



CIRRELT

Centre interuniversitaire de recherche
sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

Evolutionary Algorithms for Vehicle Routing

Jean-Yves Potvin

November 2007

CIRRELT-2007-48

Bureaux de Montréal :

Université de Montréal
C.P. 6128, succ. Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :

Université Laval
Pavillon Palasis-Prince, local 2642
Québec (Québec)
Canada G1K 7P4
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Evolutionary Algorithms for Vehicle Routing

Jean-Yves Potvin^{1,*}

¹ Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), and Department of Computer Science and Operations Research, Université de Montréal, P.O. Box 6128, Station Centre-ville, Montréal, Canada H3C 3J7

Abstract. This paper is a survey of the literature on applications of evolutionary algorithms for vehicle routing problems. It reports about genetic algorithms, evolution strategies and particle swarm optimization, when applied to the classical capacitated vehicle routing problem and many of its variants. The performance of evolutionary algorithms is also compared with the best alternative problem-solving approaches on benchmark instances.

Keywords. Evolutionary algorithms, genetic algorithms, evolution strategies, particle swarm optimization, vehicle routing.

Acknowledgements. Financial support for this work was provided by the Natural Sciences and Engineering Research Council of Canada (NSERC). This support is gratefully acknowledged.

Results and views expressed in this publication are the sole responsibility of the authors and do not necessarily reflect those of CIRRELT.

Les résultats et opinions contenus dans cette publication ne reflètent pas nécessairement la position du CIRRELT et n'engagent pas sa responsabilité.

* Corresponding author: Jean-Yves.Potvin@cirrelt.ca

Dépôt légal – Bibliothèque nationale du Québec,
Bibliothèque nationale du Canada, 2007

© Copyright Potvin and CIRRELT, 2007

1 Introduction

This paper reviews applications of evolutionary algorithms (EAs) for solving vehicle routing problems. EAs are stochastic search methods that operate on a population of solutions by simulating, at a high level of abstraction, the evolution processes observed in nature. The survey excludes the work done on the Traveling Salesman Problem (TSP), a canonical combinatorial optimization problem that has been used as a testbed for many new methodologies, including EAs. We thus focus on extensions of the TSP, where a fleet of vehicles, starting from one or a number of depots, must visit a set of customers at minimum cost, subject to side constraints. A precise definition of the classical capacitated vehicle routing problem and its variants is given in Section 2.

This review is divided along methodological lines rather than vehicle routing types. It just happens that the same researchers have often addressed different variants of the vehicle routing problem with the same basic problem-solving methodology, slightly adapted or extended to address the specific characteristics of a given variant. It is thus easier to follow the work done by these researchers. Section 6, however, provides a classification by problem type.

We also need to mention that the reference section is divided into two parts. The first part corresponds to core references that are directly related with the topic of this survey. The second part contains auxiliary references that are mentioned for completeness, but are not directly related with the survey. We did not want to mix core references with auxiliary references to allow the reader to get a complete view of the work done in the vehicle routing domain with EAs, just by considering the core references.

Although this survey is the only one that provides an exhaustive account of applications of EAs to vehicle routing problems, the reader can find overviews, surveys and classification of metaheuristics for solving combinatorial optimization problems at large in [18, 24, 39, 126], and reviews of metaheuristics for vehicle routing problems in [22, 23, 35, 36, 37, 38, 40, 87].

2 The Vehicle Routing Problem

In the Vehicle Routing Problem (VRP), least-cost delivery or collection routes must be designed for a set of customers, using a fleet of vehicles located at a central depot, subject to side constraints. More formally, let $G = (V, A)$ be a graph with $V = \{1, \dots, n\}$ the vertex set. In this graph, vertex 1 stands for the depot, while the remaining vertices are customers. Also, with every arc $(i, j) \in A$, $i \neq j$, is associated a non negative distance matrix $C = (c_{ij})$ (in some contexts, c_{ij} represents a travel cost or travel time). Note that when C is symmetrical, the set A can be conveniently replaced by a set E of undirected edges. We also assume that a fleet of m vehicles is available at the depot to visit the customers, where m is fixed or free. In the latter case, the number of vehicles in a solution can vary between 1 and $n - 1$. All vehicles are identical and have the same finite capacity Q .

Usually, the number of vehicles is free and the objective is to minimize the total distance traveled by the vehicles to visit all customers. Sometimes, a fixed cost is associated with each vehicle and the objective is then a weighted sum of fixed costs and travel costs. The problem then consists in designing a set of optimal routes such that:

- each customer is visited exactly once by exactly one vehicle;
- all vehicle routes start and end at the depot;
- some side constraints are satisfied.

This problem is NP-Hard, as it is a generalization of the TSP. Accordingly, it has given rise to a large number of heuristic problem-solving methodologies, including metaheuristics, to address instances of reasonable size. The most common types of vehicle routing problems reported in the literature are:

