Original Article

Vehicle routing in a Spanish distribution company: Saving using a savings-based heuristic

Alex Grasas[a,*], Jose Caceres-Cruz[b], Helena R. Lourenço[a], Angel A. Juan[b] and Mercè Roca[c]

[a]Department of Economics and Business, Universitat Pompeu Fabra, Ramon Trias Fargas 25–27, 08005 Barcelona, Spain.
E-mails: alex.grasas@upf.edu; helena.ramalhinho@upf.edu
[b]Department of Computer Science, Multimedia and Telecommunication, IN3-Open University of Catalonia, Roc Boronat 117, 08018 Barcelona, Spain.
E-mails: jcaceresc@uoc.edu; ajuanp@uoc.edu
[c]ESCI International Business School, Pg. Pujades 1, 08003, Barcelona, Spain.
E-mail: merce.roca@esci.upf.edu

*Corresponding author.

Abstract In this article we present a Vehicle Routing Problem (VRP) faced by a large distribution company in the Northeast of Spain. The company distributes products from its central facilities to a chain of around 400 stores all over the country. One of the peculiarities of the VRP of this company – which is common among real-life VRPs – is the presence of a heterogeneous fleet where vehicles with different capacities can make multiple trips during a single day. This variant of the problem, which we refer as Heterogeneous Fleet and Multi-trip VRP, has been barely studied in the literature. To solve the problem, we use an algorithm based on the well-known savings heuristic with a biased-randomization effect and three local search operations. Our approach is simple to implement as it needs few parameters and no fine-tuning processes, which are usually cumbersome and require experts' involvement. We obtain savings of around 12 per cent in transportation costs, which represent around €30 000 saved per week.
OR Insight (2013) **26,** 191–202. doi:10.1057/ori.2013.2; published online 13 March 2013

Keywords: heterogeneous vehicle routing problem; multi-trip vehicle routing; savings heuristic; randomized algorithms; real-life applications

Received 11 October 2012; accepted 31 January 2013 after one revision

Introduction

In the past few years, logistics and transportation companies are growingly facing demanding situations with fewer available resources. This is basically because of market instability and to an increase of competitiveness in the business environment. Road haulage represents the main mode of goods transportation in Europe and other places in the world. Since 2000, the economic and environmental impact caused by road transportation has been growing. Governments and corporations around the world have placed their attention on terrestrial logistics and the optimization of distribution processes. Such optimization has become critical to companies of any size in order to improve the quality of their service and the customer satisfaction, as well as to reduce costs.

In this article we present a Vehicle Routing Problem (VRP) faced by a large distribution company in the Northeast of Spain. In a VRP, a fleet of homogeneous vehicles is based at a single depot to serve demands for a set of geographically dispersed customers. Each vehicle has the same capacity, and it leaves from the depot, visits some customers, and returns to the depot. Each customer has a known demand and is served by exactly one visit of a single vehicle. The problem consists in finding a sequence of deliveries (routes) for each vehicle so that all customers are served and the total distance travelled by the fleet is minimized. Although the VRP, first introduced by Dantzig and Ramser (1959), has been studied for over 50 years (Toth and Vigo, 2002; Laporte, 2009), real-life applications such as the one analysed here, present several additional constraints and operation rules that must be considered when implementing solution algorithms and methods. These additional considerations may affect customers, depot(s), vehicles and so on, and may have a significant impact on the solutions.

In our study case, the main difference with respect to the original VRP is the presence of a heterogeneous fleet where vehicles with different capacities can make multiple trips on a single day. This variant of the problem, known as Heterogeneous Fleet and Multi-trip VRP, has been barely studied in the literature. Lourenço *et al* (2013) describe the problem in detail and compare it with the rest of related literature. To the best of our knowledge, only Prins (2002) proposes a similar version of the problem that differs in how some customers can (or cannot) be served by a particular type of vehicle. We use Lourenço *et al*'s algorithm which is a relatively simple to implement method based on the well-known savings heuristic (Clarke and Wright, 1964) with a biased-randomization effect and three local search processes. This approach reduces the company distribution costs on the order of €30 000 per week. In addition, the method only needs few parameters that avoid complex fine-tuning processes, which are usually time consuming and require experts' involvement. This is particularly useful as it is intended to be used by the distribution company dispatchers responsible for route planning.

