Vehicle Routing: Bridging the gap between theory and practice

Alex Van Breedam

Faculty of Applied Economics
University of Antwerp - RUCA
Middelheimlaan 1
B-2020 Antwerp (Belgium)
E-mail: alexvb@ruca.ua.ac.be

Abstract

The success of computerized vehicle routing systems depends, for the major part, on the variety and the complexity of the real-world side-constraints that can be handled. This paper gives an overview of some side-constraints important to practitioners and presents some ways to consider these side-constraints in solution methods for complex vehicle routing problems.

1 Introduction

The Vehicle Routing Problem (VRP) can be defined as the problem of finding a set of routes for a fleet of vehicles which have to service a number of stops. Vehicles depart from and arrive at a depot. In the standard VRP all vehicles are assumed to be homogeneous with respect to their capacity, whereby the demand quantity of each stop is deterministic and no single demand exceeds the vehicle capacity. However, the VRP can be extended with various side-constraints, such as mixed pick-ups (backhauls) and deliveries (linehauls), hard and soft time-windows, route duration constraints etc.

The complexity of the side-constraints dealt with in the literature to a large extent determines the applicability of academically developed methods and concepts to real-world vehicle routing and scheduling problems.

The aim of this article is to give an overview of some types of side-constraints important to industry. The nature of each type is described along with some ways to implement it in solution methods. Moreover, some relevant references to literature are provided.

The first section of this paper contains some general guidelines on how to solve the VRP. In following sections the standard VRP is extended with three major categories of side-constraints: customer-related, vehicle-related and depot-related side-constraints.

2 The Vehicle Routing Problem

The VRP was first formulated by Dantzig and Ramser (1959). Lenstra and Rinnooy Kan (1981) showed that the VRP with or without side-constraints is an NP-hard combinatorial problem. Hence, exact algorithms are only useful for tiny problems. For real life problems, heuristics are much more appropriate (from a commercial point of view). Heuristics generate suboptimal but good solutions in a computing time which is proportional to a low-order power of the number of stops.

Basically, two types of heuristics can be distinguished: initial and improvement heuristics. Initial heuristics generate a feasible solution to the VRP, given the data on customers, depot, vehicles and side-constraints.

An initial solution to the VRP can be obtained with a wide variety of heuristics. Principally, three main groups of initial heuristics can be distinguished: route-construction heuristics, two-phase heuristics and heuristics based on exact algorithms.

Route-building heuristics construct routes by iteratively adding stops to the routes, which can be built sequentially or simultaneously. Criteria used for assigning stops to routes can be based on the nearest neighbor principle (Tyagi (1968), Baker and Schaffer (1986), Solomon (1987), Balakrishnan (1993)), on the savings principle (Clarke and Wright (1964),Gaskell (1967), Yellow (1970), Paessens (1988), Webb (1972), Golden (1977), Nelson et al. (1988), Bodin (1983), Van Landeghem (1988), Knowles (1967), Tillman and Cochran (1968), Holmes and Parker (1976), McDonald (1972), Buxey (1979)), on the generalized savings principle (Altinkemer and Gavish (1991), Desrochers and Verhoog (1989)) or on the insertion and selection principle (Mole and Jameson (1976),Baker and Schaffer (1986), Solomon (1987), Savelsbergh (1990b)).

Two-phase methods belong to the second category of initial heuristics for the VRP. These heuristics generate a solution to the VRP in two distinct stages. The generalized assignment heuristic (Fisher and Jaikumar (1981), Nygard et al. (1988), Baker (1992), Koskodis et al. (1992)), the route first-cluster second heuristic (Beasley (1983)), the sweep heuristic (Gillet and Miller (1974)) and the two-phase heuristic (Christofides et al. (1979), Potvin and Rousseau (1993)) can be considered as two-phase methods.

The last category of initial heuristics are the heuristics based on exact algorithms, like the incomplete tree-search heuristic (Christofides et al. (1979)).

An initial feasible solution to the VRP can be enhanced through the application of an improvement heuristic. These procedures try to improve a feasible solution by relocating and/or exchanging stops within or between routes.

Within route improvement heuristics are typical for the TSP and are based on the k-opt and Or-opt procedures (Croes (1958), Lin (1965), Lin and Kernighan (1973), Or (1976)). These heuristics are used to optimize the route-sequence of each route of the VRP solution separately. The reader is referred to Solomon et al. (1988) and Savelsbergh (1990a) for applications of within route improvement heuristics to the VRP with side-constraints.

Between routes improvement methods are especially designed for the VRP. These procedures try to improve an initial feasible solution by moving stops between routes.

In the case of improvement heuristics, and particularly in that of between-routes improvement, a distinction can be made between local and global optimization methods.

The local optimization heuristic, the traditional "descent" method, finds a local minimum by performing only moves of stops which result in the improvement of the objective function value. As a result, local optimization heuristics are trapped in the local optimum in which they descend.

Global optimization heuristics, on the contrary, succeed in leaving the local optimum by temporarily accepting moves which cause a worsening of the objective function value. These heuristics are often called "metaheuristics" because the procedure used to generate a new solution out of the current one is embedded in a heuristic which determines the search strategy. Popular metaheuristics are Genetic Algorithms, Simulated Annealing and Tabu Search. For applications of Simulated Annealing to the VRP, the reader is referred to Robusté et al. (1990), Alfa et al. (1991), Teodorovic and Pavkovic (1992), Osman (1993), Thangiah et al. (1994a), Van Breedam (1995), Janssens and Van Breedam (1995) and Van Breedam (1994). With respect to the application of Tabu Search to the VRP, the following publications are to be noticed: Pureza and França (1991), Taillard (1992), Gendreau et al. (1992), Stewart et al. (1992), Osman (1993), Semet and Taillard (1993), Thangiah et al. (1994c) and Van Breedam (1994) and Rochat and Semet (1994).