- *Capacitated VRP*. In the capacitated VRP, a non negative demand d_i is associated with each customer $i \in V \setminus \{1\}$ and the total demand on a route should not exceed the vehicle capacity Q . In some time- or distance-constrained variants, additional constraints state that the length of a route should not exceed a given limit. Note that VRP will stand for the capacitated VRP, with or without route length limits, in the following.
- *VRP with time windows*. In the VRPTW, a vehicle must arrive at customer i within the time interval $[e_i, l_i]$. In most cases, it is possible for the vehicle to wait at customer i and thus to arrive before the lower bound e_i . Some variants also allow a vehicle to arrive late, thus after l_i , although a penalty is incurred in the objective.
- *VRP with time deadlines*. The VRPTD is a special case of the VRPTW, where the time windows are replaced by time deadlines. More precisely, there is no lower bound e_i at each customer i , only an upper bound l_i .
- *Time-dependent VRP*. The TVRP is aimed at modeling more realistic situations by considering travel times that depend both on the distance between two customers and the time of the day. For example, it takes longer to get from one customer to the next during rush hours. The TVRP takes these considerations into account.
- *Periodic VRP*. The PVRP is a variant of the VRP where the horizon extends over a number of periods. Routes are constructed for each period, and each customer is visited once or more over the horizon, depending on the customer requirements.
- *VRP with heterogeneous fleet*. The VRPHF is a variant of the VRP where the vehicles do not share the same characteristics (e.g., different capacities).
- *Multiple Depot VRP*. In the MDVRP, there are many depots and each vehicle can start or end its route from any of these depots.
- *VRP with backhauls*. In the VRPB, some customers require deliveries (linehauls) while others require pickups (backhauls). Each route is thus a mix of linehaul and backhaul customers, where the backhauls are typically visited after the linehauls. It has been quickly recognized that substantial cost savings can be achieved by allowing empty vehicles that return from deliveries to pick up inbound products. In the grocery industry, for example, supermarkets and grocery suppliers would be the linehaul and backhaul customers, respectively.
- *VRP with simultaneous delivery and pick-up*. In the VRPSDP, the same customer can ask both for goods to be delivered from the depot and for other goods to be picked up and brought to the depot.
- *Pick-up and delivery problem*. Each customer i in a PDP has both a pick-up location i^+ and a delivery location i^- . That is, the demand collected at i^+ must be delivered at i^- . Both the pick-up and delivery points of customer i must be in the same route (pairing constraint) and i^+ must be visited before i^- (precedence constraint). Different real-world applications, usually with time windows (PDPTW), are reported in the literature, like courier and dial-a-ride services.

We examine different adaptations of EAs for these vehicle routing problems. In the sections that follow, genetic algorithms (GA) and evolution strategies (ES) are reviewed. Particle swarm optimization (PSO), which is often related to EAs, is also briefly considered.

3 Genetic Algorithms

Genetic algorithms (GAs) have emerged from the work of Holland at the University of Michigan [150]. Starting from some initial population, the search mechanism of a simple GA is divided into four phases: (1) evaluation of each solution in the population, (2) selection of parent solutions, (3) application of crossover and mutation operators to parent solutions to generate offspring solutions and (4) replacement of the old population by the new population of offspring solutions. This process is repeated for a number of iterations or until the system does not improve anymore. This search mechanism is applied on solutions encoded as bit strings. Through this coding scheme, generic operators have been designed that manipulate bit strings without any knowledge of the corresponding solution (apart from the solution value). A simple GA can be sketched as follow:

1. Create an initial population of P solutions.
2. Evaluate each solution.
3. Repeat for a fixed number of iterations:
 - 3.1 Repeat until P offspring solutions are created:
 - 3.1.1 Select two parent solutions in the population (with replacement) using a randomized selection procedure based on the solution values.
 - 3.1.2 Apply crossover to the two parent solutions to create two offspring solutions.
 - 3.1.3 Apply mutation (with a certain probability) to each offspring.
 - 3.1.4 Insert the two offspring in the new population.
 - 3.2 Evaluate each offspring in the new population.
 - 3.3 Replace the old population by the new one.
4. Return the best solution found.

In Step 1, the initial population can be created randomly, but it is better to use construction heuristics to create at least a fraction of the population. Different techniques exist for selecting the parent solutions in Step 3.1.1. Typically, a randomized roulette-wheel selection assigns a selection probability to each solution which is proportional to its value. More recently, tournament selection procedures have been devised. Here a number of solutions are selected in the population and the best one becomes a parent. The procedure is repeated twice to get two parents. The crossover and mutation operators are applied in Steps 3.1.2 and 3.1.3. The classical one-point crossover creates two offspring by randomly selecting a cut point on the two parent strings and by exchanging the two substrings after the cut point. The mutation operator then processes each offspring position by position and switch the bit value at each position with a small probability. This secondary operator introduces small perturbations into the solutions and is often useful to maintain a diversity of solutions in the population.

When applied to combinatorial optimization problems in general, and to vehicle routing problems in particular, the simple GA described above exhibits a number of shortcomings.





Figure 1: One-point crossover

These shortcomings have been recognized early in the case of the TSP, a canonical combinatorial optimization problem which can be viewed as a simple vehicle routing problem. In this particular case, a single vehicle of infinite capacity, starting from any given vertex, visits all other vertices and returns to its starting point, with the objective of minimizing the total distance. The main difficulties encountered when a classical GA is applied to the TSP can be summarized as follows:

- The binary string representation is quite cumbersome. An integer string representation, where each integer stands for a customer index, is much more natural. In the literature, this is referred to as the path representation, which is simply the sequence of customer indices in a tour. In this representation, the last customer in the sequence is implicitly connected to the first one to form a tour.
- The classical one-point crossover proved inadequate, even when applied on the path representation. As illustrated in Figure 1 on a tour with six vertices, this operator leads to invalid tours with missing and duplicated vertices. This is also true for generalized version of this operator, like the 2-point and M-point crossover operators, where substrings between two consecutive cut points are exchanged.
- Special crossover and mutation operators, called order-based operators, must be designed to allow valid offspring tours to be generated. A well-known crossover operator for the TSP is the order crossover *OX* [167]. First, two cut points are randomly chosen. The substring between the two cut points on the first parent is copied to the offspring. Then, the remaining positions are filled by following the customer order on the second parent, starting at the position just after the second cross point. When the end of the sequence is reached, the procedure resumes at position 1. Figure 2 provides an example on a tour with six vertices. The substring made of vertices 3 and 4 in parent 1 is first copied to the offspring. Then, the customers are processed in the order 2, 4, 1, 6, 5, 3 on the second parent. Vertex 2 is inserted at position 5, vertex 4 is discarded because it is already present in the offspring, vertex 1 is inserted at position 6, vertex 6 at position 1, vertex 5 at position 2 and vertex 3 is discarded. This operator thus produces a valid offspring tour from two parent tours. Many order-based operators like this one are reported in the literature to handle sequencing problems, where the

Figure 2: *OX* crossover

solutions differ only by the ordering of their elements [171].

