^+ \cup W_{r^*}^- \cup V_{r^*}^+$ can be removed from all networks G^a , $a \in A$.

Fourth, note that for arcs $(i_p, j_{p'}) \in E$ that have both endpoints in N^R , at most one of the arcs $\bigcup_{q, q' \in P} \{(i_q, j_{q'})\}$ of one of the networks G^a , $a \in A$, can be present in a solution. Hence, if $\Sigma_{ij} = \sum_{a \in A} \sum_{q, q' \in P^a} \tilde{x}_{i_q, j_{q'}}^a$ is fractional, we create two branches by setting $\sum_{a \in A} \sum_{q, q' \in P^a} x_{i_q, j_{q'}}^a$ to zero and one respectively. If only one of the endpoints is in N^R , we can apply a similar branching rule. If several Σ_{ij} are fractional, we choose a pair (i, j) with value closest to $1/2$. The zero-branch

is implemented by eliminating the associated arcs from all networks G^a , $a \in A$. The one-branch first fixes $u_r = 0$ for the corresponding request(s) r associated with i_p and j_p and eliminates incompatible arcs from all networks G^a , $a \in A$.

The fifth and last rule is branching on individual arcs, which finally ensures integrality of the x_{ij}^a and u_r variables. (Recall that model (1) poses no integer requirement on the λ^{aa} variables.) If \tilde{x}_{ij}^a for an arc $(i, j) \in E^a$ is fractional, then one can branch on $x_{ij}^a = 0$ and $x_{ij}^a = 1$, where we select the combination (a, i, j) for which \tilde{x}_{ij}^a is closest to $1/2$. Branching on individual arcs $(i, j) \in E^a$ is implemented by eliminating (i, j) from the underlying network G^a for $x_{ij}^a = 0$, while for the branch $x_{ij}^a = 1$, all ingoing arcs $\delta^-(i)$ are eliminated for networks $G^{a'}$, $a' \neq a$ and all arcs $\delta^+(i) \setminus \{(i, j)\}$ are eliminated from the network G^a .

As branch-and-bound node-selection rule, we apply a best-bound-first strategy, because our primary goal is to improve the dual bound.

6. Experimental results

The results reported in this section were obtained using a standard PC with an Intel(R) Core(TM) i7-2600 3.4 GHz processor and 16 GB of main memory. The algorithms were coded in C++ with MS-Visual Studio 2010. The callable library of CPLEX 12.5 was used for solving the linear relaxations of the restricted master program in the column generation algorithm.

6.1. Test instances

Meisel and Kopfer (2014) created a set of 180 APVRP test instances (henceforth referred to as MK instances) with the following characteristics: The instances range in size from 38-1600 tasks, 2-25 active vehicles, 4-50 passive vehicles, and a planning horizon t^{max} of 2500-5000 time units. Locations are randomly placed in a 100×100 area, distances and travel times are set to the Euclidean distance, time windows are of width 1000 with random start times in $[0, t^{max} - 1000]$, and service times for loading and unloading range between 50 and 100 time units. Each passive vehicle is compatible with 2 active vehicles, each request is compatible with 3 passive vehicles. Due to their huge planning horizon t^{max} where time units can model minutes of a working week, these instances are very hard to solve with an exact approach.

Hence, we additionally created a new set of instances, in which a time unit can model 10 minutes of a working week. It is clear that our new instances are on average easier to solve than the MK instances, however, the chosen time discretization is fine enough for many practical situations. The new set contains 20 instances with 38 tasks, 2 active, and 4 passive vehicles (leading to extended networks with approximately 170 vertices and 3500 arcs), and 20 instances with 76 tasks, 4 active, and 8 passive vehicles (resulting in extended networks with ca. 650 vertices and 4000 arcs). The instances have a planning horizon of 1000 time units, and service times s_r^+ and s_r^- vary between 25 and 50. To each of the 40 instances, we assigned time windows that allow different levels of flexibility: For each request $r \in R$, we first compute the travel time between its pickup and its delivery location. Initially, each request time window $[e_r, l_r]$ (refer to Table 1 for the definition of e_r and l_r) is set so that the time window width $l_r - e_r$ is equal to the computed travel time plus the service times s_r^+ and s_r^- . By adding a *time window flexibility* of $F = 25, 50, 100$, and 200 , we create 160 new instances. A time window flexibility of F means that e_r and l_r are further shifted apart by exactly F time units (it is ensured that the new $[e_r, l_r]$ is in the overall planning horizon $[0, t^{max}]$). All other characteristics of the new instances are the same as in the MK instances.

