KATHOLIEKE UNIVERSITEIT
LEUVEN

Faculty of Business and Economics

The periodic vehicle routing problem:
a case study

S. Coene, A. Arnout and F. Spieksma

DEPARTMENT OF DECISION SCIENCES AND INFORMATION MANAGEMENT (KBI)

KBI 0828

The Periodic Vehicle Routing Problem: A Case Study *

Sofie Coene[†]   Arent Arnout   Frits C.R. Spieksma[†]

October 29, 2008

Abstract

This paper deals with a case study which is a variant of the Periodic Vehicle Routing Problem (PVRP). As in the traditional Vehicle Routing Problem (VRP), customer locations each with a certain daily demand are given, as well as a set of capacitated vehicles. In addition, the PVRP has a horizon, say T days, and there is a frequency for each customer stating how often within this T-day period this customer must be visited. A solution to the PVRP consists of T sets of routes that jointly satisfy the demand constraints and the frequency constraints. The objective is to minimize the sum of the costs of all routes over the planning horizon. We develop different algorithms solving the instances of the case studied. Using these algorithms we are able to realize considerable cost reductions compared to the current situation.

Key words: Periodic Vehicle Routing, case study

1   Introduction

In this paper we study a routing problem of a Belgian company collecting waste at slaughterhouses, butchers, and supermarkets. Planning of the routes occurs over a time period of several days (time horizon) in which customers are visited with different frequencies. For instance, supermarkets might request service every day, while for a small butcher one collection a week suffices. The resulting problem is a variant of the Periodic Vehicle Routing Problem (PVRP). As in the traditional Vehicle Routing Problem (VRP), customer locations each with a certain demand function are given, as well as a set of capacitated vehicles. In addition, the PVRP has a horizon, say T days, and there is a frequency for each customer stating how often within this T-day period this customer must be visited. A solution to the PVRP consists of T sets of routes that jointly satisfy the demand constraints and the frequency constraints. The objective is to minimize the sum of the costs of all routes over the planning horizon. Obviously, this problem is at least as hard as the VRP.

PVRP and variants

Several variants of the PVRP are described in literature. A classification of the different variants of the PVRP can be found in a survey by Mourgaya and Vanderbeck [18]. Different objective functions are distinguished, such as minimizing the distance traveled, the driving time, or total transportation cost; however also regionalization of routes, an even spread of workload over the vehicles, the number of vehicles, and service quality can be part of an optimization function. Differences also occur in the constraints which can be divided in three categories: constraints concerning (i) the planning of visits (different frequencies, restrictions on certain days, etc.), (ii) the type of demand (constant or variable; we return to this issue later), and (iii) the vehicles. Where the PVRP is mostly situated on tactical and operational level, Francis et al. [12] include strategic decisions in their model: frequencies of service are variables within the model, and not given parameters. Another variant is the PVRP with intermediate facilities which is described by Angelelli and Speranza [3], Kim et al. [17], and Alonso et al. [2]. Intermediate facilities are locations where vehicles can unload (or reload) and thus renew capacity during a route; this happens in our case, see Section 2.

Case studies

The PVRP is a relevant problem; it occurs for companies that have to carry out periodic repair and maintenance activities or that collect/deliver goords periodically. Blakely et al. [7] describe a case for periodic maintenance of elevators at different customer locations. Further case studies concerning waste collection and road sweeping can be found in Beltrami and Bodin [6] and Eglese and Murdoch [11]. Claassen and Hendriks [9] describe a milk collection problem where it is important that the goods are collected when fresh. For the collection of raw materials for a manufacturer of auto parts, on the contrary, a very long time horizon is considered, see Alegre et al. [1]. Hemmelmayr et al. [15] investigate the periodic delivery of blood products to hospitals by the Austrian Red Cross. In this case, the regularity of deliveries is of uttermost importance. Many other case studies are described in Francis et al. [13] and the references contained therein.

Solution methods

The PVRP is situated on the border between tactical and operational planning, combining the classical VRP with planning over a time horizon. That is why solution methods often consist of two phases. Beltrami and Bodin [6] consider two approaches. In a first approach, routes are developed and then assigned to days of the week; in a second approach customers are assigned to days in a first phase and in a second phase the routing problem for every single day is solved using classical techniques for solving VRPs. This second approach is used in many papers, such as Baptiste et al. [5], Tan and Beasley [19], and Christofides and Beasley [8]. Tan and Beasley [19] first solve an assignment problem to assign customers to days such that total demand in each day does not exceed demand capacity while taking pairwise distances between customers into account. After that stage they solve a VRP for each day in the planning horizon. Thus, they approach this problem as an extension of the assignment problem with a routing component. Christofides and Beasley [8] on the other hand formulate the PVRP as a routing problem with a selection decision. Customers are ordered in descending order of "importance", depending on

the demands, and then selected for a route on a certain day depending on the increase in total cost for the whole period. These approaches are the more classical solution strategies. Recent PVRP literature has focused on metaheuristic methods and mathematical based approaches to solve the problem; we refer to Francis et al for an overview. [13].

