

DRIVE: Dynamic Routing of Independent VEHICLES

Martin Savelsbergh
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205
U.S.A.

Marc Sol
Department of Mathematics and Computing Science
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
The Netherlands

Abstract

We present DRIVE (Dynamic Routing of Independent VEHICLES), a planning module to be incorporated in a decision support system for the direct transportation at Van Gend & Loos BV. Van Gend & Loos BV is the largest company providing road transportation in the Benelux with about 1400 vehicles transporting 160,000 packages from thousands of senders to tens of thousands of addressees per day. The heart of DRIVE is a branch-and-price algorithm. Approximation and incomplete optimization techniques as well as a sophisticated column management scheme have been employed to create the right balance between solution speed and solution quality. DRIVE has been tested by simulating a dynamic planning environment with real-life data and has produced very encouraging results.

April 1996
Revision July 1996

1 Introduction

Van Gend & Loos BV¹ is the largest company providing road transportation in the Benelux, which is the region consisting of Belgium, The Netherlands, and Luxembourg,

¹Van Gend & Loos BV is a member of the Nedlloyd group, a Dutch organization providing world wide transportation services by land, air and sea.

with about 1400 vehicles transporting 160,000 packages from thousands of senders to tens of thousands of addressees per day. The service of Van Gend & Loos BV can be roughly divided into two parts, the *regular transportation system* and the *direct transportation system*. In the regular transportation system, shipments, ranging from small packages to loads of up to four pallets, are picked up at the sender and then delivered at the closest distribution center. During the night, the loads are transported from that distribution center to the distribution center closest to the destination of the load, from where they are delivered at their final destination during the next day. In the direct transportation system, shipments, ranging in size from four pallets to full truck loads, are picked up by a vehicle at the sender and are delivered by that same vehicle at the destination; there is no transshipment at any distribution center.

Operating the direct transportation system of Van Gend & Loos BV is a complex task and management at Van Gend & Loos BV has realized that to handle the anticipated growth and to take advantage of the anticipated technological changes, such as global positioning systems and direct communication with vehicles, their planning department may benefit from more advanced planning tools than they currently have.

This paper discusses such a planning tool: DRIVE, Dynamic Routing of Independent Vehicles. DRIVE is to be incorporated in a decision support system for the direct transportation system at Van Gend & Loos BV. The heart of DRIVE is a branch-and-price algorithm for the general pickup and delivery problem (GPDP).

In a GPDP a set of routes has to be constructed in order to satisfy transportation requests. Each transportation request specifies the size of the load to be transported, the location where it is to be picked up (the origin), the location where it is to be delivered (the destination), and time windows defining allowable pickup and delivery times. Each load has to be transported by one vehicle from its origin to its destination without any transshipment at other locations. A heterogeneous fleet of vehicles is available to operate the routes. Each vehicle is characterized by a capacity, a depot where it is stationed, and a time window specifying when the vehicle is available.

The GPDP can naturally be formulated as a set partitioning problem in which the set of transportation requests has to be partitioned into a set of feasible routes. In recent years, set partitioning formulations have become very popular for many combinatorial optimization problems. There are two main reasons for this. First, for many problems alternative formulations are not known. This is the case, for example, in crew pairing problems [Anbil, Tanga, and Johnson 1993]. Second, for many problems where alternative formulations are known, the linear programming relaxation of the set partitioning formulation often yields a stronger bound. This is the case, for example, in cutting stock problems [Vance et al. 1992].

Many large-scale set partitioning problems have been solved successfully by branch-and-price algorithms. Branch-and-price algorithms [Barnhart et al. 1996] solve mixed

integer programming formulations with huge numbers of variables. In branch-and-price algorithms, sets of columns are left out of the linear program because there are too many columns to handle efficiently and most of them have their associated variable equal to zero in an optimal solution anyway. In order to check the optimality of a linear programming solution, a subproblem, called the *pricing problem*, is then solved in order to identify columns to enter the basis. If such columns are found, the linear program is reoptimized. Branching occurs when no columns price out to enter the basis and the linear programming solution is optimal and fractional. Branch-and-price, which is a generalization of branch-and-bound with linear programming relaxations, allows column generation to be applied throughout the branch-and-bound tree.

Our discussion of DRIVE can be broken down into three parts. First, we present a branch-and-price algorithm for the GPDP. This algorithm is very flexible and is easily turned into an approximation algorithm that is capable of producing high quality solutions for large instances in an acceptable amount of computation time. This is a nontrivial task since the underlying solution paradigm is based on the enumeration of the solution space, but essential because the algorithm is to be used in a real-life operational planning environment. The proper balance between quality and speed has been realized by using heuristics, whenever appropriate, within the overall optimization framework, by using a sophisticated column management scheme, and by using problem specific information, whenever possible, in the branching scheme, the primal heuristic, and the column selection procedure. Secondly, we present the algorithmic adjustments that had to be made to the branch-and-price algorithm for the GPDP to be able to handle the specific characteristics of the direct transportation system of Van Gend & Loos BV. Finally, we present the results of a case study. These results show the viability of our approach and illustrate that sophisticated optimization based heuristics can be used successfully in complex real-life situations. This is one of the contributions of our work, since there seems to be a tendency to resort to relatively simple heuristics to solve complex real-life decision situations.

As the results of the case study were very encouraging, the management at Van Gend & Loos BV has put forward a proposal to the board of directors that asks for permission and funds to build a decision support system for the direct transportation system with DRIVE as the core planning tool.

This paper is organized as follows. In Section 2, we discuss the problem characteristics of the direct transportation system at Van Gend & Loos BV, their current planning methodology, and the envisioned planning methodology. In Section 3, we present the branch-and-price algorithm we have developed for the GPDP. In Section 4, we discuss the algorithmic adjustments that had to be made to the branch-and-price algorithm to be able to handle the specific characteristics of the direct transportation system of Van Gend & Loos BV. In Section 5, we present the computational experiments that have

been conducted. Finally, in Section 6, we make some concluding remarks.

2 The direct transportation system

In this section we take a closer look at the direct transportation system at Van Gend & Loos BV. We discuss problem characteristics as well as current and envisioned solution methodology.

2.1 Problem characteristics

2.1.1 Requests

All requests specify one pickup location and one or more delivery locations, which have to be visited in a predefined order. A single origin-destination pair is called a *shipment*. A request consisting of multiple delivery locations is said to consist of multiple shipments. These shipments have to be picked up at the same time by a single vehicle. All requests specify time windows for the pickup and all the deliveries.

Not all the transportation requests are known in advance; some become available in real time. This implies that at the time a new request becomes available, the set of routes that is currently being executed has to be modified. At the beginning of a day, about 60% of all requests that have to be served during that day are known.

2.1.2 Vehicles

There is a heterogeneous fleet of vehicles. Vehicle capacities are specified in pallets; there are vehicles with a capacity of 14, 16, 24 and 26 pallets. Capacity is not the only characteristic that is used to differentiate between vehicles. Vehicles are also distinguished based on their physical characteristics, because some clients have specific demands on the properties of a vehicle that serves their request.

The fleet of vehicles is not stationed at a central depot; each vehicle has its own *home location*.

A *working period* of a vehicle is a period of consecutive days in which the vehicle is in use. A vehicle starts and ends a working period at its home location with no load on board. Each night of a working period a driver sleeps at one of a set of *sleeping locations*, which includes his own home location. On Friday night a vehicle has to return home. At this time the vehicle does not have to be empty, meaning that its working period has not yet ended. On Monday morning the vehicles are again available at their home location. A driver must have a 45 minute lunch break between 11 am and 2 pm each day of a working period.

The vehicles are rented by Van Gend & Loos BV on a daily basis for working periods of unknown length. When the number of transportation requests decreases, some vehicles will be sent home in order to end their working period. A number of vehicles is rented permanently, i.e. they cannot be returned. When the number of transportation requests increases, new vehicles are rented. A new vehicle can only start its working period at the beginning of the next working day. This implies that, at the end of a day, when only part of the requests that have to be served the next day are known, the vehicles for the next day have to be selected. Van Gend & Loos BV does not exchange one vehicle for another, but only returns vehicles when the total number of vehicles that is currently in use exceeds the number of vehicles that will be needed the next day.

2.1.3 Costs

The costs incurred by serving the requests are the drivers' pays, which include the lease costs of the vehicles. Each day of a working period a driver gets paid an amount proportional to the distance traveled that day unless this amount does not exceed some specified minimum, in which case this minimum is paid. Furthermore a driver gets a compensation for each night of a working period that he does not spend at his home location.

2.1.4 Problem size

The planning area, i.e., the region in which all locations are situated, is the Benelux. About 100 vehicles are used per day, of which 50 vehicles are rented permanently. The number of requests that has to be served per day is about 250 to 300. The total number of shipments per day is about 500.

2.2 Current methodology

A central dispatching office collects new requests and assigns them to vehicles. At every stop of a route the driver of a vehicle calls this central office for new instructions. At present, the office cannot directly contact the drivers, although it is anticipated that a more sophisticated communication system will be implemented in the future.

The planning is done by a team of five people. Two of them do the actual planning, the other three act as an interface between the planners and the drivers. All five people take turns on both jobs. The planning area is divided into two parts, each part covered by one planner. The planners basically use two sources of information: a *request list* with all transportation requests that have to be served within the next twenty hours, and a *vehicle list* with information about where and when the vehicles that are in use will become empty. On both lists, the items are geographically grouped. On the request

list, this grouping is done by considering the origins of the requests. Within a group, the requests are listed in order of nondecreasing earliest pickup time. On the vehicle list, vehicles within one group are listed in order of nondecreasing time when they become empty.

The planners use a three-phase approach. First they try to find combinations of requests that should be served by one vehicle. Requests are combined, based on proximity of origins and destinations in both space and time, and based on total load. In the second phase, these combinations are tentatively assigned to vehicles. Only when a driver calls for new instructions, these tentative assignments are made permanent.

During the first part of a day the planners only focus on the work that has to be done during that day. In the afternoon, the following day is also considered. At this time the planners must also decide how many vehicles they need for the next day. Though only a fraction of the requests of the next day is known at that time, usually there is some knowledge about the total amount of work that has to be done during the next day. Such knowledge is available as planner's expertise. The number of vehicles that will be used the next day is primarily based on the amount of work that has to be done during the first part of the morning, but is increased if the expected total amount of work for the next day makes this necessary.

2.3 Envisioned methodology

Because the pickup and delivery problem at Van Gend & Loos BV is a dynamic problem, we designed DRIVE to work in an iterative way. When DRIVE is invoked, it produces a plan, based on the current set of routes, the current set of known requests and some estimate on future workload. This plan should be seen as a base plan that a planner can modify. While evaluating the base plan, a planner must focus primarily on the short-term decisions proposed by DRIVE, i.e., the assignments of loads to vehicles, and the routing of the vehicles within, say, the next hour. As soon as the plan is accepted, these short-term decisions are made permanent. This means that the first parts of the routes in the plan will be executed as planned. The remaining parts may be changed when new requests become available. As soon as necessary, DRIVE is used again to produce a new plan that respects all permanent decisions.

In this way, all routes are divided into a *head*, which is the part that will be executed as planned, and a *tail*, which is the remaining part that may change in the future. In this environment a planner is always busy preparing the tails of the routes. When a driver calls for new instructions, no calculations have to be performed because these instructions have been previously stored as the head of the driver's route.

