Dynamic Multi-trip Vehicle Routing with Unusual Time-windows for the Pick-up of Blood Samples and Delivery of Medical Material

Nicolas Zufferey[1], Byung Yun Cho[2] and Rémy Glardon[2]
[1]*Geneva School of Economics and Management, GSEM - University of Geneva, Uni-Mail, 1211 Geneva 4, Switzerland*
[2]*LGPP, École Polytechnique Fédérale Lausanne, Lausanne, Switzerland*

Keywords: Dynamic Vehicle Routing, Diversion, Time-window, Heuristics.

Abstract: Given a fleet of identical vehicles and a set of n clients to be served from a single depot, the well-known vehicle routing problem (*VRP*) consists in serving each client (with a deterministic demand) once with a unique vehicle, with the aim of minimizing the total traveled distance. In this work, the basic *VRP* is extended within a medical environment, leading to *MVRP* (for medical *VRP*). Indeed, the depot is typically a laboratory for blood analysis, and a client is assumed to be a medical location at which blood samples should be picked up by a vehicle. In order to have efficient tests at the laboratory, at most 90 minutes should elapse between the release time of the blood sample and the delivery time at the laboratory. In addition, only a proportion of the demand is known in advance and the travel times depend on the traffic conditions. A fleet of non-identical vehicle is considered (with different speeds and capacities), and each location has to be visited anytime a blood sample is available. Finally, medical items should be daily delivered from the laboratory to some medical locations. A transportation cost function with three components has to be minimized. Solution methods are proposed, which are able to account for all the specific features of the problem. The experiments highlight the benefit of considering diversion opportunities (which consists in diverting a vehicle away from its planned destinations).

1 INTRODUCTION TO THE PROBLEM

This study is motivated by a real situation encountered in the city of Geneva (Switzerland). When the analysis of blood samples is required at the considered hospital, the samples are sent to an external laboratory (denoted *LABO*, which cannot be named because of a non-disclosure agreement). In order to preserve the quality of the samples, it is very important to deliver them to *LABO* as soon as possible. More precisely, no more than 90 minutes should elapse between the availability of a sample and its delivery to *LABO*. The pick-up of the samples at different locations is ensured by the vehicles managed by *LABO*. These vehicles are continuously turning around the city in order to collect and deliver all blood samples (Grasas et al., 2014). In contrast with the broad existing literature on the vehicle routing problem (*VRP*), the combination of the following features makes the considered problem new. It is denoted *MVRP* (for medical *VRP*), for which the planning horizon is a day (from 8am to 6pm).

- Time-windows: the deliveries (blood sample and medical items) have to be performed before their associated deadlines (Ciavotta et al., 2009).
- Non-identical vehicles: two fleet of vehicles are available, namely cars and scooters. In the considered city, a scooter is on average slightly faster than a car, but its capacity is lower.
- Stochastic demand: even if the static requests can be considered before the beginning of the day, there are stochastic requests all along the day.
- Dynamic planning: the travel times depend on the traffic conditions, and diversion (i.e., diverting a vehicle away from its planned route) is allowed.
- Multi-trip with pick-up/delivery: a location has to be visited when a blood sample is available. This leads to the situation where each vehicle is allowed to come back to depot as many times as decided.

The objective function f to minimize contains three types of costs (denoted f_1, f_2 and f_3), considered in a lexicographic order (i.e., a higher level objective is infinitely more important than a lower level one). Note that a lexicographic optimization is often

Zufferey, N., Cho, B. and Glardon, R.
Dynamic Multi-trip Vehicle Routing with Unusual Time-windows for the Pick-up of Blood Samples and Delivery of Medical Material.
DOI: 10.5220/0005733303660372
In *Proceedings of 5th the International Conference on Operations Research and Enterprise Systems (ICORES 2016)*, pages 366-372
ISBN: 978-989-758-171-7

used in practice (Respen et al., pted; Solnon et al., 2008; Thevenin et al., 2015; Zufferey et al., 2006).