The rest of this paper is structured as follows. In section 2 we describe the case under consideration in further detail, in section 3 we give a mathematical formulation, in section 4 we propose a solution method and in the 5th and final section we give some computational results and formulate a conclusion.

2 The case

2.1 A general description

As mentioned, we study a problem that is encountered by a Belgium transportation company, which is responsible for the collection of waste at slaughterhouses, butchers, and supermarkets. This company, which we call company A for confidentiality reasons, has clients all over Belgium and in some areas of northern France.

Legislation that originated from the BSE-epidemic (Bovine Spongiform Encephalopathy, commonly known as mad-cow disease) in the nineties, stipulates that (i) there are 3 categories of animal waste, depending on the risk of containing BSE; (ii) waste from different categories has to be collected separately. Company A only collects waste from two categories: category 1 (high-risk waste) and category 3 (low-risk waste). All high-risk waste is collected in order to be destroyed, while low-risk waste can be further processed into e.g. pet food. Vehicles assigned to collect high-risk waste cannot be used to collect low-risk waste and vice versa. In fact, this implies that company A has to solve two different instances; one instance for the periodic collection of high-risk waste and a second instance for the periodic collection of low-risk waste.

In the current planning process of company A, routes are constructed manually on a regular basis, e.g. every month, and, during that period minor modifications to the routes can be made depending on changes in the set of customers. Company A wishes to decrease the dependence upon human expertise and wants to professionalize the planning procedure. Also, the management of company A wishes to plan the routes more efficiently in order to reduce travel time and travel distance. Several opportunities need to be explored: (i) can total driving time be decreased? (ii) can the routing be done using smaller vehicles? and (iii) is it possible to decrease the vehicle fleet? All of this could be possible through more efficient planning, but obviously, the same level of service towards the customers should be retained. All this also has environmental consequences thanks to a reduction in petrol use and in vehicle use. It is clear that it would not be possible to change the vehicle fleet whenever the routing plan changes. Our routing plan, however, should allow the management to "assemble" a good vehicle fleet in the long term. In the short term, the current vehicle fleet must be respected when constructing a routing plan.

2.2 The low-risk waste instance: details

Here we describe some properties of the instance corresponding to the low-risk waste. There are 48 customers, spread out over Belgium and northern France (see Figure 1). The planning period is one week (actually 6 days), and each customer requests a certain frequency of visit over the planning period. Table 1 gives an overview of how often each frequency occurs. There are 5 different frequencies and a frequency of 4 days does not occur. Company A has 3 trucks

Frequency (nr of days)	1	2	3	5	6
nr of customers	21	15	5	5	2

Table 1: Frequencies low-risk waste

available for collecting low risk waste; their capacities are 12, 22, and 26 tons respectively. Some clients are located in the center of a city and cannot be reached by a truck of 22 tons or bigger. There is a central depot where each route starts in the morning and ends in the evening. When a truck is fully loaded, the driver can unload at a disposal facility and next continues its route. Notice that these disposal facilities can be seen as intermediate facilities, see [3]. Trucks do not need to return empty in the evening, they can also dispose of their load during the tour of the following day. Only when a truck does not drive on the following day it has to be emptied before returning to the depot. An affiliated company processes the waste and has one disposal facility where the trucks can unload 24 hours a day. Loading and unloading times depend on the volume. Legally, the maximum driving time for a driver is restricted to 90 hours within two weeks and the company restricts the daily driving hours to 10 hours. Notice that these are only driving hours, they do not include loading and unloading times.

2.3 The high-risk waste instance: details

The instance corresponding the collection of high-risk waste contains 262 customers, distributed all within Belgium (see Figure 2). The planning period is 2 weeks (10 days), and again customers have certain frequencies of visit within that period. In Table 2 we give an overview of the different frequencies for the high-risk waste instance. The capacities of the three trucks available are

