_____

Models and Algorithms for the Heterogeneous Dial-a-Ride Problem with Driver Related Constraints

Sophie N. Parragh
Jean-François Cordeau
Karl F. Doerner
Richard F. Hartl

March 2010

CIRRELT-2010-13

Models and Algorithms for the Heterogeneous Dial-a-Ride Problem with Driver Related Constraints

Sophie N. Parragh[1], Jean-François Cordeau[2,*], Karl F. Doerner[1], Richard F. Hartl[1]

[1] Departmenet of Business Administration, University of Vienna, Bruenner Strasse 72, 1210 Vienna, Austria

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), and Canada Research Chair in Logistics and Transportation, HEC Montréal, 3000 Côte-Sainte-Catherine, Montréal, Canada H3T 2A7

Abstract. This paper introduces models and algorithms for a dial-a-ride problem arising in the transportation of patients by non-profit organizations such as the Austrian Red Cross. This problem is characterized by the presence of heterogeneous vehicles and patients. In our problem, two types of vehicles are used, each providing a different capacity for four different modes of transportation. Patients may request to be transported either seated, on a stretcher or in a wheelchair. In addition, some may require accompanying persons. The problem is to construct a minimum-cost routing plan satisfying service related criteria, expressed in terms of time windows, as well as driver related constraints expressed in terms of maximum route duration limits and mandatory lunch breaks. We introduce both a three-index and a set partitioning formulation of the problem. The linear programming relaxation of the latter is solved by a column generation algorithm providing tight lower bounds on the optimal solution value. We also propose a variable neighborhood search heuristic which yields high quality solutions on realistic test instances. Finally, we integrate the heuristic and the column generation approach into a collaborative framework.

Keywords. Passenger transportation, dial-a-ride-problem, column generation, variable neighborhood search.

_____

* Corresponding author: Jean-Francois.Cordeau@cirrelt.ca

1   Introduction

Our study was motivated by a problem faced by non-profit organizations such as the Austrian Red Cross (ARC) in the field of patient transportation. These organizations have to devise daily routing plans for their ambulances to serve a set of transportation requests that have been formulated in advance. Each request is defined by origin and destination locations, a number of passengers to be transported, and a time at which the passengers should either be picked-up at the origin or dropped-off at the destination.

The ARC disposes of vehicles of two different types, each type providing a different capacity for four modes of transportation: staff seat, patient seat, stretcher and wheelchair. Each patient may demand to be transported either seated, on a stretcher, or in a wheelchair. Accompanying persons may also be present and seating room for them has to be provided on the vehicle. Accompanying persons are allowed to use a staff seat. They may also use a patient seat or sit on the stretcher. Patients that can be transported seated are not allowed to use a staff seat. They can, however, be transported on the stretcher in the case where there are no more patient seats available. Patients that have to be transported on a stretcher or in a wheelchair can only be transported in the corresponding mode. Finally, some users require the presence of a second staff member aboard the vehicle.

The ambulance dispatcher has to assign drivers to vehicles (there are usually fewer drivers than available vehicles) and driver related constraints have to be taken into account in the planning process. These constraints refer to maximum shift lengths and mandatory breaks. Non-profit organizations like the ARC have to ensure that their expenses are kept low while providing a reasonably high service level. Thus, the aim is to construct a routing plan that is of minimum routing cost while respecting service related criteria, expressed in terms of time windows, as well as labor regulations. The resulting problem can be defined as a Heterogeneous Dial-A-Ride-Problem with Driver related constraints (HDARPD).

To address this complex routing problem we propose a Variable Neighborhood Search (VNS) heuristic relying on several neighborhood operators. We also introduce a column generation algorithm that is used to compute lower bounds based on a set-partitioning formulation of the problem. These bounds are used to assess the quality of the solutions produced by the heuristic. Finally, we discuss how the heuristic and the column generation approach can be integrated into a collaborative framework.

The contributions of this work are fourfold. First, we introduce a formulation for the HDARPD that incorporates routing and driver deployment decisions. This integration is made possible by the fact that the time frame within which regular patient transportation is performed corresponds to the drivers' work shifts. Second, a column generation algorithm is proposed. The central new aspect lies in the way the maximum route duration limit together with a non-empty time window at the start depot is treated in the dynamic programming part. Third, an efficient meta-heuristic algorithm tailored to the problem is presented. Fourth, the integration of the two methods into a collaborative scheme is investigated.

The remainder of the paper is organized as follows. The next section discusses related literature. Section 3 then introduces some notation and two mathematical formulations of the problem. These are followed by the column generation algorithm in Sections 4 and 5, respectively, and by the VNS heuristic in Section 6. Computational experiments are reported in Section 7 and are followed by the conclusion.

2   Related Work

Dial-A-Ride Problems (DARP) are generalizations of Pickup and Delivery Problems with Time Windows (PDPTW). In the DARP people are being transported instead of goods. This gives rise to the issue of service quality which can be ensured either through additional constraints or with extra terms in the objective function. The first approach is followed, e.g., by Parragh et al. (2009), while we adopt the latter approach in this paper.

Like the majority of the works published on the DARP, we address here its static version. Previous publications considering heterogeneous versions of the DARP involve, e.g., the work of Toth and Vigo (1997). The heterogeneity considered by these authors refers to two modes of transportation (seated passengers and passengers in wheelchairs) and to several different types of vehicles. The authors have devised a parallel insertion heuristic and a tabu thresholding algorithm for this problem. Another heterogeneous version of the DARP was described by Melachrinoudis et al. (2007). They developed a tabu search heuristic for a problem with several different types of vehicles in terms of capacity limits but only one mode of transportation. Heterogeneous vehicles in terms of capacity were also considered by Rekiek et al. (2006) who introduced a grouping genetic algorithm.

In the context of a dynamic environment, Beaudry et al. (2009) have adapted the tabu search heuristic of Cordeau and Laporte (2003) to solve a heterogeneous DARP that arises in large hospitals. It involves transportation requests requiring three different modes of transportation (seated, on a bed, or in a wheelchair) and several different types of vehicles. Another recent implementation of a computer-based planning system for a dynamic problem in a large German hospital, considering specific constraints such as multi-dimensional capacities, was reported by Hanne et al. (2009).

For overviews of the DARP and the closely related PDPTW, we refer the reader to Berbeglia et al. (2007), Cordeau and Laporte (2007) and Parragh et al. (2008a,b).

In related research, Xu et al. (2003) have addressed a practical pickup and delivery problem with multiple vehicle types, multiple time windows and several compatibility constraints both between requests and between requests and vehicles. In addition, first-in-first-out loading rules and driver work rules were taken into account. This problem was solved by a heuristic column generation algorithm in which several pricing heuristics are employed. In their experiments, the solution obtained at the root node is often integer. If not, a Mixed Integer Program (MIP) is solved on the set of generated columns. Column generation integrated into a branch-and-cut framework has also been successfully applied to the standard PDPTW by Ropke and Cordeau (2009). This algorithm has outperformed an earlier branch-and-cut algorithm by Ropke et al. (2007) in terms of the maximum problem size that can be solved. These results indicate that the use of a column generation-based algorithm is a promising direction. Finally, a combination of column generation with a local search method has recently been studied by Danna and Le Pape (2005) in the form of a so-called collaborative scheme.

In an earlier study, Parragh (2009b) has analyzed and solved a restricted version of the HDARPD, focusing solely on heterogeneous users and fleet aspects. For this simpler variant, state-of-the-art branch-and-cut algorithms for the DARP (Cordeau, 2006, Ropke et al., 2007) were successfully adapted. In the present work, maximum user ride times are not considered explicitly but rather implicitly in terms of time windows at both the pick-up and the drop-off locations. This modification makes the application of column generation possible and allows us to solve much larger instances.

3 Problem Formulation

In the following the basic notation needed to formulate the HDARPD is given. Thereafter, two different problem formulations are presented: a three-index model and a more compact set-partitioning formulation. The latter will serve as the basis for the proposed column generation framework.

3.1 Notation

The HDARPD is modeled on a complete directed graph $G = (V, A)$, where V is the set of vertices and A the set of arcs. To each arc (i, j) are associated a non-negative travel cost c_{ij} and a non-negative travel time t_{ij}. A set of n customer requests, each consisting of a pickup and delivery vertex pair $\{i, n+i\}$, have to be served. The set of pickup vertices is denoted by $P = \{1, \ldots, n\}$ and the set of delivery vertices by $D = \{n+1, \ldots, 2n\}$.

At every pickup vertex one patient has to be transported and this patient may demand one of three different transportation modes. For every transportation mode $r \in R$, where $R = \{0, 1, 2, 3\}$ denotes the set of transportation modes, let q_i^r denote the load of passenger i with respect to mode r. Passengers may have to be transported seated ($q_i^1 = 1$), on a stretcher ($q_i^2 = 1$), or in a wheelchair ($q_i^3 = 1$). Each patient may also require an accompanying person ($q_i^0 = 1$). The demand at every delivery vertex is equal to $q_{n+i}^r = -q_i^r$ for all $r \in R$. In addition, a patient may need additional personnel (referred to as "attendants" in the following), besides the driver, to be on the vehicle. This is represented by the binary parameter a_i taking value 1 if an attendant is needed, and 0 otherwise. Furthermore, let V_a be the set of all vertices demanding an attendant on board the vehicle. Attendants occupy a staff seat when they are on the vehicle.

Every user specifies a time window $[e_i, l_i]$ either for the pick-up (origin) or for the drop-off (destination) location, and the beginning of service has to start within this time window. In case a vehicle arrives too early, it has to wait until service is possible. A maximum passenger ride time L is implicitly considered in order to provide reasonable service quality. This is done by artificially constructing a time window at the origin (resp. destination) relative to the time window given at the corresponding destination (resp. origin): in case of an outbound request (a time window is given for the destination), the time window at the origin i is set to $e_i = \max\{0, e_{n+i} - L - d_i\}$ and $l_i = \min\{l_{n+i} - t_{i,n+i} - d_i, \bar{H}\}$, where $\bar{H}$ is the end of the planning horizon. In case of an inbound request (a time window is given for the origin), the destination time window is set to $e_{n+i} = e_i + d_i + t_{i,n+i}$ and $l_{n+i} = \min\{l_i + d_i + L, \bar{H}\}$. The parameter d_i denotes the service time at vertex i.

A set K of m heterogeneous vehicles is available to serve the transportation requests. Each vehicle $k \in K$ is associated with constants $C^{r,k}$ giving the amount of resource r available on the vehicle. The ARC disposes of two basic vehicle types. Type 1 (T1) provides 1 staff seat, 6 patient seats, and 1 wheelchair place. Type 2 (T2) provides 2 staff seats, 1 patient seat, 1 stretcher, and 1 wheelchair place. Patients demanding to be transported seated may use a patient seat or the stretcher. Patients demanding a stretcher can only be transported on a stretcher. The same applies to wheelchair passengers. Accompanying persons, however, may use a staff seat or a patient seat. They may also use the stretcher if no other transportation mode is available.

Each route has to start at the depot vertex 0 within a prespecified time window and finish at the end depot $2n+2$, respecting a route duration limit T. This limit is based on Austrian labor regulations. Driver working shifts are limited to 8.5 hours per day including a (lunch) break of $H = 30$ minutes that has to start within a given time window $[e_H, l_H]$. The lunch break can be held at every vertex. In addition, only a certain number of drivers m_d (usually $m_d < m$) and

only a limited number of attendants (in Austria these are employees serving their alternative service) m_a are available. An attendant can only work during morning or afternoon periods on a vehicle. An attendant working in the morning has to return to the noon depot $2n + 1$ within a certain time window $[e_{2n+1}, l_{2n+1}]$. An attendant working in the afternoon has to be picked-up at the noon depot within the time window. In case there are more drivers available than actually needed to serve all requests, the excess drivers may be employed for attendant duties, serving only half of the day on a vehicle.