6.2. Algorithmic setup

For all experiments, we used objective function weights of $\alpha = 10$, $\beta = 1$, and $\gamma = 10,000$, i.e., a hierarchical objective of first fulfilling as many requests as possible, then minimizing traveled

distance, and then minimizing route completion time. We set a CPU time limit of 2 hours. Preliminary tests showed that bidirectional labeling was superior to the unidirectional variant, that the best results were obtained with an ng-neighborhood size of 15, applying limited discrepancy search and heuristic dominance during the pricing as described in Section 5.2.6, and stopping each labeling procedure as soon as 100 or more negative reduced cost routes have been found.

6.3. Algorithmic performance

As for the MK instances, Meisel and Kopfer (2014) could not solve any of these to optimality with the branch-and-cut algorithm presented in their paper. Our algorithm was able to solve to optimality one 38-task MK instance within the 2-hour time limit. For this solved instance the previous upper bound was improved by 10.05 %. The average gap between the upper and the lower bound after 2 hours was 6.79 % for the MK instances, with a minimum non-zero gap of 3.47 % and a maximum gap of 13.90 %. The average increase in the lower bound at the root node of the branch-and-bound tree compared to the one of the branch-and-cut by Meisel and Kopfer (2014) was a significant 18.35 %.

Tables 4 and 5 report the results on the test instances we created. Both tables have the same columns. The first column indicates the time window flexibility F of the instances considered in each row, the second specifies the number of instances solved to optimality. The next columns indicate minimal, average, and maximal values of (i) the CPU times, (ii) the gap between the value of the LP relaxation at the root node and the best known upper bound (also using upper bounds for instances with smaller F as bounds for instances with a bigger F), (iii) the percentage of the root gap closed at the end of the optimization (100.0 % for instances that were solved to optimality), and (iv) the number of branch-and-bound nodes solved. It can be seen from the tables that we are able to consistently solve instances with 38 tasks and a time window flexibility of up to 100. In general, the instances become more difficult with increasing time window flexibility, and timeout is reached for most 76-task instances, with significant gaps remaining. For the 76-task instances with a time window flexibility of 100 and 200, the number of solved branch-and-bound nodes even decreases, on average, compared to the instances with shorter time windows. This means that the time needed to solve a node increases considerably. Analyses showed that this is mostly caused by the slow convergence of the master program as well as the increasing solution times for the pricing problems, which, in turn, are due to the increasing number of labels created because of the increasing flexibility offered by longer time windows.

Table 6 shows that there is a positive impact of increasing time window flexibility on possible solution quality. The values in the table reflect the objective function improvement obtained with increasing time window flexibility. For example, for the 38-task instances, the average objective function value improves by 1.20 % if the time window flexibility is increased from 25 to 50.

TW flex.	# Solved	Time [sec]			Gap at root [%]			Gap closed [%]			# Nodes solved		
		min	avg	max	min	avg	max	min	avg	max	min	avg	max
25	20/20	36	191	931	0.00	0.43	1.66	100.00	100.00	100.00	1	5.40	23
50	20/20	43	362	2485	0.00	0.55	2.12	100.00	100.00	100.00	1	8.30	47
100	19/20	38	1076	7200	0.00	0.99	3.00	64.22	98.21	100.00	1	16.95	56
200	14/20	109	3466	7200	0.00	2.68	9.90	8.91	79.91	100.00	1	26.20	106
All	73/80	1274			1.16			94.53			14.21		