1. Taxi cost (f_1): each request which cannot be served on-time by *LABO* is treated by a taxi (or a courier service) at a high cost (involving a fixed cost as well as a variable component depending on the traveled time).
2. Car cost (f_2), formulated as the number of used cars. Obviously, if a request is served with a scooter, the cost is lower than if it is served by a car.
3. Travel cost (f_3), computed as the total travel time of all the vehicles. This cost is proportional to the fuel consumption.

Because of the difficulty of the problem, heuristic or metaheuristic approaches have to be used (see (Gendreau and Potvin, 2010; Zufferey, 2012a) for general references on such solution methods). A main contribution of this work is the design of efficient solutions methods able to account for all the specific components of the *MVRP*. Relying on (Cho, 2015), this study is organized as follows. The *VRP* and some of its extensions, with the associated literature review, are discussed in Section 2. The employed methodology is presented in Section 3. Experiments are discussed in 4. They showed the benefit of allowing diversion opportunities in the proposed solution method. Finally, a conclusion ends up the paper in Section 5.

2 THE *VRP* AND SOME EXTENSIONS

As depicted in (Zufferey, 2012a), modern methods for solving complex optimization problems are often divided into exact methods and metaheuristic methods. An exact method guarantees that an optimal solution will be obtained in a finite amount of time. Among the exact methods are branch-and-bound, dynamic programming, Lagrangian relaxation based methods, and linear and integer programming based methods (Nemhauser and Wolsey, 1988). However, for a large number of applications and most real-life optimization problems, which are typically NP-hard (Garey and Johnson, 1979), such methods need a prohibitive amount of time to find an optimal solution. For these difficult problems, it is preferable to quickly find a satisfying solution. If solution quality is not a dominant concern, then a simple heuristic can be employed, while if quality occupies a more critical role, then a more advanced metaheuristic procedure is warranted. There are mainly two classes of metaheuristics: local search and population based methods. The former type of algorithm works on a single solution (e.g., descent local search, simulated annealing, tabu search, and variable neighborhood search), while the latter makes a population of solutions evolve (e.g., genetic algorithms, scatter search, ant colonies, adaptive memory algorithms). At each iteration of a local search, a *neighbor* solution is generated from the *current* solution by performing a modification on the current solution, called a *move*. The reader interested in a recent book on metaheuristics is referred to (Gendreau and Potvin, 2010).

As exposed in (Zufferey et al., 2015), the *VRP* is one of the most popular problems in combinatorial optimization because of its obvious applications in transportation. It consists in designing the route of each of the k identical vehicles with the aim of minimizing the total traveled distance f (or the total cost or the total travel time). All vehicles are initially in a depot, where each route starts and ends. Each client v (with a known demand) has to be visited once by the collection of routes. The problem is defined in an undirected graph $G = (V,E)$, where $V = \{v_0, v_1, \ldots, v_n\}$ is the vertex set and $E = \{(v_i, v_j) \mid v_i, v_j \in V, i < j\}$ is the edge set. Note that v_0 is the depot and the other vertices are clients. The following lexicographical approach is generally used: minimize k, then the total distance f. The two most well-known constraints associated with the VRP are: (1) *capacity*: each vehicle has a limited capacity Q, thus the demand of each route cannot exceed Q; (2) *autonomy*: each vehicle has a limited autonomy A, thus the total duration of each route cannot exceed A.

For survey papers on the *VRP*, the reader is referred to (Cordeau et al., 2005; Cordeau et al., 2002; Cordeau and Laporte, 2004; Gendreau et al., 2002; Golden et al., 1998; Laporte and Semet, 2002). Many algorithms have been developed for the *VRP*. Among them, there are some successful classical heuristics such as Clarke & Wright, Two-matching, Sweep, 1-Petal and 2-Petal, as tested in (Cordeau et al., 2002). However, the best performance is achieved by metaheuristics (e.g. (Cordeau et al., 2001; Mester and Braysy, 2007; Nagata and Braysy, 2009; Rochat and Taillard, 1995; Toth and Vigo, 2003; Vidal et al., 2014)). Relying on (Zufferey et al., 2015), such competitive metaheuristics are discussed below.