3 A branch-and-price algorithm for the GPDP

The core structure of the direct transportation system at Van Gend & Loos BV is a GPDP. For a survey of the GPDP see Savelsbergh and Sol [1995]. Although the GPDP does not capture all the characteristics of the direct transportation systems at Van Gend & Loos BV, we felt that an efficient and effective algorithm for the general pickup and delivery problem would provide an invaluable tool.

We decided to develop a set partitioning based algorithm for the GPDP. The primary reasons for this decision were:

- A set partitioning approach focuses on the assignment of transportation requests to vehicles, which is the most important and most difficult aspect of multi-vehicle routing problems with relatively few stops per trip.
- A set partitioning approach is very flexible in the sense that most restrictions imposed on the routes can easily be incorporated.
- Set partitioning approaches have provided very promising results for various types of vehicle routing problems [Dumas, Desrosiers, and Soumis 1991, Desrochers, Desrosiers, and Solomon 1992].
- Set partitioning approaches benefit from the rapid advances in linear programming technology.

Our efforts have resulted in the branch-and-price algorithm for the GPDP described below. This algorithm forms the basis of DRIVE.

3.1 A set partitioning formulation

Let N be the set of transportation requests. For each transportation request $i \in N$, a load of size $q_i \in \mathbb{N}$ has to be transported from origin i^+ to destination i^- . Define $N^+ := \{i^+ \mid i \in N\}$ as the set of origins and $N^- := \{i^- \mid i \in N\}$ as the set of destinations. For each request $i \in N$ the pickup time window is denoted by $[e_{i^+}, l_{i^+}]$ and the delivery time window by $[e_{i^-}, l_{i^-}]$. Furthermore, let M be the set of vehicles. Each vehicle $k \in M$ has a capacity $Q_k \in \mathbb{N}$, is available in the interval $[e_{k^+}, l_{k^+}]$, and is stationed at depot k^+ . Define $M^+ := \{k^+ \mid k \in M\}$ as the set of depots. Let $V := N^+ \cup N^- \cup M^+$.

For all $i, j \in V, k \in M$ let d_{ij} denote the travel distance, t_{ij}^k the travel time, and c_{ij}^k the travel cost. Note that the dwell time at origins and destinations can be easily incorporated in the travel time and therefore will not be considered explicitly.

Definition 1 A pickup and delivery route R_k for a vehicle of type k is a directed cycle through a subset $V_k \subset N^+ \cup N^- \cup \{k^+\}$ such that:

1. R_k starts in k^+ not before e_{k^+} .
2. For all $i \in N$, $i^+ \in V_k$ if and only if $i^- \in V_k$.
3. If $\{i^+, i^-\} \subseteq V_k$, then i^+ is visited before i^- .
4. Each location in $V_k \setminus \{k^+\}$ is visited exactly once.
5. Each location in $V_k \setminus \{k^+\}$ is visited within its time window.
6. The vehicle load never exceeds Q_k .
7. R_k ends in k^+ not after l_{k^+} .

To formulate the pickup and delivery problem as a set partitioning problem, we define

$\Omega_k :=$ the set of all feasible pickup and delivery routes for vehicle k ,

$$\delta_{ir}^k := \begin{cases} 1 & \text{if } i \in N \text{ is served on route } r \in \Omega_k, \\ 0 & \text{otherwise,} \end{cases}$$

$c_r^k :=$ the cost of route $r \in \Omega_k$,

and introduce binary variables x_r^k ($k \in M, r \in \Omega_k$) equal to 1 if route $r \in \Omega_k$ is used and 0 otherwise. The pickup and delivery problem can now be formulated as follows:

$$\begin{aligned} & \text{minimize} && \sum_{k \in M} \sum_{r \in \Omega_k} c_r^k x_r^k \\ & \text{subject to} && \sum_{k \in M} \sum_{r \in \Omega_k} \delta_{ir}^k x_r^k = 1 && \text{for all } i \in N, && \text{(partitioning constraints)} \\ & && \sum_{r \in \Omega_k} x_r^k \leq 1 && \text{for all } k \in M, && \text{(availability constraints)} \\ & && x_r^k \in \{0, 1\} && \text{for all } k \in M, r \in \Omega_k. \end{aligned}$$

We denote this formulation by P and its linear programming relaxation by LP .

We will consider pickup and delivery problems where the primary objective is to minimize the number of vehicles needed to serve all transportation requests and the secondary objective is to minimize the total distance traveled. This is accomplished by taking the objective function

$$\sum_{k \in M} \sum_{r \in \Omega_k} (F + L_r^k) x_r^k,$$

where L_r^k denotes the length of route r for vehicle k and where $F > |N| \max_{r \in \Omega_k, k \in M} L_r^k$ is a large constant. This cost structure can be achieved by defining the travel costs $c_{ij}^k = d_{ij}$ ($i \neq k^+$), and $c_{k^+j}^k = F + d_{k^+j}$ ($j \neq k^+$).

3.2 A branch-and-price algorithm

We have developed a branch-and-price algorithm for the GPDP based on the formulation presented above. The lower bound provided by LP is usually excellent and often much better than the lower bounds provided by more traditional formulations with variables x_{ij}^k indicating whether or not a vehicle of type k travels from location $i \in V$ to location $j \in V$.

A column generation scheme has been applied to solve LP in order to handle the large number of variables that arises due to the size of the sets Ω_k . Instead of explicitly enumerating all feasible routes in order to find a variable that prices out to enter the basis, in a column generation approach the non-basic variable with the smallest negative reduced cost is found by solving an optimization problem, called the *pricing problem*. In this way the feasible routes are generated on the fly as needed and only a small fraction of all feasible routes is used to solve LP .

Dumas, Desrosiers and Soumis [1991] were the first to develop and implement a branch-and-price algorithm for the pickup and delivery problem with time windows. Our branch-and-price algorithm differs from theirs in various aspects. Most of these differences are prompted by the necessity to be able to solve large instances quickly.

The key features to accomplish a proper balance between solution speed and solution quality are

- The use of heuristics, whenever possible, to solve the pricing problem.
- The use of a sophisticated column management scheme to keep the linear programs as small as possible and to keep the number of linear programs that have to be solved as small as possible.
- The use of a column selection procedure that increases the chances of finding or constructing feasible solutions within the active set of columns.
- The use of a branching scheme that concentrates on high level decisions.
- The use of linear program based primal heuristics.

3.2.1 Solving the linear programming relaxation

Suppose that for each vehicle $k \in M$ a set $\Omega'_k \subseteq \Omega_k$ of feasible pickup and delivery routes is explicitly known. The *restricted master problem* LP' is defined as follows:

$$\text{minimize} \quad \sum_{k \in M} \sum_{r \in \Omega'_k} c_r^k x_r^k$$

$$\begin{aligned}
\text{subject to } & \sum_{k \in M} \sum_{r \in \Omega'_k} \delta_{ir}^k x_r^k = 1 && \text{for all } i \in N, \\
& \sum_{r \in \Omega'_k} x_r^k \leq 1 && \text{for all } k \in M, \\
& x_r^k \geq 0 && \text{for all } k \in M, r \in \Omega'_k
\end{aligned}$$

Suppose that LP' has a feasible solution x and let (u, v) be the associated dual solution, i.e., the dual variables u_i ($i \in N$) are associated with the partitioning constraints and the dual variables v_k ($k \in M$) are associated with the availability constraints. From linear programming duality we know that x is optimal with respect to LP if and only if for each $k \in M$ and for each $r \in \Omega_k$ the reduced cost d_r^k is nonnegative, i.e.,

$$d_r^k = c_r^k - \sum_{i \in N} \delta_{ir}^k u_i - v_k \geq 0 \quad \text{for all } k \in M, r \in \Omega_k.$$

Testing the optimality of x with respect to LP can thus be done by solving the *pricing problem*

$$\min \{c_r^k - \sum_{i \in N} \delta_{ir}^k u_i - v_k \mid k \in M, r \in \Omega_k\}.$$

Let z denote the value of the solution to the pricing problem and let k_z and r_z denote the corresponding vehicle type and route. If $z \geq 0$, then x is also optimal with respect to LP , otherwise r_z defines a column that can enter the basis and has to be added to Ω'_{k_z} . This yields the following *column generation scheme*:

1. Find initial sets Ω'_k containing a feasible solution x .
2. Solve the restricted master problem LP' .
3. Solve the pricing problem. If $z \geq 0$ then stop, otherwise set $\Omega'_{k_z} := \Omega'_{k_z} \cup \{r_z\}$ and go to Step 2.

Due to the presence of the availability constraints, it is nontrivial to find initial sets $\Omega'_k \subset \Omega_k$ containing a feasible solution to LP . However, if they exist, such sets can always be found using a two-phase method similar in spirit to the two-phase method incorporated in simplex algorithms to find an initial basic feasible solution. Define LP_+ as

$$\begin{aligned}
\text{minimize } & \sum_{k \in M} \sum_{r \in \Omega_k} c_r^k x_r^k + \sum_{i \in N} p y_i \\
\text{subject to } & \sum_{k \in M} \sum_{r \in \Omega_k} \delta_{ir}^k x_r^k + y_i = 1 && \text{for all } i \in N \\
& \sum_{r \in \Omega_k} x_r^k \leq 1 && \text{for all } k \in M \\
& x_r^k \geq 0 && \text{for all } k \in M, r \in \Omega_k \\
& y_i \geq 0 && \text{for all } i \in N
\end{aligned}$$

where y_i ($i \in N$) is an artificial variable and $p > \max_{k \in M, r \in \Omega_k} c_r^k$ is an appropriate penalty cost. Problem LP_+ can be solved by the above column generation scheme by initializing $\Omega'_k = \emptyset$ for each $k \in M$.

The artificial variables y_i are not deleted when they have all become non-basic, i.e., when a feasible solution to problem LP has been found. Because of their high cost, these variables will stay non-basic and will not interfere in the optimization process. However, during the branching process, the sets Ω'_k are restricted by the branching scheme, possibly yielding an initial infeasible LP in a node. In that case, the artificial variables will reappear in the basis, and the first phase is automatically started in order to find sets Ω'_k that do contain a feasible solution for the linear program associated with this node.

Because the route costs $c_r^k = F + L_r^k$ are constructed such that $F > |N| \max_{r \in \Omega_k, k \in M} L_r^k$, we can improve the lower bound Z_{LP} when the optimal LP solution x corresponds to a non-integral number of vehicles. More precisely, if $m = \lceil \sum_{k \in M} \sum_{r \in \Omega_k} x_r^k \rceil$, then the constraint

$$\sum_{k \in M} \sum_{r \in \Omega_k} x_r^k \geq m$$

is a valid inequality that may be added to LP , see also Desrosiers et al. [1996]. Adding this constraint does not change the pricing problems, because the dual value of this constraint appears as a constant in their objective functions.

The column generation scheme presented above can be made more flexible based on the two following observations:

- Any column with negative reduced cost is a candidate to enter the basis and can be added to the restricted master;
- If several columns with negative reduced cost exist, they can be added simultaneously to the restricted master.

