The Dynamic Multi-Period Vehicle Routing Problem

Wen, Min; Cordeau, Jean-Francois; Laporte, Gilbert; Larsen, Jesper

The Dynamic Multi-Period Vehicle Routing Problem

Min Wen
Jean-François Cordeau
Gilbert Laporte
Jesper Larsen
June 2009

The Dynamic Multi-Period Vehicle Routing Problem

Min Wen[1], Jean-François Cordeau[2], Gilbert Laporte[3], and Jesper Larsen[1]

[1] Department of Management Engineering, Technical University of Denmark,
Produktionstorvet DTU-Building 426, DK-2800 Kongens Lyngby, Denmark, {miwe@man.dtu.dk, jesla@man.dtu.dk}
[2] Canada Research Chair in Logistics and Transportation and CIRRELT, HEC Montréal,
3000 chemin de la Côte-Sainte-Catherine, Montréal, Canada H3T 2A7, cordeau@crt.umontreal.ca
[3] Canada Research Chair in Distribution Management and CIRRELT, HEC Montréal,
3000 chemin de la Côte-Sainte-Catherine, Montréal, Canada H3T 2A7, gilbert@crt.umontreal.ca

Abstract. This paper considers the *Dynamic Multi-Period Vehicle Routing Problem* which deals with the distribution of orders from a depot to a set of customers over a multi-period time horizon. Customer orders and their feasible service periods are dynamically revealed over time. The objectives are to minimize total travel costs and customer waiting, and to balance the daily workload over the planning horizon. This problem originates from a large distributor operating in Sweden. It is modeled as a mixed integer linear program, and solved by means of a three-phase heuristic that works over a rolling planning horizon. The multi-objective aspect of the problem is handled through a scalar technique approach. Computational results show that our solutions improve upon those of the Swedish distributor.
Keywords: Dynamic, Multi-Period, Multi-Objective, Vehicle Routing, Variable Neighborhood Search.

1 Introduction

The purpose of this paper is to model and solve the *Dynamic Multi-Period Vehicle Routing Problem* (DMPVRP). Our study is motivated by the case of Lantmännen, a large distributor operating in Sweden, but our contribution is of general applicability. In the DMPVRP, customers place orders dynamically over a planning horizon consisting of several periods (or days). Each request specifies a demand quantity, a delivery location and a set of consecutive periods during which delivery can take place. The distributor must plan its delivery routes over several days so as to minimize the routing cost and customer waiting, and to balance the daily workload over the planning horizon.

Lantmännen is one of the largest groups within the food, energy and agricultural industries in the Nordic countries. The company is owned by 42,000 Swedish farmers, hires 13,000 employees, and generates sales of SEK 36 billion per year. One of its activities is the distribution of fodder to the farmers at their request from one of several terminals which usually operate independently of

each other, except in periods of exceptional activity. Here we consider a single terminal, Västerås, located in southern Sweden. It is the busiest terminal in terms of number of vehicles and orders. Figure 1 shows the locations of the customers and of the terminal. Customers place orders over time and the distribution schedule of a given day is constructed for several vehicles at the beginning of that day. It serves some of the unfulfilled orders and typically leaves some for the following days. A fair amount of foresight is required so as not to create infeasible situations in the future while creating efficient routes. Unfulfilled orders after the schedule has been built and new orders accumulated during the day are considered for scheduling the following day. Because the drivers do not interact with the customers when delivering, no time windows need to be specified.
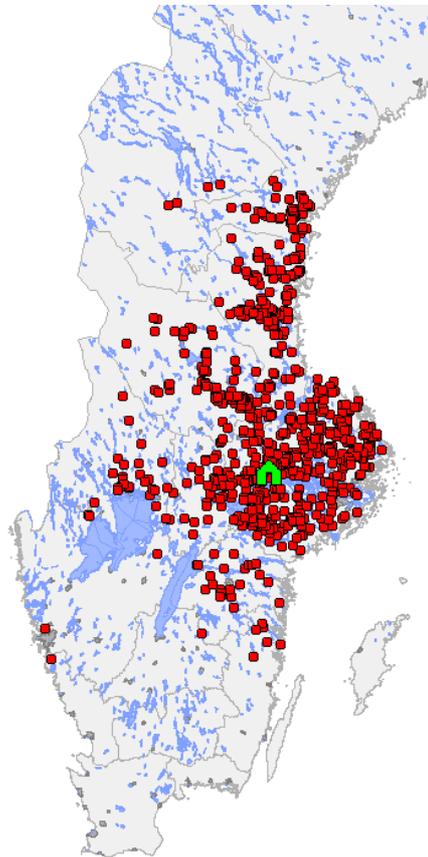

Figure 1. Locations of customers and depot (represented by a house) in the Lantmännen case study

The literature on the DMPVRP is rather scarce. To our knowledge, the closest work is that of Angelelli et al. (2007, 2009) who considered a special case of the DMPVRP with a single vehicle and a planning horizon of two days.

The DMPVRP is closely related to the *Periodic Vehicle Routing Problem* (PVRP) in which all information is available at the beginning of the planning horizon. In the PVRP, customers specify a service frequency and sets of allowable combinations of visit days. For example, if a customer specifies a frequency of 2 and the combinations {1, 3} and {2, 4}, then the customer wishes to be visited twice, on days 1 and 3, or on days 2 and 4. In the DMPVRP, visit frequencies are equal to 1 and visit combinations are made up of consecutive days. The PVRP is usually solved heuristically. The best known algorithms for this problem are those of Cordeau, Gendreau and Laporte (1997) and of Hemmelmayr, Doerner and Hartl (2009). Francis, Smilowitz and Tzur (2008) have solved a variant of the PVRP in which service frequency is a decision variable. Mourgaya and Vanderbeck (2007) have solved another variant that includes routing cost minimization and daily workload balance.