In recent years, several successful applications in industry of computerized vehicle routing systems have been reported. The reader is referred to Fisher et al. (1982), Bell et al. (1983), Evans and Norback (1985), Golden and Wasil (1987), Semet and Taillard (1993) and Rochat and Semet (1994) for some examples.

The success of most of the above-mentioned examples stems from the aptitude of the computerized vehicle planning system to cope with some specific real-world side-constraints. Some of the side-constraints, of which we feel that they are important to industry, will be considered in the next sections.

3 Customer-related side-constraints

These types of constraints deal with the nature of the demand-point, i.e. the customer. Different types of side-constraints are considered.

The first type includes the combination of linehauls and backhauls. Instead of having only drop-off or pick-up points, many real-life problems have both pick-up and delivery stops.

Other specific situations occur when the quantity to be delivered can be split over more than one route.

Some publications have been dedicated to the period routing problem. There is some industrial interest for this topic because a delivery schedule for several days or weeks can be developed. This is particularly useful for medium-term planning.

Most research attention has been directed at time-windows at the customer site because it remains a hot topic for practical applications.

3.1 Mixed pick-up and delivery

The Vehicle Routing Problem with Backhauls (VRPB) occurs frequently in some branches of industry. Examples are the distribution of beverages where some stops require a delivery and others a pick-up, the distribution of grocery industries where goods are picked-up at the suppliers' and delivered at the supermarkets.

We will assume that there is a depot and that the destination of all picked-up goods is always the depot. Dial-a-ride problems with precedence relationships, whereby each customer has a pick-up and delivery location lies beyond the scope of this paper.

Two different approaches can be considered. The first states that no deliveries are allowed after pick-ups. The second approach does not impose this sequence, that is to say, pick-ups and deliveries can be mixed as long as the route remains feasible.

As far as the first approach is concerned, a number of solution procedures are proposed.

The first to attempt to solve the VRPB were Deif and Bodin (1984), who adapted the Clarke and Wright savings heuristic by including a term that penalizes, and hence delays, the linkage of a backhaul to a delivery.

Goetschalckx and Jacobs-Blecha (1986) propose a two-phased solution methodology, composed of the generation of an initial feasible solution, which is improved in the second phase. The initial heuristic is based on the principle of spacefilling curves. Pick-up and delivery points are first projected on a line, then clustered, and subsequently routed by visiting each point in each cluster according to their position along the line. The feasible solution obtained with this heuristic is improved by means of within-route 2-opt and 3-opt improvement heuristics.

An optimal solution for small size VRPB can be obtained with the algorithm of Yano et al. (1987). This procedure is a two-step process: in the first step, a feasible solution is generated by means of an initial route-building heuristic; in the second, a set-covering is solved with branch-and-bound. Experiments were performed on problems with up to 40 pick-up and delivery points. The number of stops per route rarely exceeded four. Hence, computing time required to find an optimal solution remains acceptable. However, this algorithm is practically unusable for larger real-world problems with even more complex side-constraints.

Instead of providing specific features for including backhauls, it is also possible (and easier) to verify the backhaul constraint at the feasibility check of a route. This idea is supported by Thangiah et al. (1994b), who recommend a backhaul check at each insertion. Their solution methodology is a sequential insertion heuristic followed by a number of local search improvement heuristics.

With respect to the second approach, where backhaul points can be sequenced anywhere in a route, the following references are to be mentioned. Golden and Stewart (1985) proposed a two-phased concept. In a first stage, all deliveries are routed using an initial heuristic. Subsequently, an insertion procedure is used to insert the pick-up points into the existing routes. The insertion criterion is provided with a penalty term in order to delay pick-ups towards the end of the route. Obviously, this solution methodology is only appropriate when the number of delivery points greatly exceeds the number of pick-up points.

Casco et al. (1988) improved the above-mentioned heuristic of Golden and Stewart (1985) in a number of ways. First, instead of discarding the pick-up points from the initial solution, these points are directly linked with a two-way route to the depot. All deliveries are planned with a route-building heuristic, in which the vehicles are not filled to their full capacity. The spare capacity is useful for later insertion of backhauls. The criterion for inserting backhaul points into delivery routes during the second phase of the procedure is not only provided with a penalty to force the insertion of the backhaul towards the end of the route, but also with one which takes account of the delivery load after the pick-up. The rationale is that insertions of backhauls are appropriate if the load of the remaining deliveries in the route is small. From a practical point of view, this relieves the driver of excessively shuffling the remaining load on his truck before loading the pick-up load. This residual load-shuffling capacity can be added as an input parameter to the planning heuristic.

Van Breedam (1994) analyzed the effect of mixed pick-up and delivery on the behavior of eleven initial heuristics (8 route-construction and 3 two-phase heuristics). Backhaul constraints were taken into account at the feasibility check of a route. Backhauls were allowed to be inserted anywhere in the route. His most important conclusion was that the route-sequence imposed by combining pick-ups and deliveries tends to fade the differences between the solution of the initial heuristics and the solution when only capacity constraints are considered.