Thus, the set of all vertices is given by $V = P \cup D \cup \{0, 2n + 1, 2n + 2\}$, and the set of all arcs by $A = \{(i, j) : i \in V \setminus \{2n + 2\}, j \in V \setminus \{0\}, i \neq j\}$.

3.2 A three-index formulation

We now introduce a three-index formulation based on the following binary decision variables:

$$x_{ij}^k = \begin{cases} 1, & \text{if arc } (i, j) \text{ is traversed by vehicle } k, \\ 0, & \text{otherwise,} \end{cases}$$

$$z_0^k = \begin{cases} 1, & \text{if a driver is assigned to vehicle } k, \\ 0, & \text{otherwise,} \end{cases}$$

$$z_1^k = \begin{cases} 1, & \text{if an attendant is assigned to vehicle } k \text{ in the morning,} \\ 0, & \text{otherwise,} \end{cases}$$

$$z_2^k = \begin{cases} 1, & \text{if an attendant is assigned to vehicle } k \text{ in the afternoon,} \\ 0, & \text{otherwise,} \end{cases}$$

$$v_i^k = \begin{cases} 1, & \text{if the lunch break is held at vertex } i, \\ 0, & \text{otherwise.} \end{cases}$$

In addition, we let $u \in \{0, \ldots, m_d\}$ denote the number of drivers that serve as additional attendants, while B_i^k and $Q_i^{r,k}$ denote, respectively, the beginning of service of vehicle k at vertex i and the load of the vehicle with respect to resource r when leaving vertex i. Finally, let W_H^k represent the waiting time until the lunch break on vehicle k.

The objective considered is the minimization of total routing costs:

$$\min \sum_{k \in K} \sum_{i \in V} \sum_{j \in V} c_{ij} x_{ij}^k. \tag{1}$$

For ease of exposition, we introduce the constraints of the model in several groups. The first group contains the demand constraints and defines the basic structure of the routes:

$$\sum_{k \in K} \sum_{j \in V} x_{ij}^k = 1 \qquad \forall i \in P, \tag{2}$$

$$\sum_{j \in V} x_{ij}^k - \sum_{j \in V} x_{n+i,j}^k = 0 \qquad \forall i \in P, k \in K, \tag{3}$$

$$\sum_{i \in V} x_{ij}^k - \sum_{i \in V} x_{ji}^k = 0 \qquad \forall j \in P \cup D \cup \{2n + 1\}, k \in K, \tag{4}$$

$$\sum_{j \in V} x_{0j}^k = z_0^k \qquad \forall k \in K, \tag{5}$$

$$\sum_{i \in V} x_{i,2n+2}^k = z_0^k \qquad \forall k \in K, \tag{6}$$

$$x_{ij}^k \in \{0, 1\} \qquad \forall i, j \in V, k \in K, \tag{7}$$

$$z_0^k \in \{0, 1\} \qquad \forall k \in K. \tag{8}$$

Constraints (2) ensure that each request is served exactly once while (3) ensure that each origin-destination pair is visited by the same vehicle. Flow conservation is imposed by equalities (4). The subsequent equalities (5) and (6) guarantee that, if a driver is assigned to a vehicle, the vehicle starts at and returns to the depot.

Attendant related conditions form another block of constraints:

$$\sum_{i \in V} x_{i,2n+1}^k = \max\left\{ z_1^k, z_2^k \right\} \qquad \forall k \in K, \tag{9}$$

$$Q_0^{0,k} \geq z_1^k \qquad \forall k \in K, \tag{10}$$

$$x_{i,2n+1}^k = 1 \implies Q_{2n+1}^{0,k} \geq Q_i^{0,k} - z_1^k + z_2^k \qquad \forall i \in V, k \in K, \tag{11}$$

$$x_{i,2n+1}^k = 1 \implies Q_{2n+1}^{r,k} \geq Q_i^{r,k} \qquad \forall i \in V\, k \in K, r \in R \setminus \{0\}, \tag{12}$$

$$z_1^k + z_2^k \geq a_i \sum_{j \in V} x_{ij}^k \qquad \forall i \in V, k \in K, \tag{13}$$

$$z_1^k, z_2^k \in \{0,1\} \qquad \forall k \in K. \tag{14}$$

Equalities (9) make sure that the noon depot is used if an attendant is assigned to the corresponding vehicle for morning or afternoon periods. This is necessary to ensure either the appropriate pick-up or drop-off of the attendant at the beginning or at the end of the shift. An attendant assigned to a vehicle requires a staff seat and this is modeled in inequalities (10)-(12). Furthermore, a user demanding an attendant aboard the vehicle can only be visited if an attendant is on the vehicle, which is reflected by constraints (13).

Consistency of resource and load variables is guaranteed by the following constraints:

$$x_{ij}^k = 1 \implies Q_j^{r,k} \geq Q_i^{r,k} + q_j^r \qquad \forall i \in V, j \in V \setminus \{2n+1\}, k \in K, r \in R, \tag{15}$$

$$\sum_{r'=r}^2 Q_i^{r',k} \leq \sum_{r'=r}^2 C^{r',k} \qquad \forall i \in V, k \in K, r \in R \setminus \{3\}, \tag{16}$$

$$Q_i^{3,k} \leq C^{3,k} \qquad \forall i \in V, k \in K, \tag{17}$$

$$Q_i^{r,k} \geq 0 \qquad \forall i \in V, k \in K, r \in R. \tag{18}$$

Inequalities (15) ensure load propagation from one vertex to the next. Upgrading constraints for resources 0, 1, and 2 are given by (16). They guarantee that capacity restrictions regarding these resources are not violated and that each patient demanding resource 0, 1 or 2 can only be loaded if there is either enough capacity of the resource demanded or of another one with a higher number (0 = staff seat, 1 = patient seat, 2 = stretcher). Finally, constraints (17) guarantee that patients demanding resource 3 (wheelchair place) can only be transported if there is enough capacity of resource 3.

Suppose an empty T2 vehicle (2 staff seats, 1 patient seat, 1 stretcher, 1 wheelchair place) visits two origin locations in a row. At each location a seated patient with an accompanying person has to be picked up. The first seated passenger will fill the patient seat, the accompanying person the first staff seat. If it were not possible to have seated patients sit on the stretcher, the second origin location could not be visited. However, the given upgrading conditions allow seated passengers to sit on the stretcher. Therefore, origin location two can be visited. The seated passenger will use the stretcher and the accompanying person the second staff seat. Suppose the same two origin locations were visited by an empty T1 vehicle (1 staff seat, 6 patient seats, 1 wheelchair place). In this case the first seated patient would again use a patient seat and the accompanying person the staff seat. At the second location there is no more empty staff seat available. However, again due to the given upgrading possibilities, the accompanying

person can use a patient seat. Thus, after visiting location two, three patient seats and the staff seat will be occupied.

The next inequalities define the beginning of service for each vertex:

$$x_{ij}^k = 1 \wedge v_i^k = 0 \ \Rightarrow \quad B_j^k \geq B_i^k + d_i + t_{ij} \qquad\qquad \forall i,j \in V, k \in K, \qquad (19)$$

$$x_{ij}^k = 1 \wedge v_i^k = 1 \ \Rightarrow \quad B_j^k \geq B_i^k + d_i + t_{ij} + W_H^k + H \qquad \forall i,j \in V, k \in K, \qquad (20)$$

$$B_{n+i}^k \geq B_i^k \qquad\qquad \forall i \in P, k \in K, \qquad (21)$$

$$B_{2n+2}^k - B_0^k \leq T \qquad\qquad \forall k \in K, \qquad (22)$$

$$\sum_{i \in V} v_i^k \geq z_0^k \qquad\qquad \forall k \in K, \qquad (23)$$

$$v_i^k \in \{0,1\} \qquad\qquad \forall i \in V, k \in K. \qquad (24)$$

If vertex i is chosen for the lunch break ($v_i = 1$), in addition to the service time associated with this vertex, the vehicle will wait (see variable W_H^k) at this vertex until the lunch break time window starts and then stay until the lunch break has been concluded. Note that these constraints also take care of subtour elimination given that $t_{ij} + d_i > 0$ for all $i,j \in V : i \neq j$. Inequalities (21) ensure that every user's pick-up location is visited before the corresponding drop-off location. Total route duration is limited by (22). If a driver is assigned to a vehicle then constraints (23) ensure that the associated route contains a break.

The different limits on the beginning of service are given in the following:

$$e_i \leq B_i^k \leq l_i \qquad\qquad \forall i \in V, k \in K, \qquad (25)$$

$$z_1^k = 1 \wedge z_2^k = 0 \ \Rightarrow \quad B_0^k \leq B_i^k \leq B_{2n+1}^k \qquad \forall i \in V_a, k \in K, \qquad (26)$$

$$z_1^k = 0 \wedge z_2^k = 1 \ \Rightarrow \quad B_{2n+1}^k \leq B_i^k \leq B_{2n+2}^k \qquad \forall i \in V_a, k \in K, \qquad (27)$$

$$v_i^k = 1 \ \Rightarrow \quad e_H \leq B_i^k + d_i + W_H^k \leq l_H \qquad \forall i \in V, k \in K, \qquad (28)$$

$$W_H^k \geq 0 \qquad\qquad \forall k \in K. \qquad (29)$$

Standard time window constraints are modeled by (25). Additional time related constraints provide new bounds on the beginning of service at those vertices where an attendant is required. Bounds for morning periods are given by (26) and for afternoon periods in (27). In case an attendant is present during both periods, no additional bounds are needed. Furthermore, constraints (28) guarantee that the lunch break starts within the lunch break time window $[e_H, l_H]$.

Finally, inequalities (30) and (31) limit the number of drivers and attendants that can be assigned to vehicles. Each driver who is used for attendant duties, instead of driving, can be employed for either morning or afternoon periods, thus increasing the number of attendants available and decreasing the number of drivers by the same amount u:

$$\sum_{k \in K} z_0^k \leq m_d - u, \qquad\qquad (30)$$

$$\sum_{k \in K} z_1^k + \sum_{k \in K} z_2^k \leq m_c + u, \qquad\qquad (31)$$

$$u \in \{0, \ldots, m_d\}. \qquad\qquad (32)$$

3.3 A set-partitioning formulation

The HDARPD can be reformulated in a more compact way. Let $\mathcal{T}$ denote the set of available vehicle types and Ω_t the set of feasible routes for vehicles of type $t \in \mathcal{T}$. Let also Ω be the set

of all feasible routes, i.e., $\Omega = \bigcup_{t \in \mathcal{T}} \Omega_t$. Furthermore, let m_t denote the number of available vehicles of type t. For each route $\omega \in \Omega$, let c_ω be the cost of the route while the constants $b_{i\omega}$ and g_ω represent the number of times vertex $i \in P$ is traversed by ω and the number of attendants needed by route ω, respectively. Finally, variable y_ω takes value 1 if and only if route ω is used in the solution. The problem can thus be formulated as the following set-partitioning problem (SP):

$$\min \sum_{\omega \in \Omega} c_\omega y_\omega \tag{33}$$

subject to

$$\sum_{\omega \in \Omega} b_{i\omega} y_\omega = 1 \qquad\qquad \forall i \in P, \tag{34}$$

$$\sum_{\omega \in \Omega_t} y_\omega \le m_t \qquad\qquad \forall t \in \mathcal{T}, \tag{35}$$

$$\sum_{\omega \in \Omega} y_\omega \le m_d - u, \tag{36}$$

$$\sum_{\omega \in \Omega} g_\omega y_\omega \le m_a + u, \tag{37}$$

$$u \ge 0, \tag{38}$$

$$y_\omega \in \{0,1\} \qquad\qquad \forall \omega \in \Omega. \tag{39}$$

The objective function (33) minimizes the total cost of the selected routes. Constraints (34) guarantee that every request is served exactly once. Inequalities (35) limit the number of vehicles of type t that can be used in the solution. Constraint (36) ensures that at most as many vehicles are used as there are drivers. Drivers that are not needed in routes may be employed as additional attendants, thus increasing the number of attendants that can be used by u in constraint (37).