This flexibility can be exploited effectively to improve the overall efficiency of a column generation scheme. The basic idea is to solve the pricing problem approximately as long as this produces columns with a negative reduced cost, and only solve the pricing problem optimally when solving it approximately fails to produce columns with a negative reduced cost. In this way, the number of times the pricing problem is solved optimally, which is computationally prohibitive, is reduced considerably.

We have embedded approximation as well as optimization algorithms for the pricing problem into an effective and efficient column management system. The approximation algorithms try to generate many routes with negative reduced cost very fast. Any route with a negative reduced cost, whether generated by an approximation algorithm or an

optimization algorithm, is stored in a *column pool*. Rather than solving the pricing problem at every iteration, we first search the column pool for columns with negative reduced cost. If successful, one or more of these are selected and added to the restricted master problem. If unsuccessful, we clean the column pool and invoke the pricing algorithms in order to try to refill the pool. Note that each time the restricted master problem is reoptimized the dual variables change. Therefore, the reduced costs of the columns in the pool have to be updated after every reoptimization. Cleaning the pool consists of removing all columns with reduced cost larger than some threshold $D_{\max} \geq 0$. A positive threshold value can be useful, because the reduced costs change after every reoptimization, possibly to a negative value. If the pricing problem is solved approximately, LP cannot be solved to optimality and so P cannot be solved to optimality. Therefore, as soon as the approximation algorithms fail to produce columns with negative reduced cost, an optimization algorithm is used to solve the pricing problem optimally to either prove optimality of LP or to find new columns with negative reduced costs. The column generation scheme we propose now looks as follows:

1. Find initial sets Ω'_k containing a feasible solution x .
2. Set the column pool equal to \emptyset .
3. Solve the restricted master problem LP' .
4. If the column pool contains columns with negative reduced cost, select some of these columns, add them to the restricted master problem and go to Step 3.
5. Delete columns with reduced cost larger than D_{\max} from the column pool and start the approximation algorithms for the pricing problem. If these are successful, add the generated columns to the pool and go to Step 4.
6. Solve the pricing problem to optimality. If $z \geq 0$, then stop, otherwise add the generated columns to the pool and go to Step 4.

There are several ways to choose columns from the column pool to add to the restricted master problem. The first possibility is to select the column with the minimum reduced cost. In this way, the linear program will not grow very rapidly, but a linear program has to be solved for each added column. A second possibility is to select all columns with negative reduced costs from the pool. This will reduce the number of linear programs that have to be solved, but the linear programs will become very large. We have chosen a more adaptive greedy selection scheme that selects partial solutions to problem P . More precisely, we select a set of columns with negative reduced costs that correspond to a set of routes satisfying the following requirements:

- Each transportation request is served on at most one route;
- The number of vehicles of a certain type assigned to the routes does not exceed the available number of vehicles of that type.

The set of columns is constructed by successively choosing a column with minimum negative reduced cost such that the two requirements are still satisfied. The selection stops when no more such columns are available in the column pool.

The above column selection mechanism is motivated by two observations. First, adding columns corresponding to partial solutions to P increases the chance of encountering integral solutions during the solution of the master problem. Second, it prevents the addition of similar columns, which would happen if the columns would be selected merely based on their reduced cost and the dual variables are far from being optimal.

3.2.2 The pricing problem

The pricing problem decomposes into several independent problems, one for each vehicle, since

$$z = \min\{c_r^k - \sum_{i \in N} \delta_{ir}^k u_i - v_k \mid k \in M, r \in \Omega_k\}$$

is equal to

$$z = \min_{k \in M} \min\{c_r^k - \sum_{i \in N} \delta_{ir}^k u_i - v_k \mid r \in \Omega_k\},$$

i.e., the problem of finding a minimum cost route for vehicle k , using a modified cost structure, that serves a subset of the transportation requests. Denote these independent problems by S_k for $k \in M$. Problem S_k can be viewed as a shortest path problem with precedence constraints, capacity constraints and time windows, on the perturbed cost matrix c' .

A dynamic programming algorithm using labeling techniques to handle the precedence, capacity and time constraints, can be used to solve this shortest path problem [Dumas, Desrosiers, and Soumis 1991]. However, for instances with many transportation requests and time and capacity constraints that are not very tight, solving S_k to optimality becomes computationally prohibitive.

As pointed out in the previous section, the pricing problem may be solved approximately as long as columns with negative reduced cost are found. Solving the pricing problem approximately can be done in several ways.

One approach is to speed up the dynamic programming algorithm in earlier iterations of the column generation process by working on a reduced network [Dumas, Desrosiers,

and Soumis 1991]. Network reduction is achieved by deleting nodes corresponding to requests with a low dual value and by deleting arcs with relatively high cost. Obviously, this reduces the state space of the dynamic program and thus the computation times, but it no longer guarantees that the optimal solution is found. If no more profitable routes can be found in the reduced network, it is enlarged and the dynamic programming algorithm is started again. Note that this approach guarantees that in the end the pricing problem is solved to optimality.

Also note that the dynamic programming algorithm may encounter many columns with negative reduced costs before it identifies the one with the smallest reduced cost. Obviously all these columns could be stored in the column pool. However, since the number of columns with negative reduced cost in the first iterations of the column generation process is huge, this would lead to an unmanageable column pool. Furthermore, only a few of the columns generated in the first iterations of the column generation scheme will be actually added to the restricted master problem. Therefore, in our implementation, we have put an upper bound on the number of columns that the algorithm can create in one run for each vehicle. The algorithm will stop as soon as the upper bound is reached, which reduces computation times drastically. Note that this approach still guarantees that in the end the pricing problem is solved to optimality.

Another approach is to use fast approximation algorithms instead of the dynamic programming algorithm. The fast approximation algorithms we have used are based on construction and improvement algorithms for the single-vehicle GPDP. In their description, we use $I_r^k(j)$ to denote the minimal cost of inserting transportation request j into route r for vehicle k . If transportation request j cannot be inserted into route r , then $I_r^k(j) = \infty$.

Since computing the true insertion cost involves the solution of a single-vehicle pickup and delivery problem with time windows, i.e., a traveling salesman problem with time windows, precedence constraints and capacity constraints, we have chosen to work with an approximate insertion cost, namely the cheapest insertion cost. If r is a route consisting of n stops, our insertion algorithm first calculates τ_i ($1 \leq i \leq n$), which is the latest possible time the vehicle may arrive at stop i in order to ensure feasibility, with respect to the time windows of the remaining part of the route $(i, i+1, \dots, n)$. Because a request requires two stops to be inserted into a route, there are $\mathcal{O}(n^2)$ possible insertions for each request. By evaluating these insertions in the right order and using the values τ_i , each possible insertion can be checked for feasibility and cost in constant time. The cheapest insertion cost of request j into route r can therefore be computed in $\mathcal{O}(n^2)$ time.

Construction algorithms

Construction algorithms build routes from scratch. Define the reduced insertion cost

$D_r^k(j) := I_r^k(j) - u_j$. The construction algorithms initialize a route r and then repeatedly try to decrease the reduced cost of the route by inserting a request with $D_r^k(j) < 0$. If $D_r^k(j) \geq 0$ for all requests not in the route, and $d_r^k < 0$, then r is added to the column pool. The route can be initialized as a loop from k^+ to k^+ , or as a route serving some requests that are likely to be served by a vehicle of type k .

Improvement algorithms

Improvement algorithms modify existing routes. Note that the current LP solution provides us with a set of routes r with $d_r^k = 0$ (at least all the routes associated with basic variables). When such a route is used as a starting point for a local search algorithm, we expect to find a route with negative reduced cost very fast. The following two algorithms start with a route r with $d_r^k = 0$ and then try to decrease the reduced cost of the route by deleting requests from the route and replacing them by other requests. The first algorithm performs profitable swaps until no such swap can be found. The second algorithm performs a variable depth search. When $i \in N$ is served on route r , we denote by $r \setminus \{i\}$ the route obtained from r by deleting request i . When $j \in N$ is not served on route r , we denote by $r \cup \{j\}$ the route obtained from r by inserting request j in the cheapest way.

The first algorithm works as follows:

1. Let $d_{(r \setminus \{i_0\}) \cup \{j_0\}}^k = \min\{d_{(r \setminus \{i\}) \cup \{j\}}^k \mid i, j \in N, \delta_{ir}^k = 1, \delta_{jr}^k = 0\}$.
2. If $d_{(r \setminus \{i_0\}) \cup \{j_0\}}^k < d_r^k$, then set $r = (r \setminus \{i_0\}) \cup \{j_0\}$ and go to Step 1.
3. If $d_r^k < 0$, then add r to the column pool.

At each iteration of the variable depth search algorithm the best swap is performed, even if this increases the reduced cost. The algorithm maintains a set $F \subset N$ of requests that were deleted from the route in a previous iteration. These requests are not allowed to reenter the route. The best route found over all iterations is added to the column pool if it has negative reduced cost:

1. $\bar{r} = \emptyset$, $d_{\bar{r}}^k = \infty$, and $F^k = \emptyset$.
2. Let $d_{(r \setminus \{i_0\}) \cup \{j_0\}}^k = \min\{d_{(r \setminus \{i\}) \cup \{j\}}^k \mid i, j \in N, \delta_{ir}^k = 1, \delta_{jr}^k = 0, j \notin F^k\}$.
3. If $d_{(r \setminus \{i_0\}) \cup \{j_0\}}^k = \infty$, then go to Step 6, otherwise set $r = (r \setminus \{i_0\}) \cup \{j_0\}$ and $F^k = F^k \cup \{i_0\}$.
4. If $d_r^k < d_{\bar{r}}^k$, then set $d_{\bar{r}}^k = d_r^k$ and $\bar{r} = r$.
5. Go to Step 2.

6. If $d_{\bar{r}}^k < 0$, then add \bar{r} to the column pool.

3.2.3 Solving the Integer Program

In order to obtain integral solutions, we need a branching scheme that excludes the current fractional solution, validly partitions the solution space of the problem, and does not complicate the pricing problem too much. The third requirement almost always excludes the standard branching rules based on variable fixing. Fixing a variable to 1 does not complicate the pricing problem, because it just reduces the size of the problem. However, fixing a variable to 0 corresponds to forbidding a certain solution to the pricing problem. Deeper down the search tree this implies that a set of solutions to the pricing problem must be excluded, which is in general very complicated, if not impossible.

In order to develop branching schemes that satisfy the requirements given above, the structure of the pickup and delivery problem has to be exploited. The pickup and delivery problem can be roughly divided into two parts. First, the assignment of the requests to the available vehicles. Second, the construction of pickup and delivery routes. This suggests two types of branching schemes, one that partitions the solution space with respect to assignment decisions, the other with respect to routing decisions.