Ultimately, customers may require simultaneous delivery and pick-up services. Min (1989) was confronted with a problem of this type for the distribution of library material for a public library center and its 22 branches. The author developed a three-phase cluster first-route second heuristic to solve this problem. In a first phase, stops are clustered according to their geographical proximity using the average linkage method, which is a hierarchical clustering method. Subsequently, trucks with drivers are assigned to the clusters. The third phase consists of determining the route-sequence of a truck within its cluster by solving a TSP with mixed loadings using an adapted branch-and-bound process.

In the margin, we can mention the work of Hall (1991) which is directed at creating and assessing spatial models for the VRPB with multiple terminals and uniformly distributed stops. The proposed principals of optimal district shape and the optimal traversal of a district can be helpful for the selection of seed points used by some types of heuristics for initializing routes.

3.2 Split delivery

Sometimes the demand of a stop may be fulfilled by more than one vehicle. This can occur for instance when the demand required by a customer is larger than the capacity of a vehicle, when a customer can be serviced by more than one depot, or when it is less costly to service a customer more than once.

A solution methodology for this problem was proposed by Dror and Trudeau (1990). In their methodology, an initial solution is generated by means of a variant of the savings heuristic. Subsequently, improvement heuristics based on one-node and two-node swaps between routes are performed. After each route-construction, a 2-opt within route improvement is executed. This procedure yields a good VRP solution. Thereafter, two specific split-delivery routines are launched. The k-split interchange tries to split the demand of a stop between k-vehicles. The second routine is a route-addition heuristic which considers the elimination of split deliveries by adding a route. The four above-mentioned improvement routines all work with a best improvement strategy, i.e. all feasible moves are evaluated and only the best is selected.

Some important observations result from split delivery problems. Stops with a split demand are usually the ones with above-average load. Splits resulting in high distance savings usually take place at great distances from the depot. Important savings in distance and vehicles as compared to traditional vehicle routing solutions can be realized for problems with stops of which the demand exceeds 10% of the vehicle's capacity.

This approach can easily be adapted to problems where the demand of some stops exceeds the vehicle's capacity.

We found no other publications on the general aspects of split delivery. In this field, an important research topic could be the combination of split delivery and multiple depot, whereby each depot supplies different products to the customer who requires a mix of products.

3.3 Period routing

Quite a lot of real-world planning situation requires a weekly schedule in addition to the daily planning, for instance for fuel oil and industrial gas distribution. The Period Vehicle Routing Problem (PVRP) is the problem of finding routes for all days of a given T-day period. The number of customer visits per week is lower than or equal to the number of days (T) of the period. These types of problems are also called allocation/routing problems. The allocation part consists of the assignment of customers to days of the period, while the routing part governs for the daily planning.

Ball (1988) gives an overview of some application environments and some solution methodologies for the PVRP. Most solution methodologies are based on the cluster-first route-second principle.

Dror and Levy (1986) proposed the assignment of customers to days of the period and subsequently to solve a VRP for each day. The generalized assignment problem for assigning customers to days is solved heuristically with a linear programming based approach. The VRP for each day is solved in two stages. In the first stage, a feasible solution is obtained with an initial heuristic. This solution is further improved by moving stops between routes and between days. The authors felt that this improvement routine was sufficient to compensate for the absence of a component which takes account of the geographic dispersion of the stops.

Geographic considerations are taken into account by the approaches of Christofides and Beasley (1984) and Russell and Igo (1979). The latter associate customers with a single allowable delivery combination with their appropriate days. In the case of Christofides and Beasley (1984), a set of possible centers is given, and the optimization process makes a choice. In both approaches, each day of the week becomes associated with a customer representing the cluster center. The remaining customers are assigned to the days of the week in order of decreasing service frequency. The assignment cost of customers to days is given by the distance of the customer to the center. This approach tends to associate geographic areas with days of the period. Like with seed points, different associations of centers to days of the period can be tried. An interchange heuristic is used to evaluate better assignments. Once a solution for this p-median problem has been generated, a VRP must be solved for each day. Again, a two-stage solution process is proposed, consisting of an initial solution and an improvement heuristic.

A comparable approach was proposed by Dror and Ball (1987) and Trudeau and Dror (1992) but has been adapted to handle stochastic demands in the case of inventory routing. Inventory routing implies that each customer maintains a local inventory. By means of a real life case of the distribution of heating oil, some problems like the number of stockouts and route failures are addressed.

Tan and Beasley (1984) adapted the generalized assignment heuristic of Fisher and Jaikumar (1981) for the VRP to the PVRP. In a first step, a number KT (K vehicles, T-days period) of seed points is generated using the seed point generation method of Fisher and Jaikumar (1981). The second step consists of assigning K seeds to each day of the period by solving a linear assignment problem. The assignment cost of a stop to a seed corresponds to the minimum extra distance travelled when inserting a stop in the route from the depot to the seed. In a third step, the seed point generation method is executed with the stops of each day in order to assign stops to vehicles. This results in a new set of KT seed points and the entire procedure can be repeated from step two. Once the linear assignment problem is solved in step two, an initial heuristic combined with an improvement heuristic are used to solve the VRP for each day.