A Real Vehicle Routing Application

The distribution company of this study distributes products from its central facilities in the Northeast of Spain to a chain of around 400 stores all over the country. Orders from every store are received daily (that is, Monday through Saturday), and the distribution is then carried out from a central depot by a company-owned fleet of 169 vehicles. This fleet includes trucks of different capacities (see Table 1 for the fleet composition). At a glance, the daily distribution planning process unfolds as follows:

(i) *Order placement*: stores place orders by noon with no restriction on the number of boxes.
(ii) *Order planning*: orders are received at the central depot and may be adjusted depending on product availability.
(iii) *Route planning*: three route dispatchers plan routes to all stores by 2pm (see more details on route planning below).
(iv) *Distribution*: vehicles load the cargo at the depot and depart to the stores. Truck loading is divided into three shifts (at 2pm, 3pm and 5pm, respectively).
(v) *Delivery*: vehicles arrive at the stores between 5pm and 1am of the next day, and unload their cargo.
(vi) *Return to depot*: after the last store in the route is served, vehicles return to the depot.

The route planning step establishes the routes that vehicles must follow to deliver the products. This is, in fact, the solution to the VRP, as defined in the

Table 1: Fleet composition of the distribution company

Vehicle type	*Capacity (in boxes)*	*Number of vehicles*
A	222	8
B	414	5
C	482	139
D	550	3
E	616	6
F	676	3
G	752	4
H	1.210	1
Total	—	169

'Introduction'. This phase is obviously the crucial step in the distribution process as it determines most of the total distribution costs. Currently, this task is executed manually by three route dispatchers. They divide all stores into three geographical areas, so that each dispatcher is responsible for the routes in their region (that is, each of them solves a smaller VRP). For each region, they have sets of predetermined routes that modify slightly according to daily demand and truck availability. Trucks and stores are usually assigned to one of the loading shifts, so that routes include stores in the same shift only. In addition, there exist other specific constraints on the routing problem that make this company's VRP quite unique. These need to be considered when designing routes:

(a) The number of trucks available each day and shift may vary because of eventualities.
(b) Not all types of vehicles can visit all stores. For example, large trucks cannot access some stores for maneuverability reasons.
(c) Some stores have restrictions on their delivery times. For example, trucks may not be allowed in some urban areas before (or after) a determined time. These are known as delivery time windows.
(d) Some trucks are allowed to make multiple trips (generally two). In days of high demand, for example, the total capacity of all available vehicles may not be enough to cover all demand, so that some trucks perform two trips on that same day. This implies that some stores could be visited twice by the same truck or by a different one, having to split their order in two.
(e) Each truck is driven by a single driver, so there is an upper bound on the route duration given by the maximum number of working hours (that is, 8 hours).

With all this information and constraints, the three company dispatchers have about 1 hour every day to configure the delivery routes. The planning is done manually with some computer aid to perform simple verifications (like tracking the number of boxes yet to be assigned). As the number of stores continues to grow, the need for a scientific method to help the decision making becomes more latent. This is a very complex problem that requires more sophisticated methods to obtain better and faster solutions that allow the company to save considerable costs in transportation.

Solving the Heterogeneous VRP

Vehicle routing algorithms aim at finding sets of routes to serve all customers and minimize the total distance travelled by all vehicles. The more conventional versions of the problem tend to ignore the problem of vehicle availability, as they assume that an unlimited fleet of identical vehicles is available to perform the routes. Unfortunately, this is hardly the case in practice, and companies, such as the one studied in this article, usually have a limited fleet of vehicles that restricts route planning. Even in cases where distribution processes are outsourced to third-party logistics providers, whose fleets could be quite large, special attention must be given to vehicle assignment in order to use resources efficiently. For this reason, some variants of the VRP regarding fleet availability or composition have been analysed in the literature. We focus our attention in those variants with heterogeneous fleets (that is, containing vehicles with different capacities). In the Vehicle Fleet Mix Problem (Prins, 2009) the fleet is heterogeneous but there are unlimited number of vehicles for every type. The Fleet Size and Mix VRP (Golden *et al*, 1984) considers also an unlimited fleet but incorporates a decision on the fleet size. The reader is referred to Baldacci *et al* (2008) and Prins (2009) for an extensive survey of different heterogeneous fleet variants of the VRP.