Frequency (nr of days)	1	2	4
nr of customers	62	186	14

Table 2: Frequencies high-risk waste

9, 12, and 12 tons respectively. The collected waste must be delivered at two disposal facilities of an external company. Note that time windows apply to these facilities. Loading and unloading times are constant, 10 resp. 30 minutes. Restrictions on driving hours and depot locations are the same as for low-risk waste.
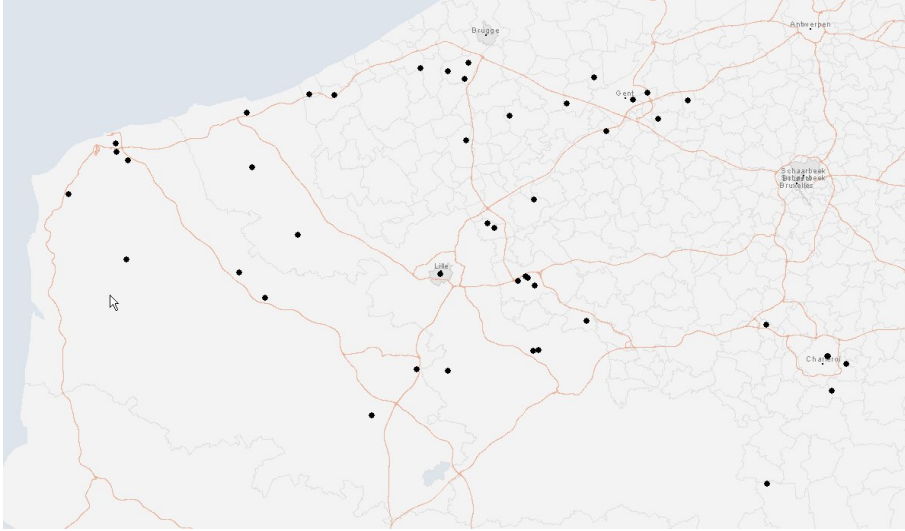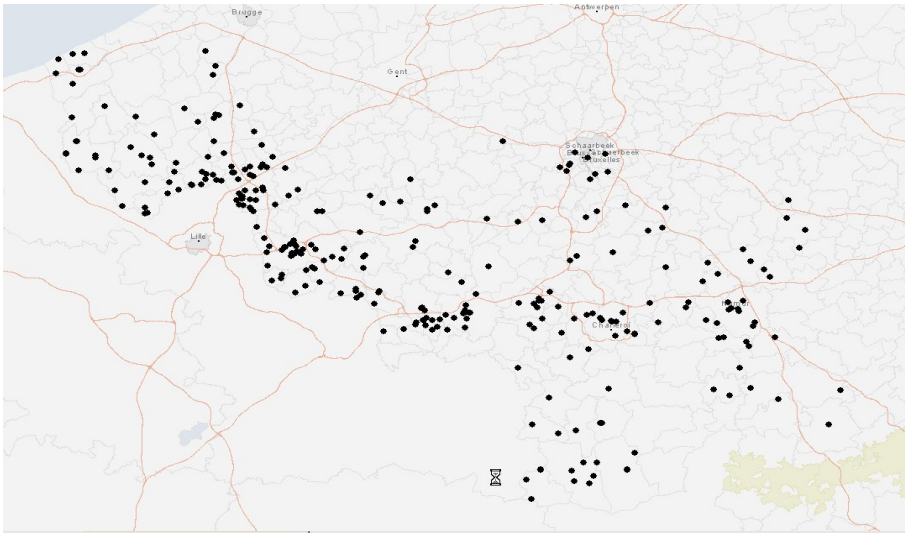
Figure 1: Clients of category 3: instance 1

Figure 2: Clients of category 1: instance 2

3 Model

An important aspect of our model are the "visit-frequencies". Based on historical data, it can be accurately predicted how much waste a customer produces per day, and how much waste can be stored at a customer's site. On the basis of these data and in consultation with the customer, a *visit-frequency* is computed for each customer: the number of times a collection is done during the planning period. We also compute a demand (q_i) depending upon this frequency. As an example, consider some customer i that generates 10 tons of waste per day, and has a storage capacity of 30 tons, while the planning period is $T = 5$. We compute the corresponding visit-frequency f_i for this client i as follows: $f_i = \lceil \frac{10T}{30} \rceil = 2$. More in particular, the visits must be spread well-balanced over the planning period. Hence, this customer can be visited on days 1 and 3, but also on days 1 and 4, 2 and 4, 2 and 5, and 3 and 5; but not on any two consecutive days. The demand of this customer equals $q_i = \frac{10T}{f_i} = 25$. Notice that the actual amount of waste collected at customer i may differ from the predicted amount q_i.

For every frequency, we define all possible combinations of days in the planning period, and we call them scenarios. Then, for every customer we need a constraint that selects exactly one scenario from the set of scenarios that correspond to the frequency associated with the customer. How exactly the scenarios are assigned to customers will become clear further in this section; we first give some notation.

3.1 Notation

We define the network $G = (V, A)$. The customers to be visited are represented by nodes 1 to N, the depot is represented by node 0 and the disposal facilities by nodes $N+1$ to $N+M$. Thus vertex set $V = \{0, 1, 2, \ldots, N, N+1, \ldots, N+M\}$ and A is the arc set with for each arc (i, j) a travel distance d_{ij} and a travel time c_{ij} $(i, j \in V)$. The planning period has a length of T days and a customer is visited within this period according to a certain scenario $c \in C$. Essentially, a scenario is a set of days within $\mathcal{T} = \{1, \ldots, T\}$; choosing a scenario for a customer means that the customer is visited during these days. We let f^c equal the number of visiting days in scenario c. We are given a number of vehicles K, each with a certain capacity Q_k. For each customer $i \in \mathcal{N} = \{1, \ldots, N\}$, a frequency f_i, a quantity q_i and a loading time l_i are given. Quantity q_i is the quantity to be collected at each visit; this number is based on the average amount of waste that is collected at each visit (see our discussion earlier). We can use this average amount to approximate reality because visits are spread evenly over the planning period. Further, a driver may not drive longer than D_d hours a day, and no longer than D_T hours in the total planning period of T days. Disposal facilities can only be visited within their time windows $[r_i, s_i]$, $i \in \mathrm{M} = \{N+1, \ldots, N+M\}$. Finally, a_{ct} is equal to 1 if day $t \in \mathcal{T}$ is visited within scenario $c \in C$ and 0 otherwise. We then define the following 5 sets of variables: x_{ijkt} is a binary variable which is equal to 1 if customer $i \in V$ is visited after customer $j \in V$ by vehicle $k \in \mathcal{K}$ at day $t \in \mathcal{T}$, 0 otherwise; y_{ic} is a binary variable which is 1 if client $i \in V$ is visited according to scenario $c \in C$ and 0 otherwise; let L_{ikt} and T_{ikt} be the total load and the time, respectively, of vehicle $k \in \mathcal{K}$ at day $t \in \mathcal{T}$ after having visited customer $i \in V$; and finally,