Table 4: Results on 38-task instances

TW flex.	# Solved	Time [sec]			Gap at root [%]			Gap closed [%]			# Nodes solved		
		min	avg	max	min	avg	max	min	avg	max	min	avg	max
25	17/20	131	3594	7200	0.00	0.80	2.39	46.60	90.54	100.00	1	22.70	49
50	8/20	282	5419	7200	0.28	1.55	4.31	21.22	67.91	100.00	3	25.90	46
100	2/20	2814	6882	7200	0.47	3.53	6.95	10.70	34.08	100.00	11	25.50	45
200	0/20	7200	7200	7200	6.10	8.48	13.47	1.77	6.28	15.39	5	11.60	19
All	27/80	5775			3.59			49.70			21.43		

Table 5: Results on 76-task instances

TW flex. increase	Objective improvement [%]	
	38 tasks	76 tasks
25 → 50	1.20	1.54
50 → 100	1.94	1.56
100 → 200	5.84	N/A

Table 6: Impact of time window flexibility

7. Conclusions

This paper has investigated the exact solution of the active-passive vehicle-routing problem (APVRP) by means of a branch-and-price method. The problem includes synchronization constraints for the operations and the movement of active and passive vehicles. It supports a flexible coupling of these transport resources in order to achieve an efficient resource utilization and high-quality transport solutions. The synchronization constraints introduce linear vertex costs in the ESPPRC pricing problem, thus significantly complicating its solution. To solve the pricing problems, we use a bidirectional labeling algorithm and apply a refinement of the ng-path relaxation approach. What is more, we propose a sophisticated procedure for merging forward and backward labels. Computational experiments show that our method delivers improved results for the APVRP benchmark instances introduced by Meisel and Kopfer (2014). For a newly created benchmark set, it is capable of solving instances with 76 tasks, 4 active, and 8 passive vehicles to optimality.

In its current form, the APVRP already covers a number of relevant applications from the areas of distribution logistics, health care management, the security industry and others. Nevertheless, there are several further features of real-world problems that are not yet supported by the presented model. This includes, for example, the possibility of a temporary drop-off of passive vehicles at intermediate locations (e.g., public parking lots) that neither belong to the set of initial or final locations nor to the set of pickup and delivery locations. Such drop-off locations could reduce the detours needed by active vehicles for the exchange of passive vehicles. Furthermore, a rendezvous of active vehicles is required in some applications where passive vehicles cannot be left behind but must be handed over directly from one active vehicle to another. This is the case if only one of the active vehicles is equipped with a certain loading equipment, if the passive vehicle cannot be left unattended for security reasons, or if cargo documents need to be exchanged among the active vehicles. Moreover, in many or even most practical applications, vehicles with a capacity of more than one are used, and performing load transshipments between such vehicles is quite common in practice. These features and their corresponding synchronization requirements represent practically relevant and mathematically non-trivial extensions of the APVRP and will be a subject of our future research.

Acknowledgement. This research was funded by the Deutsche Forschungsgemeinschaft (DFG) under grants no. IR 122/5-2 and DR 963/2-1.

Appendix: Notation

Sets

R	Set of all requests
r	A request
A	Set of classes of active vehicles
a	A class of active vehicles
K_a	Number of active vehicles of class a
P	Set of passive vehicles
p	A passive vehicle
R^p	Set of requests compatible with passive vehicle p
P^r	Set of passive vehicles compatible with request r
P^a	Set of passive vehicles compatible with active vehicle a
A^p	Set of classes of active vehicles compatible with passive vehicle p

Locations

ℓ_r^+	The pickup location of request r
ℓ_r^-	The delivery location of request r
o	Initial location of all active vehicles
d	Final location of all active vehicles
o_p	Initial location of passive vehicle p
d_p	Final location of passive vehicle p

Cost parameters

α	Weight for travel distance
β	Weight for arrival time at the final depot d
γ	Penalty for not fulfilling a request
c_{ij}	Travel distance from vertex i to vertex j