Other routing problems with a dynamic component are often encountered in the context of dynamic pickup and delivery problems (Psaraftis, 1988; Mitrović-Minić, Krishnamurti and Laporte, 2004; Branke et al. 2005; Hvattum, Løkketangen and Laporte, 2006, 2007; Pureza and Laporte, 2008), but these papers do not consider a multi-period horizon. For recent literature reviews, see Larsen, Madsen and Solomon (2008), and Berbeglia, Cordeau and Laporte (2009).

Another strand of literature relevant to our problem is about the *Multi-Objective Vehicle Routing Problem* encountered in school bus routing (Pacheco and Marti, 2006; Alabas-Uslu, 2008), waste collection (Lacomme, Prins and Sevaux, 2006), and hazardous products transportation (Dell'Olmo, Gentili and Scozzari, 2005; Zografos and Androutsopoulos, 2008; Tan, Chew and Lee, 2006). The two main solution strategies for multi-objective problems are the scalar technique, which consists in minimizing a weighted linear combination of the objectives, and the Pareto method which identifies a set of non-dominated solutions. We refer to Jozefowiez, Semet and Talbi (2008) for a recent survey of these methods in the context of vehicle routing.

In this paper we formulate the DMPVRP as a mixed integer linear program using the scalar technique. We then develop a three-phase heuristic for its solution, and we show that our results outperform those of Lantmännen. The remainder of the paper is organized as follows. The math-

ematical model is described in Section 2. The heuristic is described in Section 3, followed by computational results in Section 4 and by conclusions in Section 5.

2 Mathematical Problem Description

We start with a more detailed description of the DMPVRP. To capture the problem more precisely, we also formulate it as a mixed integer linear program.

2.1 Problem description and analysis

The DMPVRP is solved over a planning horizon divided into days. Customer orders arrive at any time and must be fulfilled within a set of consecutive service days which can start as early as the day after the order is placed. A set of homogeneous vehicles are available at the depot. These vehicles depart from the depot at 00.00 and return to the depot at the latest at 23.59 on the same day. The objectives are to minimize the total routing cost (proportional to travel time) and customer waiting, and to balance the daily workload over the planning horizon. Each customer must be visited exactly once by one vehicle within its feasible service period, each vehicle must depart from and return to the depot in the same day, and the load of each vehicle cannot exceed its capacity.

This problem is dynamic in the sense that orders are revealed incrementally over time. The daily planning must determine which orders should be fulfilled on that day and in which sequence the vehicles should visit the customers. These decisions are made without the knowledge of future orders. However, even if the problem is dynamic, the routing problem at the beginning of each particular day over the planning horizon can be viewed as a static problem since the routes for that day are planned based on the orders known so far and the routes are fixed before their execution.

Figure 2 illustrates the planning process for a small DMPVRP example consisting of two days. For simplicity, we assume that the demand of each order is one, the capacity of the vehicle is three, and two vehicles are available. Before the first day, as shown in (a), six orders are already logged in the system. Three of these, denoted by triangles, can be fulfilled either on the first day or the second day, while the other three can only be served on the first day. At the beginning of the planning horizon, the planner has to construct the routing plan for the first day, as shown in (b). Before the

second day is planned, three new orders have arrived, as shown in (c). The routes for the second day are shown in (d).
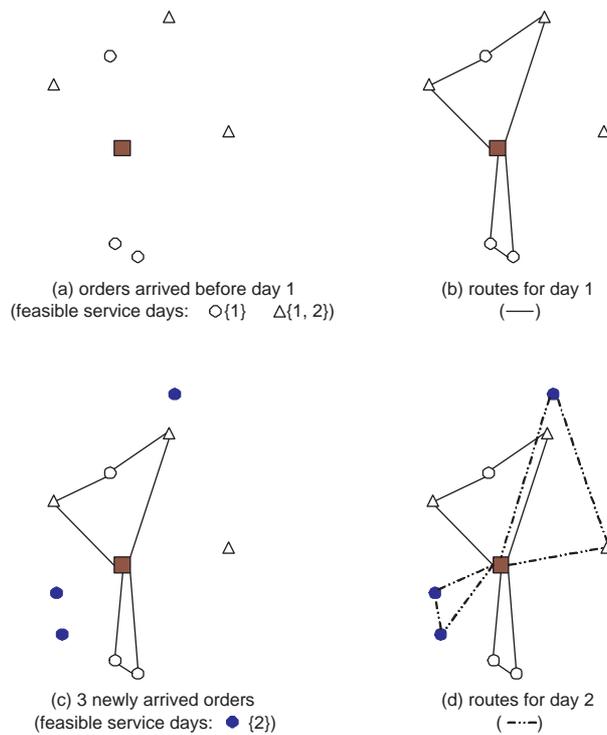
Figure 2. A small instance of the DMPVRP

This example illustrates that the challenging part of the problem is to decide on the first day whether to serve the triangle orders, or whether to postpone them until the second day without knowing which orders will arrive during the first day. On the one hand, if the new orders are destined for locations close to those of the triangle orders, it may be wise to postpone them so as to minimize the total travel time. On the other hand, if too many orders are postponed, customer waiting is prolonged and the feasibility of the next day's solution may be jeopardized due to the limited available vehicle capacity.

2.2 Mathematical formulation

The planning for each particular day can be regarded as a special case of the PVRP with unit visit frequency and consecutive allowable delivery periods. Without loss of generality, we present the formulation for the planning problem on day t ($t \in T$, where $T = \{1, 2, \ldots, r\}$ denotes the planning horizon). Denote the updated planning horizon on day t by $T' = \{t, t+1, \ldots, r\}$, the set of known but unvisited orders by $N = \{1, 2, \ldots, n\}$, and the set of vehicles by $K = \{1, 2, \ldots, m\}$. The depot is located at 0 and the set of all locations is $N_0 = N \cup \{0\}$. The parameter c_{ij} represents the travel time on arc $(i, j) \in A$, where A is the set of arcs between all the locations in N_0. Each order i specifies a demand q_i and a service time d_i. We denote the original consecutive feasible service days for order i by $\{a_i, \ldots, b_i\}$. Note that the first feasible day has to be adjusted to $a'_i = \max\{t, a_i\}$ when planning on day t. Each vehicle has a capacity Q and each route has a duration limit D. The binary variables x^t_{ijkl} denote the decisions made on day t. They are equal to 1 if and only if vehicle k travels from i to j on day l. The constraints are defined as follows:

$$\sum_{l \in \{a'_i, \ldots, b_i\}} \sum_{k \in K} \sum_{j:(i,j) \in A} x^t_{ijkl} = 1 \qquad \forall i \in N \tag{1}$$

$$\sum_{i \in N} \sum_{j:(i,j) \in A} q_i x^t_{ijkl} \leq Q \qquad \forall k \in K, l \in T' \tag{2}$$

$$\sum_{i \in N} \sum_{j:(i,j) \in A} (c_{ij} + d_i) x^t_{ijkl} \leq D \qquad \forall k \in K, l \in T' \tag{3}$$

$$\sum_{j \in N} x^t_{0jkl} = 1 \qquad \forall k \in K, l \in T' \tag{4}$$

$$\sum_{i:(i,h) \in A} x^t_{ihkl} - \sum_{j:(h,j) \in A} x^t_{hjkl} = 0 \qquad \forall h \in N, k \in K, l \in T' \tag{5}$$

$$\sum_{i \in N} x^t_{i0kl} = 1 \qquad \forall k \in K, l \in T' \tag{6}$$

$$x^t_{ijkl} \in \{0, 1\} \qquad \forall (i, j) \in A, k \in K, l \in T'. \tag{7}$$

Constraints (1) ensure that each customer is visited once by exactly one vehicle within its feasible service days. Constraints (2) guarantee the vehicle capacity limit is not exceeded. The duration limit on each route is ensured by constraints (3). Constraints (4)–(6) state that each vehicle must
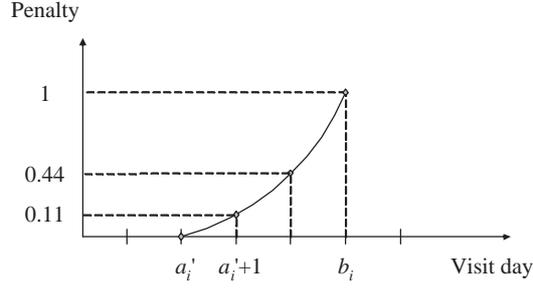
Figure 3. Penalty curve

start and end its route at the depot and that flow is conserved at each customer location. Constraints (7) define the binary variables.

The first objective, minimizing the total travel time of visiting the orders in N, can be formulated as

$$f_1^t = \sum_{l \in T'} \sum_{k \in K} \sum_{(i,j) \in A} c_{ij} x_{ijkl}^t. \tag{8}$$

To minimize the total customer waiting, for each customer having multiple feasible service days, we assign a penalty for not visiting it on the first of its feasible service days. This penalty increases quadratically with customer waiting time, and goes up to 1 if the customer is visited at the end of its feasible service days, as shown in Figure 3. This penalty function favors short waiting times for several customers, as opposed to long waiting times for a few. For example, letting three customers wait for one day is preferable to letting one customer wait for three days. Let N' denote the set of customers having multiple feasible service days, and let the integer variable y_i^t be the day when customer i is visited. The second objective can be formulated by:

$$f_2^t = \sum_{i \in N'} \left(\frac{y_i^t - a_i'}{b_i - a_i'} \right)^2, \tag{9}$$

where

$$y_i^t = \sum_{l \in \{a_i', \dots, b_i\}} \sum_{k \in K} \sum_{j:(i,j) \in A} l x_{ijkl}^t \qquad \forall i \in N'. \tag{10}$$

The third objective, balancing the daily workload over the planning horizon, is more difficult to define since future orders are unknown. In a static problem, this objective can be achieved by minimizing the total deviation of daily workload, where a single day's workload deviation is measured

7

by the absolute value of the difference between that day's workload and the average daily workload over the planning horizon. However, in the dynamic case, it is unwise to allocate the known orders evenly to all future days of the planning horizon. Instead, it seems preferable to focus on the workload of the current day, since we have the complete knowledge of the orders accumulated at the beginning of that day. Moreover, since the actual average daily workload cannot be obtained until the end of the planning horizon, we use an estimate of the average daily workload, denoted by $\tilde{w}^t$, based on historical data. The third objective is hence formulated as:

$$f_3^t = \left| \sum_{k \in K} \sum_{(i,j) \in A} c_{ij} x_{ijkt}^t - \tilde{w}^t \right|. \tag{11}$$

As mentioned, scalar techniques and the Pareto method are the two most used strategies for multi-objective optimization. However, in a dynamic context, the Pareto method is inappropriate because even if it were possible to determine a set of Pareto optimal solutions, it would be necessary to implement one of these before the next day's planning, without guidelines on how to make this selection. We have therefore opted to implement the scalar method with weights, 1, w_2 and w_3, for objective f_1^t, f_2^t and f_3^t, respectively, and we work with the aggregate objective

$$f^t = f_1^t + w_2 f_2^t + w_3 f_3^t. \tag{12}$$

3 A Three-Phase Rolling Horizon Heuristic

We propose a three-phase rolling horizon heuristic to handle the dynamic aspect of the problem. Planning on day t starts with adjusting the set of feasible service days for the yet unvisited customers, including those revealed on day $t-1$. A three-phase heuristic (TPH) is then applied to construct the delivery plan for that day. In order to minimize the total travel time over the planning horizon, instead of only planning the routes for day t, the TPH also optimizes the routes for τ days in the future. Let $T_t = \{t, \ldots, t + \tau\}$ be the planning horizon considered on day t. Phase I selects the customers to be visited within T_t. The selection is necessary because the feasible service days of the customers may not be entirely included in T_t. To this end, we perform a time-space correlation analysis on the known customers. In Phase II, given the customers selected for period T_t, routes are constructed by treating the planning problem as a PVRP with a service frequency equal to 1 over the planning horizon T_t. This routing problem is solved by means of a variable

neighborhood search heuristic. In Phase III, the routes to be executed on day t are postoptimized by means of a tabu search algorithm, and the customers visited on day t are removed from further consideration. This three-phase scheme is summarized in Algorithm 1.