All variants of the VRP are very difficult to solve optimally. The heterogeneous fleet version we study is no exception. Although some exact methods have been proposed in the literature, their use in real applications is limited as the problem size increases. Therefore, most approaches propose heuristics and metaheuristics that provide near-optimal solutions. A metaheuristic is a computational method designed to solve large optimization problems like this one. The method starts from an initial solution, and it searches iteratively for better solutions using a series of rules and conditions. Most of the heuristics and metaheuristics used in the heterogeneous VRP have been adapted from successful implementations on the classical VRP (Osman and Salhi, 1996; Gendreau *et al*, 1999; Tarantilis *et al*, 2004; Li *et al*, 2007; Prins, 2009; Euchi and Chabchoub, 2010; Brandao, 2011).

We use a method to solve the heterogeneous VRP of the distribution company that is a modified version of the algorithm developed by Juan *et al* (2010). The reader is referred to Lourenço *et al* (2013) for a complete technical description of the algorithm used in this article. The approach is a hybrid algorithm that combines the parallel version of the classical savings heuristic, originally developed by Clarke and Wright (1964), with Monte Carlo simulation. The savings heuristic iteratively combines routes while cost savings (in terms of distance or time) are obtained and problem constraints are not violated. In the parallel version, unlike the sequential version, several routes can be constructed at a time. The savings heuristic is the most referred method in the VRP literature and the most commonly used in VRP software (Rand, 2009). The method is simple to apply and provides relatively good solutions in reasonable time. The algorithm benefits from these features and introduces some variations that make it suitable for the company problem. In particular, it uses a biased randomization that transforms the heuristic into a multi-start probabilistic algorithm (step iv. below), and three local searches that improve the resulting solution (step vii. below). Next, we summarize the steps that the algorithm follows to solve the problem:

(i) It constructs an initial solution in which a virtual vehicle is assigned to each customer.
(ii) It computes savings associated with each edge connecting each pair of customers. These savings are cost reductions obtained when routes are merged.
(iii) It stores edges in a list in decreasing order of savings.
(iv) It creates a new savings list of edges by randomizing the original list. The randomization is done using a biased probability distribution such as the Geometric (this tends to place higher-saving edges to the top of the list).
(v) It uses a multi-start procedure, so that different randomizations of the previous step lead to different algorithm starts.
(vi) While the savings list is not empty, it extracts the edge at the top of the list and merges two different routes if:

    (a) The extreme points of the edge are external (that is, connected directly to the depot).
    (b) The capacity and maximum route length constraints are not violated. This validation is done as in Prins (2002). Vehicles are listed in decreasing capacity and current routes in decreasing accumulated demand. A temporary assignment between the two lists is searched. If a successful match is found, the merge is valid; otherwise, the merge is discarded.

    After each merge, a fast 2-Opt local search (Croes, 1958) is performed over the new route.
(vii) Once the savings list is empty, the resulting solution is improved using two additional local search methods (Juan *et al*, 2011b):

    (a) A cache-based local search, which updates routes by checking a table that contains best-known routes, found in previous iterations, connecting the same sets of costumers.
    (b) A splitting-based local search, which divides current solution into sets of routes with their corresponding vehicles, and applies the previous steps of the algorithm to these smaller VRPs to find better local solutions.

    Then, the resulting solution is compared with best solution found, and updated (if appropriate).
(viii) When a time-based criterion is reached, the best solution is returned.

An interesting feature of this algorithm is that it can be easily parallelized. The randomization made at the beginning of the algorithm (step iv.) allows different starts by changing the seed used to initialize the pseudo-random number generator. Therefore, it is possible to run different instances in parallel by using different seeds in different threads, cores or even computers (Juan *et al*, 2011a). This increases the possibility of finding better solutions in a given period of time, something very interesting as the company needs to plan the routes every day in less than 1 hour.

Results

Before running the algorithm to solve the problem, all necessary input data had to be compiled and prepared in the appropriate format. This basically refers to all problem parameters and constraints, which include data from all stores (demands and postal addresses), vehicle capacities, truck-store incompatibilities, delivery time windows and maximum time per route (that is, at most 8 hours per route). With all addresses, including that of the depot, we constructed distance and time matrices. These matrices contained all travel distances and times between every pair of stores, and between all stores and the depot. For about 400 stores plus a depot, this implied finding around 160 000 distances and times. To automate this quest, we developed a web application (http://vrp.upf.edu) that uses Google Maps in which the user uploads an Excel file with all addresses, and the application returns, in few seconds, a plain text file with the matrices in a format ready for our algorithm.