the load of a vehicle $k \in \mathcal{K} = \{1, 2, \ldots, K\}$ at the beginning of day $t \in \mathcal{T}$ is denoted by S_{1kt} and the load at the end of the day by S_{2kt}.

3.2 The model

$$\text{(PVRP) Minimize} \sum_{i \in V} \sum_{j \in V} \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} d_{ij} x_{ijkt} \tag{1}$$

subject to

$$\sum_{c \in C: f^c = f_i} y_{ic} = 1 \qquad \forall i \in \mathcal{N}, \tag{2}$$

$$\sum_{j \in V} \sum_{k \in \mathcal{K}} x_{ijkt} - \sum_{c \in C} a_{ct} y_{ic} = 0 \qquad \forall i \in \mathcal{N}; \forall t \in \mathcal{T}, \tag{3}$$

$$\sum_{i \in V} x_{ihkt} - \sum_{j \in V} x_{hjkt} = 0 \qquad \forall h \in V; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{4}$$

$$\sum_{j \in \mathcal{N}} x_{0jkt} \le 1 \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{5}$$

$$\sum_{i \in V} \sum_{j \in V} \sum_{t \in \mathcal{T}} d_{ij} x_{ijkt} \le D_T \qquad \forall k \in \mathcal{K}, \tag{6}$$

$$\sum_{i \in V} \sum_{j \in V} d_{ij} x_{ijkt} \le D_d \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{7}$$

$$\sum_{i \in S} \sum_{j \in S} x_{ijkt} \le |S| - 1 \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}; S \subseteq \mathcal{N}; |S| \ge 2, \tag{8}$$

$$(L_{ikt} + q_j - L_{jkt}) \le (1 - x_{ijkt})M \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}; \forall i, j \in \mathcal{N}, \tag{9}$$

$$L_{ikt} \le Q_k \qquad \forall i \in \mathcal{N}; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{10}$$

$$L_{ikt} = 0 \qquad \forall i \in \mathcal{M}; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{11}$$

$$S_{2kt} = S_{1k(t+1)} \qquad \forall k \in \mathcal{K}; \forall t \in \{1, \ldots, T-1\}, \tag{12}$$

$$L_{ikt} - S_{2kt} \le (1 - x_{i0kt})M \qquad \forall i \in V; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{13}$$

$$S_{1kt} - L_{jkt} \le (1 - x_{0jkt})M \qquad \forall j \in V; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{14}$$

$$S_{2kt} = 0 \qquad \forall t \in \{fridays\}; \forall k \in \mathcal{K}, \tag{15}$$

$$S_{2kt} \le Mz \qquad \forall k \in \mathcal{K}; \forall t \in \{1, \ldots, T-1\}, \tag{16}$$

$$1 - \sum_{j \in V} x_{0jk(t+1)} \le M(1-z) \qquad \forall k \in \mathcal{K}; \forall t \in \{1, \ldots, T-1\}, \tag{17}$$

$$T_{ikt} + l_i + c_{ij} - T_{jkt} \le (1 - x_{ijkt})M \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}; \forall i, j \in V, \tag{18}$$

$$r_i x_{ijt} \le T_{ikt} \le s_i x_{ijkt} \qquad \forall k \in \mathcal{K}; \forall t \in \mathcal{T}; \forall i, j \in \mathcal{N}, \tag{19}$$

$$x_{ijkt} \in \{0, 1\} \qquad \forall i, j \in V; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}, \tag{20}$$

$$y_{ic}, z \in \{0, 1\} \qquad \forall i \in \mathcal{N}; \forall c \in C, \tag{21}$$

$$L_{ikt}, T_{ikt}, S_{1kt}, S_{2kt} \ge 0 \qquad \forall i \in V; \forall k \in \mathcal{K}; \forall t \in \mathcal{T}. \tag{22}$$