Algorithm 1 : Rolling horizon framework

1: Input: the set N_new_t of customers revealed on each day $t \in T$
2: Output:the routing plan $\mathcal{R} = \{R_1, \ldots, R_{|T|}\}$ for horizon T
3: $N \leftarrow \emptyset$
4: **for** $t = 1$ to $|T|$ **do**
5:     AdjustVisitDays(N)
6:     $N \leftarrow N \cup N\_new_{t-1}$
7:     $N_t \leftarrow$ SelectCustomers(N)   // Phase I
8:     $\{R_t, \ldots, R_{t+\tau}\} \leftarrow$ RouteCustomer(N_t, T_t)   // Phase II
9:     $R_t \leftarrow$ Optimize(R_t)   // Phase III
10:    $N \leftarrow N \setminus \{i : i \in R_t\}$
11:    $\mathcal{R} \leftarrow \mathcal{R} \cup R_t$
12: **end for**

In the TPH, τ is a user-defined parameter. A small value of τ results in a planning problem of small size for the subsequent solution phases and hence reduces the computational burden, whereas a large value of τ helps optimize the total routing cost over the planning horizon. A sensitivity analysis on τ is conducted in Section 4.

3.1   Phase I: Customer selection

The customer selection phase attempts to determine a good set of customers to be visited in the future τ days without relying on routing information. This is achieved by analyzing the time-space correlation between the known customers, as shown in Algorithm 2. More specifically, for each customer i we define a compatibility index q_{il} for each of its allowable service days, where $l \in \{a'_i, \ldots, b_i\}$. A larger value of q_{il} corresponds to a higher visit preference for day l. The parameter is determined as follows. First set q_{il} equal to 0 for all customers and feasible service days. Now consider two customers i and j having common allowable service days. If $c_{ij} \leq \rho$ then both q_{il} and q_{jl} are increased by $1/(c_{ij} + \delta)^\varepsilon$ ($l \in \{a'_i, \ldots, b_i\} \cap \{a'_j, \ldots, b_j\}$), where ρ, δ and ε are user-defined parameters. A smaller c_{ij} results in a larger increment (see Figure 4). For each customer i, the day with the highest compatibility index is selected as the best service day. The customers whose best

service days lie within T_t are selected for visit during that horizon. This procedure is described as Algorithm 2.

Algorithm 2 : Phase I (Customer selection)

 1: Input: the set of known customers N
 2: Output: the set of customers N_t to be visited within period T_t
 3: **for** $i = 1$ to $|N|$ **do**
 4:     **for** $l = a'_i$ to b_i **do**
 5:         $q_{il} \leftarrow 0$
 6:     **end for**
 7: **end for**
 8: **for** $i = 1$ to $|N| - 1$ **do**
 9:     **for** $j = i + 1$ to $|N|$ **do**
10:         **if** $c_{ij} \leq \rho$ and $\{a'_i, \ldots, b_i\} \cap \{a'_j, \ldots, b_j\} \neq \emptyset$ **then**
11:             **for** $l \in \{a'_i, \ldots, b_i\} \cap \{a'_j, \ldots, b_j\}$ **do**
12:                 $q_{il} \leftarrow q_{il} + 1/(c_{ij} + \delta)^\varepsilon$
13:                 $q_{jl} \leftarrow q_{jl} + 1/(c_{ij} + \delta)^\varepsilon$
14:             **end for**
15:         **end if**
16:     **end for**
17: **end for**
18: **for** $i = 1$ to $|N|$ **do**
19:     $v_i \leftarrow \arg\min_{l \in \{a'_i, \ldots, b_i\}} q_{il}$
20: **end for**
21: $N_t \leftarrow \{i : v_i \in T_t\}$

3.2 Phase II: Variable neighborhood search

The aim of the Phase II is to construct routes for customers on each day of T_t. This problem is treated as a PVRP with frequency 1, where the planning horizon is $\{t, t + 1, \ldots, t + \tau\}$, and each selected customer i must be served with frequency 1 between day $\max\{t, a_i\}$ and day $\min\{t + \tau, b_i\}$. The PVRP is solved by means of a variable neighborhood search heuristic (see Algorithm 3), made up of three components: initialization, local search and shaking. An initial solution is first constructed by means of a sweep heuristic. The local search phase is based on a tabu search (TS) algorithm that uses simple insertion moves to transfer customers from their route to another route. For each customer, all possible reinsertion positions are attempted and the one leading to the
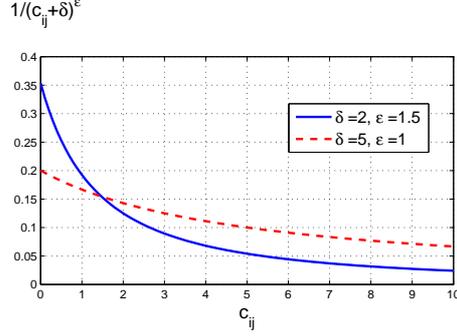
Figure 4. The increment value curve for different value of δ and ε

minimum objective value is selected. If TS fails to improve the solution within a preset number θ of iterations, it restarts from another solution provided by a shaking phase, based on a ruin and recreate approach (RRA) (Schrimpf et al., 2000; Pisinger and Ropke, 2007). This procedure is initiated from the best known solution and attempts to iteratively improve it by removing $\xi\%$ of the customers that have the largest removal costs, and reinserting them by means of the regret insertion method described in Algorithm 4. If the RRA finds a better solution or fails to improve the best solution after κ iterations, TS is reapplied to it. Phase II stops after ω iterations.