- Heuristic information must be incorporated into the GA to obtain competitive results. For example, powerful heuristic-driven local search-based operators are now commonly used within GAs and often play the role of the mutation operator. This kind of GA is also known in the literature as a memetic algorithm [163, 164].

Vehicle routing problems are more complex than the TSP since they comprise both an assignment subproblem, where each customer is assigned to a particular vehicle, and a sequencing subproblem, where the customers assigned to the same vehicle are sequenced to form a route. Clearly, the difficulties observed in the case of the TSP are even worse here. In particular, the representation, solution evaluation and problem-solving methodologies need to be adapted to solutions made of multiple routes, as it is explained below.

Representation

With regard to the representation issue, the main formalisms for vehicle routing problems (some of which inherited from the TSP) are:

- *No representation.* Since many GAs for combinatorial optimization have already drifted away from the evolutionary orthodoxy, some authors propose to forget as well about the representation and to apply the genetic operators directly on the solutions.
- *Random keys.* An encoding scheme based on random numerical keys is proposed in [8] for sequencing problems. That is, each element of a solution is tagged with a random key and a sequence is obtained by sorting these keys. By manipulating and modifying the random keys with crossover and mutation operators adapted to this representation, different sequences can be generated. For vehicle routing problems, a solution element is a customer index which is tagged with an integer that represents the assigned vehicle, plus a real number in the interval (0,1). By sorting the real numbers which are coupled with the same vehicle number, a route for this vehicle is obtained [8, 53, 123].

- *Path representation.* Here, a sequence of customer indices is used with separators to indicate where the routes end. When the number of vehicles is a variable, the number of separators also varies, thus leading to a solution string of variable length. To obtain a fixed length representation, some authors use a sufficiently large number of separators and interpret consecutive separators in the sequence as a single separator. Another approach is to add pointers at the end of the representation to indicate where each route ends. In decoder or route first-cluster second approaches (see Sections 3.2 and 3.4), a single sequence without any separator is used, thus leading to the classical path representation for the TSP.
- *Assignment representation.* In this representation, each position in the string corresponds to a customer (i.e., position 1 stands for customer 1, position 2 stands for customer 2, ...) and the value at each position is a vehicle number, namely, the vehicle assigned to the corresponding customer. Since this representation is not an ordering, classical crossover and mutation operators can be used to generate new vehicle assignments. An alternative representation is reported in [50], where the assignment is encoded as a binary matrix. In this matrix, the columns correspond to customers and the rows to vehicles. If the entry (i, j) is 1, it means that customer j is visited by vehicle i . Accordingly, all entries set to 1 in a row correspond to a group of customers visited by the same vehicle. These representations are appropriate to address the assignment subproblem, but not the sequencing subproblem.
- *Cluster representation.* In this representation, each position contains a cluster identifier, that is, a group of customers visited by the same vehicle. This representation is typically used at some upper level, where different partitions of the set of customers are explored before the sequencing issue is addressed.
- *Multi-part representation.* In this more complex representation, each part is a different string that corresponds to a vehicle route (this is called a Genetic Vehicle Representation (GVR) in [80]). A path representation is then used within each part to encode the corresponding route. Note that this is pretty close to approaches where the GA works directly on the solutions (see *No representation*, above).

Problem-solving methodology

Obviously, solutions with multiple routes also impact the problem-solving methodologies. For one thing, the genetic operators need to be adapted to the chosen representation. Also, the GA must account for the side constraints that lead to these multiple routes. This is usually done in the following ways:

- The design of genetic operators that ensure solution feasibility.
- The use of a decoder that constructs a feasible solution from an ordering of customers produced by the GA (see Section 3.2).
- The introduction of penalties in the evaluation of an infeasible solution.
- The use of repair operators to restore feasibility.

Solution evaluation

Multiple routes impact the way a solution is evaluated. A pervasive issue with regard to vehicle routing problems, particularly in the case of the VRPTW, concerns the number of vehicle routes in a solution versus the total distance or travel time of the vehicles. Most problem-solving approaches for the VRPTW first minimize the number of vehicles and, for

the same number of vehicles, minimize the total distance or travel time. This issue appears in the proposed methodologies, either via (1) the use of operators specifically designed for each objective, (2) the application of distinct search phases for each objective, (3) the co-evolution of two populations, one for each objective and (4) the development of true bi-objective approaches.

The reader will note that the considerations mentioned above emerge at different places in this survey. In the following, we first address GAs, which are by far the most widely used EAs for solving vehicle routing problems. The various methodological lines proposed in the literature are described, starting with extensions of methodologies previously developed for the TSP.

3.1 TSP extensions

Since a lot of work has been done on the TSP, it is not surprising that problem-solving approaches developed for the TSP have been adapted to vehicle routing problems. A good example is the edge assembly crossover *EAX*, a powerful operator for the TSP [165], which was later adapted to the VRP in [70]. In the case of the VRP, this operator can be described as follows:

1. Construct a graph by combining the edges of the two parent solutions.
2. Partition the set of edges in this graph with cycles created by alternately selecting an edge from the first and second parents.
3. Select a subset of cycles.
4. Generate an intermediate solution as follow. Take one parent and remove all edges that are in the subset of cycles. Then, the edges in the subset of cycles from the other parent are added. At this point, the intermediate solution is a collection of routes connected to the depot plus subtours that are not connected to the depot.
5. Create a complete solution. A greedy heuristic is applied where, at each iteration, a subtour is merged at least cost to a route or to another subtour. The procedure is repeated until a set of routes is obtained that cover all customers.
6. Eliminate capacity violations. Using a penalty function for route overcapacity, restricted 2-opt exchanges [160] and customer exchanges are applied until a feasible solution is obtained. These modifications are said to be restricted because an infeasible route is necessarily involved.