To compute a lower bound, the linear programming relaxation of SP (LSP) can be solved, where LSP is obtained by replacing (39) with

$$y_\omega \ge 0 \qquad\qquad \forall \omega \in \Omega. \tag{40}$$

Due to the large size of Ω the above formulation will not be solved directly. Instead, a restricted version of this problem, considering only a small subset of columns $\Omega' \subset \Omega$, will be solved. The set Ω' is generated by solving LSP using column generation. In column generation LSP decomposes into a (restricted) master problem and $|\mathcal{T}|$ subproblems, one for each vehicle type. Let π_i, σ_t, λ, and ϕ be the dual variables associated with constraints (34) for index i, with constraints (35) for index t, and with constraint (36) and (37), respectively. The reduced cost of column y_ω corresponding to route $\omega \in \Omega_t$ is given by

$$\bar{c}_\omega = c_\omega - \sum_{i \in P} b_{i\omega} \pi_i - \sigma_t - \lambda - g_\omega \phi. \tag{41}$$

Subproblem t corresponds to finding a single vehicle route $\omega \in \Omega_t$ for a vehicle of type t such that its reduced cost $\bar{c}_\omega$ is minimum. It is subject to constraints (3)-(29), omitting superscript k, setting $z_0 = 1$, and replacing k by t in the case of the capacities $C^{r,t}$. The index t refers to the vehicle type.

4 Solving the Column Generation Subproblem

In order to find negative reduced cost columns we implemented a label setting shortest path algorithm and several heuristics. The use of heuristics, aiding the exact procedure in finding negative reduced cost paths, may yield significant run time reductions (see Savelsbergh and Sol, 1998). Following the findings of Ropke and Cordeau (2009) in the context of the PDPTW, our label setting algorithm considers only elementary paths, i.e. every vertex can be visited at most once in a path. This requires storing and processing additional information during the solution of the subproblem, but usually yields better results than the use of a non-elementary shortest path problem because it provides stronger lower bounds.

In the following the elementary shortest path algorithm will first be described in detail. Thereafter, the heuristics will be briefly discussed.

4.1 The label setting shortest path algorithm

The labeling algorithm implemented is based on the one described by Ropke and Cordeau (2009) for solving the elementary shortest path problem with time windows, capacity, and pickup and delivery. The subproblem we have to solve is also a constrained shortest path problem. It can be described as an elementary shortest path problem with time windows, multiple capacities, pickup and delivery, and route duration constraints.

4.1.1 Lunch break and attendant requirements

To properly treat the lunch break requirements and the presence of attendants on the vehicle, three additional artificial vertices are introduced in the graph: one denoted by $2n + 3$ for the morning attendant, one denoted by $2n + 4$ for the afternoon attendant, and one denoted by $2n + 5$ for the lunch stop. If a path contains the morning attendant vertex, an attendant is aboard the vehicle during the morning shift. Similarly, if the afternoon attendant vertex is part of the path, an attendant is assigned to the vehicle in the afternoon. Every path has to contain the lunch vertex, with one exception: if the vehicle returns to the end depot before the end of the lunch time window and the lunch has not yet taken place, it is assumed that it is held at the end depot. The vertex that is followed by the lunch vertex on the constructed graph is the one where the lunch break will be held.

In the graph, vertex $2n + 3$ can only be visited from the origin depot 0 and $2n + 4$ only from the noon depot $2n + 1$. Travel times for these vertices are set to $t_{0,2n+3} = t_{2n+1,2n+4} = 0$, $t_{2n+3,j} = t_{0,j}$ and $t_{2n+4,j} = t_{2n+1,j}$ for all j. In the case of the lunch vertex, travel times from all vertices to this vertex are set to $t_{i,2n+5} = 0$ for all i. Let now i_{2n+5} denote the vertex visited directly before the lunch vertex. The travel times from the lunch vertex are dynamically set to $t_{2n+5,j} = t_{i_{2n+5},j}$. The time windows of the attendant pickup vertices are set to the beginning and the end of the planning horizon, respectively. In the case of the lunch vertex a time window is given, i.e., $e_{2n+5} = e_H$ and $l_{2n+5} = l_H$ (see above). The service times at the attendant vertices are set to $d_{2n+3} = d_{2n+4} = 0$ and to $d_{2n+5} = H$ in case of the lunch vertex. The load is set to $q^0_{2n+3} = q^0_{2n+4} = 1$ at the attendant vertices and to zero for all other resources of artificial vertices.

4.1.2 Label management

For each label the following data is stored: η - the vertex of the label, δ - the departure time at η, Q^r_{cum} - the cumulative load of resource r when leaving vertex η, c_{cum} - the accumulated cost up to vertex η, $b \in \{0, 1\}$ - whether a lunch stop has already taken place or not, $\alpha \in \{0, 1\}$

- whether an attendant is aboard the vehicle or not, $o \in \{0, 1\}$ - whether the noon depot has already been visited or not, $\mathcal{V} \subseteq \{0, \ldots, 2n + 5\}$ - the set of vertices visited along the path, $\mathcal{O} \subseteq \{1, \ldots, n\}$ - the set of open requests, f - the forward time slack, w_{cum} - the accumulated waiting time, and a pointer to its parent label. The resources f and w_{cum} are needed to check whether the route duration limit can be respected; f gives the maximum time the departure at the noon depot can be shifted forward in time. Section 4.1.4 will explain these resources in further detail.

The extension of a label κ along an arc $(\eta(\kappa), j)$ is only possible if the following holds:

$$\delta(\kappa) + t_{\eta(\kappa),j} \leq l_j, \tag{42}$$

$$Q_{cum}^r(\kappa) + q_j^r + \sum_{r'=r+1}^{2} q_j^{r'} \leq C_{cum}^{r,t} \quad \forall r \in R, \tag{43}$$

$$\alpha(\kappa) \geq a_j, \tag{44}$$

$$(1 - b(\kappa))(\max\left\{\delta(\kappa) + t_{\eta(\kappa),j}, e_j\right\} + d_j) \leq l_H, \tag{45}$$

$$\max\left\{\max\left\{\delta(\kappa) + t_{\eta(\kappa),j}, e_j\right\} + d_j, (1 - b(\kappa))e_H\right\} + (1 - b(\kappa))H - e_0 - F_j^0 \leq T, \tag{46}$$

$$j \notin \mathcal{V}(\kappa). \tag{47}$$

Here, $C_{cum}^{r,t}$ is set to $C^{r,t} + \sum_{r'=r+1}^{2} C^{r',t}$. The final time slack F_j^0 at vertex j is given by $F_j^0 = \min\left\{\min[f(\kappa), l_j - (\delta(\kappa) + t_{\eta(\kappa),j}) + w_{cum}(\kappa)], w_{cum}(\kappa) + \max(0, e_j - (\delta(\kappa) + t_{\eta(\kappa),j}))\right\}$, i.e. the minimum over all forward time slacks and the total waiting time until vertex j (see Section 4.1.4 for further details). Condition (42) ensures time window feasibility: the departure from the previous vertex $\eta(\kappa)$ plus the travel time from $\eta(\kappa)$ to j has to be smaller than or equal to the end of the time window at vertex j. According to condition (43) a path can only be extended along arc $(\eta(\kappa), j)$ if all loading restrictions are satisfied. Condition (44) states that if vertex j demands an attendant aboard the vehicle ($a_j = 1$), it can only be visited if an attendant is currently on the vehicle ($\alpha(\kappa) = 1$). In the case where the lunch node is not yet part of the path, extension along arc $(\eta(\kappa), j)$ is only possible if it can still be feasibly inserted after vertex j. This is taken care of by (45). Feasibility with respect to route duration is guaranteed by condition (46). For further details on this issue we refer to Section 4.1.4. Finally, elementarity is ensured by (47).

Moreover, κ and j must comply with the following conditions:

$$j \in D \Rightarrow j - n \in \mathcal{O}(\kappa), \tag{48}$$

$$j = 2n + 2 \Rightarrow \mathcal{O}(\kappa) = \emptyset \wedge \alpha - o \leq 0. \tag{49}$$

Condition (48) ensures that if j is a delivery, it can only be visited if the request is open, i.e. the corresponding pickup has already been visited. Condition (49) ensures that a label can only be extended to the end depot $2n + 2$ if there are no more open requests and if the noon depot was visited, in case an attendant is currently on the vehicle.

If a label can feasibly be extended along arc $(\eta(\kappa), j)$, a new label κ' is generated at vertex j:

$$\eta(\kappa') = j, \tag{50}$$

$$\delta(\kappa') = \max\left\{\delta(\kappa) + t_{\eta(\kappa),j}, e_j\right\} + d_j, \tag{51}$$

$$Q_{cum}^0(\kappa') = \begin{cases} Q_{cum}^0(\kappa) - 1 & \text{if } j = 2n + 1 \wedge \alpha(\kappa) = 1, \\ Q_{cum}^0(\kappa) + q_j^0 + \sum_{r'=1}^{2} q_j^{r'} & \text{otherwise,} \end{cases} \tag{52}$$

$$Q_{cum}^r(\kappa') = Q_{cum}^r(\kappa) + q_j^r + \sum_{r'=r+1}^{2} q_j^{r'} \quad \forall r \in R \setminus \{0\}, \tag{53}$$

$$c_{cum}(\kappa') = c_{cum}(\kappa) + \bar{c}_{\eta(\kappa),j}, \tag{54}$$

$$b(\kappa') = \begin{cases} 1 & \text{if } j = 2n+5, \\ b(\kappa) & \text{otherwise}, \end{cases} \tag{55}$$

$$\alpha(\kappa') = \begin{cases} 1 & \text{if } j \in \{2n+3, 2n+4\}, \\ 0 & \text{if } j = 2n+1, \\ \alpha(\kappa) & \text{otherwise}, \end{cases} \tag{56}$$

$$o(\kappa') = \begin{cases} 1 & \text{if } j = 2n+1, \\ o(\kappa) & \text{otherwise}, \end{cases} \tag{57}$$

$$\mathcal{V}(\kappa') = \mathcal{V}(\kappa) \cup \{j\}, \tag{58}$$

$$\mathcal{O}(\kappa') = \begin{cases} \mathcal{O}(\kappa) \cup \{j\} & \text{if } j \in P, \\ \mathcal{O}(\kappa) \setminus \{j-n\} & \text{if } j \in D, \\ \mathcal{O}(\kappa) & \text{otherwise}, \end{cases} \tag{59}$$

$$w_{cum}(\kappa') = w_{cum}(\kappa) + \max\left\{0, e_j - (\delta(\kappa) + t_{\eta(\kappa),j})\right\} \tag{60}$$

$$f(\kappa') = \min\left\{f(\kappa), w_{cum}(\kappa) + l_j - (\delta(\kappa) + t_{\eta(\kappa),j})\right\}. \tag{61}$$

At the origin depot (the first vertex along each path) these labels are initialized by setting $\delta(0) = e_0$, $Q_{cum}^r(0) = 0$ for all $r \in R$, $c_{cum} = 0$, $b(0) = 0$, $\alpha(0) = 0$, $o = 0$, $f(0) = l_0 - e_0$, and $w_{cum}(0) = 0$.