- Adaptive Memory (*AM*). *AM* (Rochat and Taillard, 1995) has been proved to be a good algorithm for the *VRP* and introduces a very innovative approach. At each generation of *AM*, an offspring solution s is built route by route from a central memory M (which contains routes), then s is improved with a local search, and the resulting solution is used to update M (i.e. routes of M are replaced with routes of s).

- Unified Tabu Search (*UTS*). *UTS* (Cordeau et al., 2001) has been proved to be a very flexible algorithm (easily adapted to variations of the *VRP*) with competitive quality and speed. *UTS* relies on a tabu search using an objective function which dynamically penalizes the constraint violations (the penalty component is likely to be increased if the last iterations violate the constraints).
- Granular Tabu Search (*GTS*). *GTS* (Toth and Vigo, 2003) has proved to be a very balanced algorithm in terms of speed and quality. It uses a tabu search framework and relies on the use of *granular* neighborhoods to discard the edges that rarely would belong to a competitive solution. *GTS* uses a granularity threshold which is dynamically adjusted.
- Active Guided Evolution Strategies (*AGES*). *AGES* (Mester and Braysy, 2007) has been proved to be very efficient (it is one of the best *VRP* method), with a reasonable speed. *AGES* is a combination of several procedures (including local search techniques), but an important drawback is its significant number of parameters.
- Edge Assembly-based Memetic Algorithm (*EAMA*). *EAMA* (Nagata and Braysy, 2009) combines an edge-assembly crossover with well-known local search procedures.
- Unified Solution Framework for Multi-Attribute *VRP* (*USFMA*). *USFMA* (Vidal et al., 2014) is able to tackle a wide range of *VRP* variants. Using a diversity management process, the proposed method is a hybrid genetic algorithm relying on problem-independent local search and genetic operators.

Several extensions of the *VRP* can be found in the literature (e.g., time-windows, pick-up and delivery, multi-trips). As indicated in (Lorini et al., 2011; Respen et al., 2014), the recent developments observed in communication facilities have led to the consideration of dynamic vehicle routing problems where new customer requests must be inserted in the currently planned routes as soon as they occur (Gendreau and Potvin, 1998; Psaraftis, 1995). A good survey about methodologies for solving different types of dynamic vehicle routing problems can be found in (Ichoua et al., 2000). The travel times can also be time-dependent to account for rush hours (Fleischmann et al., 2004; Horn, 2000; Ichoua et al., 2003; Kaufman and Smith, 1993).

As detailed in (Respen et al., 2014), dynamic vehicle routing, where a part of the information about the customers to be visited is not known in advance, is attracting a growing attention from transportation companies (see, for example, (Gendreau and Potvin, 1998; Psaraftis, 1995)). An interesting survey on this topic can be found in (Pillac et al., 2013), where the problem is first discussed, applications are reviewed, and solution methods are presented. In particular, tabu search led to impressive results on different dynamic vehicle routing problems. Nowadays, new possibilities offered by localization devices such as global positioning systems (GPS) can be exploited to improve vehicle routing management. In (Potvin et al., 2006), the authors are interested in a vehicle routing problem with time windows and dynamic travel times. The travel times include three different components: long-term forecasts, such as those based on long-term trends (time-dependency), short-term forecasts, where the travel time on a link is modified with a random uniform value to account for any new information available when a vehicle is ready to depart from its current location, and dynamic perturbations, which represent any unforeseen events that might occur while traveling on a link (e.g., accident causing sudden congestion). A modification to a planned route is only possible when the vehicle is at a customer location. That is, a planned route cannot be reconsidered while a vehicle is traveling on a link. An extension of this model is proposed in (Lorini et al., 2011), where the position of each vehicle can be obtained when a vehicle reaches some lateness tolerance limit or when a new customer request occurs. Based on this information, the planned route of each vehicle is reconsidered, including the possibility of diversion (i.e., redirecting a vehicle en route to its current destination). The results show that the setting of an appropriate lateness tolerance limit can provide substantial improvements. In line with (Respen et al., 2014), in the second proposed heuristic of this study, we present a further extension by assuming that the position of each vehicle is known at all time, thanks to accurate GPS devices. This assumption allows the system to react appropriately.