The approach of Gaudioso and Paletta (1992) is somewhat different, in so far that stops are assigned to delivery combinations one at a time. Stops are arranged in decreasing order of their delivery frequency. Within a same delivery sequence, stops are ranked in decreasing order of their demand. The first stop of this list is selected for allocation to a delivery combination. The feasible delivery combinations for each customer are ordered lexicographically with respect to possible days of service. If there exists a delivery combination for which the assignment of the customer does not increase the vehicle fleet size, then that combination is selected. Otherwise, any other delivery combination is chosen. This approach is especially designed for a VRP whose objective is to minimize the number of vehicles. Moreover, vehicles are allowed to complete more than one route a day.

3.4 Time-windows

The Vehicle Routing Problem with time-windows (VRPTW) has received a lot of attention in the literature. This is probably due to the wide applicability of these types of side constraints in real-world cases.

Basically, two types of time-windows can be distinguished: hard and soft time-windows. Hard time-windows restrict the delivery time at the customer's site. Violation of this window implies that the customer cannot be serviced. In the case of soft time-windows, violation induces a penalty but the customer can be serviced anyway. Hard time-windows most often correspond with the opening and closing times of a customer, for instance 9.00 am to 5.00 pm. Examples of soft time-windows are the lunch break between 12.30 pm and 2.00 pm or the time-interval preferred by the customer for being serviced.

Much work in vehicle routing has centred on time-windows. Almost all types of heuristics have been provided to handle hard time-windows easily. The addition of a simple time-window feasibility check to the other route feasibility checks are a minimal requirement. The criteria of some heuristics which consider time-windows explicitly are more elaborated. For a sequential nearest neighbor heuristic, Baker and Schaffer (1986), Solomon (1987), Balakrishnan (1993) and Van Breedam (1994) developed a time-window oriented nearest neighbor criterion. Van Landeghem (1988) extended the savings criteria of the parallel savings heuristic with a time-oriented part in order to take account of time-windows. The most interesting time-windows extensions can be made by means of the insertion and selection criteria of a sequential insertion heuristic. This is demonstrated by the work of Baker and Schaffer (1986), Solomon (1987) and Van Breedam (1994). Potvin and Rousseau (1993) and Van Breedam (1994) used some of these criteria in a parallel version of the insertion heuristic.

A problem associated with the generalized assignment heuristic of Fisher and Jaikumar (1981) is that the clustering and routing phases are separated. Consequently, time aspects are only considered during the routing phase, which can result in unfeasible routes in the case of difficult time-windows. Koskodis et al. (1992) proposed treating the time-windows as being soft in a first stage. Through subsequent iterations of the generalized assignment heuristic using an updated cost matrix, the time-windows violation penalties are minimized. If no penalty remains, then the soft time-windows solution is a feasible hard time-windows solution.

Some work has also been done on the speeding up of the feasibility checks for time-windows. This is particularly valuable for within and between routes improvement procedures. Savelsbergh (1988), Solomon et al. (1988) and Baker and Schaffer (1986) presented some work in this field.

Comparative studies on the effect of hard time-windows on the behavior of heuristics are those by Solomon (1987) and Van Breedam (1994). Both authors experienced the quality of the time-window oriented sequential insertion heuristic as compared to a number of other heuristics. An important conclusion of Van Breedam (1994) was that non-sequential heuristics using seed points for initializing routes give bad solutions in the case of binding time-windows.

Very little research has been done on the problem of soft time-windows for the VRP (VRPSTW). Min (1991) proposes a mixed-integer goal programming model with a problem-size reduction procedure to solve a real-world case of moderate size.

Balakrishnan (1993) presents a nearest neighbor, a savings and a space-time heuristic for the VRPSTW. The main objective was to show that by allowing violation of certain customer time-windows, considerable savings in total costs could be made, considering that soft time-windows are maybe more realistic than hard time-windows for real-world applications.

4 Vehicle-related side-constraints

Less attention has been directed at vehicle-related side-constraints than to customer-related side-constraints. Nevertheless, a heterogeneous fleet of vehicles, most often in accordance with site-dependencies, are minimal requirements for practical applications.

Another side-constraint is the maximal route-time for the truck-driver, most often according to specific industry regulations.

The complexity of the VRP is increased when multicompartment vehicles are considered. Such vehicles are used in some important branches of industry.

Finally, the attention is focused on time-dependent travel time. The importance of this topic grows with the increasing traffic saturation of highways and metropolitan areas.

4.1 Heterogeneous fleet and site-dependencies

Most real-world cases are confronted with heterogeneous vehicles with respect to their capacity. When a company has to decide whether to lease or buy vehicles, it is faced with the fleet size and mix vehicle routing problem (FSMVRP): a VRP in which the composition and the size of the fleet has to be determined so as to minimize the sum of fixed costs and routing costs.

This problem was extensively treated by Golden et al. (1984) and Gheysens et al. (1984). They proposed some adaptations of the savings heuristic in order to take advantage of the different vehicle types. The savings formula is extended to consider, in addition to savings in routing costs, savings in fixed costs of a vehicle and opportunity savings. Opportunity savings can be defined in various ways: proportional to the unused capacity of the vehicle servicing the resulting combined tour of the savings heuristic; the cost of the smallest vehicle that can service the entire unused capacity of the vehicle assigned to the combined tour. Other versions of the savings heuristic use opportunity savings to encourage the use of a larger vehicle when it is profitable to do so, i.e. when a combined tour requires a larger vehicle than that of the two tours to be combined.

The authors also presented a route first-cluster second heuristic, whereby first a giant tour with all customers is built with a TSP-heuristic. This tour is then partitioned by finding the shortest path through a set of nodes with an aggregated demand exceeding the largest vehicle capacity. The solution obtained with this heuristic is further improved with a within-route TSP-heuristic.