4.1.3 Dominance

The dominance criterion used here is similar to the one denoted by DOM1' by Ropke and Cordeau (2009). Let $\mathcal{U}(\kappa)$ denote the set of unreachable requests of label κ, where $\mathcal{U}(\kappa) = \mathcal{V} \cup \left\{i \in P : \delta(\kappa) + t_{\eta(\kappa),i} > l_i\right\}$. According to this criterion a label κ dominates another label κ' if

$$\eta(\kappa) = \eta(\kappa'), \delta(\kappa) \leq \delta(\kappa'), c_{cum}(\kappa) \leq c_{cum}(\kappa'), \mathcal{U}(\kappa) \subseteq \mathcal{U}(\kappa'), \mathcal{O}(\kappa) \subseteq \mathcal{O}(\kappa'). \tag{62}$$

In addition, in our case, the following has to hold:

$$\min\left\{f(\kappa), w_{cum}(\kappa)\right\} \geq \min\left\{f(\kappa'), w_{cum}(\kappa')\right\}, \tag{63}$$

$$b(\kappa) = b(\kappa'), \alpha(\kappa) = \alpha(\kappa'), o(\kappa) = o(\kappa'). \tag{64}$$

This implies that the amount of time by which the departure from the origin depot can be shifted forward in time in label κ has to be at least as large as in label κ'. If the lunch node or the noon depot were already visited by the partial path represented by label κ the same has to be true for κ'. The same applies to whether an attendant is aboard the vehicle or not. Furthermore, we only apply the dominance check to labels with $\eta \in V$.

4.1.4 Time windows at the start depot and minimum route duration

Imposing a route duration limit together with a non-empty time window at the origin depot in a label setting algorithm requires further adjustments. One option, introduced by Desaulniers

and Villeneuve (2000) and reviewed by Irnich (2008), consists in appending two resources to each label, coupled by a max-term, denoted as q and z. These are extended as follows:

$$q(\kappa') = t_{\eta(\kappa),j} + \max\left\{q(\kappa) - (t_{\eta(\kappa),j} + d_{\eta(\kappa)}), z(\kappa) - l_j\right\} \tag{65}$$

$$z(\kappa') = t_{\eta(\kappa),j} + \max\left\{z(\kappa), q(\kappa) - (t_{\eta(\kappa),j} + d_{\eta(\kappa)}) + e_j\right\}. \tag{66}$$

Here, we show that this is equivalent to using the forward time slack notion developed by Savelsbergh (1992).

The idea of Desaulniers and Villeneuve (2000) can be explained as follows. Let $U = \{0, 1, 2, \ldots, \tilde{n}\}$ denote a feasible path. The earliest possible departure time $\tilde{e}_j$ at node j when traveling along U can be calculated as

$$\tilde{e}_0 = e_0, \tag{67}$$

$$\tilde{e}_j = \max\left\{\tilde{e}_{j-1} + (d_{j-1} + t_{j-1,j}), e_j\right\} \qquad \forall j \in \{1, \ldots, \tilde{n}\}. \tag{68}$$

Similarly, the latest arrival time $\tilde{l}_j$ at node j for which waiting can be avoided can be computed by setting

$$\tilde{l}_0 = l_0, \tag{69}$$

$$\tilde{l}_j = \min\left\{\tilde{l}_{j-1} + (d_{j-1} + t_{j-1,j}), l_j\right\} \qquad \forall j \in \{1, \ldots, \tilde{n}\}. \tag{70}$$

Let now U be a feasible (partial) path starting at the origin depot 0 and ending at node i. Then, let $s_i = \sum_{(k,l) \in U}(t_{kl} + d_k)$ denote the sum of actual travel times along U (including service time but excluding waiting time), the following two parameters can be defined:

$$\tilde{q}_i = s_i - \tilde{l}_i, \tag{71}$$

$$\tilde{z}_i = \max\left\{s_i, \tilde{q}_i + \tilde{e}_i\right\}. \tag{72}$$

Then the minimum duration of path U (ending at i) is equal to $\max\{\tilde{z}_i, \tilde{q}_i + \delta_i\}$, where δ_i denotes the departure time from vertex i.

Now it can be shown that $\tilde{q}_i$ is equivalent to the forward time slack defined by Savelsbergh (1992). Let $\tilde{f}_i^0$ denote the forward time slack from the origin depot to the end of the path U (here node i). It is the maximum amount of time by which the departure at the depot can be shifted forward without violating any other time window constraint. It is computed as follows (where B_j denotes the beginning of service at node j):

$$\tilde{f}_i^0 = \min_{0 \leq j \leq i}\left\{l_j - \left[B_0 + \sum_{p=1}^{j}(d_{p-1} + t_{p-1,p})\right]\right\}. \tag{73}$$

Let us assume that $e_0 = 0$ and therefore $B_0 = 0$. We obtain,

$$\tilde{f}_i^0 = \min_{0 \leq j \leq i}\left\{l_j - \sum_{p=1}^{j}(d_{p-1} + t_{p-1,p})\right\}. \tag{74}$$

Let us now rewrite $\tilde{l}_i$ as

$$\tilde{l}_i = \min_{0 \leq j \leq i}\left\{l_j + \sum_{p=j+1}^{i}(d_{p-1} + t_{p-1,p})\right\}, \tag{75}$$

and replace $\tilde{l}_i$ by this formulation in $\tilde{q}_i$. We obtain

$$\tilde{q}_i = \sum_{p=1}^{i}(d_{p-1} + t_{p-1,p}) - \min_{0 \leq j \leq i}\left\{l_j + \sum_{p=j+1}^{i}(d_{p-1} + t_{p-1,p})\right\} \tag{76}$$

$$= -\min_{0 \leq j \leq i}\left\{l_j - \sum_{p=1}^{j}(d_{p-1} + t_{p-1,p})\right\}. \tag{77}$$

This shows that $\tilde{q}_i = -\tilde{f}_i^0$. Furthermore, let us examine the minimum route duration time as defined by Desaulniers and Villeneuve (2000) and given by $\max\{\tilde{z}_i, \tilde{q}_i + \delta_i\}$. We still assume that $B_0 = e_0 = 0$. By substituting $\tilde{z}_i$ this term can be rewritten as $\max\{s_i, \tilde{q}_i + \tilde{e}_i, \tilde{q}_i + \delta_i\} = \max\{s_i, \tilde{q}_i + \max(\tilde{e}_i, \delta_i)\}$. If δ_i is computed correctly it will always be greater than or equal to $\tilde{e}_i$ in our case. Therefore, $\max\{s_i, \tilde{q}_i + \delta_i\}$ is equivalent to the above expression. Now if we denote by $\tilde{w}_i$ the accumulated waiting time until node i, it is easy to see that $\delta_i - \tilde{w}_i = s_i$. Thus, $\max\{s_i, \tilde{q}_i + \delta_i\} = \delta_i - \min(-\tilde{q}_i, \tilde{w}_i) = \delta_i - \min(\tilde{f}_i^0, \tilde{w}_i)$. This corresponds to what Cordeau and Laporte (2003) use to compute the minimum duration of a given route.

Thus, to handle a time window at the start depot together with a route duration limit, the notion of forward time slack defined by Savelsbergh (1992) can be used. In order to do so, we use the additional resources w_{cum} and f (see above). To generalize what has been shown to the case of $e_0 > 0$, these resources are initialized, as already pointed out, with $f(0) = l_0 - B_0$ ($B_0 = e_0$) and $w_{cum}(0) = 0$.

4.1.5 Label elimination

Labels that are currently in the queue of labels to be processed can be eliminated if the deliveries of open requests cannot be reached in a feasible path. As in Ropke (2005) and Ropke and Cordeau (2009) we consider sets of one and two deliveries and one set of three deliveries. The last set consists of the delivery that is furthest away from the current vertex, the delivery that is furthest away from these two, and the delivery that is furthest away from the current vertex and the two previously selected deliveries. If for one of these sets of deliveries no path can be found that serves all deliveries in the set in a feasible way, the label can be eliminated.

In addition to these sets of deliveries, we check whether there is an attendant on the vehicle and whether the noon depot has not been visited yet. If this is the case and the noon depot cannot be reached in a feasible way from the current vertex, the label can also be eliminated.

4.2 Heuristic algorithms

To accelerate the column generation process several heuristic procedures are used to generate negative reduced cost columns. These can be divided into two classes; those based on the labeling algorithm, and those that rely on simple construction and improvement principles.

4.2.1 Heuristics based on the labeling algorithm

Two of the heuristic algorithms used to generate columns are based on the exact labeling algorithm. Both were also used by Ropke and Cordeau (2009) in the context of the PDPTW. The first heuristic (*LimLabels*) simply limits the number of labels that can be in the queue of unprocessed labels at any time. At first the limit is set to 500. If no negative reduced cost column can be found with this limit, it is increased to 1000. If again no negative reduced cost columns are generated the limit is set to 2000. The second heuristic (*LimGraph*) applies the exact labeling algorithm on a reduced graph. Two reduced graphs are used. In Graph G_5 every

pickup and delivery vertex is only connected to the five closest pickup vertices and the five closest delivery vertices. In addition, if not already present, connections between each pickup vertex and its corresponding delivery vertex are added, and the start depot is connected to all pickup vertices and the morning attendant node; the noon depot is connected to the afternoon attendant node; all delivery vertices are connected to the end depot. All vertices except the end depot are connected to the lunch node and the noon depot. The morning attendant node, the afternoon attendant node and the lunch node are connected to all vertices except the start depot. The second reduced Graph G_{10} is constructed in the same way. However, instead of the five closest pickups and deliveries, it considers the ten closest.

4.2.2   Construction/improvement based heuristics

The four remaining pricing heuristics use construction or improvement algorithms. Like the labeling heuristics described above, they are based on those used by Ropke and Cordeau (2009) to solve the PDPTW. *ConstrHeur* is a simple construction heuristic that starts from every pickup and delivery vertex pair and iteratively adds requests by means of a best insertion criterion regarding reduced cost. *RandConstrHeur* is also a construction heuristic but here randomized best insertion is used to insert additional requests. The randomization process favors requests that least increase the reduced costs of the partial route. In both construction heuristics, every time a new request is inserted, the resulting route is checked to see whether it has a negative reduced cost. If this is the case, the corresponding column is generated. After the check, the route undergoes local search based improvement (see Section 6.3), minimizing actual routing cost, and considering only moves that yield a feasible route. In the case where this results in another negative reduced cost route, the corresponding column is again added to the pool. Heuristic *LNSCurrBasis* applies Large Neighborhood Search (LNS) (Shaw, 1998) on the routes in the current basis. It works as follows. In a removal step, up to 50% of the requests forming the route are randomly removed from it. In an insertion step, requests are reinserted using randomized best insertion as described above. These two steps are repeated until no further improvement can be found. Between 15 and 20 non-improving steps are performed. Finally, *LNSRandConstr* simply improves all solutions obtained by the randomized construction algorithm with LNS.

5   The Column Generation Framework

The overall column generation framework that we use to solve the HDARPD is summarized in Algorithm 1. During the initialization phase initial columns are generated (see Section 5.1) and added to Ω', and a number of pre-processing steps (see Section 5.2) are applied. Then, LSP is solved on Ω' and the dual variable values associated with the different constraints are retrieved. Based on these values, negative reduced cost columns are generated using the different heuristics in a certain sequence (see Section 5.3) and if all fail, then the exact dynamic programming procedure is called. All new negative reduced cost columns are added to Ω' and the LSP is solved again. This is repeated until no new negative reduced cost column can be identified. In this case, the optimal solution to LSP has been found. Additionally, every 10 iterations the proposed VNS heuristic (see Section 6) is applied to the current solution of LSP. This process is represented in Figure 1. The so-called collaborative scheme is described in further detail in Section 5.4.