3 METHODOLOGY

In this section, additional information is first given, and the proposed heuristics are then presented. Because the planning horizon is a day (from 8am to 6pm) and the concerned territory is a rather small city, the autonomy constraint (i.e., the refueling of vehicles) can be ignored, because fuel can be provided to the vehicles before or after the planning horizon.

3.1 Additional Assumptions

The graph G is built based on the Geneva network, which allows to deduce the distance matrix D. Two types S (for scooters) and C (for cars) of vehicles are considered. Each vehicle has its own speed coefficient ρ (in km/h), and it is assumed that $\rho_S = 1.1 \cdot \rho_C$, where ρ_S (resp. ρ_C) is the speed coefficient for scooters (resp. cars). Indeed, a scooter generally moves faster than a car within a city environment. The standard travel time $\hat{t}_{ij}$ between two locations i and j is computed as D_{ij}/ρ (where $\rho \in \{\rho_S, \rho_C\}$). A coefficient T is used to model the traffic situation: $T = 1$ in standard conditions, $T > 1$ if there is traffic congestion (e.g., beginning and end of the day), and $T < 1$ if we have low traffic conditions (e.g., second part of the morning, first part of the afternoon). The actual travel time t_{ij} between locations i and j is computed as $\hat{t}_{ij} \cdot T$. Based on the situation of the city of Geneva, two coefficients are considered: low traffic condition between 10am and 3pm (with $T = 0.5$) and high traffic condition for the remaining part of the planning horizon (with $T = 1.25$).

Two types of request are considered: (1) the blood samples represent 80% of the demand; (2) freight (i.e., the delivery of medical material from the depot) constitutes 20% of the demand. The previous requests are all static (i.e., their release times are known before the planning horizon), whereas 40% of the blood requests are dynamic (i.e., they appear randomly during the day). Each request j is associated with: a unique location, a volume q_j, a release time r_j (at the customer location), and a deadline d_j (at the *LABO* depot). In other words, a time window $[r_j, d_j]$ is associated with each request j, where in contrast with the classical *VRP* literature, d_j is not a due date (or time) at location j, but a deadline at the depot. Note that formally, a due date can be exceeded but penalized, whereas a deadline cannot be exceeded. For each request j, its release time r_j is generated based on a uniform distribution during the whole day (but not within the 90 last minutes of the day for the blood request). The deadline d_j of any blood sample is assumed to be 90 minutes after its release time (i.e., $d_j = r_j + 90$), whereas the deadline for the freight is randomly generated with a uniform distribution in interval $\min(r_j + 60; d_0)$, where d_0 is the closing time of *LABO* (i.e., 6pm).

3.2 Heuristics

Because of the dynamic nature of *MVRP*, quick reactions have to be taken during the day. For this reason, sophisticated metaheuristics cannot be employed, thus a straightforward but fast solution method is designed.

The proposed heuristic relies on a fixed fleet $F = (S, C)$, where S is the set of scooters and C the set of cars. In a first phase, before the beginning of the planning horizon, a *static solution* is generated with a greedy insertion procedure *GR*, followed by a descent local search method *DLS* based on the well-known CROSS-exchanges (as in (Lorini et al., 2011)). *GR* starts from scratch and at each step, it inserts a request j to a route R (of a scooter or a car) in order to minimize the augmentation of f_2 (ties are broken with f_3, and then it tries to balance the request load over all the vehicles), while satisfying the capacity and the deadline constraints. If it is not possible to find a feasible insertion, a taxi is used for request j (i.e., the value of f_1 augments). Because the static solution will be significantly modified with the occurrence of random events (e.g., dynamic travel times, stochastic requests), there is no need to use more advanced methods to generate it.