To illustrate the performance of the algorithm we summarize the results obtained for 10 real instances from 10 business days in 2012 (see Table 2). We imposed a time bound of 10 minutes to the algorithm (step viii. of the algorithm). For each instance, the table shows the number of stores visited that day as well as the total demand (in boxes) delivered. The table compares the solution obtained by the company dispatchers with our solution, showing the total logistics cost and the number of routes used. Notice that more routes does not necessarily imply higher costs as can be seen in some instances. Multi-trips appear in almost all solutions, that is, those in which the number of routes exceeds the total number of vehicles (that is, 169). The algorithm was executed using 10 different seeds per instance (step v.). The table reports both the average cost of the 10 runs and the best solution found. The best solution for each instance was obtained in few seconds and the average cost reduction was around 12 per cent, which represents savings of around €5000 per day.

Figure 1 shows the 20 longest routes obtained by Lourenço *et al*'s algorithm in instance 10, providing a picture of the territorial extension supplied by the

0953-5543 **OR Insight** Vol. 26, 3, 191–202

Table 2: Comparison of results for 10 real instances

Instances	*Stores visited*	*Total demand delivered*	*Company solution*		*Our solution*					
			Cost (1)	*Routes*	*Best cost (2)*	*Routes*	*Time (sec)*	*Gap (2-1) %*	*Average 10 seeds (3)*	*Gap (3-1) %*
1	372	77 913	45 302.40	173	39 534.11	180	63.22	−12.73	39 841.99	−12.05
2	366	79 130	47 184.35	182	41 072.46	183	55.20	−12.95	41 399.65	−12.26
3	371	91 901	53 941.43	218	49 669.31	219	66.48	−7.92	50 082.32	−7.15
4	364	63 078	36 897.29	136	31 378.63	135	58.94	−14.96	31 543.09	−14.51
5	372	83 571	50 872.80	197	45 485.83	200	29.73	−10.59	45 836.63	−9.90
6	373	85 773	51 315.40	207	45 275.62	206	7.67	−11.77	45 681.39	−10.98
7	372	84 023	50 492.74	200	45 165.12	197	28.53	−10.55	45 493.28	−9.90
8	374	85 539	51 427.10	208	44 386.64	200	65.94	−13.69	44 909.39	−12.67
9	370	89 596	54 446.01	215	49 053.97	212	59.57	−9.90	49 354.83	−9.35
10	372	76 846	45 056.40	172	38 973.19	175	29.33	−13.50	39 252.86	−12.88
Average	—	—	48 693.59	190.80	42 999.49	190.70	46.46	−11.86	43 339.54	−11.17

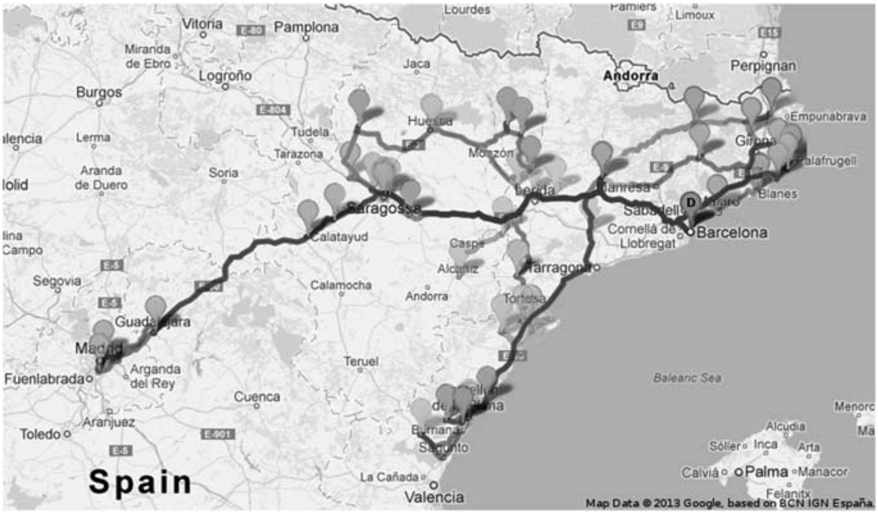

Figure 1: Geographical situation of the 20 longest routes of instance 10, using Google Maps.

company. These routes represent a total distance of 9057 km and 68 customers. Depot is marked with a 'D' pinpoint while customers are the remaining pinpoints.