This model finds a scenario for every customer and a set of routes for each day of the planning period such that total travel distance is minimized. The first constraints (2) make sure that exactly one scenario is selected for every customer

and in such a way that within this scenario the customer is visited according to its frequency. A customer then will be visited on the days of the selected scenario; this is ensured by constraints (3). Constraints (4) make sure that when a vehicle arrives at a customer, it also leaves from that customer. Constraints (5) impose that each vehicle can be used at most once every day. Constraints (6) and (7) keep the number of driving hours for every vehicle within the restrictions for the whole planning period, and within the daily restrictions, respectively. Constraints (8) are subtour elimination constraints. Correct counting of vehicle loads is ensured by constraints (9), with M a big number, and constraints (10) keep the amount of vehicle load within the capacity. Constraints (11) impose that vehicles are empty when they have visited a disposal facility. A vehicle does not need to dispose of all of its load at the end of the day. It follows that the load of a vehicle at the end of a day needs to be equal to the load of that vehicle at the start of the following day, this is insured by constraints (12) to (14). At certain moments in the week though vehicles do need to unload at the end of the day, e.g at the end of the week (15) and when a vehicle is not used on the next day, see constraints (16) and (17). Equations (18) and (19) make sure that time windows for the disposal facilities are not violated. Finally, constraints (20) through (22) impose binary conditions and nonnegativity conditions on the variable set. Solving this model in order to obtain an optimal solution for a small data set of less than 10 customers already takes a very long time. That is why in the following section we develop a heuristic solution approach.

4 Solution approach

We study methods that consist of two phases: in one phase customers are assigned to days, and in another phase a VRP-instance is solved. In subsection 4.1 we investigate methods that first assign customers to days, and next solve a VRP-instance. We consider two ways of assigning customers to days, namely striving for an "even spread" of the number of visits on a day (subsection 4.1.1), or use a geographically based clustering approach (subsection 4.1.2). Finally, we describe how we use ILOG Dispatcher to solve a VRP-instance (subsection 4.1.3). In subsection 4.2 we describe a method that first solves a VRP-instance, and then assigns customers to days.

4.1 First assign customers to days, then route

4.1.1 Assigning customers to days: algorithm ES

Here, we focus on the spread of visits over the planning horizon. We solve the following problem (problem ES). We are given a set V of customers, and each customer i has an associated frequency f_i. Further we have a set C of scenario's, where a scenario consists of a set of f^c days, meaning that in that scenario c visits are performed on f^c different days. These scenario's need to be assigned to customers such that the frequency of each customer equals the frequency of the scenario assigned to the customer and the maximum number of customers visited on each day of the planning horizon is minimized. We refer to this problem as problem ES and we claim the following.

Fact 1 *Problem ES is NP-hard.*

Proof. We prove NP-hardness of problem ES by a reduction from Exact Cover by 3-sets (X3C). In X3C a finite set X containing $3n$ elements and a collection C of 3-element subsets of X are given. The question is whether there exists a subset C' of C such that every element of X occurs in exactly one triple of C'. This problem is proven to be NP-complete by Garey and Johnson [14]. For every instance of X3C we can create an instance of problem ES as follows. There are n customers, each with frequency f_i equal to 3. The planning period lasts $3n$ days and each triple from C corresponds to a scenario consisting of three days corresponding to the elements of the triple. Now, if an assignment of scenario's to customers can be found such that exactly 1 customer is visited on each day (even spread), then also a solution for X3C exists, and vice versa.□

Thus, this proves that in general problem ES is a hard problem to solve. However, for our purposes, due to the relatively small size of the problem (we have $T = 5$ (10), $|C| = 15$ (20), $N = 48$ (262)) we can solve problem ES in reasonable time using the following integer program (IP). We define parameter a_{ct}, which is equal to 1 if day t is visited in scenario c, and 0 otherwise. Variable y_{ic} is equal to 1 if customer i is assigned scenario c and 0 otherwise. z_t is an integer variable representing the number of customers visited on day t. Notice that this variable is not strictly needed; we add it to make the model more clear.

$$\text{Minimize } w \tag{23}$$

subject to

$$w \geq z_t \qquad \forall t \in \mathcal{T}, \tag{24}$$

$$\sum_{c \in C: f^c = f_i} y_{ic} = 1 \qquad \forall i \in \mathcal{N}, \tag{25}$$

$$\sum_{i \in V} \sum_{c \in C: f^c = f_i} a_{ct} y_{ic} = z_t \qquad \forall t \in \mathcal{T}, \tag{26}$$

$$y_{ic} \in \{0, 1\} \qquad \forall i \in \mathcal{N}; \forall c \in C, \tag{27}$$

$$z_t, w \in \mathbb{Z} \qquad \forall t \in \mathcal{T}. \tag{28}$$

The goal of this IP is to assign a scenario to every customer such that visits are spread evenly over the planning period. We enforce this by minimizing the maximal number of customers visited on a day during the planning period (24). In that way we will visit more or less the same amount of customers on each day. Constraint (25) makes sure that every customer is visited according to exactly one scenario that matches its frequency. We can solve this IP optimally using ILOG Cplex 10.2.

Notice that the location of the customers is completely ignored in this approach. Clearly, this implies that there is a risk that customers positioned close to each other are scheduled on different days. In the next section, we assign scenario's to customers such that this potential disadvantage is taken into account.