Algorithm 1 The column generation framework

    initialization generate initial columns by VNS, introduce artificial columns and add them to Ω', do pre-processing (graph pruning, time window tightening)

    repeat

      solve LSP on Ω'

      every 10 iterations apply VNS to the current LSP solution (collaborative scheme)

      generate new negative reduced cost columns heuristically (*ConstrHeur*, *LNSCurrBasis*, *RandConstrHeur*, *LNSRandConstr*, *LimLabels*, *LimGraph*)

      if no negative reduced cost columns are found **then**

        run exact label setting algorithm to find negative reduced cost columns

      end if

      add new negative reduced cost columns to Ω'

    until no more negative reduced cost columns can be found

    return optimal solution of LSP

Figure 1: The collaborative scheme

5.1 Initial columns

An initial set of columns is generated by means of a heuristic algorithm, namely a VNS. The VNS is described in detail in Section 6. A limit of 2×10^4 iterations is applied regardless of whether a feasible solution can be found within this time limit or not. Here, ascending moves are not allowed. In addition, one artificial column for each $i \in P$ is generated, having a coefficient of 1 in the row corresponding to the demand constraint for request i and zeros in all other rows. These columns are given a sufficiently large cost of M.

5.2 Pre-processing

Before starting the column generation process, several pre-processing steps are performed. These are mostly based on time window tightening and graph pruning techniques. They are based on those described in detail by Cordeau (2006). In addition, we use the cyclic time window tightening steps described by Desrochers et al. (1992) and Kallehauge et al. (2005):

$$e_k := \max \left\{ e_k, \min[l_k, \min_{(i,k)}(e_i + (t_{ik} + d_i))] \right\}, \tag{78}$$

$$e_k := \max \left\{ e_k, \min[l_k, \min_{(k,j)}(e_j - (t_{kj} + d_k))] \right\}, \tag{79}$$

$$l_k := \min\left\{l_k, \max[e_k, \max_{(i,k)}(l_i + (t_{ik} + d_i))]\right\}, \tag{80}$$

$$l_k := \min\left\{l_k, \max[e_k, \max_{(k,j)}(l_j - (t_{kj} + d_k))]\right\}. \tag{81}$$

These steps are repeated until no further time window tightening is possible. Note that the only arcs considered are those that have not been eliminated in previous pre-processing steps. Arcs that would lead to an infeasible solution regarding vehicle capacity restrictions are also removed from the graph.

5.3 Sequence of pricing heuristics

The first pricing heuristic invoked to find new negative reduced cost columns is selected using a roulette wheel mechanism. Initially, only *ConstrHeur* can be chosen. Its score is set to one, while the score of all others is set to zero. Thereafter, in every column generation iteration the score of the heuristic that obtained one or more new negative reduced cost columns is increased by one. Thus, its probability of being chosen as the first heuristic to be tried is increased. The heuristics are ordered as in Ropke and Cordeau (2009): *ConstrHeur - LNSCurrBasis - RandConstrHeur - LNSRandConstr - LimLabels - LimGraph*. If *ConstrHeur* is chosen as the first heuristic but it fails to yield a new negative reduced cost column for either vehicle types, we switch to *LNSCurrBasis*. If *LNSCurrBasis* fails we switch to *RandConstrHeur* and so on. If *LNSCurrBasis* is chosen as the first heuristic and fails to obtain negative reduced cost columns, only the heuristics following *LNSCurrBasis* in the list are tried in the above order. This applies to all heuristics. If also *LimGraph* (the last one on the list) fails to generate negative reduced cost columns, the exact label setting algorithm is started. If this procedure also fails to generate negative reduced cost columns, the optimal solution of the relaxed problem (LSP) has been found. In the case where it is also integer the optimal solution of SP has been found. In all labeling algorithms we stop as soon as 50 negative reduced cost columns have been generated. All labeling algorithms use a sorted queue of unprocessed labels. Labels are ordered according to increasing reduced cost.

LSP is solved again on Ω' every time at least one new negative reduced cost column for either vehicle type could be generated. Every heuristic tries to find negative reduced cost columns for T1 and T2 vehicles in alternating order. If new negative reduced cost columns for T1 vehicles are generated, before re-solving LSP on the updated Ω', their validity for T2 vehicles is checked. In case they are valid, they are added to the column pool for T2 vehicles. Only then LSP is solved. This is not done if one of the heuristics yields new negative reduced cost columns for T2 vehicles; it is not very likely that columns for T2 vehicles are also valid for T1 vehicles.

5.4 Collaborative scheme

Every 10 iterations (one iteration corresponds to finding one or more negative reduced cost columns, adding these columns to Ω', and solving LSP) the optimal solution of the current LSP is passed to the VNS. If there is still an artificial column in the basis the VNS resumes the search using the last incumbent of the previous run. If there are no more artificial columns in the basis, two different scenarios may occur: either the solution is integer and it is passed to the VNS as is, or the solution is fractional. In the latter case, all duplicate requests are removed before passing it on to the VNS. Duplicate requests are kept on the route that is associated with the y_ω closest to one. Empty routes are eliminated. In this case there can be more vehicles in use of a certain type, due to the fractionality of the solution, than actually available. However, it is still passed to the VNS keeping all increased limits on the number of routes per vehicle type

and setting the number of drivers available to $\bar{m}_d = \max(m_d, m_{cg} - 1)$ (m_{cg} gives the number of routes used after having removed duplicate requests). Thus, at most as many drivers as there are currently needed minus one or, in case this number is smaller than the original limit, at most the original number of drivers can be used. A new best solution generated by the VNS might be infeasible regarding the number of vehicles of a certain vehicle type employed and in terms of the number of drivers, but no other constraint. This entails that the resulting columns might not be combinable as such but they are all feasible. If the solution obtained by the VNS is of lower cost (actual costs not reduced costs) than the current solution of the master problem after being converted into an integer solution, the corresponding routes are transformed into columns and added to the master. Here the VNS is run for 10^4 iterations and ascending moves are not considered, i.e. a deteriorating solution cannot become a new incumbent solution (see Section 6.4). This collaborative scheme is inspired by Danna and Le Pape (2005). In contrast to Danna and Le Pape (2005), however, we do not only use the collaborative local search method to improve the best integer solution found so far. We use it to improve the current optimal solution, regardless of whether it is integer of fractional.

6 Variable Neighborhood Search

This section introduces the VNS heuristic that we have developed to solve the HDARPD. VNS was introduced by Mladenović and Hansen (1997) and, as summarized in Algorithm 2, it consists of four major design elements: the initialization phase, the shaking operators, the iterative improvement mechanism, and the decision whether to move to the new solution or not.

Algorithm 2 Variable neighborhood search

    initialization {determine an initial solution s and set $h \leftarrow 1$}
    repeat
        shaking {determine a solution s' in the neighborhood h of s}
        iterative improvement {apply local search to s' yielding s''}
        move or not {if s'' meets the acceptance requirements the incumbent solution s is replaced
        by s'' and $h \leftarrow 1$, otherwise $h \leftarrow (h \mod h_{max}) + 1$; if s'' is feasible and better than s_{best},
        set $s_{best} \leftarrow s''$}
    until some stopping criterion is met
    return s_{best}

In the initialization phase a first incumbent solution s is generated. In the repeat loop the algorithm iterates through neighborhoods of different sizes. At each iteration a solution s' is generated at random in the neighborhood of the current incumbent solution s (shaking phase). The new solution s' may be improved by means of an iterative improvement procedure (local search) yielding s''. If s'' meets the acceptance criteria (move or not) it replaces s and becomes the new incumbent. In this case or whenever the largest neighborhood size is reached, the search continues with the smallest neighborhood. If s'' does not constitute a new incumbent, the size of the neighborhood is increased.

Infeasibilities are allowed during the search and they are penalized through the following evaluation function:

$$\hat{f}(s) = c(s) + \sum_{r \in R} \hat{\alpha}_r q_r(s) + \hat{\beta}d(s) + \hat{\gamma}w(s) + \hat{\zeta}a(s). \tag{82}$$

The term $c(s)$ gives the routing cost of solution s. The terms $q_r(s)$, $d(s)$, $w(s)$, and $a(s)$ represent load violations ($\forall r \in R$), duration violations, time window violations, and attendant

violations (if there are more attendants needed in s than there are available), respectively. The associated penalty parameters $\hat{\alpha}_r$, $\hat{\beta}$, $\hat{\gamma}$, and $\hat{\zeta}$ are dynamically adjusted throughout the search (see Parragh et al. (2010) and Parragh (2009a) for details). Note that a solution s can only become a new best solution s_{best} if $q_r(s) = d(s) = w(s) = a(s) = 0$ for all $r \in R$.

This VNS is based on the ones developed by Parragh et al. (2009, 2010) and Parragh (2009a) for the standard DARP and the multi-objective DARP. Its different design elements are described in further detail in the following.

6.1 Initialization

The initialization procedure generates an initial solution which will often be infeasible. It is constructed as follows. In a first step the average number of requests per vehicle is computed and rounded to the next integer. Then, all requests are inserted into routes in the order in which they appear in the instance file, starting with the first route of the first type, opening the next route as soon as the average number of requests per vehicle has been reached. If a request does not fit into a route due to a lack of the demanded resource on the current vehicle, it is inserted into the next vehicle route with this resource available. The procedure ends as soon as all requests have been inserted into some route. Finally all routes are checked to see whether a request demands an attendant aboard the vehicle. If this is the case, the noon depot is inserted at the best possible position and an attendant is assigned to the corresponding vehicle for the corresponding shift (morning or afternoon).

6.2 Shaking

During the shaking phase four different neighborhood operators are employed: the first swaps two sequences of vertices, the next two are based on the move operator, and the last one makes use of the ejection chain idea described by Glover (1996). A swap, a move, and an ejection chain neighborhood operator in combination with a "zero split" neighborhood have been successfully employed by Parragh et al. (2009) in a solution framework for the multi-objective DARP, and, in a refined version, by Parragh et al. (2010) in the context of the standard DARP. The neighborhood operators employed in the current work are based on those. They had to be adapted, however, in order to cope with the special characteristics of the HDARPD.

Swap neighborhood   In the swap neighborhood (S) two sequences of vertices are exchanged. Two routes are chosen at random from all routes currently in use (vehicles not assigned to a driver are not considered). Then, the starting points and the lengths of the sequences to be moved are randomly selected. Insertion into the new routes is conducted in a one by one fashion: considering each request at a time and inserting it into its best position, using the notion of *critical vertices* as described by Cordeau and Laporte (2003). Note that if a sequence of vertices is moved their corresponding origins (destinations) have to be moved as well. The neighborhood size in this context refers to the maximum length (number of vertices) of the two sequences.

Move neighborhoods   The first move neighborhood (M) consists in moving requests from their current routes to other routes. First, the requests to be moved are randomly chosen from the set of all requests. Second, for each request to be moved an insertion route is determined across all routes, excluding the request's original route. At most one additional route that is currently not assigned to a driver can be considered. As described in further detail in Parragh (2009a), insertion routes are either selected randomly or the "closest" route in terms of spatial distance is taken. Random and "closeness" selection are chosen with a probability of 0.5 each.

Third, the selected requests are moved to the best position in their new routes (again using the notion of critical vertices). The size of this neighborhood structure is defined by the maximum number of requests moved.

The second move neighborhood (Mx) distinguishes itself from the first one in the way the insertion routes are selected. Here only random selection is employed. Furthermore, all routes, also those lacking a driver, are eligible for selection.

Chain neighborhood   The third neighborhood is referred to as chain neighborhood (C). It works in a similar way as the ejection chain neighborhood defined by Glover (1996). In contrast to its original version, the number of routes affected is a parameter. From the first route a sequence of requests (selected as in the swap neighborhood) is moved to the second route. In a second step, the sequence that decreases the evaluation function value of the second route by the most is moved to a third route (it may also be the first route). The second step is repeated until the maximum number of sequences moved has been reached. All insertions of sequences into their new routes are done one by one in the best possible way using the notion of critical vertices. All routes are selected randomly. Here the neighborhood size represents at the same time the number of sequences moved and the maximum sequence length. Again, empty routes may be selected.