The *EAX* crossover is illustrated in Figure 3. From two parents, a combined graph is obtained. Then, based on the particular set of cycles shown in this example, an offspring is produced. The latter is obtained from parent 1, by removing the edges of parent 1 and by adding the edges of parent 2 that are in the set of cycles. Clearly, the offspring is not a valid solution since it contains two subtours that are not connected to the depot. These subtours are thus merged with the routes to obtain a valid solution, as mentioned above.

This GA with *EAX* reached the best solution on a subset of instances taken from Christofides, Mingozzi and Toth's VRP set [135] and found ten new best solutions on Golden et al.'s set [149]. This work has also been extended to the VRPTW in [71]. In this case, an additional term for time window violations is included in the penalty function when feasibility is restored with exchanges.

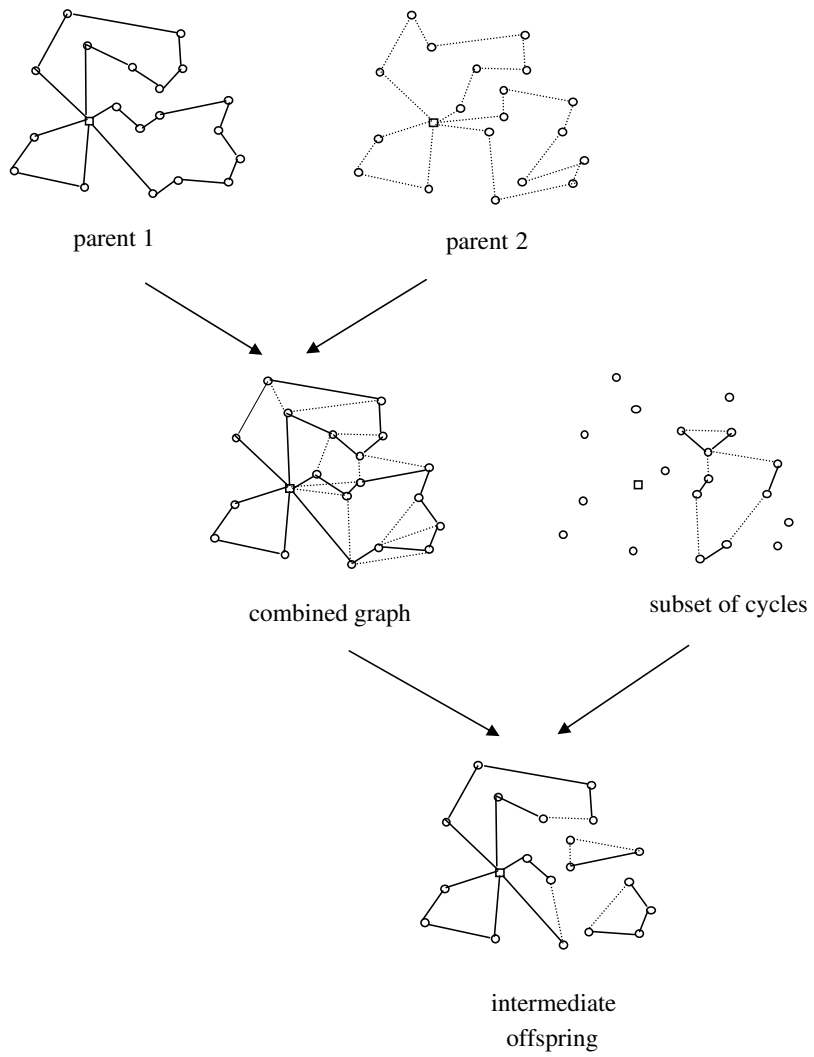


Figure 3: *EAX* crossover

An adaptation of the natural crossover, previously developed for Euclidean TSPs [153], is reported in [55] for the VRPTW. This operator, works on a two-dimensional graphical representation of the problem. It partitions the set of customers into two different classes by drawing one or more curves or geometric figures, like rectangles and ellipses, over the graphical representation. That is, customers are on one side or the other of the curve or customers are inside or outside of the geometric figure. Then, arcs from the first (respectively, second) parent solution with both endpoints in the first (respectively, second) class are transferred to the offspring solution. At this stage, disconnected segments from both parents are obtained. A repair algorithm is then applied to connect these segments to form feasible routes. Basically, routes are constructed one by one and the last customer on the current partial route is connected at least cost to the first customer of a segment, where the cost takes into account spatial and temporal issues. A local optimization is also applied on the resulting solution, based on intra- and inter-route customer moves and on 2-opt [160]. An example of the natural crossover is shown in Figure 4. In this example, a rectangle is drawn to divide the customers into two classes, namely those that are inside or outside of the rectangle. The customers inside the rectangle are linked according to the edge pattern of parent 1, while the others are linked according to the edge pattern of parent 2. The intermediate offspring contains four segments that need to be reconnected to produce a valid solution.

In [56, 100], the GENITOR algorithm [191] is adapted to the VRP by modifying its edge recombination crossover *ER*. This operator, originally designed for the TSP, progressively extends a tour by adding edges found in one of the two parents. These edges are stored in a data structure called an edge table, where all parental edges incident to a given customer node are grouped together. In the case of the VRP, *ER* is modified in the following ways: (1) the depot is always used as the starting point of a route, (2) the edge table includes the depot for which the number of incident edges is two times the number of vehicles, (3) only edges leading to the closest nodes are considered and (4) capacity constraints are taken into account to stop the construction of the current route and starts a new route. In this work, the mutation operator is a 2-opt [160] which is applied at every G iterations, where G is a parameter. In a follow-up paper [29], an open loop power distribution problem is modeled as a VRP and solved using the same approach, except that the 2-opt is replaced by a 3-opt. In [28, 30], the same modified *ER* operator is used for a multi-depot vehicle routing problem where each route can start and end at a different depot and where the number of vehicles at each depot is known. After a clustering phase that assigns each customer to the closest depot, and a routing phase at each depot that creates open half routes (by dividing the capacity in two) with a greedy heuristic, the GA with modified *ER* is applied at each depot to optimize the half routes. When this is done, another GA is applied to find the best way to link pairs of half routes together to form complete routes. Basically, the crossover operator in this GA combines half routes taken from two different solutions. A similar approach to this problem, but where a decoder is used to form the half routes, is discussed in Section 3.2.