4.1.2 Assigning customers to days: algorithm CL

In a second algorithm, algorithm CL (CLuster), we partition the customers using the algorithm "k-medoids clustering" before assigning scenario's to customers. This clustering algorithm partitions the locations into k clusters and attempts to minimize the squared error, i.e. the squared distances between points labeled

to be in a cluster and a point designated as the center of that cluster. k-medoids chooses data points as centers. We use this particular clustering method because it is suited for cases where the distance matrix is given (instead of the locations), as is the case here. We use the k-medoid method as it is defined in the C Clustering Library by de Hoon et al. [10]. The algorithm starts with an arbitrary selection of k data points that will act as centers (medoids) and then tries to improve by swapping the medoids with other points.

We apply this method for several values of k, and then we solve (29)-(34) (which is a modification of (23)-(28)), which takes clusters into account when assigning customers to days. Define parameter o_{il}, which is equal to 1 if customer $i \in \mathcal{N}$ is in cluster $l \in 1, \ldots, k$, with k the number of clusters, and 0 otherwise. z'_{tl} is a variable which is equal to 1 if and only if a customer of cluster l is visited on day t. Now, our goal is to minimize the total sum of the number of clusters visited each day. A cluster is visited if a customer of that cluster is visited. Customers that belong to the same cluster will thus be as much as possible assigned to the same day.

$$\text{Minimize } w' \tag{29}$$

subject to

$$\sum_{t \in \mathcal{T}} \sum_{l \in \{1, \ldots, k\}} z'_{tl} \leq w' \tag{30}$$

$$\sum_{c \in C: f^c = f_i} y_{ic} = 1 \qquad \forall i \in \mathcal{N}, \tag{31}$$

$$\sum_{c \in C: f^c = f_i} \sum_{i \in V} a_{ct} y_{ic} o_{il} \leq M z'_{tl} \qquad \forall t \in \mathcal{T}; \forall l \in \{1, \ldots, k\}, \tag{32}$$

$$y_{ic}, z'_{tl} \in \{0,1\} \qquad \forall i \in \mathcal{N}, \forall c \in C, \forall t \in \mathcal{T}, \forall l \in \{1, \ldots, k\} \tag{33}$$

$$w' \in \mathbb{Z}, \tag{34}$$

with M a large number.
Notice that a potential risk associated to the outcome of model (29)-(34) is that the driving time needed to visit all the customers assigned to a certain day can exceed D_d. To prevent this, a constraint on the number of customers visited on each day is added; we come back to this issue in section 5.2.

4.1.3 Routing using ILOG Dispatcher

Having assigned all customers to scenario's, either using ES or CL, we know which customers need to be visited every day. Thus, for each day in the planning horizon we now need to solve a, more or less standard, VRP. Some specific constraints do need to be taken into account, such as vehicles that have different capacities, vehicles that can dispose of their load during the day and then continue the route (they can collect more than their capacity on one day), some customers can only be reached by small vehicles, vehicles do not need to unload at the end of the day, and the number of driving hours per vehicle per day is limited. A solution algorithm for this VRP is implemented using ILOG Dispatcher 4.4, which is a C++ library based on ILOG Solver and that offers features especially adapted to solving problems in vehicle routing. We

implement all the standard VRP constraints, as well as the problem-specific constraints mentioned above.

In the resulting routing algorithm, an initial solution to the VRP is constructed using a standard savings heuristic, which makes a trade off between more vehicles with shorter routes and fewer vehicles with longer routes. Then, ILOG's dispatcher performs neighborhood search to improve this solution. Both intra-route and inter-route neighborhoods are considered. As intra-route neighborhoods, 2-Opt and Or-Opt are implemented.ILOG's dispatcher also interchanges between routes: the "relocate" neighborhood (inserting a customer in another route); the "exchange" neighborhood (swapping two customers from different routes); and the "cross" neighborhood (exchanging the ends of two routes). We refer to ILOG Dispatcher user's manual [16] for a more elaborate description. We have restricted ourselves here to these methods to improve the routes.

4.2 First route, then assign customers to days: algorithm MR

In this approach, called algorithm MR (Mega Route), we first construct large routes visiting all the customers and the disposal facilities. To accomplish this, we use the same VRP heuristic as described in the previous section. Then, on the basis of these routes, customers are assigned to days using model (29)-(34). Customers belonging to the same route are then visited as much as possible on the same day. Then, we resolve a VRP for each day in the planning horizon in order to obtain better routes. Algorithm MR is similar to algorithm CL, where we first cluster the customers, since geographically close customers tend to end up being visited on a same day. Routing can yield a very different grouping of the customers though.

5 Computational results

The different algorithms were implemented in Microsoft Visual C++ 2005, in combination with ILOG CPLEX 10.2 and ILOG Dispatcher 4.4. All algorithms are run on a personal computer with a 2.80 GHz Intel Pentium IV processor and 504 MB of RAM. Notice that we have not been concerned with running times of our approach. This is mostly because the application did not enforce strong limitations on the amount of computing time used. In subsections 5.1 and 5.2 computational results for the low-risk waste instance and for the high-risk waste instance, respectively, are given. In section 5.3 we discuss the performance of the different methods for the two instances.