In order to obtain a *full solution*, a discrete event simulator (e.g., (Silver and Zufferey, 2005)) is needed to generate the random events occurring all along the day. The used time bucket is a minute. Anytime a request appears, it is greedily assigned to a route (as in *GR*), and *DLS* is then directly performed. There is no need to discuss the computing time of the proposed overall method, because the insertion of a new request only requires a small fraction of a second.

In the basic version $H1$ of the heuristic, each insertion can only be performed after the current destination of each vehicle, whereas in the enhanced version $H2$ of the heuristic, it is allowed to divert a vehicle away from its current destination. Of course, $H2$ is only possible if there is an information system allowing an efficient communication between the vehicles and the dispatching office, as described in (Lorini et al., 2011). Giving an opportunity to divert a vehicle away from its initially planned route is a significant advantage: it gives more flexibility to the planner. Note that the insertion procedure (involving *GR* and *DLS*) works with *expected* travel times, which are different from the *actual* (i.e., simulated) travel times. In such a context, the use of diversion actions is more than relevant.

Let $s^\star$ be the best solution encountered during the search process, and let $f^\star = (f_1^\star, f_2^\star, f_3^\star)$ be its associated simulated values. We have now all the ingredients to formulate a generalizable approach in Algorithm 1, where the stopping condition can be the non-reduction of $f^\star$ at the end of the main loop.

Algorithm 1: General solution method for *MVRP*.

Initialization

1. Choose an initial fleet $F = (S,C)$ of vehicles (e.g., the one associated with the involved company).
2. Set $f^\star = (f_1^\star, f_2^\star, f_3^\star) = (+\infty, +\infty, +\infty)$.

While a stopping condition is not met, do

1. Set s_F as the *empty solution* (i.e., it does not contain any request).
2. Before the planning horizon, generate a *static solution* s_F while considering only the known static requests. For this purpose, *GR* and *DLS* are sequentially used to insert the requests one by one.
3. Use the discrete event simulator to extend s_F as a *full solution*. Anytime a request appears, insert it greedily to the current solution s_F (i.e., to a route) with *GR*, followed by *DLS*. In the variant $H1$ of the method, each insertion can only be performed after the current destination of a vehicle, whereas in the variant $H2$, diversion is allowed.
4. Update $s^\star$: if $f(s_F) < f^\star$ (under the lexicographic approach), set $s^\star = s_F$ and $f^\star = f(s_F)$.
5. Modify the fleet $F = (S,C)$ by augmenting or reducing S or C (but not both) by one unit (it is forbidden to consider again an already investigated fleet).

Return $s^\star$ with value $f^\star = (f_1^\star, f_2^\star, f_3^\star)$.

4 EXPERIMENTS

In addition to the above provided information, the following data is also assumed to be given.

- Instance size: 300 static requests, 200 dynamic requests, 25 material requests, $n = 20$ locations.
- The distances (in km) between the medical locations belong to $[5, 30]$.
- *LABO* fleet: 15 cars and 10 scooters, car capacity = 900 (liters), scooter capacity = 60 (liters).
- Volume q_j (integer) of a blood sample request: uniformly generated in interval $[1, 10]$ (liters).
- Volume q_j of a medical material request: uniformly generated in set $\{10, 20, 30, 40, 50, 60\}$ (liters).
- For each location, a blood request j is generated every t minutes, where t is an integer uniformly generated in interval $[23, 33]$ (if j is static) or in interval $[40, 50]$ (if j is dynamic).
- The car base speed (i.e., ρ_C) is fixed to 17km/h (based on the provided practical information).