The work in [26, 44] address the VRPTW through an extension of a matrix-based crossover operator for the TSP [143]. Basically, an entry (i, j) in the matrix is set to a large value if customers i and j are in the same route in both parents and share the same precedence relationship. A small value is used when i and j are not in the same route. Intermediate values are used when customers i and j are in the same route in both parents but do not share the same precedence relationship or when i and j are in the same route in one parent only. Offspring solutions are then generated by applying Solomon's I1 insertion heuristic [181], using the values in the matrix to modify the insertion costs. For example, if customer i precedes customer j in both parent solutions, so that entry (i, j) in the matrix has a large value, the insertion of j just after i is greatly favored by dividing its insertion cost by this large value. The GA is used in conjunction with a tabu search heuristic [145]

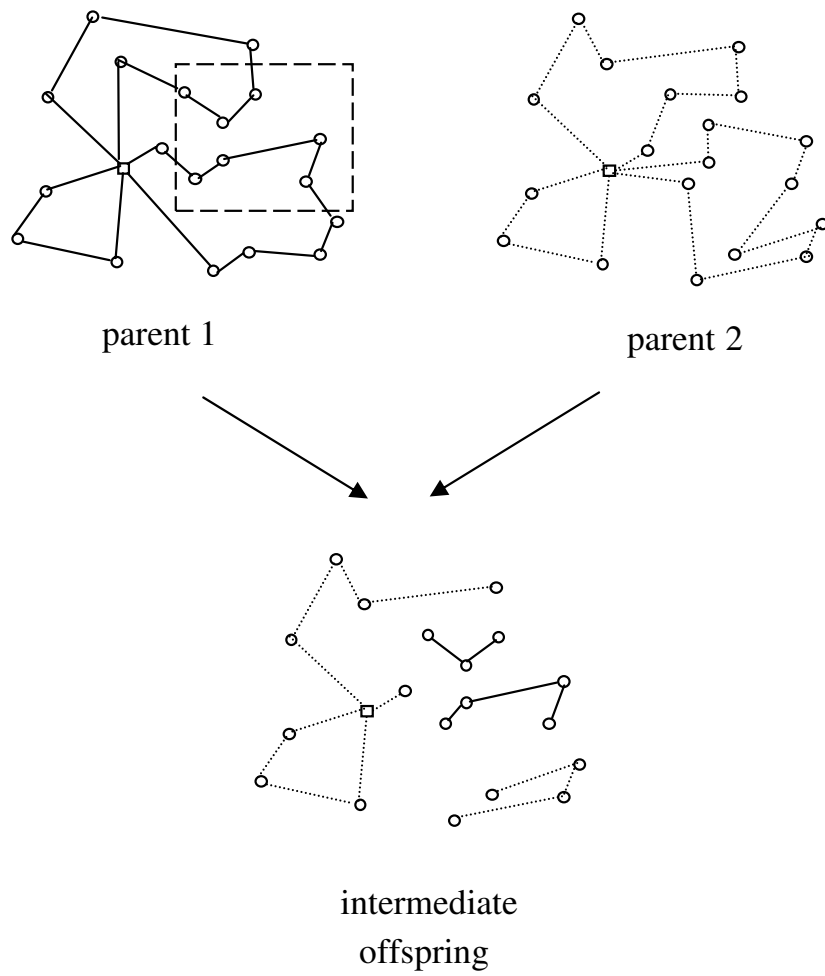


Figure 4: Natural crossover

that further improves a fraction of the solutions in the current population. The authors propose to run the GA without tabu search in the first iterations to favor diversification. Then, the intensity of the tabu search is progressively increased, by applying it on a larger fraction of the population and by running it longer on each solution. With this approach, a graceful convergence of the population is observed. A completely different approach for designing a crossover operator for the VRPTW is also proposed in this work, based on the concept of a dominant and a recessive parent, where the dominant parent is the one with better evaluation. The operator first makes the offspring a duplicate of the dominant parent. Then, the path representation of the recessive parent is swept position by position to get the current customer i and its immediate successor j . These customers are located in the offspring and an attempt is made to put customer j immediately after customer i in the offspring, if this is feasible.

The VRPB is solved in [109] with a number of order-based crossover operators for the TSP that are adapted to take into account backhaul customers. Unfortunately, there is little information about how this adaptation is realized. A problem that exhibits features of both the VRPB and VRPSDP is solved in [31] with an algorithm called CLOVES. In this application, each customer is either a pure delivery (D), a pure pickup (P) or a delivery and pickup (DP). Furthermore, the customers must be visited in the following order: pure delivery, delivery and pickup, pure pickup. In the CLOVES algorithm, only the last stage is based on a GA. The first stages construct a solution for this problem by (1) clustering customers for each of the D, P and DP categories, (2) ordering the customers within each cluster to create segments and (3) generating complete routes for the vehicles by appropriately assigning D, P and DP segments to each vehicle. The GA is then called on a population made of the solution obtained plus additional ones produced by applying Or-opt exchanges on this solution [168]. Two variants of the *OX* crossover [167] are proposed. The first variant is the classical operator, where a segment between two crossover points is chosen on one parent and transferred to the offspring. The empty locations in the offspring are then filled by sweeping the second parent from the first to the last position (to preserve, as much as possible, the required D-DP-P ordering). The second variant does the inverse: all customers outside of the crossover points are copied in the offspring and the empty locations between the two crossover points are filled by sweeping the second parent. Due to the unavailability of benchmarks for the problem considered, the algorithm was tested on VRP and VRPB benchmark instances. The average gap with the best-known solutions in these two cases was typically under 1%.