5.1 The low-risk waste instance: results

In Table 3 some computational results are shown. Applying algorithm ES, we find a set of routes requiring 3 vehicles. On several days only 2 of them are being used. Notice that visits are well spread over the week, as each day a comparable distance is traveled. Between vehicles though, there is a rather uneven workload. Total traveling time is 5859 minutes in 6 days.

As can be seen in Table 3, a solution by algorithm CL is obtained using all three vehicles. Comparing this solution with the previous one we can see that the variation in route length over the days is higher and that total traveling time decreased to only 5551.

The last columns in Table 3 show the results of algorithm MR. Here also three vehicles are used and total traveling distance is rather high (5934). Thus, routing the customers before assigning them to the different days in the planning horizon does not yield a good solution for this instance.

Concluding, this suggests that using the geographic structure yields a better routing plan. Running times vary from about 1 hour for ES, to a few hours for CL and up to 48 hours for MR. We experience a large variation in time needed by the VRP solver. Finally notice that the current routes used in practice have a total traveling time of 6565; using algorithm CL we could improve them with 15.5%.

5.2 The high-risk waste instance: results

Let us now apply the same algorithms to the instance for collection of high-risk waste. The results are summarized in Table 4. Using algorithm ES, every day the same amount of customers is visited, unloading only occurs once in two days and on Fridays and only two vehicles are required. Total travel time is then 10272 for two weeks.

Applying algorithm CL, two large clusters are obtained. Based on these clusters we assign all the customers to the days of the planning horizon. This

Algorithm	ES				CL				MR			
Vehicle	12T	22T	26T	TOT	12T	22T	26T	TOT	12T	22T	26T	TOT
Day												
1	452	391	240	1083	495	413	476	1384	339	235	567	1141
2	509		551	1060	409		279	688	339	252	220	811
3	511	378	240	1129	416	439	237	1092	439	288	588	1315
4	427	280	482	1189	512	374	214	1100	354	364	276	994
5	561		482	1140	460	163	406	1029	584	510	321	1415
6			258	258			258	258			258	258
				5859				5551				5934

Table 3: Results low-risk waste (traveling time in minutes)

Algorithm	ES			CL			
Vehicle	9T	12T	TOT	9T	9T	12T	TOT
Day							
1	558	350	908	311		385	696
2	662	495	1157	387		495	882
3	596	342	938	0			0
4	587	570	1157	579	583	398	1560
5	450	549	999	551			551
6	308	599	907	445		423	868
7	641	434	1075	494	619	115	1228
8	599	334	933	0			0
9	592	607	1199	585	577	356	1518
10	450	549	999	311		116	427
			10727				7730

results in a very uneven spread of the customers; on some days 100 customers are visited, and as a result 3 vehicles are required, and on other days no customers are visited. Total traveling time is then 7730, which is clearly much better than ES; however, this leads to the usage of 3 vehicles. In order to see whether the instance could be solved using two vehicles only, we add a constraint to (29)-(34) bounding the number of customers assigned to one day. Solving this model, we obtain routes with a total travel time of 8013 and using only 2 vehicles on most days.

Finally, we also apply algorithm MR to this instance. We create 2 giant routes, visiting all the customers and the disposal facilities. Again, when dividing the customers over the days, we obtain a highly uneven spread of customers. Visits are only performed on 4 of the 10 days, with up to 171 customers on one day. In fact, as mentioned as a potential risk in section 4.1.2, this is not feasible given the available driving time; if we relax these constraints total distance traveled is 6591. Again, we need to limit the number of customers visited per day, and we set this limit equal to 75, which is more or less the amount of customers that can be served on one day by 2 vehicles. This yields a total traveling time of 7655 and utilizes 2 vehicles. More details are shown in Table 4.

Concluding, for this instance it is important to take into account geography. Running times are limited to seconds for all three algorithms. Compared to the original routes which take 8434, algorithm ES performs badly; MR improves the currently used routes with 9% and uses one vehicle less.

5.3 Discussion

Depending on the instance, the methods used above perform differently. Clustering the customers yields a better result for both instances compared to algorithm ES, but the effect is larger for the high-risk waste instance. Algorithm MR performs comparably as CL for this instance but a lot worse than ES and CL for the low-risk waste instance. Explanation of this phenomenon is that there are customers in the low-risk instance with a very high frequency (see Table 1) and these clients are rather spread over the country (see Figure 1). Ignoring the geographic locations has then only a small impact on the solutions, as on each day the vehicles have to travel in different directions anyway. There is a

Algorithm	CL-limited clients per day				MR-limited clients per day		
Vehicle Day	9T	9T	12T	TOT	9T	12T	TOT
1	540		279	819	527		527
2	510		655	1165	452	673	1125
3	506		225	731	504		504
4	0			0	538	675	1213
5	544		675	1219	0		0
6	541		289	830	577		577
7	551	203	519	1273	455	673	1128
8	542		228	770	546		546
9	0			0	514	686	1200
10	591		615	1206	257	578	835
				8013			7655

Table 4: Results high-risk waste (traveling time in minutes)

large difference in the frequencies of the high-risk waste instance compared to the low-risk waste instance (see Tables 1 and 2). Many customers require the same frequency, and no customer has a very high frequency (maximally 4 visits in 10 days). Together with the fact that the instance is larger, this explains why clustering has a vast effect on the solutions.