3.3 Phase III: Postoptimization

Phase III aims to minimize the total travel time on day t. This problem is a *Capacitated Vehicle Routing Problem* which is solved by the TS heuristic of Cordeau, Gendreau and Laporte (1997). In this algorithm, intermediate infeasible solutions are allowed during the search and are controlled by means of a penalized objective $f'(s,t) = c(s,t) + \alpha q(s,t) + \beta d(s,t)$, where $c(s,t)$ is the total travel time by all vehicles on day t, and $q(s,t) = \sum_{k \in K} (\sum_{(i,j) \in A} q_i x_{ijkt}^t - Q)^+$ and $d(s,t) = \sum_{k \in K} (\sum_{(i,j) \in A} (c_{ij} + d_i) x_{ijkt}^t - D)^+$ are the total violations of the capacity and duration constraint on day t, where $(x)^+ = \max\{0, x\}$. The coefficients α and β are positive self-adjusting penalties. Since we only optimize the routes on day t, the last two objectives, minimizing customer waiting and balancing the daily workload, are not considered in this phase.

11

Algorithm 3 : Phase II (Variable neighborhood search)

 1: Input: the set of customers N_t to be visited within period T_t
 2: Output: the solution $s^\star$
 3: $s \leftarrow \text{SweepHeuristic}(N_t)$
 4: $s^\star \leftarrow s$
 5: $iteration \leftarrow 0$
 6: **while** $iteration < \omega$ **do**
 7:     $counter \leftarrow 0$
 8:     **while** $counter < \theta$ **do**
 9:         $s \leftarrow \text{TabuSearch}(s)$
10:         $iteration \leftarrow iteration + 1$
11:         **if** $s < s^\star$ **then**
12:             $counter \leftarrow 0$
13:             $s^\star \leftarrow s$
14:         **else**
15:             $counter \leftarrow counter + 1$
16:         **end if**
17:     **end while**
18:     $counter \leftarrow 0$
19:     $s \leftarrow s^\star$
20:     **while** $counter < \kappa$ **do**
21:         $s \leftarrow \text{RRA}(s, \xi)$
22:         $iteration \leftarrow iteration + 1$
23:         **if** $s < s^\star$ **then**
24:             $s^\star \leftarrow s$
25:             break
26:         **else**
27:             $counter \leftarrow counter + 1$
28:         **end if**
29:     **end while**
30: **end while**
31: return $s^\star$
---

4 Computational Results

The heuristic just described was implemented in C and executed on a Linux computer with lx24-amd64 architecture and two Gbytes of RAM. The data and parameters used in our tests are first described. Sensitivity analyses on the parameters used in the heuristic are then performed. Finally we provide the results of our tests on the Lantmännen data.

Algorithm 4 : Phase II (Ruin and recreate heuristic)

1: $numToRemove$ is the number of customers to be removed and reinserted
2: N_t is the set of customers in the solution s
3: Input: current solution s
4: Output: updated solution s
5: $N_{Rem} \leftarrow \emptyset$
6: **while** $|N_{Rem}| < numToRemove$ **do**
7:    **for** $i \in N_t$ **do**
8:       $RC_i \leftarrow$ CalculateRemovalCost(i, s)
9:    **end for**
10:    $i^\star \leftarrow \arg\min_{i \in N_t} RC_i$
11:    $s \leftarrow$ RemoveCustomer($i^\star$, s)
12:    $N_{Rem} \leftarrow N_{Rem} \cup \{i^\star\}$
13:    $N_t \leftarrow N_t \setminus \{i^\star\}$
14: **end while**
15: **while** $N_{Rem} \neq \emptyset$ **do**
16:    **for** $i \in N_{Rem}$ **do**
17:       $bestIC_i \leftarrow$ CalculateBestInsertionCost(i, s)
18:       $secondIC_i \leftarrow$ CalculateSecondBestInsertionCost(i, s)
19:    **end for**
20:    $i^\star \leftarrow \arg\max_{i \in N_{Rem}} (secondIC_i - bestIC_i)$
21:    $s \leftarrow$ InsertCustomer($i^\star$, s)
22:    $N_{Rem} \leftarrow N_{Rem} \setminus \{i^\star\}$
23: **end while**

4.1 Data and parameters

Real-world data were collected from Lantmännen. There are altogether 11 data sets, five of which involve a 10-day planning horizon and six involve a 15-day planning horizon. On average 80 orders are received every day. The number of feasible service days ranges from one to 15 and is equal to 2.5 on average. Figure 5 shows the distribution of the number of days elapsed between the day at which an order is placed and the first feasible service day. Most customers order two or three days before the start of the service period. The average demand of the orders is 6,306kg, and the vehicles have a capacity of 40,000kg. We use Euclidian distances and assume the vehicle speed is 45km/hour.

Based on preliminary tests, parameters ρ, δ and ε in Phase I of the TPH were set to 60, 2 and 1.5, respectively. The maximum numbers of non-improving iterations for the TS and the RRA of Phase II, i.e., parameters θ and κ, were set to 10^2 and 10^4, respectively. In the RRA, between 25% and 35%
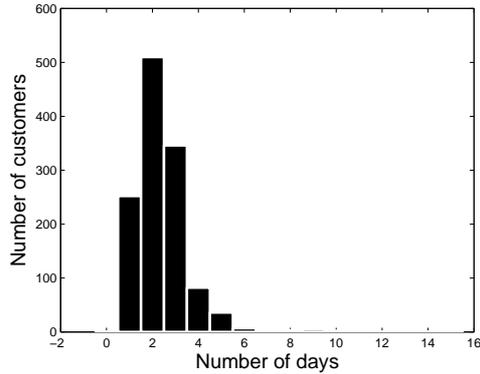
Figure 5. Distribution of the number of days before the start of the service period when customers call in

of the customers are removed and reinserted in each iteration. The estimated daily workload for objective function f_3^t in Equation (11) is obtained from the workload of the previous five days and is updated adaptively for each planning day.

4.2 Sensitivity analyses

This section describes the sensitivity analyses that were performed to assess the behaviour of the TPH.