A problem motivated from a real-world application, where a single vehicle has to perform daily tours to visit on-shore oil wells, is addressed in [93]. The tour starts and ends at the oil treatment station where separation of oil from water occurs. As the vehicle cannot visit all wells in a single day, the maximum amount of oil should be collected within a given time horizon. Although this problem is more closely related to the prize collecting TSP [141], it is considered here due to the presence of a physical depot where the vehicle must start and end its route. The solution representation is made of a variable number of wells, as the authors have chosen not to include unvisited wells in it. The crossover operator thus works on a variable-length representation and is also extended to apply to a varying number of parents. It creates an offspring, starting with a randomly chosen node. Then, the next node j is probabilistically chosen based on a measure that takes into account how often node j is found just after node i in the parent solutions. Two different types of local search heuristics are then applied in sequence with a given probability associated with each one. The first one tries to insert a well not currently visited in the solution. The other one exchanges a visited well with an unvisited one. At regular interval, a so-called data mining module is applied to discover relevant subsequences in the best solutions found thus far. These subsequences are then used to construct new solutions. A crossover mechanism similar to the one mentioned

above is applied, except that when the next node is selected, frequent subsequences that involve that node as the starting node are considered for inclusion in the solution.

A similar, but more complex, multi-vehicle problem is reported in [95, 96, 97]. This problem is in fact a PDPTW, motivated from a less-than-truckload application, but where it is still not possible for the fleet of vehicles to visit all customer requests. Accordingly, some requests are subcontracted to a common carrier. The problem is thus to select the customer requests to be visited, to assign these requests to the available vehicles and to sequence them to form routes. This should be done at least cost, which is the travel cost to visit the selected customers plus the subcontracting cost for the remaining customers. A path representation is used (with unserved requests at the end), along with pointers to indicate the end of each route. The GA starts with an initial population of solutions generated with a construction heuristic reported in [95]. A uniform crossover operator is then applied to create each offspring route. That is, a bit mask is generated to decide if each customer request in the offspring route is taken from parent 1 or parent 2. When the entry in the bit mask is 0, the request is taken from parent 1; when the entry is 1, the request is taken from parent 2. Once the current route is completed, the procedure is repeated with the next route. A mutation operator is also used, where a customer request is reinserted at another location. The request can be taken from the unvisited ones, in which case the operator injects a new request in the solution. Although the crossover and mutation operators are designed to guarantee that the pairing and precedence constraints are satisfied, this is not necessarily the case for the capacity and time window constraints. Consequently, a 2-opt [160] is applied to reduce the travel costs and constraint violations. Afterward, if the solution is still infeasible, requests are removed and added to the unvisited ones to make the solution feasible.

3.2 Decoders

In a decoder, an ordering of customers (here, generated by a GA) is given as input to a route construction heuristic that decodes the ordering into a feasible solution to the problem. Basically, the customers are considered one by one, based on the ordering indicated by the GA, and feasibly inserted into the current solution. The GA is thus free to explore the search space of customer orderings and does not need to account for side constraints, which are handled by the construction heuristic. This approach can be easily adapted to vehicle routing problems with different side constraints, due to the generic nature of the insertion heuristics that are used to construct a solution.

Blanton and Wainwright [17] were the first to propose a decoder approach for the VRPTW. In their GA, the Merge crossover operators *MX1* and *MX2* exploit a global precedence relationship among customers, based on their time window, which is independent of any particular solution. Namely, the crossover operators rely on the assumption that it is desirable for customer i to appear before customer j in the ordering if the time window of i is earlier than the time window of j . In the *MX1* crossover, the two parent orderings are processed position by position and, at each position, the customer with the earlier time window is transferred to the offspring. The resulting ordering is thus biased toward the one obtained by sorting the customers in non-decreasing order of their time window. The *MX2* operator works similarly and differs only in the way the two parent orderings are processed.

Each ordering produced by the GA is then decoded into a solution of the VRPTW and evaluated by adapting an approach reported in [152] for a partitioning problem. Assuming a fixed number of vehicles m , the first m customers in the ordering are used to initialize the vehicle routes, where each route visits a single customer. The remaining customers are then inserted one by one into any of these routes (so that all routes grow in parallel), based on

a least-cost insertion measure. A customer who cannot be feasibly inserted in the solution remains unvisited. Such infeasible solutions are kept in the population but are highly penalized. Tests on problems with 15, 30, 75 and 99 customers were performed by integrating the two operators into the GENITOR package [191]. These operators were shown to outperform general purpose order-based crossover operators like *ER* [191] and *OX* [167]. An improvement is proposed in [63] for instances where customers are geographically clustered (as instance classes C1 and C2 in Solomon's data set [181]). Rather than using the first m customers in the sequence to initialize m routes, the authors take care to select customers from different clusters to initialize the routes.

In [86], the authors argue that *MX1* and *MX2* are too strongly biased by the customers' time windows. Accordingly, they propose a crossover operator that stands between general purpose order-based crossover operators and the *MX1* and *MX2* operators. The *1X* crossover is similar to the classical one-point crossover and thus generates invalid orderings. These orderings are then repaired by eliminating duplicate customers and by inserting all customers that are not yet part of the ordering. When inserting customers into the ordering, their insertion order is determined by the time occurrence of their time window (i.e., the customer with earliest time window is inserted at the first available position, etc.). As opposed to [17], the construction heuristic for decoding the ordering into a solution is Solomon's I1 insertion heuristic [181], where the routes are built one by one. That is, a first route is initialized with the first customer in the sequence, the next customers are inserted in this route until it is full; then, a second route is initialized, etc. The new crossover operator was shown to outperform *MX1* and *MX2* on Solomon's VRPTW instances. The same GA was then applied to the VRPB with time windows in [85], by slightly modifying Solomon's insertion heuristic to account for the precedence relationship among linehaul and backhaul customers. Similar decoder approaches for the VRPTW are also reported in [69, 104, 127]. In [104, 127], Solomon's I1 sequential insertion heuristic is used, while a time-oriented parallel savings heuristic, also proposed by Solomon in [181], is used to construct routes in parallel in [69].