As mentioned before, we obtain routes that might yield a rather uneven workload for the different vehicles. Company A does not consider this as being a problem. We propose routes using three vehicles for the small instance. Not every day, though, all vehicles are needed, thus it might be interesting to consider renting a vehicle and driver at an external company for the days necessary. Changing the volume of the vehicles in order to diminish the number of vehicles does not seem really useful. For the low-risk waste instance a vehicle with capacity of 12 tons must be available for customers not reachable with a larger vehicle. For the high-risk waste instance, the volumes to be picked up are rather small and unloading is only necessary after two days. It is not the volume but the time restrictions that have the highest impact on travel times. That is also the reason why running times differ a lot between the low-risk waste instance and high-risk waste instance. In the low-risk waste instance the volumes to be picked up are higher such that vehicles unload several times during the day. This complicates the VRP model and increases running times.

6 Conclusion

We consider a problem occurring in practice, and we modeled it as a PVRP. Using different approaches (including ILOG's dispatcher), we were able to improve the current routes of Company A, using 1 vehicle less for one of the instances. This not only means a reduction in cost due to the gain in travel time, but also a reduction in wage costs and material costs. We use rather simple algorithms to assign the customers to the days and to solve the VRP's. The use of clustering customers depends highly on the dispersion of the customers and on their frequencies. We deal with this instance of the PVRP by considering the problems of assigning customers to days and routing the customers independently.

For those interested to further study the problem, the instances can be found on the following website: http://www.econ.kuleuven.ac.be/public/N05012/.

References

[1] J. Alegre, M. Laguna, and J. Pacheco. Optimizing the periodic pick-up of raw materials for a manufacturer of auto parts. *European Journal of Operational Research*, 179:736–746, 2007.

[2] F. Alonso, M. J. Alvarez, and J. E. Beasley. A tabu search algorithm for the periodic vehicle routing problem with vehicle trips and accessibility restrictions. *Journal of the Operational Research Society*, 59:963–976, 2008.

[3] E. Angelelli and M. G. Speranza. The periodic vehicle routing problem with intermediate facilities. *European Journal of Operational Research*, 137:233–247, 2002.

[4] A. Arnout. Het periodieke routeplanningsprobleem: een case studie. Master Thesis, 2007.

[5] S. Baptiste, R. C. Oliviera, and E. Zúquete. A period vehicle routing case study. *European Journal of Operational Research*, 139:220–229, 2002.

[6] E. J. Beltrami and L. D. Bodin. Networks and vehicle routing for municipal waste collection. *Networks*, 4:65–94, 1974.

[7] F. Blakely, B. Bozkaya, B. Cao, W. Hall, and J. Knolmajer. Optimizing periodic maintenance operations for schindler elevator corporation. *Interfaces*, 33:67–79, 2003.

[8] N. Christofides and J. E. Beasley. The period routing problem. *Networks*, 14:237–256, 1984.

[9] G. D. H. Claassen and Th. H. B. Hendriks. An application of special ordered sets to a periodic milk collection problem. *European Journal of Operational Research*, 180:754–769, 2007.

[10] M. de Hoon, S. Imoto, and S. Miyano. *The C Clustering Library*, 2008.

[11] R. W. Eglese and H. Murdock. Routing road sweepers in a rural area. *The Journal of the Operational Research Society*, 42:169–194, 1991.

[12] P. Francis, K. Smilowitz, and M. Tzur. The period vehicle routing problem with service choice. *Transportation Science*, 40:439–454, 2006.

[13] P. Francis, K. Smilowitz, and M. Tzur. The period vehicle routing problem and its extensions. In B. Golden, S. Raghavan, and E. Wasil, editors, *The Vehicle Routing Problem*, pages 239–261. Springer, 2008.

[14] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. Freeman, San Francisco, 1979.

[15] V. Hemmelmayr, K. F. Doerner, R. F. Hartl, and M. W. P. Savelsbergh. Delivery strategies for blood products supplies. *OR Spectrum*, available online, 2008.

[16] ILOG. *ILOG Dispatcher 4.1, User's manual*, 2005.

[17] B. Kim, S. Kim, and S. Sahoo. Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, 33:3624–3642, 2006.

[18] M. Mourgaya and F. Vanderbeck. Problème de tournées de véhicules multipériodiques: Classification et heuristique pour la planification tactique. *RAIRO Operations Research*, 40:169–194, 2006.

[19] C. C. R. Tan and J. E. Beasley. A heuristic algorithm for the period vehicle routing problem. *Omega*, 12:497–504, 1984.