Time parameters

t_{ij}	Travel and service time from vertex i to vertex j
$[e_r, l_r]$	Time window for fulfilling request $r \in R$
s_r^+	The service duration to process an empty passive vehicle at request r
s_r^-	The service duration to process a loaded passive vehicle at request r
$[0, t^{max}]$	Planning horizon

Network parameters

$G^a =$	
(V^a, E^a)	Extended graph for active vehicle class $a \in A$ with vertex set V^a and arc set E^a
$\delta^+(i)$	$\{j \in V^a : (i, j) \in E^a\}$; forward star of vertex $i \in V^a$
$\delta^-(i)$	$\{j \in V^a : (j, i) \in E^a\}$; backward star of vertex $i \in V^a$
N^R	$\bigcup_{r \in R} (V_r^- \cup W_r^+ \cup W_r^- \cup V_r^+)$; set of request vertices
For each passive vehicle $p \in P^a$ and each compatible request $r \in R^p$:	
v_{rp}^-	Vertex where empty passive vehicle is delivered
w_{rp}^+	Vertex where loaded passive vehicle is picked up
w_{rp}^-	Vertex where loaded passive vehicle is delivered
v_{rp}^+	Vertex where empty passive vehicle is picked up

(Continued on next page)

(Continued from previous page)

N^a $\cup_{p \in P^a, r \in R^p} \{v_{rp}^-, w_{rp}^+, w_{rp}^-, v_{rp}^+\}$; set of request vertices for each passive vehicle $p \in P^a$ and each compatible request $r \in R^p$

For each request $r \in R$:

V_r^- $\{v_{rp}^- : p \in P^r\}$; set of vertices where empty passive vehicle may be delivered for pickup of r

W_r^+ $\{w_{rp}^+ : p \in P^r\}$; set of vertices where passive vehicle loaded with r may be picked up

W_r^- $\{w_{rp}^- : p \in P^r\}$; set of vertices to where passive vehicle loaded with r may be delivered

V_r^+ $\{v_{rp}^+ : p \in P^r\}$; set of vertices where empty passive vehicle may be picked up after delivery of r

Set-partitioning formulation parameters

Ω^a Set of all routes that can be defined in graph $G^a = (V^a, E^a)$, $a \in A$

q A route in $\bigcup_{a \in A} \Omega^a$

X_{ij}^q Number of times arc (i, j) is traversed on route q

T_i^q Service start time at request vertex i on route q

b_i^q Number of times vertex i is visited on route q

c^q Cost of route q ; weighted sum of arrival time at the destination d and length of the route

λ^{aq} Continuous variables measuring the flow of active vehicles of class $a \in A$ along route $q \in \Omega^a$

x_{ij}^a Binary variables indicating the number of times arc $(i, j) \in E^a$ is traversed by an active vehicle of class $a \in A$

u_r Binary variable indicating whether or not request $r \in R$ remains unfulfilled

Pricing problem parameters

π_r Unrestricted dual variable of constraint (1b)

ϕ_{rp}^+ Unrestricted dual variable of constraint (1c)

ϕ_{rp}^- Unrestricted dual variable of constraint (1d)

τ_r^+ Nonnegative dual variable of constraint (1e)

τ_r^- Nonnegative dual variable of constraint (1f)

μ_p^P Unrestricted dual variable of constraint (1g)

μ_a^A Nonpositive dual variable of constraint (1h)

\tilde{c}_{ij} Reduced costs of arc (i, j)

\tilde{c}_i Linear vertex costs of vertex i

x_{ij} Binary variable indicating whether or not arc $(i, j) \in E^a$ is traversed

T_i Continuous variable indicating the point in time when the service at vertex $i \in V^a$ begins

Labeling algorithm parameters

τ A task

\mathcal{N}_i ng-neighborhood at vertex i ; $\mathcal{N}_i \subset N^R$; \mathcal{N}_i contains i and the ν closest request vertices j for which a cycle (j, \dots, i, \dots, j) would be feasible when taking time windows and travel and service times into account