A messy GA [146] for the VRPTW is described in [103]. Messy GAs can be seen as a relaxation of classical GAs, because they are more flexible with regard to the solution representation. In particular, each element is a couple (position, value) that can appear at any particular location along the string. Messy GAs also allow for variable-length strings that may be under or overspecified with regard to the problem under study. Overspecification occurs when two elements share the same position but have different values. When the string is swept to decode it into a solution, overspecification is handled by selecting the first element encountered and by ignoring the following ones. Underspecification occurs when a particular position is absent. In that case, the missing value is filled using the current template, which is the best solution from the previous iteration. As in classical GAs, there is a parent selection phase followed by the application of a cut-and-splice operator, which is similar to the one-point crossover, but adapted to variable-length strings. In the case of the VRPTW, an element is a couple (customer index, vehicle). Overspecification occurs when the same customer is served by two or more vehicles. Underspecification occurs when a customer is not served by any vehicle. Decoding the string into a solution of the problem is obtained by considering each couple in turn along the string and by applying Solomon's I1 heuristic [181] to insert the corresponding customer in the vehicle's route. As the authors minimize total distance only, they often obtain better solutions on Solomon's benchmark instances than other approaches, but they use much more vehicles. It is worth noting that the GA developed by the authors only approximately follows the original design of messy genetic algorithms, and the theoretical justifications at the core of this design are more or less ignored.

A VRP with time deadlines is addressed in [65, 74]. One interesting feature of the pro-

posed GA is a route crossover *RC* based on a bit mask where the number of bits corresponds to the number of routes in the first parent solution. If the bit value is 1, the corresponding route is transferred to the offspring solution. If the bit value is 0, the customers in the corresponding route are added to a temporary list. At the end, the customers in the temporary list are sorted, based on their relative order in the second parent and inserted, one by one, into the routes of the current offspring solution. Using the final ordering obtained, including the addition of infeasible customers at the end, a decoder then produces the final solution with a sequential construction heuristic.

A multi-depot VRP, where each route can originate and terminate at a different depot and where the number of vehicles at each depot is known, is addressed in [98, 99]. In this particular application, the number of vehicles that start their route from a given depot and end their route at another depot is specified for each pair of depots. The problem-solving methodology first assigns a number of closest customers to each depot. This number corresponds to the number of vehicles that goes in and out of the corresponding depot. Each assigned customer is used to initialize a route that goes in (first customer of the route) or out (last customer of the route) of the depot. Afterward, a GA is applied to handle the remaining customers and enlarge the routes until half of the vehicle capacity is reached. A decoder is used for this purpose, that is, an ordering of unassigned customers is provided as input to a construction heuristic that links each customer in turn to the closest end node of a half route. When all customers are routed, the route halves are linked together in a heuristic way by sequentially and randomly selecting a half route and connecting it in the cheapest possible way to another half route to form a complete route (while taking care to satisfy the specifications about the number of vehicles that travels from one depot to another).

Finally, a GA for solving a dynamic vehicle routing problem is described in [43]. Here, new customer requests are received over the day and must be integrated in real-time into the current vehicle routes. In this work, the horizon is divided into time slices. A new vehicle routing problem is then defined and solved by the GA at the end of each time slice by incorporating any new request that has been received during that time. A variable-length string representation is proposed that encodes the new customer requests that have not been assigned yet, plus pointers to existing routes. A decoder then inserts these new requests one by one into existing routes, if this is feasible, or creates new routes for them, based on their ordering in the string.

3.3 Cluster First-Route Second

The well known cluster first-route second problem-solving methodology is based on the natural partition of a vehicle routing problem into two subproblems, namely (1) an assignment subproblem where groups or clusters of customers that are visited by the same vehicle are identified and (2) a sequencing subproblem where the customers within each cluster are sequenced to form routes.

Sam R. Thangiah was one of the first researchers to apply GAs to vehicle routing problems with this framework. Basically, clusters of customers are first identified by a GA, while the routing is done afterward with standard operations research techniques. The GIDEON system [114, 116, 118] was first developed in the context of the VRPTW. Its clustering and routing phases work as follows. In the clustering phase, a genetic sectoring algorithm called GENSECT partitions the set of customers into m clusters, one for each vehicle, by identifying $m - 1$ rays originating from the central depot. All customers lying between two consecutive rays are then grouped together to form a cluster. As each ray is defined by its polar angle, $m - 1$ polar angles are encoded on a bit string (to be more precise, each angle

is encoded as an offset from a fixed increment angle). To evaluate the clusters produced by GENSECT, a route is formed within each cluster. This is done through a least-cost insertion heuristic, where the cost accounts for the total travel distance, as well as penalties for time window and capacity violations. At the end, a local search heuristic is applied to the best solution produced by GENSECT for further improvement. The neighborhood structure is based on the λ -interchanges of Osman [169], with $\lambda = 2$. In this case, one or two customers are moved to another route or exchanged with one or two customers in another route. This final phase is particularly important to reduce or eliminate violations of the capacity or time window constraints. GIDEON was used to solve a school bus routing problem in [117]. It was also applied to VRPTDs [119, 122], where the time window at each customer is replaced by a time deadline.