An important issue is to determine the fleet $F = (S,C)$ of vehicles. One can deduce that if F is too small (resp. C too large, S too large), f_1 will be too high (resp. f_2 too high, f_3 too high). Initially, F is defined as the current situation encountered by *LABO* (i.e., 15 cars and 10 scooters). For each fleet F, 10 runs of the heuristic (either $H1$ or $H2$) on 10 scenarios are performed, and average costs (i.e., over 100 experiments) are computed for each component (f_1, f_2 and f_3). Then, other fleets are tested by adding a single vehicle (either a scooter or a car), until bad results (i.e., solutions with high costs) are obtained. It was first observed that the initial fleet is significantly understaffed. As a goal of *LABO* consists in assigning at most 10% of the total number of requests to the taxis, then a fleet of 18 scooters and 12 cars was found to be the most efficient. We have performed various sensitivity analysis, which are summarized as follows.

- With the initial fleet of vehicles, if ρ_C increases from 17km/h to 27km/h, the number of taxi requests averagely decreases from 19.5% to 7.3%. The most sensitive interval for ρ_C is $[19, 22]$, corresponding to taxi requests in [18.8%, 11.2%].
- With 15 cars, the average number of requests assigned to taxis decreases from 84 to 62 if the number of scooters increases from 10 to 24. The most sensitive interval for the number of scooters is $[10, 15]$, corresponding to taxi requests in [84, 66].
- With 18 scooters, the average number of requests assigned to taxis decreases from 79 to 62 if the number of cars increases from 8 to 18. The most sensitive interval for the number of cars is $[8, 12]$, corresponding to taxi requests in [79, 64].

We have observed the benefit of the proposed ingredients introduced in the heuristics. Firstly, the use of the CROSS-exchanges (Taillard et al., 1997) leads to a reduction of the total travel time (i.e., f_3) by roughly 15%. More precisely, the total travel time averagely decreases from 114.8 hours to 101.9 hours. Secondly, considering $T \in \{0.5, 1.25\}$ for two distinct time periods (versus $T = 1$ for the whole day) leads to an augmentation of 22% on the number of assigned requests to F. It means that no attention should be paid to standard traffic conditions (i.e., with $T = 1$) when designing and tuning a solution method, because standard traffic conditions are far away from real conditions. Thirdly, the use of diversion opportunities has an important impact, as $H2$ is able to roughly assign the double of dynamic requests when compared to $H1$.

5 CONCLUSION

In this paper, we have introduced *MVRP*, a new problem relying on the well-known vehicle routing problem. In *MVRP*, blood samples have to be picked-up (when available) at some medical locations and then delivered on-time (in order to preserve the quality of the blood samples) to the depot, which is a laboratory denoted *LABO*. The planning horizon is a day. Two fleets of vehicles are managed by *LABO*: cars and scooters. If *LABO* is not able to assign a request to one of its vehicle, it can call an external taxi to treat the request (but at a higher cost). For *LABO*, the involved transportation functions to minimize are the taxi costs, the number of employed cars, and the total traveled distance of its vehicles. Because of the dynamic nature of the problem (indeed, the demand is stochastic and the travel times depend on the traffic), a quick solution method has to be employed.

The performance of a solution method can be evaluated according to several criteria (Zufferey, 2012a). The most relevant criteria are presented below.

- Quality: value of the obtained results, according to a given objective function.
- Speed: time needed to get good results.
- Robustness: sensitivity to variations in problem characteristics and data quality.
- Ease of adaptation of the method to a problem, because, as mentioned in (Woolsey and Swanson, 1975), ”people would rather live with a problem they cannot solve than accept a solution they cannot understand”.
- Possibility to incorporate properties of the problem. It is admitted that an efficient metaheuristic should incorporate knowledge from the considered problem (Grefenstette, 1987).

The second heuristic, able to divert away a vehicle from its current destination, seems to perform well according to all the above criteria. Future works might include: the consideration of maintenance constraints with an extended planning horizon (Hertz et al., 2009), the use of other constructive algorithms with a learning process (Zufferey, 2012b), and the development of exact methods (e.g., based on linear programming) to benchmark the heuristics on deterministic cases.