Finally, the authors propose an adaptation of the generalized assignment heuristic of Fisher and Jaikumar (1981). First, a good vehicle mix is obtained using a lower bound procedure. Subsequently, as many seeds as there are vehicles are generated. The first seed corresponds to the customers farthest from the depot, while each next seed is the farthest from the set of previously selected seeds and the depot. The largest vehicle is assigned to the farthest seed, as larger routes are expected to satisfy a larger demand. Stops are assigned to seeds by solving a generalized assignment problem. A TSP is solved within each group of stops assigned to a seed.

Results showed that the performance of the adapted savings heuristic was rather poor as compared to the route first-cluster second and the adapted generalized assignment heuristic.

The problem of site-dependencies is even more complex because certain customers cannot be serviced by certain types of vehicles. This problem was addressed by Nag et al. (1988). Basically, these authors considered three approaches to the problem, based on the generalized assignment heuristic of Fisher and Jaikumar (1981).

The input of the first approach is the assignment of stops to vehicles obtained with a sweep-like heuristic. The sweep heuristic performs a sweep procedure for each vehicle type sequentially, starting with the smallest type. An artificial load is associated with each vehicle during the sweep procedure in order to get a more balanced allocation. Routes are then formed within each group obtained by the sweep procedure. If some routes are not feasible, the customers which can be serviced by vehicles of a larger type are removed. Their inclusion is postponed until the sweep procedure for larger vehicle types is implemented. The allocation obtained with this process is then used for the first generalized assignment heuristic. For each vehicle type, seeds are selected and a generalized assignment problem is solved. A TSP-heuristic is used to sequence each route.

The second approach uses only the solution obtained by the sweep-like procedure to determine seed points for each type separately. Subsequently, one generalized assignment problem is solved for the entire fleet of vehicles. To determine the route-sequence within each group, a TSP-heuristic is used.

The third approach involves a seed selection for each type of vehicle, starting with the largest type. Subsequently, customers which coincide with seeds are eliminated and the seed selection is performed for the next largest vehicle type. Next, one generalized assignment problem is solved and a TSP-heuristic is executed to determine the route-sequence.

Experiments showed that approaches two and three outperformed approach one.

4.2 Multicompartment vehicles

Multicompartment vehicles are typical of some branches of industry, for instance the simultaneous delivery of different crude oil derivatives.

Basically, it may be sufficient to extend the check for the capacity constraint to multiple capacity constraints per vehicle. However, in order to take advantage of possible delivery combinations, specific criteria for stop selection must be provided. Route-construction heuristics will be much more easily adaptable to the multicompartment principle than two-phase heuristics.

We are not aware of publications reporting a general approach in order to cope with multicompartment vehicles.

4.3 Maximal route-time

Maximal route-time is frequently imposed on truck-drivers in accordance with regulations.

These types of side-constraints can be easily added to the feasibility check embedded in heuristics of the route-construction type.

As far as heuristics with a separate clustering and routing phase (e.g. the generalized assignment heuristic of Fisher and Jaikumar (1981)) are concerned, the problem is much more complicated. During the clustering phase, one cannot consider temporal aspects because no route-sequence can be determined at that time. Hence, the resulting route in a cluster can be infeasible with respect to the route-time.

Two ways can be considered to cope with this problem. The first is to discard a stop of the infeasible route as long as the route remains infeasible. This stop is reinserted in another existing route or, if not possible, in a new route. Different criteria may be used for choosing the stop to be discarded: the stop with the smallest demand, the stop closest to the depot, etc...

The other approach consists of making the procedure iterative. This means that if a solution of the general assignment heuristic is unfeasible, the heuristic is restarted with alternative seed locations. An attempt is made to move the seeds in response to the magnitude of the violation of the maximal route-time. This process is repeated until a feasible solution is found or a number of meaningful seed locations have been tried out. A comparable procedure was proposed by Nygard et al. (1988) for the deadline VRP, which is a VRP with stops having a one-sided time-window, i.e. a latest allowable arrival time.

4.4 Time-dependent travel speed

The interest for time-dependent travel speed for vehicles has grown proportionally with the increasing traffic congestion problems. In addition, the importance of time-dependent travel speed is largely dependent on the scale of the VRP. The smaller the scale on which to perform the routing, the more important it is to obtain an accurate planning. Discarding time-dependent travel speed in metropolitan areas can result in an underestimation of the total routing time. For nationwide routing, the bias is more limited.

The time-dependent VRP (TDVRP) is a VRP for which the travel time between two nodes depends on the distance between the points and the time of the day. It is to be noticed that the triangle inequality does not hold anymore for the TDVRP.

The most obvious approach consists of making the cost of each link between two stops dependent on the time of the day. This approach was proposed by Malandraki and Daskin (1992). Heuristics which can easily be adapted for this concept are the ones where stops can only be added at the end of the route in construction. The authors present an sequential and a parallel nearest neighbor heuristic which satisfy this principle.

The adaptation of heuristics where stops are not only added at the end of a route, for instance insertion and improvement heuristics, results in a substantial increase in computation time of the heuristic. To test route feasibility, the entire route next to the inserted stop have to be recomputed, including the time-dependent costs which may have changed due to the changing time period. Ahn and Shin (1991) partially bypassed this problem and conceived some efficient feasibility checks for insertion and within-route improvement heuristics, based on the monotonicity property of the arrival time function.