Repair function   Both of the move operators as well as the chain neighborhood operators may construct a solution with more vehicles in use than drivers available. Such a solution has to be "repaired". The repair procedure employed here simply chooses one route at random out of all non-empty routes and redistributes the requests forming this route to other non-empty routes. Request insertion is done one by one in the best possible way, again using the notion of critical vertices. The repair procedure is repeated until the number of vehicles in use meets the number of available drivers.

Neighborhood sequence   The shaking operators described above are applied in the following order: S1 – M1 – C1 – S2 – M2 – C2 – S3 – M3 – C3 – S4 – M4 – C4 – Mx4 (the number given in addition to the neighborhood abbreviation indicates the neighborhood size). This means that the first neighborhood ($h = 1$) corresponds to applying the shaking operator S1, whereas, e.g. in case of $h = 6$ operator C2 will be used. In Parragh et al. (2009, 2010), the last neighborhood is the parameterless zero split neighborhood which distributes requests forming a natural sequence to other routes. Here, Mx4 is the last neighborhood. In combination with the repair function it is the strongest diversification mechanism in place, regarding the number of request relocated together with the amount of change possible in terms of the number of vehicles of each vehicle type employed.

6.3   Iterative improvement

At every iteration, after moving requests and repairing the resulting solution, the obtained solution may undergo local search based improvement. As in Parragh et al. (2010), if $c(s') < 1.02c(s)$, s is automatically subject to local search based improvement, otherwise it has a 1% chance to undergo local search. If s' meets the acceptance criteria to become the new incumbent solution it is also subject to local search based improvement, given that $c(s') > 1.05c(s)$. In any case only those routes that are affected by the shaking step are improved. Here again the notion of critical vertices is used. The search moves along each of these routes. At every origin it removes the corresponding request and reinserts it, starting at the beginning of the route, at every possible position. If an improving position is found, the request is kept at its new position

and the search restarts at the first origin on the route (first improvement). If no improving position is encountered, the request is kept at its original position and the search continues with the next origin on the route. The improvement heuristic terminates as soon as the last origin on the route has been reached and no improving position has been found.

6.4 Move or not

The decision as to whether the search moves to the new solution s'' or not is based on a simulated annealing acceptance criterion (Kirkpatrick et al., 1983, Černy, 1985). As in Parragh et al. (2010), in the case where $\hat{f}(s'') < \hat{f}(s)$, s'' is always accepted. In the case where $\hat{f}(s'') \geq \hat{f}(s)$, s is accepted with a probability equal to $\exp(-\frac{\hat{f}(s'')-\hat{f}(s_{best})}{\hat{t}})$. The temperature $\hat{t}$ is set such that as soon as the first feasible solution has been identified, a solution that is 0.5% worse than the current best solution is accepted with probability 0.2. Before the first feasible solution has been identified every solution that has an evaluation function value $\hat{f}(s'') \leq 1.05\hat{f}(s)$ is accepted with a probability of 1%.

6.5 Route evaluation

Every time a route is modified its routing cost and constraint violations have to be evaluated. Here the notion of forward time slack, as for the label setting algorithm in Section 4.1.4, is applied. In a first step the lunch location is determined. Thereafter the forward time slack is calculated. In a separate procedure the noon depot is either inserted in the best possible way or removed, in case there are no more requests on the route that demand an attendant. In the local search step, the noon depot insertion/removal procedure is only invoked at the very end. The same applies to the request insertion routine. Whenever a request is removed from a route it is also invoked.

7 Computational Experiments

All programs were implemented in C++. In the column generation framework, the LP solver of CPLEX 11.1 was used together with Concert Technology 2.6. All experiments were carried out on a 3.2 GHz Pentium D computer with a memory of 4 GB. Both solution procedures have been tested on three artificial data sets and on real-world data. In the following we first describe the characteristics of the different instances and we then discuss the results obtained.

7.1 Artificial instances

For each instance of the data set "A" introduced by Cordeau (2006), containing between 16 and 48 requests and between 2 and 4 vehicles, three instances with different degrees of heterogeneity have been generated. The characteristics of these instances are summarized in Table 1. Setting "X" is the most homogeneous one: 50% of the original users have been converted into seated passengers; 25% into patients on stretchers; and 25% into persons needing a wheelchair. The probability that an accompanying person (AP) is present has been set to zero. The number of attendants available has been randomly set to between 0.5 of and 1.5 times the number of drivers available (rounded up to the next integer). The probability for an attendant to be demanded by a patient has been set to 25%. For setting "Y", 25% of the original users have been transformed into seated patients, 25% into patients on stretchers and 50% into wheelchair patients; 10% of all patients are assumed to be accompanied by someone. The number of attendants has been randomly set to at least the number of drivers available and at most to twice the number of drivers. The probability for an attendant to be demanded by a patient has been set to 50%.

In the third setting, denoted by "Z", 83% of the patients are assumed to be seated, 11% are assumed to be on a stretcher, and 6% are assumed to be in a wheelchair. This setting is based on the data provided by the ARC. Finally, the probability for a patient to be accompanied by someone has been set to 50%. All other settings are equal to those of data set "Y". In all instances the vehicle fleet consists of 2 T1 and 4 T2 vehicles and the number of available drivers has been set to the original number of vehicles. At the start depot a 60-minute time window has been set, and the maximum route duration has been reduced by 60 minutes with respect to the original data. As before it is equal for all vehicles. The time window at the noon depot has been set to $e_{2n+1} = e_0 + T/2$ and $l_{2n+1} = e_{2n+1} + 15$.

Table 1: Artificial instances - data

Data set	Probability for patient to be			Probability for AP	Probability AT demanded	m_a	fleet
	seated	on stretcher	in wheelchair				
X	0.50	0.25	0.25	0.00	0.25	$\lceil m(0.5 + \rho) \rceil$	2 T1, 4 T2
Y	0.25	0.25	0.50	0.10	0.50	$\lceil m(1 + \rho) \rceil$	2 T1, 4 T2
Z	0.83	0.11	0.06	0.50	0.50	$\lceil m(1 + \rho) \rceil$	2 T1, 4 T2

AP = accompanying person, AT = attendant
ρ randomly chosen in $[0, 1]$
T1: 1 staff seat, 6 patient seats, 1 wheelchair place
T2: 2 staff seats, 1 patient seat, 1 stretcher, 1 wheelchair place

7.2 Real-world instances

We have also applied the different solution algorithms to 15 real-world instances from the ARC. They possess the following characteristics. As in data set "Z" 83% of the passengers are seated patients, 11% have to be transported on a stretcher, and 6% in a wheelchair; 50% of all these passengers take an accompanying person with them and about 40% demand additional personnel (an attendant) on the vehicle. Three T1 vehicles and 31 T2 vehicles are available. The maximum route duration (driver working hours) is 510 minutes. The lunch break has to start between 11 am and 2 pm. It lasts 30 minutes. Every driver starts working between 6:30 and 8:30 am. These two points in time give the time window at the start depot. The time window at the noon depot lasts from 12:30 to 1:00 pm. Users specify a time window for either the pick-up or the drop-off location. Time window width is equal to 30 minutes. Maximum user ride times have been set to $L_i = t_{i,n+i} + 30$ for all $i \in P$. As mentioned above, ride time limits are not explicitly considered; depending on which time window has been specified by the user, the time window for the corresponding location without time window is set relative to the existing time window; in the case of a time window at the destination, it is set to $e_i = e_{n+i} - (L_i + d_i)$ and $l_i = l_{n+i} - (t_{i,n+i} + d_i)$ at the origin; in the case of a time window at the origin, it is set to $e_{n+i} = e_i + d_i + t_{i,n+i}$ and $l_{n+i} = l_i + d_i + L_i$ at the corresponding destination.

7.3 Column generation results

In a first step two versions of the column generation framework were tested on the artificial data sets. The first version only uses pure column generation, while in the second version the collaborative scheme is employed. Table 2 gives the results for pure column generation while Table 3 gives the results for the collaborative scheme. The following information is provided. First, the time needed to compute the initial VNS solution and the total run time, excluding the initial VNS, of the respective program is given. Then, the lower bounds and the best integer solutions found throughout the search are given. In the case where the time limit was reached,

no lower bound can be given. Within a time limit of one hour, both procedures solve 26 out of 36 instances. Within an increased time limit of five days, four instances remain unsolved, in the case of pure column generation. With the collaborative scheme in place, three instances remain unsolved. Tables 2 and 3 contain the results obtained with the increased limit. The best integer solution found throughout the search is either the optimal integer solution, in the case the obtained lower bound is integer, or the solution obtained from solving SP on the set of generated columns within a run time limit of 10 minutes. Furthermore, the status of the obtained lower bound is given (integer (int.), fractional (frac.), or infeasible (inf.)). This is followed by the total number of columns generated, and by the number of times the different pricing procedures found at least one new negative reduced cost column. Rows $\bar{X}$, $\bar{Y}$, and $\bar{Z}$ give the average values for the respective data set. Row $\overline{XYZ}$ gives the total average values across all data sets.

In both cases, pure column generation and the collaborative approach, the two heuristic pricing procedures which prove to be the most useful are *ConstrHeur* and *LimLabels*. Also, all other heuristic pricing procedures contribute a number of negative reduced cost columns and thus should not be left out. When comparing the two tables, the following differences can be observed. In the case of the smaller instances, pure column generation is faster than the collaborative approach. This is due to the fact that in the case of the smaller instances, the initial VNS quite often already finds the optimal solution and thus, intermediate calls to the VNS do not improve this solution. They only increase computation times. This relation changes in the case of the larger instances. Here, in some cases the collaborative approach is faster while in others, lower computation times are due to pure column generation. The intuition is that calls to the intermediate VNS are not always useful although in many cases they are. When comparing the best integer solution found during the search, the collaborative approach outperforms pure column generation on average. When comparing the results of each instance individually, the collaborative scheme is better or equal in all but two cases. The most remarkable difference is certainly the fact that with the collaborative scheme in place instance a4-48 of data set "X" can be solved within the time limit, while in the case of pure column generation this instance cannot be solved.

Given these results, it seems fair to say that the integration of a collaborative local search method into the column generation framework has a positive impact on the performance of the whole method. As the original real-world instances are too large to be solved with either of the two exact procedures, results for these will only be provided for the heuristic method.

7.4 Heuristic results

In Table 4 the results obtained by means of the proposed VNS for the artificial data sets are presented. A limit of 10^5 iterations is used as stopping criterion. For each data set the following information can be taken from the table: the name of the instance; the average solution values over five random runs; the best solution values out of these five runs; and the respective percentage deviations from the lower bounds (see above). The average percentage gap between the lower bound and the obtained average solution value is less than 2% for all three data sets. Computing times are not really low, but acceptable (less than 9 minutes on average). The rather long computing times with respect to the low total iteration limit are due to the complex evaluation procedure, including a possible repositioning of the noon depot and the appropriate choice of the lunch break location.

As mentioned above, all real-world instances were too large to be solved by means of column generation. Therefore, in order to get an idea of how difficult practical instances are, we generated two additional real-world based data sets. In these two data sets, the number of

requests in every instance has been reduced by a factor of 5 or 3 with respect to the original data, respectively, and the time window length has been decreased to 15 minutes. For those instances that could be solved by the column generation framework, the obtained lower bound will serve to assess the solution quality of the proposed VNS.

Table 5 contains the results obtained for 8 instances that have been reduced by the factor 5 with respect to the original data. Overall, 7 out of these 8 instances could be solved by means of column generation. The following information is provided: the name of the instance; the size of the instance in terms of the total number of requests n; the number of drivers m_d; the number of attendants m_a available; the lower bound, where known; and for each iteration limit, the average and the best solution value out of 5 random runs, the corresponding deviation from the lower bound, and the total average run time in seconds. All lower bound solutions are fractional, with one exception. In the case of instance aug1108a the obtained lower bound is integer. On average, the proposed VNS yields solutions within 3.17 percent of the lower bound, using 10^5 iterations. With an increased limit of 2×10^5 iterations, the average percentage gap is reduced to 2.49. In the case of instance mai1805a, the lower bound solution is highly fractional. This explains why the obtained gap between the heuristic upper bound and the column generation lower bound is so large. It can be assumed that it is at least partly due to a larger integrality gap with respect to the other instances. Total run times are below 16 minutes for all instances.