ν Size of ng-neighborhood

τ_p^o Task of picking up the passive vehicle $p \in P$ at its initial location

τ_r^1 Task of delivering an empty passive vehicle for request $r \in R$

(Continued on next page)

(Continued from previous page)

τ_r^2	Task of picking up and delivering a passive vehicle loaded with request $r \in R$
τ_r^3	Task of picking up an empty passive vehicle after delivering request $r \in R$
τ_p^d	Task of delivering the passive vehicle $p \in P$ at its final location
L	$L = ((t^p, s^p)_{p=1}^{n+1}, c^1, S, (I^d)_{d=1}^k)$; a label
n	Number of time-slope pairs/pieces in tradeoff curve
$(t^p, s^p)_{p=1}^n$	The n pieces of the tradeoff curve; $s^p < 0, p = 1, \dots, n-1$ and $s^n \leq 0$
t^{n+1}	End time of last piece n in tradeoff curve
$c(T)$	$= c^1 + \sum_{p=0}^{q-1} s^p(t^{p+1} - t^p) + s^q(T - t^q)$ for $t^q \leq T \leq t^{q+1}$; the piecewise linear cost function/tradeoff curve of a label, i.e., a label's associated partial path; provides the minimal costs $c(t)$ incurred by the path when the service at the last vertex of the path starts at time t
c^1	Reduced costs at start time of first piece in tradeoff curve
S	Tasks that have already been performed or are unreachable along the route and which are taken into account for checking elementarity in the label extension step according to the ng-path relaxation
k	Number of intervals on which the associated label is dominated
$(I^d)_{d=1}^k$	The k intervals on which the associated label is dominated
c^*	$= c^1 - \sum_{p=1}^n s^p(t^{p+1} - t^p)$; optimal (reduced) costs
t^*	$= t^n$ for $s^n = 0$, t^{n+1} otherwise; earliest time to obtain costs c^*
s^{n+1}	$= 0$; slope of $(n+1)$ th piece; defined for convenience
$\mathcal{T}_j^{\text{test}}$	Before extending a label L at vertex i to vertex j , tasks in $\mathcal{T}_j^{\text{test}}$ are tested and L is extended only if none of these tasks has already been fulfilled along the partial path represented by L
$\mathcal{T}_j^{\text{set}}$	Upon extending a label L at vertex i to vertex j , tasks in $\mathcal{T}_j^{\text{set}}$ are marked as fulfilled
f	Index of first linear slope piece to be kept when extending a label
g	Index of last linear slope piece to be kept when extending a label
δ	Piece indicator; one if a new linear slope piece must be created upon extension of a label, zero otherwise
$f_{ij}(t, s)$	$= (\{\max\{e_j, t + t_{ij}\}\}, \min\{0, s + \tilde{c}_j\})$; function for extending pieces during label extension