Number of days to plan in TPH

As mentioned in Section 3, the TPH not only plans the routes for day t, but also for τ days in the future. We tested the TPH with different values of τ on 11 instances. Figure 6 illustrates the convergence of the TPH for three different values of τ. When τ equals 1 or 2, Phase I selects approximately 33% or 50% of the customers, respectively. The results show that $\tau = 1$ is not sufficient but $\tau = 2$ works very well. With a short running time (less than four minutes), $\tau = 2$ even provides better results than $\tau = \infty$. This is because within a given short running time, the problem of smaller size can be better optimized due to a more thorough search, and the correlation analysis provides good candidates for the customers that should be visited within the next two days.
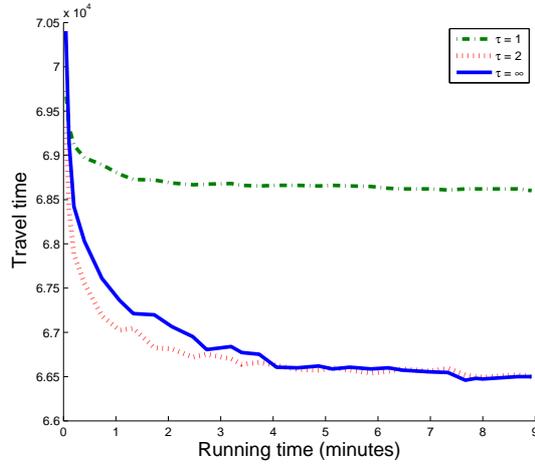
14

Figure 6. Sensitivity analysis of number of planning days (τ)

Effectiveness of correlation analysis

To further demonstrate the effectiveness of correlation analysis, we compare the results obtained with correlation analysis to those using a random scheme. In the random scheme, we assume each customer is randomly and uniformly assigned to one of its feasible service days, and customers assigned to the first τ days are selected. Figure 7 shows the comparison between the two schemes. The horizontal axis is the instance index and the vertical axis gives the total travel time over the planning horizon by using a random selection scheme or correlation analysis. The running time is set at four minutes. For all 11 instances, the solutions provided by the correlation analysis are consistently better than those obtained by the random selection scheme.

Results for the multi-objective function

In this experiment, we assess the effectiveness of the TPH to handle the multiple objectives. Table 1 shows the values of the first objective, i.e., total travel time (denoted by F_1), and of the second objective, i.e., total customer waiting (denoted by F_2), with different values of w_2 ranging from 0 to 20. Column 'F_1' and 'F_2' are the total travel time over the planning horizon and the total number of waiting days for all customers over the planning horizon, respectively. The last row 'Average' shows the average values for the 11 instances. As w_2 increases from 0 to 20, the total customer waiting is reduced by half on average, whereas the total travel time increases only slightly, by less
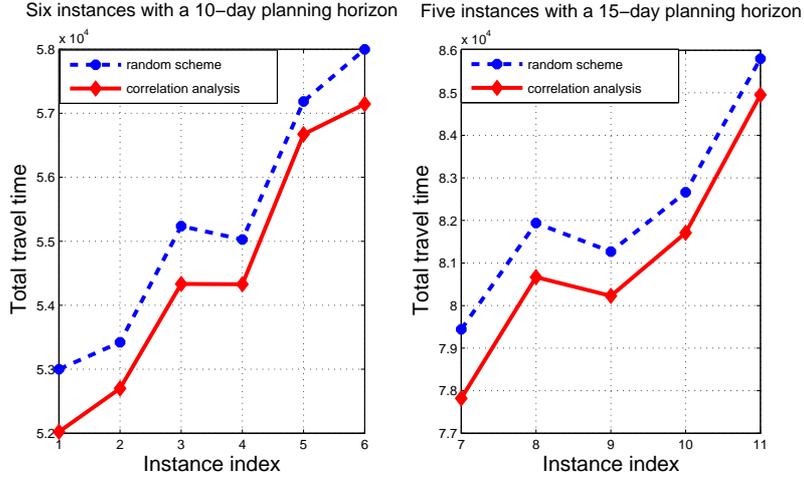
Figure 7. Comparison between correlation analysis and random selection scheme

than 1 %. Figure 8 depicts the relative changes of total travel time and total customer waiting as a function of w_2.

Similar results are obtained for the total travel time and for the total workload deviations with increasing values of w_3, as shown in Table 2 and Figure 9. In Table 2, column 'F_3' is the sum of deviations between each day's duration and the average daily duration, over the planning horizon. The last row shows the average values for the 11 instances. As can be seen from the results, when w_3 increases from 0 to 0.6, the average workload deviation decreases by more than 70%, whereas the total travel time only increases by approximately 0.5%. We also note that the rate of deviation reduction decreases as w_3 increases. In Figure 9, within the interval $0.4 \leq \beta \leq 0.6$, the deviation reduction curve is nearly flat and the deviation reduction is insignificant. This is because the objective function f_3^t used in the TPH minimizes the difference between the workload on day t and an estimation of the average daily workload instead of the actual average workload.

4.3 Comparison between TPH solutions and the company's solutions

Lantmännen already works with high quality solutions obtained by running their vehicle routing software 12 minutes each day on their latest platform. In order to establish a fair comparison, we have also run our algorithm for 12 minutes on a similar computer, but the improvement obtained

w_2	0		2		4		8		12		16		20	
Data set	F_1	F_2	F_1	F_2	F_1	F_2	F_1	F_2	F_1	F_2	F_1	F_2	F_1	F_2
101	52024	380	52271	317	52060	296	52295	251	52306	225	52386	203	52546	171
102	52700	391	52786	327	52919	297	52974	267	53001	236	53039	198	53222	176
103	54333	479	54322	412	54401	389	54439	332	54605	296	54802	272	54939	254
104	54328	449	54222	381	54305	343	54471	310	54798	263	54641	250	55038	241
105	56672	404	56522	343	56591	302	56607	252	57012	218	57063	203	57164	185
106	57144	414	57279	365	57346	320	57381	276	57684	254	57582	228	57872	199
151	77820	708	78103	593	78057	528	78079	477	78341	419	78377	385	78564	351
152	80669	658	80628	578	80485	516	80671	451	80825	403	81077	356	81232	328
153	80226	658	80213	587	79985	525	80373	459	80339	415	80690	362	80631	350
154	81710	620	81809	532	81879	492	82159	417	82248	379	82203	350	82431	318
155	84950	634	85015	540	85031	474	84948	422	85142	399	85420	342	85465	331
Average	66597	526	66651	452	66641	407	66763	355	66936	318	67025	286	67191	264