The main disadvantage of a time-dependent inter-node cost is its enormous number of parameters to be estimated and the prohibitive data storage requirements, especially when the number of time-periods is high. For these reasons, Hill and Benton (1992) proposed a more practical framework to handle time-dependent travel speeds. Instead of using a time-varying inter-node cost, a time-dependent travel speed is assigned to every location. The speed between two locations is then the average of the speeds assigned to each location. The speed at the destination location is that of the time-period corresponding with the estimated arrival time of the vehicle. It is even possible to provide different speeds when arriving at and departing from the location. By using this approach, the amount of data to be stored is limited. Moreover, most initial heuristics can easily be adapted to this approach by using a forward-scheduling process.

5 Depot-related constraints

A major topic with respect to depot-related constraints is the case of multiple depots.

5.1 Multiple depots

A lot of companies have more than one depot from which vehicles are scheduled. The difficulty in handling a multi-depot problem depends on the degree of interaction between the depots. If there is no interaction, the problem is easily solved: assign stops to each depot and solve a VRP for each depot.

This strategy was used by Gillett and Johnson (1976). First, customers are assigned to the depots and subsequently a VRP is solved for each depot using the sweep heuristic.

Tillman and Cain (1972) extended the savings approach to the multiple depot problem. For this purpose, the savings formula was slightly adapted. Golden et al. (1977) further elaborated and refined this savings heuristic in order to handle large problems.

Much more difficult to solve are multi-depot problems with interactions between the depots. Examples are cases where trucks leave one depot to make deliveries at different stops and pass by an other depot for replenishment before continuing their deliveries. Other applications require that stops must be delivered from different depots, depending on the type of commodity that is required.

It is clear that a lot of different interaction types can be distinguished in practice. In these situations, a case-sensitive tackling of the problem is required. The large variety of interaction types is probably responsible for the shortage of general publications on this subject.

6 Conclusions

In this article we have tried to give an overview of a set of side-constraints which we feel are important to industrial applications. Moreover, their complexity and the way they can be handled by solution methodologies largely determine the applicability of computerized vehicle routing systems for real-world applications.

On the one hand, a lot of research has already focused on some types of side-constraints, e.g. time-windows. On the other hand, side-constraints like multiple depots, multicompartment vehicles and site-dependencies have attracted very little attention.

The analyst system builder has a powerful set of tools at his disposal to conceive elaborated computerized vehicle routing systems. Nevertheless, a lot of time is lost to provide specific procedure to handle side-constraints for which no or insufficient theoretical framework exists in the literature. Hence, further research is necessary, especially for this category of side-constraints.

References

Ahn, B.-H. and J.-Y. Shin (1991). Vehicle-routeing with time windows and time-varying congestion. *Journal of the Operational Research Society 42*(5), 393–400.

Alfa, A., S. Heragu, and M. Chen (1991). A 3-opt based simulated annealing algorithm for vehicle routing problems. *Computers and Industrial Engineering 21*, 635–639.

Altinkemer, K. and B. Gavish (1991). Parallel savings based heuristic for the delivery problem. *Operations Research 39*(3), 456–469.

Baker, B. (1992). Further improvements to vehicle routeing heuristics. *Journal of the Operational Research Society 43*(10), 1009–1012.

Baker, E. and J. Schaffer (1986). Solution improvement heuristics for the vehicle routing and scheduling problem with time-window constraints. *American Journal of Mathematical and Management Sciences 6*(3-4), 261–300.

Balakrishnan, N. (1993). Simple heuristics for the vehicle routeing problem with soft time-windows. *Journal of the Operational Research Society 44*(3), 279–287.

Ball, M. O. (1988). Allocation/routing: models and algorithms. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 199–221. Elsevier Science Publishers.

Beasley, J. (1983). Route first - cluster second methods for the vehicle routing. *Omega 11*(4), 403–408.

Bell, W., L. Dalberto, M. Fisher, A. Greenfield, R. Jaikumar, P. Kedia, R. Mack, and P. Prutzman (1983). Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces 13*, 4–23.

Bodin, L. (1983). Solving large vehicle routing and scheduling problems in small core. In *Proceedings of the ACM '83 Conference*, pp. 27–37.

Buxey, G. (1979). The vehicle scheduling problem and monte carlo simulation. *Journal of the Operational Research Society 30*, 563–573.

Casco, D., B. Golden, and E. Wasil (1988). Vehicle routing with backhauls: models, algorithms and case studies. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 127–147. North-Holland.

Christofides, N. and J. Beasley (1984). The period routing problem. *Networks 14*, 237–256.

Christofides, N., A. Mingozzi, P. Toth, and C. Sandi (1979). *Combinatorial Optimization*. Wiley, Chicester.

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

Croes, A. (1958). A method for solving traveling salesman problems. *Operations Research 5*, 791–812.

Dantzig, G. and J. Ramser (1959). The truck dispatching problem. *Management Science 6*, 81–91.

Deif, I. and L. Bodin (1984). Extension of the clarke and wright algorithm for solving the vehicle routing problem with backhauling. In A. Kidder (Ed.), *Proceedings of the Babson Conference on Software Uses in Transportation and Logistics Management*, pp. 75–96.

Desrochers, M. and T. Verhoog (1989). A matching-based savings algorithm for the vehicle routing problem. Technical report, Cahier du Gerad G-89-04, Ecole des Hautes Etudes Commerciales de Montréal.