References

- Affi, S., Dang, D.-C., and Moukrim, A. (2015). Heuristic solutions for the vehicle routing problem with time windows and synchronized visits. *Optimization Letters*, in press.
- Andersson, H., Duesund, J. M., and Fagerholt, K. (2011). Ship routing and scheduling with cargo coupling and synchronization constraints. *Computers & Industrial Engineering*, **61**(4), 1107–1116.
- Baldacci, R., Mingozzi, A., and Roberti, R. (2011). New route relaxation and pricing strategies for the vehicle routing problem. *Operations Research*, **59**(5), 1269–1283.
- Bredström, D. and Rönnqvist, M. (2008). Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. *European Journal of Operational Research*, **191**(1), 19–31.
- Buijs, P., Vis, I. F., and Carlo, H. J. (2014). Synchronization in cross-docking networks: A research classification and framework. *European Journal of Operational Research*, **239**(3), 593–608.
- Cheung, R. K., Shi, N., Powell, W.-B., and Simao, H. P. (2008). An attribute-decision model for cross-border drayage problem. *Transportation Research Part E: Logistics and Transportation Review*, **44**(2), 217–234.
- Christiansen, M. and Nygreen, B. (1998). A method for solving ship routing problems with inventory constraints. *Annals of Operations Research*, **81**(0), 357–378.
- Desaulniers, G., Desrosiers, J., Ioachim, I., Solomon, M., Soumis, F., and Villeneuve, D. (1998). A unified framework for deterministic time constrained vehicle routing and crew scheduling problems. In T. G. Crainic and G. Laporte, editors, *Fleet Management and Logistics*, pages 57–93. Kluwer, Boston.
- Desaulniers, G., Desrosiers, J., and Solomon, M., editors (2005). *Column Generation*. Springer, New York.
- Desaulniers, G., Madsen, O., and Ropke, S. (2014). The vehicle routing problem with time windows. In P. Toth and D. Vigo, editors, *Vehicle Routing: Problems, Methods, and Applications*, MOS-SIAM Series on Optimization, pages 119–159. SIAM.
- Dohn, A., Kolind, E., and Clausen, J. (2009). The manpower allocation problem with time windows and job-teaming constraints: A branch-and-price approach. *Computers & Operations Research*, **36**, 1145–1157.
- Dohn, A., Rasmussen, M., and Larsen, J. (2011). The vehicle routing problem with time windows and temporal dependencies. *Networks*, **58**, 273–289.
- Drexler, M. (2007). *On Some Generalized Routing Problems*. Ph.D. thesis, Faculty of Business and Economics, RWTH Aachen University.
- Drexler, M. (2012). Synchronization in vehicle routing—a survey of VRPs with multiple synchronization constraints. *Transportation Science*, **46**(3), 297–316.
- Drexler, M. (2013). Applications of the vehicle routing problem with trailers and transshipments. *European Journal of Operational Research*, **227**(2), 275–283.
- Drexler, M. (2014). Branch-and-cut algorithms for the vehicle routing problem with trailers and transshipments. *Networks*, **63**(1), 119–133.
- Drexler, M., Rieck, J., Sigl, T., and Press, B. (2013). Simultaneous vehicle and crew routing and scheduling for partial- and full-load long-distance road transport. *BuR - Business Research*, **6**(2), 242–264.
- Dror, M. (1994). Note on the complexity of the shortest path models for column generation in vrptw. *Operations Research*, **42**, 977–978.
- Feillet, D., Gendreau, M., and Rousseau, L.-M. (2007). New refinements for the solution of vehicle routing problems with branch and price. *INFOR*, **45**(4), 239–256.
- Hachemi, N. E., Gendreau, M., and Rousseau, L.-M. (2013). A heuristic to solve the synchronized log-truck scheduling problem. *Computers & Operations Research*, **40**(3), 666–673.
- Hollis, B., Forbes, M., and Douglas, B. (2006). Vehicle routing and crew scheduling for metropolitan mail distribution at Australia post. *European Journal of Operational Research*, **173**(1), 133–150.
- Huber, S. and Geiger, M. (2014). Swap body vehicle routing problem: A heuristic solution approach. In R. González-Ramírez, F. Schulte, S. Voß, and J. Ceroni Díaz, editors, *Computational Logistics*, volume 8760 of *Lecture Notes in Computer Science*, pages 16–30. Springer International Publishing.
- Ioachim, I., Gélinas, S., Soumis, F., and Desrosiers, J. (1998). A dynamic programming algorithm for the shortest path problem with time windows and linear node costs. *Networks*, **31**(3), 193–204.
- Irnich, S. and Desaulniers, G. (2005). Shortest path problems with resource constraints. In G. Desaulniers, J. Desrosiers, and M. Solomon, editors, *Column Generation*, pages 33–65. Springer, New York.
- Irnich, S., Toth, P., and Vigo, D. (2014). The family of vehicle routing problems. In D. Vigo and P. Toth, editors, *Vehicle Routing*, chapter 1, pages 1–33. Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Jans, R. (2010). Classification of Dantzig-Wolfe reformulations for binary mixed integer programming problems. *European Journal of Operational Research*, **204**(2), 251–254.
- Kergosien, Y., Lenté, C., Piton, D., and Billaut, J.-C. (2011). A tabu search heuristic for the dynamic transportation of patients between care units. *European Journal of Operational Research*, **214**(2), 442–452.
- Kergosien, Y., Lenté, C., Billaut, J.-C., and Perrin, S. (2013). Metaheuristic algorithms for solving two interconnected vehicle routing problems in a hospital complex. *Computers & Operations Research*, **40**(10), 2508–2518.
- Kim, B.-I., Koo, J., and Park, J. (2010). The combined manpower-vehicle routing problem for multi-staged services. *Expert Systems with Applications*, **37**(12), 8424–8431.