REFERENCES

Cho, B. Y. (2015). Management and Optimization of Vehicle Routing Problem with Medical Challenge. Master’s thesis, École Polytechnique Fédérale de Lausanne, Switzerland.

Ciavotta, M., Meloni, C., and Pranzo, M. (2009). Scheduling dispensing and counting in secondary pharmaceutical manufacturing. *AIChE Journal*, 55 (5):1161 – 1170.

Cordeau, J.-F., Gendreau, M., Hertz, A., Laporte, G., and Sormany, J.-S. (2005). *Logistics Systems: Design and Optimization*, chapter New Heuristics for the Vehicle Routing Problem, pages 270–297. Springer.

Cordeau, J.-F., Gendreau, M., Laporte, G., Potvin, J.-Y., and Semet, F. (2002). A Guide to Vehicle Routing Heuristics. *Journal of the Operational Research Society*, 53 (5):512–522.

Cordeau, J.-F. and Laporte, G. (2004). *Metaheuristic Optimization via Memory and Evolution: Tabu Search and Scatter Search*, chapter Tabu search heuristics for the vehicle routing problem, pages 145–163. Kluwer, Boston.

Cordeau, J.-F., Laporte, G., and Mercier, A. (2001). A Unified Tabu Search Heuristic for Vehicle Routing Problems with Time Windows. *Journal of the Operational Research Society*, 52:928–936.

Fleischmann, B., Grietz, M., and Gnutzmann, S. (2004). Time Varying Travel Times in Vehicle Routing. *Transportation Science*, 38:160–173.

Garey, M. and Johnson, D. (1979). *Computer and Intractability: a Guide to the Theory of NP-Completeness*. Freeman, San Francisco.

Gendreau, M., Laporte, G., and Potvin, J.-Y. (2002). *The Vehicle Routing Problem*, chapter Metaheuristics for the VRP, pages 129–154. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia.

Gendreau, M. and Potvin, J.-Y. (1998). *Fleet Management and Logistics*, chapter Dynamic Vehicle Routing and Dispatching, pages 115 – 126. Kluwer.

Gendreau, M. and Potvin, J.-Y. (2010). *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research & Management Science*. Springer.

Golden, B. L., Wasil, E. A., Kelly, J. P., and Chao, I.-M. (1998). *Fleet Management and Logistics*, chapter Metaheuristics in vehicle routing, pages 33–56. Kluwer, Boston.

Grasas, A., Ramalhinho, H., Pessoa, L., Resende, M., Caballé, I., and N. Barba N. (), . (2014). On the Improvement of Blood Sample Collection at Clinical Laboratories. *BMC Health Services Research*, 14 (12).

Grefenstette, J. (1987). *Genetic Algorithms and Simulated Annealing*, chapter Incorporating Problem Specific Knowledge into Genetic Algorithms, pages 42–60. Morgan Kaufmann Publishers.

Hertz, A., Schindl, D., and Zufferey, N. (2009). A solution method for a car fleet management problem with maintenance constraints. *Journal of Heuristics*, 15 (5):425 – 450.

Horn, M. (2000). Efficient Modelling of Travel in Networks with Time-Varying Link Speeds. *Networks*, 36:80 – 90.

Ichoua, S., Gendreau, M., and Potvin, J.-Y. (2000). Diversion Issues in Real-Time Vehicle Dispatching. *Transportation Science*, 34 (4):426 – 438.

Ichoua, S., Gendreau, M., and Potvin, J.-Y. (2003). Vehicle Dispatching with Time-Dependent Travel Times. *European Journal of Operational Research*, 144:379 – 396.

Kaufman, D. and Smith, R. (1993). Fastest Paths in Time-Dependent Networks for Intelligent Vehicle-Highway Systems Application. *Intelligent Vehicle-Highway Systems Journal*, 1:1 – 11.

Laporte, G. and Semet, F. (2002). *The Vehicle Routing Problem*, chapter Classical heuristics for the capacitated VRP, pages 109–128. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia.