Table 6 provides similar information to Table 5. However, for this medium size data set, which is again based on the available real-world data, considering only every third request, and reducing the number of available drivers and attendants accordingly, lower bounds cannot be computed. Thus, we only provide average and best solution values for the proposed VNS within a limit of 10^5 and 2×10^5 iterations. The percentage deviations presented in Table 6 give the deviations from the best solutions encountered with the two iteration limits. As expected, more iterations lead, on average, to better solution values.

Finally, Table 7 provides the results for those data sets that are based on ARC data from 15 days in the city of Graz. As mentioned above, all of them are too large to be solved by means of column generation. Therefore, only the proposed VNS is applied. Because of the large size of these instances, two different iteration limits have been used. First, VNS is run for 5×10^5 iterations. This configuration results in solution values that are on average 3.77% worse than the best solution found during both experimental settings. On average, slightly less than 3 hours of run time is needed. Then, the limit is increased to 10^6 iterations. In less than twice the time (on average 5.5 hours), the average gap from the best known solution is reduced to 1.47%.

For the smallest instances containing up to 50 requests, 10^5 iterations will lead to solutions of high quality when compared to the lower bound. For medium size real-world instances, however, more iterations are necessary to yield solutions of acceptable quality. In the case of the largest real-world instances, a limit of 5×10^5 or even 10^6 iterations, if time allows, should be used.

8 Conclusion

This paper has described a heterogeneous dial-a-ride problem with driver related constraints and has introduced two new formulations of the problem. A column generation approach has been proposed to compute lower bounds based on the set-partitioning formulation. A variable neighborhood search heuristic has also been developed. Computational experiments show that on the artificial instances, high quality solutions are obtained within rather short run times. A collaborative scheme, integrating the VNS heuristic into the column generation framework, has also been developed. Comparisons with the pure column generation show that the collaborative

scheme improves the efficiency of the original method. We hope that the methods proposed in this paper will serve as the basis for the development of a computer aided routing tool that will support ambulance dispatchers at the Austrian Red Cross in their day-to-day work.

Acknowledgements

References

Beaudry, A., Laporte, G., Melo, T., and Nickel, S. (2009). Dynamic transportation of patients to hospitals. *OR Spectrum*, 32:77–107.

Berbeglia, G., Cordeau, J.-F., Gribkovskaia, I., and Laporte, G. (2007). Static pickup and delivery problems: A classification scheme and survey. *TOP*, 15:1–31.

Cordeau, J.-F. (2006). A branch-and-cut algorithm for the dial-a-ride problem. *Oper Res*, 54:573–586.

Cordeau, J.-F. and Laporte, G. (2003). A tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transport Res B-Meth*, 37:579–594.

Cordeau, J.-F. and Laporte, G. (2007). The dial-a-ride problem: Models and algorithms. *Ann Oper Res*, 153:29–46.

Danna, E. and Le Pape, C. (2005). Branch-and-price heuristics: a case study on the vehicle routing problem with time windows. In Desaulniers, G., Desrosiers, J., and Solomon, M. M., editors, *Column Generation*, pages 99–129.

Desaulniers, G. and Villeneuve, D. (2000). The shortest path problem with time windows and linear waiting costs. *Transport Sci*, 34(3):312–319.

Desrochers, M., Desrosiers, J., and Solomon, M. (1992). A new optimization algorithm for the vehicle routing problem with time windows. *Oper Res*, 40:342–354.

Glover, F. (1996). Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Appl Math*, 65:223–253.

Hanne, T., Melo, T., and Nickel, S. (2009). Bringing robustness to patient flow management through optimized patient transports in hospitals. *Interfaces*, 39:241–255.

Irnich, S. (2008). Resource extension functions: properties, inversion, and generalization to segments. *OR Spectrum*, 30:113–148.

Kallehauge, B., Larsen, J., Madsen, O. B. G., and Solomon, M. M. (2005). Vehicle routing problems with time windows. In Desaulniers, G., Desrosiers, J., and Solomon, M. M., editors, *Column Generation*, New York. Springer.

Kirkpatrick, S., Gelatt Jr., C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220:671–680.

Melachrinoudis, E., Ilhan, A. B., and Min, H. (2007). A dial-a-ride problem for client transportation in a health-care organization. *Comput Oper Res*, 34:742–759.

Mladenović, N. and Hansen, P. (1997). Variable neighborhood search. *Comput Oper Res*, 24:1097–1100.

Parragh, S. N. (2009a). *Ambulance routing problems with rich constraints and multiple objectives*. PhD thesis, Faculty of Business, Economics and Statistics, University of Vienna.

Parragh, S. N. (2009b). Introducing heterogeneous users and vehicles into models and algorithms for the dial-a-ride problem. Technical report, Faculty of Business, Economics and Statistics, University of Vienna.

Parragh, S. N., Doerner, K. F., Gandibleux, X., and Hartl, R. F. (2009). A heuristic two-phase solution method for the multi-objective dial-a-ride problem. *Networks*, 54:227–242.

Parragh, S. N., Doerner, K. F., and Hartl, R. F. (2008a). A survey on pickup and delivery problems. Part I: Transportation between customers and depot. *J Betriebswirtschaft*, 58:21–51.

Parragh, S. N., Doerner, K. F., and Hartl, R. F. (2008b). A survey on pickup and delivery problems. Part II: Transportation between pickup and delivery locations. *J Betriebswirtschaft*, 58:81–117.

Parragh, S. N., Doerner, K. F., and Hartl, R. F. (2010). Variable neighborhood search for the dial-a-ride problem. *Comput Oper Res*, 37:1129–1138.

Rekiek, B., Delchambre, A., and Saleh, H. A. (2006). Handicapped person transportation: An application of the grouping genetic algorithm. *Eng Appl Artif Intel*, 19:511–520.

Ropke, S. (2005). *Heuristic and exact algorithms for vehicle routing problems*. PhD thesis, Computer Science Department, University of Copenhagen.

Ropke, S. and Cordeau, J.-F. (2009). Branch-and-cut-and-price for the pickup and delivery problem with time windows. *Transport Sci*, 43:267–286.

Ropke, S., Cordeau, J.-F., and Laporte, G. (2007). Models and branch-and-cut algorithms for pickup and delivery problems with time windows. *Networks*, 49:258–272.

Savelsbergh, M. W. P. (1992). The vehicle routing problem with time windows: Minimizing route duration. *ORSA J Comput*, 4:146–154.

Savelsbergh, M. W. P. and Sol, M. (1998). DRIVE: Dynamic routing of independent vehicles. *Oper Res*, 46:474–490.

Shaw, P. (1998). Using constraint programming and local search methods to solve vehicle routing problems. In *Proceedings CP-98 (Fourth International Conference on Principles and Practice of Constraint Programming)*.

Toth, P. and Vigo, D. (1997). Heuristic algorithms for the handicapped persons transportation problem. *Transport Sci*, 31:60–71.

Černy, V. (1985). Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *J Opt Theory Appl*, 45:41–51.

Xu, H., Chen, Z.-L., Rajagopal, S., and Arunapuram, S. (2003). Solving a practical pickup and delivery problem. *Transport Sci*, 37:347–364.

Table 2: Artificial instances - pure column generation

	CPU				Stat.	Cols	Number of times new best cols found by					
	Init. VNS	All	LB	Best int.			Constr-Heur	LNS-Curr-Basis	Rand-Constr-Heur	LNS-Rand-Constr	Lim-Labels	Lim-Graph
X												
a2-16	17.50	5.77	299.37	299.37	int.	662	6	1	6	3	5	0
a2-20	79.78	94.93	376.70	376.70	int.	2282	44	9	32	5	27	0
a2-24	133.54	3518.32	461.66	461.66	int.	4217	30	53	13	3	47	0
a3-18	17.80	14.83	291.68	291.68	int.	1232	10	5	9	5	3	0
a3-24	47.53	166.43	351.19	361.39	frac.	2786	51	9	24	5	32	1
a3-30	112.62	4668.83	510.79	510.79	int.	6336	87	39	50	24	93	12
a3-36	-	-	-	602.52	-	14800	123	216	5	6	50	47
a4-16	20.42	4.12	-	-	inf.	711	12	8	4	1	2	0
a4-24	34.63	41.92	388.95	393.57	frac.	2275	23	4	3	1	10	0
a4-32	101.56	557.36	504.79	509.23	frac.	5136	41	40	5	6	28	0
a4-40	-	-	-	610.15	-	12448	147	113	1	6	41	32
a4-48	-	-	-	670.06	-	18725	109	72	103	19	89	47
$\overline{X}$	62.82	1008.06	398.14	462.46		5968	57	47	21	7	36	
Y												
a2-16	54.42	5.79	-	-	inf.	1111	15	4	10	2	2	0
a2-20	77.59	27.98	371.62	371.62	int.	1968	52	10	33	9	21	0
a2-24	234.72	1742.20	-	-	inf.	3389	18	54	27	13	35	0
a3-18	25.30	4.63	295.39	295.39	int.	596	9	2	6	1	3	0
a3-24	92.05	69.11	353.13	359.75	frac.	2806	37	14	37	9	26	0
a3-30	110.69	3278.17	487.13	487.13	int.	7752	86	11	30	26	38	24
a3-36	-	-	-	620.27	-	13328	186	49	40	19	61	47
a4-16	20.46	2.18	-	-	inf.	544	8	2	4	0	4	0
a4-24	59.90	82.75	-	-	inf.	2019	28	24	16	4	14	0
a4-32	96.58	645.16	504.91	510.73	frac.	4751	46	35	3	1	21	4
a4-40	148.65	6287.23	587.56	590.21	frac.	8161	70	25	21	14	61	30
a4-48	198.22	276470.00	717.95	717.95	int.	22932	140	52	107	37	227	122
$\overline{Y}$	101.69	26237.75	473.96	494.13		5780	58	24	28	11	43	19
Z												
a2-16	41.77	7.98	308.30	308.30	int.	1091	24	2	9	3	6	0
a2-20	97.43	449.47	398.65	398.65	int.	2877	39	15	52	12	40	0
a2-24	125.06	1293.53	423.05	423.05	int.	3965	95	31	34	11	47	0
a3-18	26.15	4.40	297.24	297.24	int.	783	13	6	2	1	2	0
a3-24	54.97	104.19	354.19	355.15	frac.	2159	34	10	28	7	26	0
a3-30	122.81	863.53	495.70	498.43	frac.	5322	91	42	27	5	33	1
a3-36	174.07	100628.00	569.36	570.60	frac.	12123	103	18	46	30	93	95
a4-16	25.77	3.30	-	-	inf.	697	11	6	10	0	2	0
a4-24	39.24	47.68	388.69	388.69	int.	1903	24	8	13	6	11	0
a4-32	121.32	769.89	498.34	513.98	frac.	4040	67	47	13	8	22	0
a4-40	119.20	58136.70	581.86	604.28	frac.	9632	72	52	45	12	20	33
a4-48	248.59	57015.70	679.21	679.21	int.	12740	100	39	93	30	163	55
$\overline{Z}$	99.70	18277.03	454.05	457.96		4778	56	23	31	10	39	15
$\overline{XYZ}$	90.01	16156.63	456.34	469.26		5508	57	31	27	10	39	15

Cols = columns, frac. = fractional, inf. = infeasible, Init. = initial, int. = integer, LB = lower bound, Stat. = status.