- Labadie, N., Prins, C., and Yang, Y. (2014). Iterated local search for a vehicle routing problem with synchronization constraints. In *ICORES 2014 - Proceedings of the 3rd International Conference on Operations Research and Enterprise Systems, Angers, Loire Valley, France, March 6-8, 2014*, pages 257–263.
- Lahyani, R., Khemakhem, M., and Semet, F. (2015). Rich vehicle routing problems: From a taxonomy to a definition. *European Journal of Operational Research*, **241**(1), 1–14.
- Laporte, G. (2013). Scheduling issues in vehicle routing. *Annals of Operations Research*.
- Liberatore, F., Righini, G., and Salani, M. (2011). A column generation algorithm for the vehicle routing problem with soft time windows. *JOR*, **9**(1), 49–82.
- Lübbecke, M. and Desrosiers, J. (2005). Selected topics in column generation. *Operations Research*, **53**, 1007–1023.
- Mankowska, D. S., Meisel, F., and Bierwirth, C. (2013). The home health care routing and scheduling problem with interdependent services. *Health Care Management Science*, **17**(1), 15–30.
- Meisel, F. and Kopfer, H. (2014). Synchronized routing of active and passive means of transport. *OR Spectrum*, **36**(2), 297–322.
- Morais, V. W., Mateus, G. R., and Noronha, T. F. (2014). Iterated local search heuristics for the vehicle routing problem with cross-docking. *Expert Systems with Applications*, **41**(16), 7495–7506.
- Mues, C. and Pickl, S. (2005). Transshipment and time windows in vehicle routing. In *Proceedings of the 8th International Symposium on Parallel Architectures, Algorithms and Networks*, pages 113–119, Piscataway. IEEE.
- Salani, M. (2005). *Branch-and-Price Algorithms for Vehicle Routing Problems*. Ph.D. thesis, Faculty of Mathematical, Physical and Natural Sciences, University of Milan.
- Salazar-Aguilar, M. A., Langevin, A., and Laporte, G. (2013). The synchronized arc and node routing problem: Application to road marking. *Computers & Operations Research*, **40**(7), 1708–1715.
- Smilowitz, K. (2006). Multi-resource routing with flexible tasks: An application in drayage operations. *IIE Transactions*, **38**(7), 555–568.
- Spliet, R. and Gabor, A. (2014). The time window assignment vehicle routing problem. *Transportation Science*.
- Xue, Z., Zhang, C., Lin, W.-H., Miao, L., and Yang, P. (2014). A tabu search heuristic for the local container drayage problem under a new operation mode. *Transportation Research Part E: Logistics and Transportation Review*, **62**, 136–150.
- Zhang, R., Yun, W. Y., and Kopfer, H. (2010). Heuristic-based truck scheduling for inland container transportation. *OR Spectrum*, **32**(3), 787–808.
- Zhang, R., Yun, W. Y., and Kopfer, H. (2013). Multi-size container transportation by truck: modeling and optimization. *Flexible Services and Manufacturing Journal*, **27**(2-3), 403–430.
- Zhang, R., Lu, J.-C., and Wang, D. (2014). Container drayage problem with flexible orders and its near real-time solution strategies. *Transportation Research Part E: Logistics and Transportation Review*, **61**, 235–251.