Lorini, S., Potvin, J.-Y., and Zufferey, N. (2011). Online vehicle routing and scheduling with dynamic travel times. *Computers & Operations Research*, 38:1086 – 1090.

Mester, D. and Braysy, O. (2007). Active-guided evolution strategies for large-scale capacitated vehicle routing problems. *Computers & Operations Research*, 34 (10):2964 – 2975.

Nagata, Y. and Braysy, O. (2009). Edge assembly-based memetic algorithm for the capacitated vehicle routing problem. *Networks*, 54 (4):205 – 215.

Nemhauser, G. and Wolsey, L. (1988). *Integer and Combinatorial Optimization*. John Wiley & Sons.

Pillac, V., Gendreau, M., Guéret, C., and Medaglia, A. L. (2013). A Review of Dynamic Vehicle Routing Problems. *European Journal of Operational Research*, 225:1 – 11.

Potvin, J.-Y., Xu, Y., and Benyahia, I. (2006). Vehicle Routing and Scheduling with Dynamic Travel Times. *Computers & Operations Research*, 33:1129 – 1137.

Psaraftis, H. (1995). Dynamic Vehicle Routing: Status and Prospects. *Annals of Operations Research*, 61:143 – 164.

Respen, J., Zufferey, N., and Amaldi, E. (2015 (accepted)). Metaheuristics for a job scheduling problem with smoothing costs relevant for the car industry. *Networks*.

Respen, J., Zufferey, N., and Potvin, J.-Y. (2014). Online vehicle routing and scheduling with continuous vehicle tracking. In *Proceedings of the 15th Annual Congress of the French Operations Research Society, ROADEF 2014*, Bordeaux, France.

Rochat, Y. and Taillard, E. (1995). Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics*, 1:147–167.

Silver, E. A. and Zufferey, N. (2005). Inventory control of raw materials under stochastic and seasonal lead times. *International Journal of Production Research*, 43:5161–5179.

Solnon, C., Cung, V., Nguyen, A., and Artigues, C. (2008). The car sequencing problem: Overview of state-of-the-art methods and industrial case-study of the ROADEF 2005 challenge problem. *European Journal of Operational Research*, 191 (3):912 – 927.

Taillard, E., Badeau, P., Gendreau, M., Guertin, F., and Potvin, J.-Y. (1997). A Tabu Search Heuristic for the Vehicle Routing Problem with Soft Time Windows. *Transportation Science*, 31:170 – 186.

Thevenin, S., Zufferey, N., and Widmer, M. (2015). Metaheuristics for a scheduling problem with rejection and tardiness penalties. *Journal of Scheduling*, 18 (1):89 – 105.

Toth, P. and Vigo, D. (2003). The Granular Tabu Search and Its Application to the Vehicle-Routing Problem. *INFORMS Journal on Computing*, 15 (4):333 – 346.

Vidal, T., Crainic, T. G., Gendreau, M., and Prins, C. (2014). A unified solution framework for multi-attribute vehicle routing problems. *European Journal of Operational Research*, 234:658 – 673.

Woolsey, R. and Swanson, H. S. (1975). *Operations Research for Immediate Applications*. Harper and Row.

Zufferey, N. (2012a). Metaheuristics: some Principles for an Efficient Design. *Computer Technology and Applications*, 3 (6):446 – 462.

Zufferey, N. (2012b). Optimization by ant algorithms: Possible roles for an individual ant. *Optimization Letters*, 6 (5):963 – 973.

Zufferey, N., Farres, J., and Glardon, R. (2015). Ant metaheuristics with adapted personalities for the vehicle routing problem. *Lecture Notes in Computer Science*, 9335:1 – 13.

Zufferey, N., Studer, M., and Silver, E. A. (2006). Tabu search for a car sequencing problem. In *Proceedings of the 19th International Florida Artificial Intelligence Research Society Conference*, pages 457 – 462, Melbourne, USA, May 11 – 13.