Concluding Remarks

Vehicle Routing is a very common problem faced by many companies today. As such, countless models have been proposed in the Operations Research literature for more than 50 years. Despite the extensive research effort, many of these theoretical approaches have had little impact on real applications. Unrealistic assumptions or implementation difficulties may explain this lack of coordination. In this article we try to shorten this gap. We present a real VRP of a distribution company in the Northeast of Spain. The company distributes products daily to around 400 stores. One of the main differences of this application with respect to other VRP studies is the presence of a heterogeneous fleet of vehicles, in which some are allowed to perform multiple trips on a single day.

We use a novel approach for solving a real-life VRP. The algorithm used is based on a randomized version of the popular savings heuristic that contains three local search processes that improve the final solution. One of the advantages of this method is its easy implementation with no complex fine-tuning required. This makes it very suitable for companies. The results we obtained reduced the company distribution costs significantly with little computational effort, as solutions were obtained in just a few seconds.

Acknowledgements

This work has been partially supported by the Spanish Ministry of Science and Innovation (ECO2009-11307, TRA2010-21644-C03) and the CYTED-HAROSA Network (CYTED2010-511RT0419), in the context of the IN3-ICSO program (http://dpcs.uoc.edu) and the Business Analytics Research Group (http://www.upf.edu/barg/). We acknowledge the personnel of the distribution company for their cooperation in this project.

About the Authors

Alex Grasas is an Assistant Professor of Operations Management in the Department of Economics and Business at Universitat Pompeu Fabra, Barcelona. He holds an MS and a PhD in Industrial and Systems Engineering from University of Florida (USA) and a BS degree in Industrial Engineering from Universitat Politècnica de Catalunya (Spain). He has been involved in different research projects and consulting for firms in the area of Operations Management. He is also a Researcher of the Business Analytics Research Group at Universitat Pompeu Fabra. His research interests include retail operations management, product assortment planning and vehicle routing.

Jose Caceres-Cruz is a PhD student of the Information and Knowledge Society Doctoral Programme in the Open University of Catalonia. He holds a Specialization on Information Systems from the Central University of Venezuela and a Computer Engineering degree from the Simón Bolívar University. His main research interests are related to optimization of vehicle routing on realistic scenarios and simulation solutions.

Helena R. Lourenço is an Associate Professor at the Economics and Business Department at the Universitat Pompeu Fabra, Barcelona, Spain. She has a Master in Statistics and Operations Research from the University of Lisbon, Portugal, and a PhD in Operations Research from Cornell University, New York, USA. She has participated in several research projects and consulting for public and private companies in the telecommunications, healthcare, transportation, logistics and production area. She has published several articles in prestigious international scientific journals and has presented her work at international congresses and conferences. Helena teaches at various Masters and Postgraduate courses in Spain and Portugal. She is currently a Researcher at the Economics Healthcare Research Centre and Business Analytics Research Group at Universitat Pompeu Fabra and also at the Centre for Operational Research at the University of Lisbon.

Angel A. Juan is an Associate Professor of Applied Optimization and Simulation in the Computer Science Department at the IN3-Open University of Catalonia (Barcelona, Spain). Dr Juan holds a PhD in Applied Computational Mathematics, an MS in Information Systems & Technology and an MS in Applied Mathematics. His research interests include Applied Optimization and Randomized Algorithms. He has published over 100 peer-reviewed papers in international journals, books and proceedings. Currently, he is the Coordinator of the CYTED-IN3-HAROSA@IB Network and Editorial Board Member of both the Int. J. of Data Analysis Techniques and Strategies and the Int. J. of Information Systems & Social Change.

Mercè Roca i Puigvert is a full time Professor and Academic Coordinator of Marketing and Research techniques at Escola Superior de Comerç Internacional (ESCI) from Universitat Pompeu Fabra, Barcelona, Spain. She has a Master degree in Economics and Management from the Economics and Business Department at the Universitat Pompeu Fabra and a PhD in Decision Sciences from Leeds University Business School. She has participated in several research and consulting projects for companies and public institutions related to international business and marketing research, acting as an expert in research techniques and statistical analysis. Her work has been published in prestigious international academic journals and presented at international congresses and conferences. She is the Academic Director of the Master of Sciences in International Business of Universitat Pompeu Fabra and is currently a member of the Research Group in International Economics and Management of ESCI, the Business Analytics Research Group (BARG) and the Experimental Economics Laboratory (LeeX) of Universitat Pompeu Fabra.