Branching on routing decisions

Dumas, Desrosiers and Soumis [1991] present a branching scheme that focuses on routing decisions. The branching scheme is based on the one proposed for the asymmetric traveling salesman problem by Carpaneto and Toth [1980]. Let $x_r^k > 0$ be a fractional variable in the current LP solution. Suppose that the corresponding route r serves n requests $\{i_1, i_2, \dots, i_n\} \subset N$ in such a way that i_p is picked up before i_q if $p < q$. Binary order variables O_{ij} ($i, j \in N^+ \cup M^+$) are introduced, where O_{ij} is equal to 1 if no requests are picked up between locations i and j , and 0 otherwise. The current subset of solutions is now divided into $n + 2$ subsets as follows. Define $i_0^+ = i_{n+1}^+ = k^+$. The first subset is characterized by the constraints $O_{i_0^+ i_1^+} = O_{i_1^+ i_2^+} = \dots = O_{i_n^+ i_{n+1}^+} = 1$. For each $j \in \{0, 1, \dots, n\}$ another subset is characterized by $O_{i_0^+ i_1^+} = O_{i_1^+ i_2^+} = \dots = O_{i_{j-1}^+ i_j^+} = 1$ and $O_{i_j^+ i_{j+1}^+} = 0$. The dynamic programming algorithm to solve the pricing problem can be easily modified such that the routes found by the algorithm satisfy the constraints on the order variables.

Branching on assignment decisions

We present an alternative branching scheme that focuses on assignment decisions rather than routing decisions. Assignment decisions constitute higher level decisions and have a greater impact on the structure of the solution. Therefore, we feel that assignment

decisions are more important than routing decisions and should be made first.

Let x be the current fractional solution to LP . Now define for each $i \in N$ and $k \in M$ the *assignment value* $z_i^k = \sum_{r \in \Omega_k} \delta_{ir}^k x_r^k$, indicating what fraction of request i is served by vehicle k in the current LP solution.

Proposition 1 *Let x be an optimal solution of LP , and let $z_i^k = \sum_{r \in \Omega_k} \delta_{ir}^k x_r^k$ ($k \in M, i \in N$). Then x is integral if and only if z is integral.*

Proof If x is integral, then trivially z is integral. Now suppose z is integral. Let $k \in M$ and $i \in N$ satisfy $z_i^k = 1$. Because $\sum_{r \in \Omega_k} \delta_{ir}^k x_r^k = 1$ and $\sum_{r \in \Omega_k} x_r^k \leq 1$, we have that $\delta_{ir}^k = 1$ for all $r \in \Omega_k$ with $x_r^k > 0$. Therefore all routes $r \in \Omega_k$ with $x_r^k > 0$ serve the same requests. Because x is an optimal solution to LP , these routes all have the same cost and so they are identical. Since all routes in Ω_k are distinct, this implies that x is an integral solution. \square

Based on the above Proposition, we propose the following branching scheme: *When x is fractional, find a request $\hat{i} \in N$ and a vehicle $\hat{k} \in M$ with $0 < z_{\hat{i}}^{\hat{k}} < 1$ and create two subsets characterized by $z_{\hat{i}}^{\hat{k}} = 0$ and $z_{\hat{i}}^{\hat{k}} = 1$ respectively.*

This branching scheme can be viewed as a special case of the branching scheme proposed by Ryan and Foster [1981] and does not complicate the pricing problem in the subsets. The restriction $z_{\hat{i}}^{\hat{k}} = 0$ is easily satisfied by ignoring request \hat{i} when solving pricing problem $S_{\hat{k}}$. It is less obvious that the restriction $z_{\hat{i}}^{\hat{k}} = 1$, i.e., requiring that request \hat{i} is served by vehicle \hat{k} , can also be satisfied when solving pricing problem $S_{\hat{k}}$. In this case, it seems to depend on the specific algorithm used. For example, if a construction heuristic is used it is easy and if a dynamic programming algorithm is used it is doable. Sol [1994] shows that, in fact, any algorithm for the pricing problem $S_{\hat{k}}$ can be used with an appropriate choice of dual variables.

3.2.4 Primal solutions

In order to keep the search tree small, we need good lower and upper bounds. To obtain good upper bounds, we have developed a primal heuristic that, in each node of the search tree, tries to construct a feasible solution starting from the current fractional solution and, if successful, tries to improve this solution.

The constructive algorithm is based on the assignment values defined in the previous section. Let x be the fractional solution to the LP. Then we define for each request $i \in N$ and each vehicle $k \in M$ the fractional assignment value $z_i^k = \sum_{r \in \Omega_k} \delta_{ir}^k x_r^k$. Note that if this value is large, the LP solution indicates that it is likely that transportation request i is served by vehicle k . The following algorithm now tries to construct a feasible solution:

1. $N_0 = N$.
Sort the pairs $(k, i) \in M \times N$ such that $z_{i_1}^{k_1} \geq z_{i_2}^{k_2} \geq z_{i_3}^{k_3} \geq \dots$
 $t = 1$.
For each vehicle $k \in M$ set r_k to be the empty route of k .
2. If $i_t \notin N_0$ or i_t cannot be added to route r_{k_t} , then go to 4.
3. Add i_t to route r_{k_t} .
Remove i_t from N_0 .
4. $t = t + 1$.
If $N_0 \neq \emptyset$ and $t \leq |M||N|$, then go to 2, otherwise stop.

Checking whether i_t can be added to route r_{k_t} in Step 2, is done with our cheapest insertion algorithm.

If a solution is found in this way, it is subjected to three local search algorithms. The first algorithm considers a single route and *reinserts* each request. Because the routes were constructed by sequential insertion, the cheapest insertion of a request in its route can differ from its current positions. The other two algorithms consider two routes and try to decrease the total cost by moving requests from one route to the other, or by exchanging two requests between routes. The cost of both operations is approximated by insertion and deletion algorithms.

The quality of the upper bound may be further improved by incorporating more sophisticated iterative improvement algorithms, such as those described in Kindervater and Savelsbergh [1992].

3.3 Computational experiments

The ultimate goal of our research is the development of a high quality approximation algorithm for the GPDP that can solve moderate size instances in an acceptable amount of computation time. We believe that, with the appropriate choices, the branch-and-price algorithm described above satisfies these requirements, and the results of the various computational experiments that we have conducted support that claim.

The computational experiments are designed to answer the following questions:

- Do we solve the linear relaxations faster when we use the column management scheme described in Section 3.2.1?
- If we never solve the pricing problem to optimality, i.e., if we only use the heuristics to solve the pricing problem, can we still get high quality solutions?

- Does generating columns at every node of the search tree, as opposed to only generating columns in the root node, lead to higher quality solutions?

We have implemented several versions of our algorithm to be able to answer these questions and start their description with a discussion of the implementation issues that are common to all of them.

The cheapest insertion algorithm that we have developed for the pricing problem is also used to construct a starting solution and an initial set of columns. The limit on the number of columns that the dynamic programming algorithm can generate for each vehicle in a single execution is set to 200. We have experimented with limits of 50, 200 and 1000 columns per vehicle, but the computation times hardly varied. Because the limit of 200 columns per vehicle produced slightly better results, this value has been chosen. The threshold value used to decide when the column pool is cleaned up has been set to 10. The branching pair (\hat{k}, \hat{i}) is chosen such that $z_{\hat{i}}^{\hat{k}} = \max\{z_i^k \mid z_i^k < 1, i \in N, k \in M\}$. The search tree is explored according to a best bound search.

We have implemented three approximation algorithms, A_1, A_2 and A_3 . All of them use the column generation scheme of Section 3.2.1. The algorithms differ either in the way they solve the LPs in the root node or in the way they solve the LPs in the other nodes of the search tree. Algorithm A_1 always uses the heuristics to solve the pricing problem and therefore never solves the pricing problem to optimality. In the root node, algorithm A_2 uses the heuristics to solve the pricing problem as long as they produce profitable columns, but reverts to the dynamic programming algorithm when the heuristics fail to produce profitable columns. Consequently, in the root node A_2 solves the LP to optimality and therefore produces a valid lower bound on the optimal solution value. In all the other nodes A_2 only uses the heuristics to solve the pricing problem. In the root node algorithm A_3 is identical to algorithm A_2 , but algorithm A_3 does not generate new columns in any of the other nodes of the search tree. As a consequence, algorithm A_3 circumvents the complications introduced by generating columns throughout the search tree and can use standard branching rules based on variable dichotomy with the given set of columns.

We have implemented two optimization algorithms, O_1 and O_2 . Both of them use the column generation scheme of Section 3.2.1, but O_1 always uses the dynamic programming algorithm for column generation, whereas O_2 uses the heuristics to solve the pricing problem as long as they produce profitable routes.

All versions have been implemented using MINTO, a Mixed INTegeR Optimizer [Nemhauser, Savelsbergh, and Sigismondi 1994]. MINTO is a software system that solves mixed-integer linear programs by a branch-and-bound algorithm with linear programming relaxations. It also provides automatic constraint classification, preprocessing, primal heuristics and constraint generation. Moreover, the user can enrich the basic al-

gorithm by providing a variety of specialized application routines that can customize MINTO to achieve maximum efficiency for a problem class. All our computational experiments have been conducted with MINTO 2.0/CPLEX 3.0 and have been run on an IBM/RS6000 model 550.

3.3.1 Test problems

Because the ability to solve the pricing problem to optimality, and thus the ability to solve an instance to optimality, strongly depends on the size of the set of all feasible solutions, and this size strongly depends on the number of transportation requests that can be in a vehicle at the same time and the width of the pickup and delivery time windows, we have developed a random problem generator that allows us to vary these instance characteristics.

Instances are constructed as follows: Generate a set of 100 points randomly within a square of size 200×200 . The distance between two points is the Euclidean distance. The travel time between two points is equal to the distance between these points. Origins (i^+), destinations (i^-) and vehicle home locations (k^+) are now chosen from this set of points. The load of a request is selected from an interval $[q^{\min}, q^{\max}]$. The capacity of all vehicles is equal to Q . The time windows of the requests are constructed in the following way: The planning period has length $L = 600$. Each window has width W . For each request i choose e_{i^+} randomly within the interval $[0, e_i^{\max}]$, where $e_i^{\max} = L - t_{i^+ i^-}$. The time windows for request i are now calculated as $[e_{i^+}, e_{i^+} + W]$ for the pickup and $[e_{i^+} + t_{i^+ i^-}, e_{i^+} + t_{i^+ i^-} + W]$ for the delivery. A time unit can be interpreted as a minute. In this way the length of the diagonal of the square corresponds to approximately half a planning period. We choose the number of available vehicles $|M| = |N|/2$.

As indicated earlier, the objective is to minimize the number of vehicles used and the total distance traveled. We have taken the fixed cost F to be $F = 10000$.

Table 1 lists the problem classes that we have used in our first experiments in order of anticipated increasing difficulty. We have randomly generated 10 instances in each problem class. Although the number of transportation requests is not large, it is large enough to enable us to analyze the characteristics of the different versions of the algorithms we have proposed.

We have also tried to obtain the problem instances used by Dumas, Desrosiers, and Soumis [1991], but unfortunately, these are no longer available.

3.3.2 Quality of the lower bound

We first consider the quality of the lower bound Z_{LP} obtained in the root of the branch and bound tree. Note that this includes the addition of the constraint $\sum_{k \in M} \sum_{r \in \Omega_k} x_r^k \geq$

Class	$ N $	$ M $	q^{\min}	q^{\max}	Q	W
A30	30	15	5	15	15	60
B30	30	15	5	20	20	60
C30	30	15	5	15	15	120
D30	30	15	5	20	20	120

Table 1: Problem classes

m (see Section 3.2.1). For all instances, the optimal number of vehicles equals m . We therefore focus on the quality of the lower bound with respect to the total distance traveled. Table 2 shows the linear programming bound at the root (Z_{LP}), the value of the optimal solution (Z_{OPT}), and the integrality gap (as a percentage of the optimal value) with respect to distance traveled only: $100(Z_{OPT} - Z_{LP})/(Z_{OPT} - mF)$. For 17 out of 40 instances this gap equals 0, indicating that the problem was solved without any branching, and only for 7 out of 40 instances the gap exceeds 1%.