Table 3: Artificial instances - collaborative scheme

	CPU						Number of times new best cols found by					
	Init. VNS	All	LB	Best int.	Stat.	Cols	Constr-Heur	LNS-Curr-Basis	Rand-Constr-Heur	LNS-Rand-Constr	Lim-Labels	Lim-Graph
X												
a2-16	17.50	17.65	299.37	299.37	int.	660	7	3	6	0	7	0
a2-20	79.51	149.48	376.70	376.70	int.	1919	24	7	23	3	18	0
a2-24	132.04	4892.11	461.66	461.66	int.	4453	29	93	28	7	50	0
a3-18	17.70	15.43	291.68	291.68	int.	1143	9	3	6	0	3	0
a3-24	47.82	322.04	351.19	358.29	frac.	2886	52	11	23	16	30	2
a3-30	112.75	1617.41	510.79	510.79	int.	5897	103	21	58	17	77	11
a3-36	-	-	-	600.83	-	16178	129	171	15	10	75	80
a4-16	20.47	5.94	-	-	inf.	665	9	4	7	1	2	0
a4-24	34.33	33.94	388.95	392.37	frac.	2157	17	0	6	0	8	0
a4-32	100.33	675.10	504.79	509.23	frac.	4829	38	15	10	5	30	1
a4-40	-	-	-	608.55	-	14533	181	44	57	32	83	46
a4-48	159.00	192018.00	665.11	665.74	frac.	19313	114	43	63	42	110	86
$\overline{X}$	72.15	19974.71	459.97	461.38		6211	59	35	25	11	41	19
Y												
a2-16	54.31	24.04	-	-	inf.	1132	13	3	4	5	5	0
a2-20	77.34	107.12	371.62	371.62	int.	1649	28	7	40	4	11	0
a2-24	234.83	709.20	-	-	inf.	3415	19	48	29	4	32	0
a3-18	24.96	6.29	295.39	295.39	int.	716	9	1	3	0	4	0
a3-24	91.86	223.08	353.13	355.47	frac.	2863	39	21	46	15	25	0
a3-30	107.60	2540.69	487.13	487.13	int.	7400	61	8	30	5	46	30
a3-36	-	-	-	598.62	-	13803	175	40	37	9	79	58
a4-16	20.40	3.06	-	-	inf.	557	10	1	2	1	3	0
a4-24	59.76	110.69	-	-	inf.	2113	26	25	17	4	10	0
a4-32	95.43	689.41	504.91	507.34	frac.	4990	47	44	3	2	21	5
a4-40	145.27	8329.59	587.56	590.21	frac.	8080	76	24	8	4	62	28
a4-48	197.55	218190.00	717.95	717.95	int.	21355	206	47	113	50	254	99
$\overline{Y}$	100.85	20993.92	489.45	490.46		5673	59	22	28	9	46	18
Z												
a2-16	41.80	32.49	308.30	308.30	int.	972	24	4	14	2	3	0
a2-20	97.05	247.69	398.65	398.65	int.	2415	57	22	48	12	19	0
a2-24	124.48	696.34	423.05	423.05	int.	4178	81	32	35	12	38	0
a3-18	26.11	22.23	297.24	297.24	int.	842	13	8	4	6	2	0
a3-24	54.03	217.16	354.19	355.15	frac.	2275	40	21	46	11	22	0
a3-30	121.22	1370.22	495.70	501.60	frac.	5716	73	40	53	21	54	1
a3-36	169.32	191169.00	569.36	579.61	frac.	12599	99	8	65	54	99	96
a4-16	25.67	8.07	-	-	inf.	630	11	11	2	1	3	0
a4-24	38.77	61.82	388.69	388.69	int.	1837	19	6	7	5	10	0
a4-32	121.41	680.31	498.34	504.26	frac.	4027	50	35	10	3	20	1
a4-40	118.14	29156.20	581.86	603.86	frac.	9511	76	38	23	23	29	26
a4-48	246.36	55650.80	679.21	679.21	int.	12386	81	37	68	37	145	60
$\overline{Z}$	98.70	23276.03	454.05	458.15		4782	52	22	31	16	37	15
$\overline{XYZ}$	91.37	21514.93	465.66	467.95		5555	57	26	28	12	41	18

Cols = columns, frac. = fractional, inf. = infeasible, Init. = initial, int. = integer, LB = lower bound, Stat. = status.

Table 4: Artificial instances - VNS (10^5 iterations, 5 runs)

	Data set X					Data set Y					Data set Z				
	Avg.	%	Best	%	CPU	Avg.	%	Best	%	CPU	Avg.	%	Best	%	CPU
a2-16	299.37	0.00	299.37	0.00	201.34	-	-	-	-	-	308.30	0.00	308.30	0.00	270.89
a2-20	377.14	0.12	376.70	0.00	345.76	371.62	0.00	371.62	0.00	680.22	398.65	0.00	398.65	0.00	436.34
a2-24	461.66	0.00	461.66	0.00	685.8	-	-	-	-	-	426.08	0.72	425.30	0.53	633.21
a3-18	291.68	0.00	291.68	0.00	89.72	295.39	0.00	295.39	0.00	132.68	297.24	0.00	297.24	0.00	128.09
a3-24	359.72	2.43	356.30	1.45	247.14	361.66	2.42	360.78	2.17	314.13	361.00	1.92	355.15	0.27	280.76
a3-30	512.34	0.30	510.79	0.00	516.5	496.02	1.82	487.13	0.00	527.24	506.05	2.09	498.29	0.52	555.64
a3-36	614.31	-	604.35	-	987.97	617.55	-	596.61	-	1108.15	580.01	1.87	572.55	0.56	942.36
a4-16	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
a4-24	402.44	3.47	394.34	1.38	169.8	-	-	-	-	-	396.92	2.12	393.13	1.14	180.79
a4-32	516.21	2.26	513.10	1.64	367.19	508.30	0.67	507.29	0.47	340.34	506.36	1.61	504.44	1.22	353.16
a4-40	623.57	-	613.33	-	753.95	595.90	1.42	592.55	0.85	644.27	600.02	3.12	595.72	2.38	617.07
a4-48	689.15	3.61	680.27	2.28	862.78	734.42	2.29	728.74	1.50	1122.68	697.64	2.71	684.56	0.79	1164.09
Avg.	467.96	1.35	463.81	0.75	475.27	497.61	1.23	492.51	0.71	608.71	461.66	1.47	457.58	0.67	505.67

Avg. = average

Table 5: Real world based instances (5 times smaller) - VNS (5 runs)

					VNS 10^5 iterations					VNS 2×10^5 iterations				
	n	m_d	m_a	LB	Avg.	%	Best	%	CPU	Avg.	%	Best	%	CPU
aug0508a	29	4	6	377.32	384.96	2.02	383.10	1.53	274.54	383.64	1.68	378.89	0.42	519.25
aug1108a	30	4	6	449.74	458.75	2.00	451.03	0.29	229.47	459.30	2.13	453.15	0.76	445.85
aug1208a	34	4	6	426.01	431.95	1.39	428.99	0.70	313.55	436.23	2.40	428.99	0.70	671.22
aug1308a	33	4	6	474.38	480.80	1.35	479.66	1.11	316.00	483.59	1.94	479.66	1.11	651.54
mai0605a	41	5	7	471.47	488.11	3.53	485.94	3.07	287.16	483.61	2.58	473.38	0.41	590.89
mai0705a	42	5	7	-	664.30	-	654.11	-	471.91	640.08	-	627.32	-	769.51
mai1805a	54	6	9	582.38	623.57	7.07	603.78	3.67	495.44	622.99	6.97	615.81	5.74	937.00
mai2105a	28	4	5	390.65	409.49	4.82	391.62	0.25	190.51	399.35	2.23	391.10	0.12	375.15
Avg.				453.14	492.74	3.17	484.78	1.52	322.32	488.60	2.49	481.04	1.16	620.05

Avg. = average

Table 6: Real world based instances (3 times smaller) - VNS (5 runs)

					VNS 10^5 iterations					VNS 2×10^5 iterations				
	n	m_d	m_a	All best	Avg.	%	Best	%	CPU	Avg.	%	Best	%	CPU
aug0508b	49	6	9	667.90	676.06	1.22	667.90	0.00	461.64	676.76	1.33	670.34	0.36	833.46
aug1108b	51	6	9	693.31	712.22	2.73	693.31	0.00	428.37	706.65	1.92	696.80	0.50	802.34
aug1208b	58	7	10	718.02	733.79	2.20	718.02	0.00	404.47	735.34	2.41	730.46	1.73	788.95
aug1308b	55	7	10	643.30	653.10	1.52	643.30	0.00	355.91	650.33	1.09	644.15	0.13	731.29
feb0402b	60	7	11	684.51	703.89	2.83	695.09	1.55	444.74	697.84	1.95	684.51	0.00	847.73
mai0605b	69	8	11	764.79	785.08	2.65	764.79	0.00	498.42	777.71	1.69	769.87	0.66	969.59
mai1805b	91	10	15	1107.04	1167.19	5.43	1140.29	3.00	743.60	1129.74	2.05	1107.04	0.00	1362.91
nov0411b	79	9	13	867.97	902.28	3.95	878.21	1.18	598.93	882.35	1.66	867.97	0.00	1173.83
Avg.				768.36	791.70	2.82	775.11	0.72	492.01	782.09	1.76	771.39	0.42	938.76

Avg. = average

Table 7: Real world instances - VNS (5 runs)

	n	m_d	m_a	All best	VNS 5×10^5 iterations					VNS 10^6 iterations				
					Avg.	%	Best	%	CPU	Avg.	%	Best	%	CPU
aug0508	147	17	26	1374.33	1406.03	2.31	1379.42	0.37	6797.44	1392.12	1.29	1374.33	0.00	12241.06
aug1108	154	18	27	1415.55	1447.50	2.26	1422.90	0.52	5121.00	1429.40	0.98	1415.55	0.00	10179.36
aug1208	174	19	29	1632.16	1688.72	3.47	1673.86	2.55	7854.12	1666.13	2.08	1632.16	0.00	15394.96
aug1308	166	19	29	1552.40	1600.94	3.13	1576.00	1.52	6729.45	1565.90	0.87	1552.40	0.00	12795.56
feb0202	233	20	30	2087.65	2144.81	2.74	2103.59	0.76	20555.08	2113.63	1.24	2087.65	0.00	37997.24
feb0402	182	21	32	1677.70	1722.32	2.66	1701.25	1.40	7484.74	1696.75	1.14	1677.70	0.00	14200.60
feb1002	186	18	27	1804.26	1854.35	2.78	1829.44	1.40	12483.68	1839.04	1.93	1804.26	0.00	24718.90
mai0605	208	22	33	1760.82	1851.03	5.12	1833.69	4.14	10917.44	1779.48	1.06	1760.82	0.00	20268.88
mai0705	210	23	35	1964.68	2077.48	5.74	2046.29	4.15	10119.45	1997.39	1.66	1964.68	0.00	18350.90
mai1805	273	30	45	2570.25	2677.39	4.17	2629.02	2.29	11725.66	2601.40	1.21	2570.25	0.00	22719.08
mai2105	140	16	24	1457.36	1494.56	2.55	1476.78	1.33	5745.56	1470.55	0.91	1457.36	0.00	10813.82
nov0411	239	25	38	2155.53	2225.80	3.26	2186.96	1.46	11984.40	2171.58	0.74	2155.53	0.00	23501.62
nov0911	247	24	36	2207.46	2303.18	4.34	2269.98	2.83	15720.66	2248.11	1.84	2207.46	0.00	29666.94
nov1211	192	21	32	1711.99	1805.97	5.49	1784.38	4.23	9152.52	1749.12	2.17	1711.99	0.00	16999.46
nov1611	219	21	32	2097.98	2236.66	6.61	2212.37	5.45	15687.24	2158.05	2.86	2097.98	0.00	30054.52
Avg.				1831.34	1902.45	3.77	1875.06	2.29	10538.56	1858.58	1.47	1831.34	0.00	19993.53

Avg. = average