Table 1. Total travel time and total customer waiting with different values of the wight w_2 assigned to customer waiting time

after four minutes is insignificant. Comparative results are presented in Table 3. Ten random runs for each instance are performed for our algorithm to obtain the average value of the total driving time, total customer waiting, and daily workload deviation. These statistics are provided in the columns 'Average total duration', 'Average total customer waiting' and 'Average total workload deviation', respectively. The best value of the total travel time within the 10 random runs are also presented in column 'Best total duration'. The results provided by the company are for a single run. The average values for all the instances are given in the last row. Regarding the total duration, our average value for 10 runs is slightly better (by 0.2%) than that of Lantmännen, probably because their solutions and ours are both very close to optimality. However, the TPH significantly improves customer waiting and workload deviation by up to 24% and 35%, compared with the company's solutions. We also found the best solutions for all instances. This is a clear sign of the effectiveness of our heuristic. One should bear in mind, however, that customer waiting and daily workload balance may not have been optimized by the company.

5 Conclusion

We have considered a real-life dynamic multi-period and multi-objective routing problem encountered by a large distributor operating in Sweden. The planning horizon consists of several periods
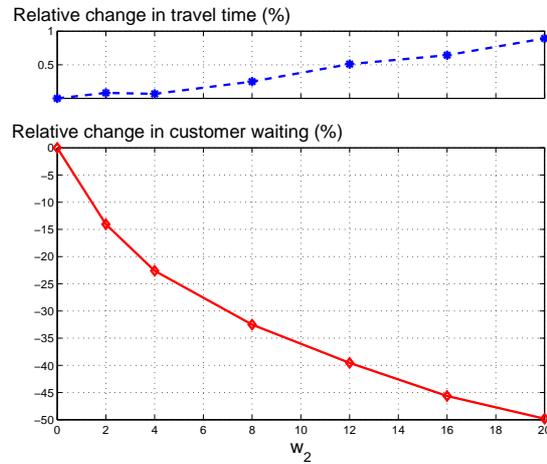
Figure 8. Relative changes in total travel time and total customer waiting as a function of the weight w_2 assigned to customer waiting time

and the problem considers three objectives, including minimization of the total travel time, minimization of customer waiting, and daily workload balancing over the planning horizon. We have presented a mixed integer linear programming formulation for the problem, and we have proposed a three-phase heuristic embedded within a rolling horizon scheme. The main idea of the heuristic is to wisely select the customers to be visited in the near future, and to route these customers so that the overall travel time can be minimized efficiently. The choice of customers to be routed on a given day is performed rather effectively through a time-space correlation analysis. The multiple objectives are handled by the scalar technique. The method was implemented and tested on real-life data. Results show that the proposed TPH provides very high quality solutions within a reasonable running time. It improves upon the company's solutions in terms of travel time, customer waiting and daily workload balance, with gains of 0.2%, 24% and 35%, respectively. Our method is general and applies to other contexts.

Acknowledgements

18

w_3	0		0.1		0.2		0.3		0.4		0.5		0.6	
Data set	F_1	F_3	F_1	F_3	F_1	F_3	F_1	F_3	F_1	F_3	F_1	F_3	F_1	F_3
101	52024	1221	51972	597	52284	344	52174	169	52303	131	52447	139	52247	123
102	52700	1146	52791	664	52676	340	52878	214	52863	154	52876	87	52786	123
103	54333	1285	54183	572	54274	430	54157	266	54107	239	54248	207	54207	190
104	54328	1603	54463	835	54566	494	54851	491	54812	399	54877	411	55063	369
105	56672	1430	56711	1043	56888	599	57045	527	56994	461	56908	405	57211	488
106	57144	1372	57310	1009	57396	752	57496	631	57522	610	57576	673	57682	645
151	77820	1218	77946	808	78022	638	78264	603	78357	552	78466	555	78400	500
152	80669	1032	80718	654	80832	499	80596	416	80747	363	80603	373	80701	336
153	80226	1170	80128	605	80199	404	80336	305	80503	276	80401	245	80483	267
154	81710	1526	81945	848	81826	726	82021	551	82318	510	82253	512	82447	494
155	84950	1460	84848	989	84971	723	84921	693	85378	631	85393	697	85386	683
Average	**66597**	**1314**	**66637**	**784**	**66721**	**540**	**66794**	**442**	**66900**	**393**	**66913**	**391**	**66964**	**383**

Table 2. Total travel time and total workload deviation for the tests with different values of the weight w_3 assigned to balance daily workload
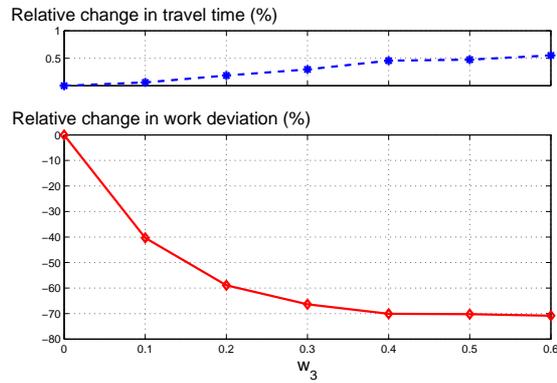

Figure 9. Relative changes in total travel time and total workload deviation as a function of the weight w_3 assigned to daily workload deviation