Based on this experiment, we conclude that the lower bound computed in the root node provides a good indication of the value of an optimal solution.

3.3.3 The optimization algorithms

To analyze the effect of using heuristics in the column generation scheme on the time required to solve the linear relaxations, we have solved the initial linear programming relaxation by both algorithm O_1 and O_2 for all instances in the problem classes A30, B30, C30 and D30. Table 3 shows the CPU time, the number of columns generated, and the number of columns added by O_1 and O_2 . The number of columns generated is the total number of columns that have been stored in the column pool during the solution process. The last column shows the quotient $\text{CPU}(O_2)/\text{CPU}(O_1)$.

Algorithm O_2 clearly outperforms O_1 . Over all 40 instances, we have observed an average decrease in computation time of 35% when using the approximation algorithms for the pricing problem. For problem class D30 the average computation time was almost halved. The number of columns in the optimal master problem never becomes very large and this number does not differ drastically for O_1 and O_2 . There is however a big difference in the number of generated columns. This is due to the fact that the dynamic programming algorithm stores all columns with negative reduced cost it encounters (up to 200 per vehicle per execution) in the column pool. Although this results in a larger pool size for O_1 than for O_2 , the larger pool size does not cause the differences in computation times. The differences in computation times can be fully attributed to the

Problem	Z_{opt}	Z_{LP}	gap	Problem	Z_{opt}	Z_{LP}	gap
A30.1	104291	104263.0	0.65	B30.1	124655	124655.0	0
A30.2	114451	114451.0	0	B30.2	94341	94332.0	0.21
A30.3	135415	135415.0	0	B30.3	124882	124872.0	0.21
A30.4	115148	115148.0	0	B30.4	104431	104431.0	0
A30.5	114535	114535.0	0	B30.5	94434	94403.0	0.70
A30.6	114967	114939.8	0.55	B30.6	104264	104264.0	0
A30.7	104247	104247.0	0	B30.7	104263	104262.0	0.02
A30.8	115154	115146.0	0.16	B30.8	135332	135332.0	0
A30.9	114883	114883.0	0	B30.9	145466	145466.0	0
A30.10	115208	115208.0	0	B30.10	114685	114685.0	0
C30.1	84664	84604.0	1.29	D30.1	94460	94440.5	0.44
C30.2	84674	84674.0	0	D30.2	73855	73832.0	0.60
C30.3	84011	83961.7	1.23	D30.3	84737	84737.0	0
C30.4	74205	74168.7	0.86	D30.4	84765	84765.0	0
C30.5	84426	84391.2	0.79	D30.5	84593	84498.5	2.06
C30.6	94547	94500.5	1.02	D30.6	84257	84247.5	0.22
C30.7	84270	84187.4	1.93	D30.7	84302	84278.5	0.55
C30.8	94778	94767.0	0.23	D30.8	74293	74097.0	4.57
C30.9	84246	84231.2	0.35	D30.9	84110	84110.0	0
C30.10	84230	84213.2	0.40	D30.10	74499	74402.5	2.14

Table 2: Optimal solutions and integrality gaps

fact that O_1 only uses optimization algorithms to solve the pricing problem whereas O_2 also uses approximation algorithms.

Based on the above observations, we have chosen algorithm O_2 to solve all the problem instances to optimality. Table 4 shows the total CPU time and the number of nodes evaluated in the search tree.

The results presented in Table 4 indicate that for loosely constrained instances the time required to solve an instance may vary significantly and in some cases may become prohibitively large.

3.3.4 The approximation algorithms

The computational experiments with the optimization algorithm have demonstrated that there can be significant differences in solution times even for instances in the same problem class. This type of behavior is generally unacceptable in practical planning situations. An analysis of the distribution of the total computation time over the various components of the branch-and-price algorithm revealed that most of the time was spent

on optimally solving the pricing problems. This suggests that we may be able to reduce the computation times and improve the robustness by reducing the number of times the pricing problem is solved to optimality. This observation motivated the three approximation algorithms.

Table 5 shows the CPU time, the number of evaluated nodes, and the relative error $(Z_{BEST} - Z_{OPT}) / (Z_{OPT} - mF)$ for the approximation algorithms A_1 , A_2 , and A_3 . When a problem is solved in the root, A_2 and A_3 are equivalent to O_2 . This event is indicated by an asterisk (*) in the last column.

Algorithm A_2 clearly outperforms the others with respect to quality of the solutions. It solves 36 out of 40 problems to optimality. For 17 of these problems this is due to the fact that no branching was required and for 19 problems A_2 found the optimal solution even though branching was required and the pricing problems were only solved to optimality in the root node. For the four problems that were not solved to optimality, the relative error was at most 0.33 %. By comparing the results of A_2 and A_3 , we conclude that it pays to use column generation during branch and bound. However, as the relative errors of A_3 are still small, it is clear that creating a good set of columns in the root is the most important issue in finding good approximate solutions.

Algorithm A_1 , which never solves the pricing problem to optimality, outperforms the others with respect to speed. For all problems the optimal number of vehicles was obtained, and though only 9 out of 40 problems are solved to optimality, the average relative error over all 40 problems is only 1.07%. These observations indicate that A_1 might be a good algorithm for practical situations where problem sizes are bigger and time and capacity constraints are less restrictive. In fact, in such situations it may not be possible to use the other algorithms because of the computation times of the dynamic programming algorithm.

Larger instances

To determine whether the conclusions drawn based on the experiments with instances with 30 transportation requests remain valid when larger instances are solved, we have generated several sets of larger instances and tested algorithms A_1 and A_2 on these sets. We did not include algorithm A_3 since its computational requirements do not differ substantially from those of algorithm A_2 and A_2 generally produces higher quality solutions. The characteristics of these problem classes are shown in Table 6.

The set DAR50 is a set of instances of the dial-a-ride problem, which is a well-known special case of the pickup and delivery problem in which loads represent people. In dial-a-ride problems the capacity restrictions are fairly loose.

As our ultimate goal is the development of a high quality approximation algorithm for the GPDP that can solve moderate size instances in an acceptable amount of com-

putation time, we have imposed an upper bound of 10 minutes of cpu time. When this bound causes the algorithm to stop, we report the best solution found so far. We have done this to mimic practical situations in which there exists a limit on the time available to construct a set of routes.

The results for classes A50 and B50 are shown in Table 7. The results for classes C50 and D50 are shown in Table 8. For these more loosely constrained problem classes, we can clearly observe the computational disadvantages associated with solving the pricing problem to optimality, since algorithm A_2 terminated prematurely for all but 2 instances, i.e., it ran into the time limit of 10 minutes. In fact, for all these instances, the algorithm has not been able to completely evaluate the root node! On the other hand, algorithm A_1 terminated normally for most instances and produces significantly better solutions. To determine the affect of the time limit on the quality of the solutions, we have run both algorithms on the same set of instances with a time limit of 30 instead of 10 minutes. The results can also be found in Table 8 and indicate that allowing more time does result in better solutions.

The results presented in Tables 7 and 8 indicate that algorithm A_2 performs well on tightly constrained instances, but that computational requirements become a bottleneck for loosely constrained instances. Algorithm A_1 on the other hand produces high quality solutions in an acceptable amount of computation time regardless of the problem class.

The results presented in Table 9 indicate that both algorithms are capable of producing high quality solutions for instances of the dial-a-ride problem.

4 DRIVE

Since the general pickup and delivery problem does not capture all the characteristics of the direct transportation system at Van Gend & Loos BV, the branch-and-price algorithm cannot be applied directly. In this section, we discuss the modeling tricks and the algorithmic adjustments that had to be made to be able to handle the specific characteristics of the direct transportation system at Van Gend & Loos BV.

4.1 Handling the dynamics

DRIVE uses an iterative approach to handle the dynamics of the direct transportation system. At each iteration, a static general pickup and delivery problem, which we call the reoptimization problem, is solved. Each reoptimization problem, which is characterized by a set of vehicles, a set of transportation requests, and a set of initial routes, depends on the solution obtained in the previous iteration. In the following subsections, we describe how the reoptimization problem is constructed when DRIVE is invoked at time τ .

4.1.1 Planning horizon

Sometimes, requests are known long before they can be served. It is not necessary to include such requests into the planning process as soon as they become available. We therefore introduce a planning horizon $H \geq 0$, and we only consider those requests that can be picked up before time $\tau + H$. The parameter H also provides a means to control the size of the problem instance.

4.1.2 Requests

Let the active set of requests consist of all requests that are known at time τ , that have not yet been completed, and that can be picked up before time $\tau + H$. We distinguish two types of active requests: permanently assigned requests and non-assigned requests. An active request that is served on an existing route and that has been picked up before τ is labeled as a permanently assigned request, since it has to be assigned to the vehicle associated with the route.

For each existing route, we introduce a virtual request representing all permanently assigned (but not yet completed) requests, i.e., representing the loads that are on board the vehicle at τ . The origin of the virtual request is the first location on the existing route visited after time τ and the destinations of the virtual request are the locations where the loads that are on board the vehicle must be delivered. The set of virtual requests replaces the set of permanently assigned requests. For the virtual requests we relax the constraint that deliveries have to be made in a predefined order, because the virtual request may contain destinations of various different requests.

We now define the set of requests for the reoptimization problem to be the set of virtual requests and the set of non-assigned requests, and we enforce each that vehicle associated with an existing route starts its new route by picking up the virtual request.

4.1.3 Route costs

The cost of a vehicle on a given day depends on the distance traveled and on the sleeping location of the driver on that day (recall that routes typically extend over several days). Let $L_r^k(t)$ be the distance traveled by vehicle k on route r on day t , and let $s_r^k(t)$ indicate whether the driver sleeps at his home location k^+ at the end of day t ($s^k(t) = 0$) or somewhere else ($s^k(t) = 1$). The amount that has to be paid to the driver is then equal to

$$\sum_t (\pi_k \max\{L^k(t), L_{\min}^k\} + \mu s^k(t)),$$

where π_k is the price per unit of distance traveled, L_{\min}^k is the guaranteed minimum traveling distance that has to be paid, and μ is the compensation that has to be paid

when a driver does not sleep at his home location. Note that a driver's pay includes the lease price of the vehicle.

These route costs cannot be used directly for the following reasons:

- The last part of a route may change in the future and therefore does not provide a reliable cost estimate.
- For a vehicle that is currently in use, the distance $L_r^k(0)$ will not exceed L_{\min}^k during a significant part of the day, since initially only a subset of the requests is known. Similarly, for all vehicles, $L_r^k(t)$ will hardly ever exceed L_{\min}^k for $t \geq 1$. Therefore, the route costs do not differentiate the various routes.