Dror, M. and M. Ball (1987). Inventory routing: reduction from an annual to short period problem. *Naval Research Logistics 34*, 891–905.

Dror, M. and L. Levy (1986). A vehicle routing improvement algorithm comparison of a greedy and a matching implementation for inventory routing. *Computers and Operations Research 13*(1), 33–45.

Dror, M. and P. Trudeau (1990). Split delivery routing. *Naval Research Logistics 37*, 383–402.

Evans, S. and J. Norback (1985). The impact of a decision support system for vehicle routing in a food service supply situation. *Journal of the Operational Research Society 36*, 467–472.

Fisher, M., A. Greenfield, R. Jaikumar, and J. Lester (1982). A computerized vehicle routing application. *Interfaces 12*(4), 42–52.

Fisher, M. and R. Jaikumar (1981). A generalized assignment heuristic for the vehicle routing problem. *Networks 11*, 109–124.

Gaskell, T. (1967). Bases for vehicle fleet scheduling. *Operational Research Quarterly 18*, 281–295.

Gaudioso, M. and G. Paletta (1992). A heuristic for the periodic vehicle routing problem. *Transportation Science 26*(2), 86–92.

Gendreau, M., A. Hertz, and G. Laporte (1992). A tabu search heuristic for the vehicle routing problem. Technical report, Working Paper ORWP 92/14, Département de Mathématiques, Ecole Polytechnique Fédérale de Lausanne, Suisse.

Gheysens, F., B. Golden, and A. Assad (1984). A comparison of techniques for solving the fleet size and mix vehicle routing problem. *OR Spektrum 6*, 207–216.

Gillet, B. and L. Miller (1974). A heuristic algorithm for the vehicle-dispatch problem. *Operations Research 22*, 340–349.

Gillett, B. E. and J. G. Johnson (1976). Multi-terminal vehicle dispatch algorithm. *Omega 4*(6), 711–718.

Goetschalckx, M. and C. Jacobs-Blecha (1986). The vehicle routing problem with backhauls. *European Journal of Operational Research 42*, 39–51.

Golden, B. (1977). Evaluating a sequential vehicle routing algorithm. *AIIE Transactions 9*, 204–208.

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

Golden, B., T. Magnanti, and H. Nguyen (1977). Implementing vehicle routing algorithms. *Networks 7*, 113–148.

Golden, B. and W. Stewart (1985). Empirical analysis of heuristics. In E. Lawler, J. Lenstra, A. Rinnooy Kahn, and D. Shmoys (Eds.), *The Traveling Salesman Problem. A Guided Tour of Combinatorial Optimization*, pp. 207–246. Wiley, Chicester.

Golden, B. and E. Wasil (1987). Computerized vehicle routing in the soft drink industry. *Operations Research 35*(1), 6–17.

Hall, R. W. (1991). Characteristics of multi-stop/multi-terminal delivery routes, with backhauls and unique items. *Transportation Research B 25B*(6), 391–403.

Hill, A. V. and W. Benton (1992). Modelling intra-city time-dependent travel speeds for vehicle scheduling problems. *Journal of the Operational Research Society 43*(4), 343–351.

Holmes, R. and R. Parker (1976). A vehicle scheduling procedure based upon savings and a solution perturbation scheme. *Operational Research Quarterly 27*(1), 83–92.

Janssens, G. and A. Van Breedam (1995). A simulated annealing post-processor for the vehicle routing problem. In V. Rayward-Smith (Ed.), *Applications of Modern Heuristics*, Chapter Case studies, pp. 175–191. Alfred Waller Ltd.

Knowles, K. (1967). The use of a heuristic tree-search algorithm for the vehicle routing and scheduling. In *Operational Research Conference, Exeter, England*.

Koskodis, Y., W. Powell, and M. Solomon (1992). An optimization-based heuristic for vehicle routing and scheduling with soft time-windows. *Transportation Science 26*, 69–85.

Lenstra, J. and A. Rinnooy Kan (1981). Complexity of vehicle routing and scheduling problems. *Networks 11*, 221–227.

Lin, S. (1965). Computer solutions to the traveling salesman problem. *Bell System Technical Journal 44*, 2245–2269.

Lin, S. and B. Kernighan (1973). An effective heuristic algorithm for the traveling salesman problem. *Operations Research 21*, 498–516.

Malandraki, C. and M. S. Daskin (1992). Time dependent vehicle routing problems: formulations, properties and heuristic algorithms. *Transportation Science 26*(3), 185–200.

McDonald, J. (1972). Vehicle scheduling - a case study. *Operational Research Quarterly 23*(4), 433–444.

Min, H. (1989). The multiple vehicle routing problem with simultaneous delivery and pick-up points. *Transportation Research A 23A*(5), 377–386.

Min, H. (1991). A multiobjective vehicle routing problem with soft time windows: the case of a public library distribution system. *Socio-Economic Planning Science 25*(3), 179–188.

Mole, R. and S. Jameson (1976). A sequential route-building algorithm employing a generalised savings criterion. *Operational Research Quarterly 27*, 503–511.

Nag, B., B. L. Golden, and A. A. Assad (1988). Vehicle routing with site dependencies. In B. L. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 149–159. Elsevier Science Publishers.

Nelson, M., K. Nygard, J. Griffin, and W. Shreve (1988). Implementation techniques for the vehicle routing problem. *Computers and Operations Research 12*(3), 273–283.

Nygard, K., P. Greenberg, W. Bolkan, and E. Swenson (1988). Generalized assignment methods for the deadline vehicle routing problem. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 107–126. North-Holland, Amsterdam.