References

Baldacci, R., Battarra, M. and Vigo, D. (2008) Routing a heterogeneous fleet of vehicles. In: B. Golden, S. Raghavan and E.E. Wasil (eds.) *The Vehicle Routing Problem: Latest Advances and New Challenges*. New York: Springer, pp. 3–27.

Brandao, J. (2011) A tabu search algorithm for the heterogeneous fixed fleet vehicle routing problem. *Computers and Operations Research* 38(1): 140–151.

Clarke, G. and Wright, J. (1964) Scheduling of vehicles from a central depot to a number of delivering points. *Operations Research* 12(4): 568–581.

Croes, G. (1958) A method for solving traveling salesman problems. *Operations Research* 6(6): 791–812.

Dantzig, G. and Ramser, J. (1959) The truck dispatching problem. *Management Science* 6(1): 80–91.

Euchi, J. and Chabchoub, H. (2010) A hybrid tabu search to solve the heterogeneous fixed fleet vehicle routing problem. *Logistics Research* 2(1): 3–11.

Gendreau, M., Laporte, G., Musaraganyi, C. and Taillard, E.D. (1999) A tabu search heuristic for the heterogeneous fleet vehicle routing problem. *Computers & Operations Research* 26(12): 1153–1173.

Golden, B., Assad, A., Levy, L. and Gheysens, F. (1984) The fleet size and mix vehicle-routing problem. *Computers & Operations Research* 11(1): 49–66.

Juan, A.A., Faulin, J., Jorba, J., Caceres, J. and Marques, J. (2011a) Using parallel & distributed computing for solving real-time vehicle routing problems with stochastic demands. *Annals of Operations Research*, http://link.springer.com/article/10.1007/s10479-011-0918-z.

Juan, A.A., Faulin, J., Jorba, J., Riera, D., Masip, D. and Barrios, B. (2011b) On the use of Monte Carlo simulation, cache and splitting techniques to improve the Clarke and Wright savings heuristics. *Journal of the Operational Research Society* 62(6): 1085–1097.

Juan, A.A., Faulin, J., Ruiz, R., Barrios, B. and Caballé, S. (2010) The SR-GCWS hybrid algorithm for solving the capacitated vehicle routing problem. *Applied Soft Computing* 10(1): 215–224.

Laporte, G. (2009) Fifty years of vehicle routing. *Transportation Science* 43(4): 408–416.

Li, F., Golden, B. and Wasil, E. (2007) A record-to-record travel algorithm for solving the heterogeneous fleet vehicle routing problem. *Computers and Operations Research* 34(9): 2734–2742.

Lourenço, H.R., Juan, A.A., Caceres-Cruz, J. and Grasas, A. (2013) A savings-based randomized heuristic for the heterogeneous fleet multi-trip vehicle routing problem. Working paper, Universitat Pompeu Fabra.

Osman, I.H. and Salhi, S. (1996) Local search strategies for the mix fleet routing problem. In: V. J. Rayward-Smith, I. H. Osman, C. R. Reeves and G. D. Smith (eds.). *Modern Heuristics Search Methods*. Chichester, UK: Wiley, pp. 131–153.

Prins, C. (2002) Efficient heuristics for the heterogeneous fleet multitrip VRP with applications to a large-scale real case. *Journal of Mathematical Modelling and Algorithms* 1(2): 135–150.

Prins, C. (2009) Two memetic algorithms for heterogeneous fleet vehicle routing problems. *Engineering Applications of Artificial Intelligence* 22(6): 916–928.

Rand, G.K. (2009) The life and times of the savings method for vehicle routing problems. *Orion* 25(2): 125–145.

Tarantilis, C.D., Kiranoudis, C.T. and Vassiliadis, V.S. (2004) A threshold accepting metaheuristic for the heterogeneous fixed fleet vehicle routing problem. *European Journal of Operational Research* 152(1): 148–158.

Toth, P. and Vigo, D. (eds.) (2002) *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Applications. Philadelphia, PA: SIAM.