We use the following route costs. For a vehicle currently in use

$$c_r^k = \pi_k \sum_{0 \leq t \leq T_r} (\alpha \max\{L_r^k(t), L_{\min}^k\} + (1 - \alpha)L_r^k(t)) + \mu s_r^k(0)$$

and for all other vehicles

$$c_r^k = \pi_k \sum_{1 \leq t \leq T_r} (\alpha \max\{L_r^k(t), L_{\min}^k\} + (1 - \alpha)L_r^k(t)) + F,$$

for some $F \geq 0$ and $0 \leq \alpha \leq 1$, and where T_r denotes the last day that the vehicle is in use. Note that we do not take the sleeping locations of days $t \geq 1$ into account.

When we choose $\alpha = 1$, we obtain a cost structure that, during most of the current day, satisfies $c_r^k \approx (T_r + 1)\pi_k L_{\min}^k + \mu s_r^k(0)$ if the vehicle is currently in use, and $c_r^k \approx T_r \pi_k L_{\min}^k + F$ for all the other vehicles, which are almost independent of r . When we choose $0 < \alpha < 1$, we reflect that we want to make good short-term decisions, but still use some of the real cost structure in the model.

In practice, the values of L_{\min}^k are much smaller than the average distance traveled per day per vehicle. If \bar{L} is this average, then we have $\bar{L}/L_{\min}^k \approx 4/3$. The part of the objective function that deals with the costs associated with travel distance does not distinguish between a solution in which four vehicles each travel a distance L_{\min}^k and a solution in which three vehicles each travel a distance \bar{L} . The first solution, however, is unacceptable for Van Gend & Loos BV, because it would decrease the average income of the drivers, which would lead to an increase of the values π_k . Therefore, we have included the constant F in the cost for a vehicle currently in use. By taking the constant $F \gg \pi_k \alpha L_{\min}^k$, we discourage using a new vehicle. A new vehicle is introduced only when a new request becomes available that cannot be served by a vehicle currently in use.

4.1.4 The reoptimization problem

As indicated before, we can distinguish two phases in the planning process of each day. In the morning, the planners focus on the work that has to be done during that day. In the afternoon, the planners also consider the work that has to be done the next day and decide how many vehicles will be required to satisfy the (anticipated) transportation request for the next day.

We can easily accomplish that the known requests for the next day are being considered in the afternoon, by properly selecting the planning horizon H . However, since the objective function minimizes the number of vehicles and only a fraction of the requests that have to be served the next day is known, we must force a solution to use enough vehicles for the next day, so that new requests can be added when they become available.

We accomplish this by adding the following constraint to the set partitioning formulation

$$\sum_{k \in M} \sum_{r \in \Omega_k} Q_k x_r^k \geq \bar{Q},$$

where Ω_k is the set of feasible routes for vehicle k , Q_k is the capacity of vehicle k , x_r^k the variable that indicates whether route r for vehicle k is selected ($x_r^k = 1$) or not ($x_r^k = 0$), and \bar{Q} an estimate of the total vehicle capacity that will be needed during the next day.

In the morning \bar{Q} is set to a small value, say $\bar{Q} = 0$, effectively making the constraint redundant. At this time, the emphasis will be on short-term decisions.

4.2 Route generation

The branch-and-price algorithm uses construction and improvement algorithms to generate additional routes. In this section, we discuss the modifications that had to be made to these algorithms to handle the specific characteristics of the direct transportation system: lunch breaks and night breaks.

A driver must have two breaks every day of a working period: a night break and a lunch break. Let Δ be the day length. So day t is the time interval $(t\Delta, (t+1)\Delta]$. A break on day t for vehicle k is characterized by a time window $[e, l] \subset (t\Delta, (t+1)\Delta]$ in which it must start, a duration σ , and a set V^k of locations at which the break can take place.

A night break has to be either at the vehicle's home location k^+ or at one of a set of common sleeping locations V_0 . On Friday the night break has to be at the vehicle's home location. On Monday through Thursday the night break location depends on the predecessor and the successor in the route. Suppose we want a night break to take place between locations u and v , and suppose that vehicle k departs from u at time D_u . The

night break takes place at a location $w \in V^k$ for which $D_u + t_{uw} \leq l$ and that minimizes the d-tour $d_{uw} + d_{wv}$.

A lunch break takes place immediately after arriving at a location on the vehicle's route, immediately before leaving a location on the vehicle's route, or at some location $j \in V_0 \cup \{k^+\}$ when traveling between two locations on the vehicle's route. The lunch location depends on the predecessor and the successor in the route. Suppose we want the lunch break to take place between location u and v , and suppose that vehicle k departs from u at time $D_u \leq l$. If $D_u \geq e$, lunch takes place at u . Otherwise, let $A_v = D_u + t_{uv}$. If $A_v \leq l$, then lunch takes place at v . Otherwise lunch takes place at a location $w \in V_0 \cup \{k^+\}$ for which $D_u + t_{uw} \leq l$ and that minimizes the detour $d_{uw} + d_{wv}$. If $D_u + t_{uw} > l$ for all $w \in V_0 \cup \{k^+\}$, then lunch takes place in u .

The locations of the break stops are not a priori determined. They depend on the predecessor and the successor stops in the route. This implies that the insertion of a request into route r may lead to the relocation of some of these break stops. This obviously complicates computing the feasibility and cost associated with inserting a new request.

Suppose that route r for vehicle k ends at the end of day j . Inserting a new request into route r may require that the vehicle is also in use on day $j + 1$. When we add a day at the end of a route, we must also introduce a lunch break and a night break for that day. This complicates the insertion algorithm. We therefore assume that, for some $T \geq 1$, all required breaks on days $0, 1, \dots, T$, are already present in the route, regardless of whether the vehicle will be used on those days or not. We do not allow inserting a request after the last night break. T should be chosen large enough to ensure that this is not restrictive. In this way, for example, a route for a new vehicle that serves no requests consists of $2T$ stops: T lunch breaks and T night breaks.

4.3 Feasible solutions

In a dynamic environment, a planning tool must be able to provide a good solution within a short amount of computation time. With an LP based branch-and-bound algorithm, we cannot guarantee that a good integral solution is found fast, if we only rely on the branching phase to produce these integral solutions. Therefore DRIVE frequently applies its primal heuristics. In fact, the primal heuristics are not only invoked when an LP has been solved to optimality, but each time a set of new routes is about to be generated, i.e., immediately before the pricing heuristics are activated.

Furthermore, DRIVE uses a heuristic construction algorithm in order to produce a starting solution very fast. This algorithm takes the solution of the previous reoptimization process as a starting point, and then sequentially assigns each new request j to the vehicle for which the marginal insertion cost is minimal. If it is infeasible to insert the

new request in an existing route, then a new vehicle is introduced. After all new requests have been inserted into a route, we try to improve the solution by applying three types of improvement algorithms. The first algorithm reinserts each request into its route. When, after request j has been inserted, new requests have been inserted into the same route, the current positions of the pickup and deliveries may no longer be optimal. Reinserting request j may therefore decrease the route cost. The other two algorithms take two routes r_1 and r_2 , and try to find r'_1 and r'_2 such that $c_{r'_1}^{k_1} + c_{r'_2}^{k_2} < c_{r_1}^{k_1} + c_{r_2}^{k_2}$, either by moving requests from r_1 to r_2 , or by exchanging requests between r_1 and r_2 .

An advantage of using a set partitioning model in a dynamic environment is the fact that many of the routes that have been generated during the solution of the previous reoptimization problem are useful for the current reoptimization problem. Suppose that r is some route for vehicle k that has been generated in the previous reoptimization process, and suppose that r serves requests $j \in N_r$. We now create a new route for vehicle k as follows. Let r' be the route for vehicle k that only serves the virtual request i_k . Now sequentially insert all requests $j \in N_r$ that have not been permanently assigned to another vehicle, i.e., all requests $j \in N_r \cap N_0$, into route r' . The resulting routes provided a good set of initial routes for the branch-and-price algorithm.

5 Case Study

We have tested DRIVE by simulating a dynamic planning environment with real-life data. These data contained all requests that had been served by Van Gend & Loos BV in a given period. These simulations were based on a stand-alone methodology. Solutions presented by DRIVE were not modified by a planner, but were considered as being executed as proposed.

During the development of DRIVE, we repeatedly compared the results of the simulations to the results of the planners at Van Gend & Loos BV over the same period. These comparisons were needed for various reasons. First, we had to make sure that the data was complete and that we were aware of all constraints that a route should satisfy. Second, we had to compare global solution characteristics, such as the number of vehicles used and the average distance traveled per vehicle per day, in order to assure that these were acceptable to Van Gend & Loos BV. Finally, we compared the total distance traveled and the total cost over the entire planning period, in order to show that DRIVE is capable of providing good solutions.

5.1 Organization of the tests

The test data covered a period of 14 working days, starting with a Thursday and ending with a Tuesday. The actual evaluation period covered 10 working days, starting with

a Monday, and ending with a Friday. The two additional days at the beginning of the simulation were introduced in order to make sure that the vehicles are not empty on Monday morning, because that would not be realistic. The two additional days at the end of the simulation prevented DRIVE from deferring requests in order to obtain better results for the last days of the evaluation period. We only present results for the evaluation period.

Besides the usual request information, such as addresses of origins and destinations, time windows and load sizes, the test data also contained the times at which the requests became available, i.e., the times at which the requests were called in by the client. We used this data to simulate the process of clients calling in new requests.

The tests have been organized such that we have simulated the invocation of DRIVE once every hour, from 6:00 to 18:00, and once at midnight, for each day of the planning period. We assumed that the solution presented by DRIVE is then executed until DRIVE is invoked again. At 15:00 we let DRIVE select the vehicles for the next day.

In practice, DRIVE must be able to produce a solution within a reasonable amount of time, such that a planner has enough time to evaluate the proposed solution. Therefore we have put an upper bound on CPU time of 5 minutes for each time we invoked DRIVE, with an exception for the run at 15:00, where the upper bound was 10 minutes.

5.2 Characteristics of the reoptimization problems

Table 10 shows the sizes of the reoptimization problems at times 0:00, 9:00, 12:00, 15:00 and 18:00 for each day of the two planning weeks. For each reoptimization problem, we list the number of active requests $|N|$, the number of non-assigned requests $|N_1|$, and the number of vehicles $|M_0|$ currently in use. The set of active requests N has been defined by taking the planning horizon H equal to 19 hours. In this way we consider requests with an earliest pickup time of 8:00 the next day, in the planning process starting at 13:00 this day. When we take H much larger, DRIVE spends much CPU time in calculating assignments that will almost certainly have to be changed when new requests become available.

5.3 The contribution of column generation

From our test results we observed that it is sometimes difficult to construct a feasible solution to a reoptimization problem, i.e., a solution in which all requests $j \in N_0$ are served. Note that we cannot always serve a new request by introducing a new vehicle. A new vehicle can only start working at the beginning of the next working day, so it cannot serve a new request that has to be picked up during the current day. The construction algorithm that we used to produce a starting solution could not always insert all new

requests. At the beginning of the column generation algorithm, all requests that could not be inserted in any route during the creation of the starting solution, get a very high dual value. Therefore the pricing heuristics will automatically try to construct routes that serve these requests.

5.4 Quality of solutions

Tables 11 and 12 show the results of DRIVE compared to the results of the planners at Van Gend & Loos BV (VGL). For each day of the planning period we show the number of vehicles used, the total distance traveled, and the total cost of the solutions. The travel distances and the costs have been scaled, such that the total travel distance and the total cost per week equal 100.0 for the solution of Van Gend & Loos BV.