		Lantmännen solution			TPH solution			
Data set	$|T|$	Total duration	Total customer waiting	Total workload deviation	Average total duration	Best total duration	Average total customer waiting	Average total workload deviation
101	10	51958	382	729	52064	51679	330	588
102	10	52702	369	844	52663	52163	312	693
103	10	53727	560	1231	54128	53711	408	530
104	10	54299	444	1117	54439	54089	372	721
105	10	56476	440	535	56429	56009	344	812
106	10	57729	411	606	57211	56740	352	716
151	15	77362	759	1170	77706	77320	632	710
152	15	81135	828	978	80303	79963	618	627
153	15	79600	767	312	79730	79342	604	518
154	15	82164	713	1315	81711	81321	526	756
155	15	85270	685	1060	84540	83992	552	740
Average		**66584**	**578**	**900**	**66447**	**66029**	**459**	**673**

Table 3. Comparison between the Lantmännen solutions and the TPH solutions

References

[1] Alabas-Uslu, C., 2008. A self-tuning heuristic for a multi-objective vehicle routing problem, *Journal of the Operational Research Society*, 59(7), 988–996.

[2] Angelelli, E., Speranza, M.G., Savelsbergh, M.W.P., 2007. Competitive analysis for dynamic multi-period uncapacitated routing problems, *Networks*, 49(4), 308–317.

[3] Angelelli, E., Bianchessi, N., Mansini, R., Speranza, M.G., 2009. Short term strategies for a dynamic multi-period routing problem, *Transportation Research Part C*, 17(2), 106–119.

[4] Berbeglia, G., Cordeau, J.-F., Laporte, G., 2009. Dynamic pickup and delivery problems, forthcoming in *European Journal of Operational Research*.

[5] Branke, T., Middendorf, M., Noeth, G., Dessouky, M., 2005. Waiting strategies for dynamic vehicle routing, *Transportation Science*, 39(3), 298–312.

[6] Cordeau, J.-F., Gendreau, M., Laporte, G., 1997. A tabu search heuristic for periodic and multi-depot vehicle routing problems, *Networks*, 30(2), 105–119.

[7] Dell'Olmo, P., Gentili, M., Scozzari, A., 2005. On finding dissimilar Pareto-optimal paths, *European Journal of Operational Research*, 162(1), 70–82.

[8] Francis, P.M., Smilowitz, K.R., Tzur, M., 2008. The period vehicle routing problem and its extensions, In Golden, B.L., Raghavan, S., Wasil, E.A., editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, 73–102, Springer, New York.

[9] Hemmelmayr, V.C., Doerner, K.F., Hartl, R.F., 2009. A variable neighborhood search heuristic for periodic routing problems, *European Journal of Operational Research*, 195(3), 791–802.

[10] Hvattum, L.M., Løkketangen, A., Laporte, G., 2006. Solving a dynamic and stochastic vehicle routing problem with a sample scenario hedging heuristic, *Transportation Science*, 40(4), 421–438.

[11] Hvattum, L.M., Løkketangen, A., Laporte, G., 2007. A branch-and-regret heuristic for stochastic and dynamic vehicle routing problems, *Networks*, 49(4), 330–340.

[12] Jozefowiez, N., Semet, F., Talbi, E.-G., 2008. Multi-objective vehicle routing problems, *European Journal of Operational Research*, 189(2), 293–309.

[13] Lacomme, P., Prins, C., Sevaux, M., 2006. A genetic algorithm for a bi-objective capacitated arc routing problem, *Computers & Operations Research*, 33(12), 3473–3493.

[14] Larsen, A., Madsen, O.B.G., Solomon, M.M., 2008. Recent developments in dynamic vehicle routing systems, In Golden, B.L., Raghavan, S., Wasil, E.A., editors, *The Vehicle Routing Problem: Latest Advances and New Challenges*, 199–218, Springer, New York.

[15] Mitrović-Minić, S., Krishnamurti, R., Laporte, G., 2004. Double horizon based heuristics for the dynamic pickup and delivery problem with time windows, *Transportation Research Part B*, 38(8), 669–685.

[16] Mourgaya, M., Vanderbeck, F., 2007. Column generation based heuristic for tactical planning in multi-period vehicle routing, *European Journal of Operational Research*, 183(3), 1028–1041.

[17] Pacheco, J., Martí, R., 2007. Tabu search for a multi-objective routing problem, *Journal of the Operational Research Society*, 57(1), 29–37.

[18] Psaraftis, H.N., 1988. Dynamic vehicle routing problems, In Golden, B.L., Assad, A.A., editors, *Vehicle Routing: Methods and Studies*, 223–248, North-Holland, Amsterdam.

[19] Pisinger, D., Ropke, S., 2007. A general heuristic for vehicle routing problems, *Computers & Operations Research*, 34(8), 2403–2435.

[20] Pureza, V., Laporte, G., 2008. Waiting and buffering strategies for the dynamic pickup and delivery problem with time windows, *INFOR*, 46(3), 165–175.

[21] Schrimpf, G., Schneider, J., Stamm-Wilbrandt, H., Dueck, G., 2000. Record breaking optimization results using the ruin and recreate principle, *Journal of Computational Physics*, 159(2), 139–171.

[22] Tan, K.C., Chew, Y.H., Lee, L.H., 2006. A hybrid multi-objective evolutionary algorithm for solving truck and trailer vehicle routing problems, *European Journal of Operational Research*, 172(3), 855–885.

[23] Zografos, K.G., Androutsopoulos, K.N., 2008. A decision support system for integrated hazardous materials routing and emergency response decisions, *Transportation Research Part C*, 16(6), 684–703.

This paper considers the Dynamic Multi-Period Vehicle Routing Problem which deals with the distribution of orders from a depot to a set of customers over a multi-period time horizon. Customer orders and their feasible service periods are dynamically revealed over time. The objectives are to minimize total travel costs and customer waiting, and to balance the daily workload over the planning horizon.

This problem originates from a large distributor operating in Sweden. It is modeled as a mixed integer linear program, and solved by means of a three-phase heuristic that works over a rolling planning horizon. The multi-objective aspect of the problem is handled through a scalar technique approach. Computational results show that our solutions improve upon those of the Swedish distributor.