Or, I. (1976). *Traveling salesman-type combinatorial optimization problems and their relation to the logistics of regional blood banking*. Ph. D. thesis, Northwestern University, Evanston, IL.

Osman, I. (1993). Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research 41*, 421–451.

Paessens, I. (1988). The savings algorithm for the vehicle routing problem. *European Journal of Operational Research 34*, 336–344.

Potvin, J.-Y. and J.-M. Rousseau (1993). A parallel route building algorithm for the vehicle routing and scheduling problem with time windows. *European Journal of Operational Research 66*, 331–340.

Pureza, V. and P. França (1991). Vehicle routing problems via tabu search metaheuristic. Technical report, Technical Report, CRT-747, Centre de Recherche sur les Transports, Université de Montréal, Montréal, Canada.

Robusté, F., C. Daganzo, and R. Souleyrette (1990). Implementing vehicle routing models. *Transportation Research B 24*, 263–286.

Rochat, Y. and F. Semet (1994). A tabu search approach for delivering pet food and flour in switzerland. *Journal of the Operational Research Society 45*(11), 1233–1246.

Russell, R. and W. Igo (1979). An assignment routing problem. *Networks 9*, 1–16.

Savelsbergh, M. (1988). *Computer Aided Routing*. Ph. D. thesis, Centrum voor Wiskunde en Informatica, Amsterdam.

Savelsbergh, M. (1990a). An efficient implementation of local search algorithms for constrained routing problems. *European Journal of Operational Research 47*, 75–85.

Savelsbergh, M. (1990b). A parallel insertion heuristic for vehicle routing with side-constraints. *Statistica Neerlandica 44*(3), 139–148.

Semet, F. and E. Taillard (1993). Solving real-life vehicle routing problems efficiently using taboo search. *Annals of Operations Research 41*, 469–488.

Solomon, M. (1987). Algorithms for the vehicle routing and scheduling problems with time-window constraints. *Operations Research 35*(2), 254–265.

Solomon, M., E. Baker, and J. Schaffer (1988). Vehicle routing and scheduling with time-window constraints: efficient implementations of solution improvement procedures. In B. Golden and A. Assad (Eds.), *Vehicle Routing: Methods and Studies*, pp. 85–105. North-Holland, Amsterdam.

Stewart, W., J. Kelly, and M. Laguna (1992). Solving vehicle routing problems using generalized assignments and tabu search. Technical report, College of Business Administration, University of Colorado, Boulder.

Taillard, E. (1992). Parallel iterative search methods for vehicle routing problems. Technical report, Working paper ORWP 92/03, Département de Mathématiques, Ecole Polytechnique Fédérale de Lausanne.

Tan, C. and J. Beasley (1984). A heuristic algorithm for the period vehicle routing problem. *Omega 12*(5), 497–504.

Teodorovic, D. and G. Pavkovic (1992). A simulated annealing technique approach to the vehicle routing problem in the case of stochastic demand. *Transportation Planning and Technology 16*, 261–273.

Thangiah, S., I. Osman, and T. Sun (1994c). Hybrid genetic algorithm, simulated annealing and tabu search methods for vehicle routing problem with time windows. Technical report, Technical Report SRU-CpSc-TR-94-27, Computer Science Department, Slippery Rock University, Slippery Rock, PA, U.S.A.

Thangiah, S., I. Osman, R. Vinayagamoorthy, and T. Sun (1994a). Algorithms for vehicle routing problems with time deadlines. *American Journal of Mathematical and Management Sciences 13*(3-4).

Thangiah, S., T. Sun, and J.-Y. Potvin (1994b). Heuristic approaches to vehicle routing with backhauls and time windows. Technical report, Technical Report SRU-CpSc-TR-94-28, Computer Science Department, Slippery Rock University, Slippery Rock, PA, U.S.A.

Tillman, F. and T. Cain (1972). An upperbound algorithm for the single and multiple terminal delivery problem. *Management Science 18*(11), 664–682.

Tillman, F. and H. Cochran (1968). A heuristic approach for solving the delivery problem. *Journal of Industrial Engineering 19*, 354–358.

Trudeau, P. and M. Dror (1992). Stochastic inventory routing: route design with stockouts and route failures. *Transportation Science 26*(3), 171–184.

Tyagi, M. (1968). A practical method for the truck dispatching problem. *Journal of the Operations Research Society of Japan 10*, 76–92.

Van Breedam, A. (1994). *An analysis of the behavior of heuristics for the vehicle routing problem for a selection of problems with vehicle-related, customer-related and time-related constraints*. Ph. D. thesis, University of Antwerp - RUCA, Belgium.

Van Breedam, A. (1995). Improvement heuristics for the vehicle routing problem based on simulated annealing. *European Journal of Operational Research 86*(3).

Van Landeghem, H. (1988). A bi-criteria heuristic for the vehicle routing problem with time-windows. *European Journal of Operational Research 36*, 217–226.

Webb, M. (1972). Relative performance of some sequential methods of planning multiple delivery journeys. *Operational Research Quarterly 23*(2), 361–372.

Yano, C., T. Chan, L. Richter, T. Cutler, K. Murty, and D. McGettigan (1987). Vehicle routing at quality stores. *Interfaces 17*(2), 52–63.

Yellow, P. (1970). A computational modification to the savings method of vehicle scheduling. *Operational Research Quarterly 21*(2), 281–283.