We observe that for the first planning week, DRIVE obtained an improvement of 4.7% of the total cost compared to the planners at Van Gend & Loos BV. For the second week, this improvement was 3.7%. The number of vehicles used by DRIVE is not substantially smaller than the number used by the planners at Van Gend & Loos BV. For the second week, it is even 7.7% higher. Apparently, DRIVE obtains the cost decrease by constructing better assignments of requests to vehicles.

The characteristics in Table 10 indicate that the first planning week is much busier than the second week. This is reflected in the difference between the solutions of DRIVE and the solutions of the planners. In a busy period, there are so many active requests that the planners must reduce the planning horizon in order to keep the problem manageable. At this time, DRIVE provides much better solutions, because it is able to look further into the future. In a quiet period, a planner has more time to make decisions, so he can better evaluate the effect of various assignments. In such a period, the cost decrease that DRIVE obtains is smaller but can still be significant.

6 Concluding remarks

Although we tested DRIVE only in a simulated environment, the results indicate that DRIVE will provide a good basis for developing a decision support system at Van Gend & Loos BV. Clearly, the simulated environment in which we tested DRIVE was not entirely realistic. In practice, it sometimes happens that a vehicle breaks down, or that a load is not yet available at an origin at its indicated earliest pickup time. Such situations, which increase total cost, did not occur in our simulations. On the other hand, our tests have indicated that DRIVE, when implemented as a stand-alone system, is capable of providing solutions that are better than those provided by the planners at Van Gend & Loos BV. When DRIVE is embedded in a DSS, its solutions serve as a starting point for these planners. We may therefore expect that the cost decrease that we observed when

DRIVE is used as a stand-alone system, will be even more significant when DRIVE is embedded in a DSS. Furthermore, the quality of the solutions produced by DRIVE may be increased by allowing more running time, switching to a faster computer, and using improved linear programming optimizers.

During the development of DRIVE, we have identified several research topics related to dynamic pickup and delivery problems that deserve more attention in the future. First, we believe that there is not yet a good insight in objective functions for dynamic routing problems. Usually it is possible to define a cost structure that can be used to evaluate the cost of a route *after* it has been executed, but that cost structure does not always provide a good objective function for a reoptimization problem. A second research topic is the development of techniques for predicting requests that are not yet available. If we have some idea about the origin and load of future requests before they are actually called in, this might help in producing better plans.

Acknowledgment

We thank P. Maier, B. Muusers, and T. Steysiger of Van Gend & Loos BV and J. Hendriks of Logion BV for their help and support throughout the project. We also thank the Dutch Technology Foundation for their financial support.

References

- R. ANBIL, R. TANGA AND E.L. JOHNSON (1993). A global approach to crew-pairing optimization, *IBM Systems Journal* 31, 71-78.
- C. BARNHART, E.L. JOHNSON, G.L. NEMHAUSER, M.W.P. SAVELSBERGH, AND P.H. VANCE (1996). Branch-and-Price: Column Generation for Solving Integer Programs. *Operations Research*, to appear.
- G. CARPANETO AND P. TOTH (1980). Some new branching and bounding criteria for the asymmetric travelling salesman problem. *Management Science* 26, 736-743.
- M. DESROCHERS, J. DESROSIERS, AND M. SOLOMON (1992). A new optimization algorithm for the vehicle routing problem with time windows. *Operations Research* 40, 342-354.
- J. DESROSIERS, Y. DUMAS, M.M. SOLOMON, AND F. SOUMIS (1996). Time constrained routing and scheduling. M.O. BALL, T.L. MAGNANTI, C.L. MONMA, AND G.L. NEMHAUSER (eds.). *Network Routing*, North-Holland, 35-139.
- Y. DUMAS, J. DESROSIERS, AND F. SOUMIS (1991). The pickup and delivery problem

with time windows. *European Journal of Operations Research* 54, 7-22.

G.A.P. KINDERVATER AND M.W.P. SAVELSBERGH (1992). Local search in physical distribution management. Memorandum COSOR 92-30, Eindhoven University of Technology.

G.L. NEMHAUSER, M.W.P. SAVELSBERGH AND G.C. SIGISMONDI (1994). MINTO, a Mixed INTeger Optimizer, *Operations Research Letters* 15, 47-58.

D.M. RYAN AND B.A. FOSTER (1981). An integer programming approach to scheduling. A. WREN (ed.) *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*. North-Holland, Amsterdam, 269-280.

M.W.P. SAVELSBERGH (1996). A Branch-and-Price Algorithm for the Generalized Assignment Problem. *Operations Research*, to appear.

M.W.P. SAVELSBERGH AND M. SOL (1995). The General Pickup and Delivery Problem. *Transportation Science* 29, 17-29.

M. SOL (1994). *Column generation techniques for pickup and delivery problems*. Ph.D. Thesis, Eindhoven University of Technology.

P.H. VANCE, C. BARNHART, E.L. JOHNSON AND G.L. NEMHAUSER (1994). Solving binary cutting stock problems by column generation and branch and bound. *Computational Optimization and Applications* 3, 111-130.

Problem	O_1			O_2			
	CPU	generated	added	CPU	generated	added	
A.30.1	83.76	7035	247	48.05	1217	246	0.57
A.30.2	20.33	5039	200	11.05	885	204	0.54
A.30.3	7.29	2006	153	6.83	961	190	0.94
A.30.4	20.32	3146	196	12.86	956	217	0.63
A.30.5	24.15	4863	207	12.57	791	180	0.52
A.30.6	10.93	3603	181	8.54	944	202	0.78
A.30.7	17.07	5261	161	10.76	920	190	0.63
A.30.8	9.64	3087	210	9.89	822	218	1.03
A.30.9	9.80	3579	186	5.65	880	219	0.58
A.30.10	12.31	4213	214	9.25	837	206	0.75
avg.	21.56	4183.2	195.5	13.55	921.3	207.2	0.70
B.30.1	17.67	3164	188	8.65	593	160	0.49
B.30.2	31.69	4655	199	25.15	1093	237	0.79
B.30.3	15.96	3469	166	10.63	691	170	0.67
B.30.4	27.13	5320	191	13.47	949	216	0.50
B.30.5	37.76	6749	226	37.07	938	214	0.98
B.30.6	24.22	5460	191	15.27	950	233	0.63
B.30.7	31.81	5934	238	17.40	975	229	0.55
B.30.8	9.47	1978	185	5.50	801	173	0.58
B.30.9	13.42	4502	165	11.00	856	211	0.82
B.30.10	19.33	3886	191	15.11	934	193	0.78
avg.	22.85	4511.7	194.0	15.93	878.0	203.6	0.68
C.30.1	60.39	8599	250	47.79	1429	251	0.79
C.30.2	179.35	10183	289	110.68	1450	272	0.62
C.30.3	290.67	12427	314	183.90	1484	285	0.63
C.30.4	336.70	13924	306	184.41	1550	302	0.55
C.30.5	177.41	9818	290	110.27	1627	286	0.62
C.30.6	74.61	7330	238	40.17	1204	239	0.54
C.30.7	344.94	11440	283	197.02	1484	277	0.57
C.30.8	126.96	7588	265	83.86	1352	284	0.66
C.30.9	194.27	12490	312	118.65	1501	287	0.61
C.30.10	370.98	13131	278	277.60	1852	323	0.75
avg.	215.63	10693.0	282.5	135.44	1493.3	280.6	0.63
D.30.1	174.21	10646	264	101.39	1438	257	0.58
D.30.2	606.85	15879	292	371.42	1349	278	0.61
D.30.3	121.16	6607	238	84.29	1198	244	0.70
D.30.4	138.02	7271	249	78.17	1399	280	0.57
D.30.5	104.64	7299	259	61.72	1143	298	0.59
D.30.6	165.40	8188	289	73.03	1332	263	0.44
D.30.7	168.99	8653	287	103.01	1518	280	0.61
D.30.8	324.02	12782	279	207.53	1218	252	0.64
D.30.9	485.28	12800	336	242.84	1865	306	0.50
D.30.10	121.66	7958	257	67.27	1507	244	0.55
avg.	241.02	9808.3	275.0	139.07	1396.7	270.2	0.58

Table 3: Performance of optimization algorithms for Z_{LP}

Class	A30		B30		C30		D30	
Nr	CPU	#N	CPU	#N	CPU	#N	CPU	#N
1	106.95	15	8.61	1	60.82	3	177.99	7
2	11.08	1	39.07	5	110.54	1	2552.56	23
3	6.83	1	14.87	3	684.49	27	83.93	1
4	12.85	1	13.28	1	1115.42	35	77.82	1
5	12.54	1	81.61	9	1082.63	57	6402.77	873
6	78.23	47	15.05	1	410.75	91	345.78	23
7	10.75	1	19.68	3	7148.49	329	205.44	9
8	19.11	5	5.41	1	183.12	9	55266.79	1555
9	5.67	1	10.91	1	457.93	17	242.40	1
10	9.23	1	15.09	1	1035.06	23	897.31	51
avg.	27.32	7.4	22.36	2.6	1228.93	59.2	6625.28	254.4

Table 4: Performance of optimization algorithm O_2

Problem	A_1			A_2			A_3			
	CPU	#N	error	CPU	#N	Err	CPU	#N	error	
A.30.1	22.69	35	1.00	50.67	13	0.26	49.34	17	0.02	
A.30.2	10.64	10	2.99	11.01	1	0.00	11.20	1	0.00	*
A.30.3	2.73	1	2.79	6.84	1	0.00	6.92	1	0.00	*
A.30.4	5.26	1	1.07	12.99	1	0.00	12.98	1	0.00	*
A.30.5	5.50	1	0.04	12.47	1	0.00	12.94	1	0.00	*
A.30.6	13.48	17	0.00	15.48	27	0.00	9.37	15	0.46	*
A.30.7	4.70	1	4.87	10.78	1	0.00	10.76	1	0.00	*
A.30.8	3.18	1	1.24	10.41	3	0.00	10.49	11	0.06	*
A.30.9	4.94	1	0.33	5.69	1	0.00	5.56	1	0.00	*
A.30.10	4.14	1	0.31	9.19	1	0.00	9.24	1	0.00	*
avg.	7.73	6.9	1.46	14.55	5.0	0.03	13.88	5.0	0.00	
B.30.1	4.12	1	2.77	8.61	1	0.00	8.80	1	0.00	*
B.30.2	8.10	1	1.08	26.10	5	0.00	25.48	5	0.00	
B.30.3	4.08	1	2.03	10.90	3	0.20	10.68	7	0.20	
B.30.4	5.16	1	0.14	13.38	1	0.00	13.62	1	0.00	*
B.30.5	16.94	15	0.14	39.79	9	0.00	38.35	23	0.00	*
B.30.6	6.38	1	2.18	15.05	1	0.00	15.33	1	0.00	*
B.30.7	7.72	1	0.84	17.69	3	0.28	17.66	3	0.00	*
B.30.8	3.40	1	0.41	5.45	1	0.00	5.52	1	0.00	*
B.30.9	4.30	3	1.43	11.20	1	0.00	11.10	1	0.00	*
B.30.10	3.70	1	1.71	15.51	1	0.00	15.46	1	0.00	*
avg.	6.39	2.6	1.27	16.37	2.6	0.05	16.20	4.4	0.02	
C.30.1	17.89	11	1.11	47.26	3	0.00	48.80	43	2.53	*
C.30.2	22.16	8	1.24	110.90	1	0.00	112.49	1	0.00	*
C.30.3	17.58	5	0.77	213.00	51	0.05	190.02	49	1.35	
C.30.4	67.94	81	1.17	205.84	41	0.00	188.24	15	0.19	
C.30.5	31.14	18	0.16	123.01	33	0.38	113.05	25	0.45	
C.30.6	28.35	18	0.33	61.56	53	0.00	41.09	21	0.00	
C.30.7	103.44	84	0.00	209.86	27	0.00	200.71	47	0.94	
C.30.8	8.07	1	1.88	84.70	3	0.65	86.05	21	0.50	
C.30.9	17.58	3	0.31	131.99	23	0.02	120.87	9	0.02	
C.30.10	61.34	55	0.09	299.27	43	0.00	279.37	15	0.00	
avg.	37.55	28.40	0.71	148.74	27.8	0.11	138.07	24.6	0.60	
D.30.1	11.27	3	1.61	102.39	7	0.22	102.05	3	0.00	*
D.30.2	26.34	10	0.29	371.77	7	0.00	376.67	7	0.00	*
D.30.3	6.50	1	3.31	84.06	1	0.00	86.52	1	0.00	*
D.30.4	10.08	1	3.44	77.33	1	0.00	79.27	1	0.00	*
D.30.5	317.57	469	0.00	232.34	389	0.00	72.64	183	0.89	*
D.30.6	9.26	1	2.37	97.31	53	0.45	74.14	15	0.63	*
D.30.7	11.06	1	0.93	105.90	9	0.00	103.88	5	0.00	*
D.30.8	52.84	37	1.05	211.94	9	0.16	208.98	21	0.00	*
D.30.9	16.69	1	2.85	242.50	1	0.00	246.81	1	0.00	*
D.30.10	8.17	1	3.73	115.86	73	0.00	69.60	7	0.27	*
avg.	46.98	52.5	1.96	164.14	55.0	0.08	142.06	24.4	0.18	

Table 5: Performance of approximation algorithms

Class	$ N $	$ M $	q^{\min}	q^{\max}	Q	W
A50	50	25	5	15	15	60
B50	50	25	5	20	20	60
C50	50	25	5	15	15	120
D50	50	25	5	20	20	120
DAR50	50	25	1	1	5	60

Table 6: Larger and less restricted problem classes

Problem	A_1			A_2		
	Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU
A.50.1	187978	35	110.46	187851	1	105.32
A.50.2	157456	1	32.44	157412	1	142.33
A.50.3	167980	23	99.37	167934	3	80.44
A.50.4	157287	15	75.63	157250	39	555.44
A.50.5	158382	1	23.72	158142	1	94.57
A.50.6	157251	1	23.38	157102	5	121.72
A.50.7	167955	19	84.54	167918	1	87.28
A.50.8	157176	194	395.64	157141	311	600.00
A.50.9	167899	22	86.73	167880	91	150.21
A.50.10	157089	27	104.61	157019	27	547.41
avg.	163645.3	33.8	103.65	163564.9	48.0	248.47
B.50.1	177130	1	21.18	177123	1	88.95
B.50.2	177683	9	43.63	177647	31	183.20
B.50.3	167961	1	26.45	167762	1	68.28
B.50.4	148001	1	27.77	147736	3	138.55
B.50.5	208177	1	11.83	208126	1	45.55
B.50.6	156872	1	29.39	156870	49	287.61
B.50.7	177116	9	30.25	177065	3	117.52
B.50.8	157008	5	45.86	156875	1	251.83
B.50.9	146801	2	40.78	146732	3	301.04
B.50.10	158516	1	32.32	158231	19	190.63
avg.	167526.5	3.1	30.95	167416.7	11.2	167.32

Table 7: Performance of algorithms A_1 and A_2 for the problem classes A50 and B50

Problem	$A_1 (\leq 600)$			$A_1 (\leq 1800)$			$A_2 (\leq 600)$			$A_2 (\leq 1800)$		
	Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU
C.50.1	148086	57	292.59	148086	57	285.93	148036	21	550.61	148036	21	550.02
C.50.2	126554	143	393.16	126554	143	392.77	177398	1	600.00	177398	1	1800.00
C.50.3	137507	2	58.42	137507	2	57.13	137482	11	577.48	137482	11	577.66
C.50.4	127379	99	398.46	127379	99	392.56	177926	1	600.00	127335	762	1800.00
C.50.5	116377	118	600.00	116190	522	1800.00	166648	1	600.00	166648	1	1800.00
C.50.6	167358	57	192.05	167358	57	191.85	188190	1	600.00	188132	1	1800.00
C.50.7	126459	75	272.66	126459	75	272.72	177388	1	600.00	177388	1	1800.00
C.50.8	126726	11	125.81	126726	11	125.85	167227	1	600.00	167227	1	1800.00
C.50.9	156769	66	600.00	106351	577	1761.26	156795	1	600.00	156795	1	1800.00
C.50.10	116707	33	251.97	116707	33	252.02	147619	1	600.00	116742	69	1670.85
avg.	134992.2	66.1	318.51	129931.7	157.6	553.21	164470.9	4.0	592.81	156318.3	86.9	1539.85
D.50.1	116078	116	600.00	116062	184	774.10	166919	1	600.00	166919	1	1800.00
D.50.2	157109	150	600.00	126852	669	1800.00	147106	1	600.00	126864	436	1800.00
D.50.3	187815	179	600.00	127579	1173	1483.97	177738	1	600.00	177738	1	1800.00
D.50.4	117038	111	487.48	117038	111	476.97	156560	1	600.00	156560	1	1800.00
D.50.5	126648	183	528.16	126648	183	516.28	167170	1	600.00	167170	1	1800.00
D.50.6	137037	4	83.92	137037	4	81.97	198132	1	600.00	137016	45	1467.27
D.50.7	116453	161	600.00	116300	997	1446.23	166710	1	600.00	166678	1	1800.00
D.50.8	136878	147	337.80	136878	147	330.74	167463	1	600.00	167463	1	1800.00
D.50.9	116303	37	228.60	116303	37	222.38	146590	1	600.00	146590	1	1800.00
D.50.10	116134	129	466.23	116134	129	457.30	167029	1	600.00	167029	1	1800.00
avg.	132749.3	121.7	453.22	123683.1	363.4	758.99	166141.7	1.0	600.00	158002.7	48.9	1766.73

Table 8: Performance of algorithms A_1 and A_2 for the problem classes C50 and D50

Problem	Z_{LP}	$A_1 (\leq 600)$			$A_1 (\leq 1800)$		
		Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU
DAR.1	125735.05	125851	3	57.17	125851	3	58.37
DAR.2	136073.81	136180	296	600.58	136145	673	1201.60
DAR.3	146219.00	146345	1	25.38	146345	1	25.05
DAR.4	146658.50	146882	275	601.02	146867	537	617.00
DAR.5	135954.90	136070	182	601.21	136047	321	733.72
DAR.6	146411.50	146520	5	38.44	146520	5	38.80
DAR.7	136152.70	136208	61	152.87	136208	61	155.93
DAR.8	125905.50	126007	273	600.39	125999	1013	1800.58
DAR.9	126790.50	127048	1	32.72	127048	1	32.96
DAR.10	135964.75	136036	26	144.85	136036	26	146.47
avg.	136186.62	136314.7	112.3	285.46	136306.6	264.1	481.05

Problem	Z_{LP}	$A_2 (\leq 600)$			$A_2 (\leq 1800)$		
		Z_{BEST}	#N	CPU	Z_{BEST}	#N	CPU
DAR.1	125735.05	146068	1	606.45	125839	9	898.71
DAR.2	136073.81	136151	73	407.96	136151	73	415.19
DAR.3	146219.00	146219	1	188.97	146219	1	191.33
DAR.4	146658.50	146779	109	333.57	146779	109	349.99
DAR.5	135954.90	136050	74	600.18	136045	1042	1800.69
DAR.6	146411.50	146420	11	218.89	146420	11	224.10
DAR.7	136152.70	136195	127	364.19	136195	127	368.19
DAR.8	125905.50	126020	228	600.26	126012	1343	1800.70
DAR.9	126790.50	126956	30	601.16	126896	87	643.18
DAR.10	135964.75	135987	19	336.73	135987	19	343.48
avg.	136186.62	138284.5	67.3	425.84	136254.3	282.1	703.56

Table 9: Performance of algorithms A_1 and A_2 for the problem class DAR50

Week 1					Week 2				
day	time	$ N $	$ N_1 $	$ M_0 $	day	time	$ N $	$ N_1 $	$ M_0 $
Mon	0:00	328	197	92	Mon	0:00	266	153	90
	9:00	300	146	92		9:00	261	111	90
	12:00	285	134	92		12:00	237	90	90
	15:00	322	170	96		15:00	227	104	87
	18:00	329	175	96		18:00	257	135	87
Tue	0:00	349	204	96	Tue	0:00	272	154	87
	9:00	332	161	96		9:00	249	118	87
	12:00	325	162	96		12:00	257	103	87
	15:00	335	193	99		15:00	271	137	90
	18:00	333	178	101		18:00	292	166	90
Wed	0:00	354	208	101	Wed	0:00	304	182	90
	9:00	345	172	101		9:00	272	118	90
	12:00	306	140	101		12:00	254	97	90
	15:00	349	186	100		15:00	256	122	86
	18:00	341	174	100		18:00	261	132	86
Thu	0:00	354	195	100	Thu	0:00	282	182	86
	9:00	313	148	100		9:00	264	118	86
	12:00	292	125	100		12:00	248	97	86
	15:00	285	133	96		15:00	261	144	92
	18:00	285	140	96		18:00	277	154	92
Fri	0:00	306	166	96	Fri	0:00	295	175	92
	9:00	275	122	96		9:00	278	113	92
	12:00	264	99	96		12:00	279	114	92
	15:00	256	123	89		15:00	281	147	95
	18:00	244	125	90		18:00	279	151	95

Table 10: Characteristics of the reoptimization problems for the test set

	Nr. of vehicles		Travel distance		Cost	
	VGL	DRIVE	VGL	DRIVE	VGL	DRIVE
Mon	92	92	20.1	18.2	20.4	18.6
Tue	97	96	19.7	19.9	20.1	20.5
Wed	100	101	20.1	20.1	20.4	20.7
Thu	100	100	20.8	18.8	21.0	19.2
Fri	97	96	19.3	16.1	18.1	16.3
Total	486	485	100.0	93.1	100.0	95.3

Table 11: Results for the first planning week

	Nr. of vehicles		Travel distance		Cost	
	VGL	DRIVE	VGL	DRIVE	VGL	DRIVE
Mon	80	90	21.1	18.1	21.3	19.1
Tue	77	87	17.5	17.5	18.0	18.5
Wed	85	90	21.0	21.1	21.3	21.3
Thu	84	86	19.1	17.0	19.4	17.9
Fri	87	92	21.3	20.0	20.0	19.5
Total	413	445	100.0	93.7	100.0	96.3

Table 12: Results for the second planning week