In [113], an alternative approach is proposed to cluster customers. The GA, called GENCLUST, adapts geometric shapes (more specifically, circles) to cluster customers. Each bit string encodes the origin and radius of a set of circles. The number of circles corresponds to the number of vehicles, and is estimated by first solving the VRPTW with Solomon's II insertion heuristic [181]. All customers found in a given circle are grouped into the same cluster. Customers that are in-between two or more circles are associated with the closest circle, while customers that are found inside two or more circles (assuming that circles can overlap) are associated with the circle that appears first along the bit string. Thus, when a small circle is included within another larger circle, and customers are found only within the small one, the circle that appears first takes all customers. The other circle becomes empty and is removed, thus saving one vehicle. As in GENSECT, the circles or clusters produced by the GA are evaluated by sequencing the customers within each cluster with a least-cost insertion heuristic to form routes.

GENCLUST was also adapted to solve the multi-depot vehicle routing problem, where each vehicle starts and ends its route at the same depot [92, 121]. Basically, when a set of circles is produced by the GA, a solution to the multi-depot problem is obtained by associating each circle or cluster of customers with the nearest depot. The interested reader will find a review of Thangiah's work with GAs for these different types of vehicle routing problems in [120].

In [81], the cluster first-route second approach is implemented using a different GA in each phase. For the clustering phase, an assignment representation is exploited by the GA to assign customers to vehicles, that is, to form clusters. Another GA then form routes within each cluster by applying the *PMX* order-based crossover [147] on a classical path representation. This problem-solving approach is applied on a PDP, more precisely a dial-a-ride problem with no time windows. In this particular application, a repair operator is needed to swap the pick-up and delivery associated with a given customer, when the precedence constraint is not satisfied.

A dial-a-ride problem with time windows is addressed in [50]. Here, the assignment representation is based on a binary matrix, where entry (i, j) is 1 if customer j is visited by vehicle i , and 0 otherwise. Thus, each row contains the group of customers assigned to a given vehicle. A crossover operator then explores different ways to perform this assignment. First, one row is selected in the matrix of each parent. A binary mask is then used to create a new row for the offspring. Each row element is inherited from parent 1 or parent 2, depending if the bit value in the mask is 0 or 1, respectively. The other rows are inherited from parent 1. A repair operator is then applied to make sure that each customer is assigned to exactly one vehicle (i.e., each column has exactly one entry equal to 1). Routes are then created from the obtained assignment using a nearest neighbor heuristic that accounts for both spatial and temporal proximity.

In [77], each element in the representation of a PDPTW corresponds to a group of

requests that are assigned to the same vehicle, rather than a single request or a single service point (pickup or delivery). A variant of the 2-point crossover operator is then applied to combine clusters from two parents on a single offspring. The substring of clusters between the two cut points on the second parent is inserted just after the first cut point on the first parent to create an offspring. Then, a repair operator removes duplicate customers (which can lead to the removal of whole clusters). Finally, unassigned requests are inserted, allocating a new vehicle if necessary. The mutation operator randomly removes a cluster and reinserts its customers into other clusters, also allocating a new vehicle if necessary. While clusters are worked on, the corresponding routes are created and updated. For example, when a customer request is reinserted in another cluster, the insertion is also performed in the corresponding route to ensure feasibility. An adaptation of this GA for a dynamic version of the PDPTW, where new requests are incorporated in real-time into the current routes, is reported in [78]. Basically, a series of static subproblems over known requests are solved under a rolling time horizon. A static subproblem is defined each time a new request occurs. At this point, each solution in the current population is updated (synchronized) to reflect the current state of the system. The population obtained becomes the initial population for the next execution of the GA. During synchronization, care is taken to maintain as much as possible the integrity of the request groupings.

In [7], the VRP is addressed with a GA using an assignment representation. Although this representation is not meant to specify the sequence of customers visited by each vehicle, the routes are implicitly inferred. That is, the customers are indexed in such a way that it makes sense for two customers with consecutive indices to be visited one after the other by the same vehicle. These indices are determined by sorting the customers according to a nearest neighbor TSP solution, when the customers are clustered, or based on their polar angle, when the customers are randomly located. The route associated with a given vehicle is then obtained by taking all customers with the same vehicle number, and by sequencing these customers based on their relative order in the representation. Local optimization heuristics, like 2-opt and 3-opt are then applied on each route. Since the representation encodes groups of customers that are assigned to the same vehicle as well as their sequence, it is not a true cluster-first, route-second approach. However, these routes are only the starting point for the local optimization heuristics. They do not correspond to the real routes, which are those obtained after the local optimization.

In the case of the work in [27, 125], the cluster first-route second paradigm is extended to a periodic VRP. Here, *cluster first* corresponds to the selection of customers that are visited during the same period over a given multi-period horizon, while *route second* corresponds to the construction of routes for each period. A binary representation is used to encode the periods where each customer is visited. For example, if an horizon of three periods is considered, there are seven possibilities: 001, 010, 100, 110, 011, 101, 111. That is, if 101 is associated with a given customer, this customer is visited in periods 1 and 3. Based on a particular selection of visits for each customer, a VRP is then solved in each period using a cost saving heuristic. The proposed crossover operator creates two offspring by exchanging the selection of visits of a given customer on the two parents (if feasible), while the mutation operator randomly changes the selection of visits of a number of customers (if feasible). A parallel implementation of this GA is reported, where subpopulations evolve in parallel and periodically exchange solutions.

3.4 Route First-Cluster Second

Route First-Cluster Second is an alternative paradigm for tackling vehicle routing problems. Basically, the path representation encodes a (likely infeasible) unique giant tour that covers all customers. Using such a representation means, in particular, that classical order-based

crossover and mutation operators previously developed for the TSP can be used, given that a single sequence of customers is worked on. The giant tour then needs to be partitioned into individual feasible routes. By working on a single sequence of customers and by coupling this representation with a procedure that transforms this sequence into a solution of the problem, an approach very similar to the decoders of Section 3.2 is obtained. However, the final routes are usually more closely tied to the original sequence through the use of decoders that can only break the sequence into subsequences. This approach can thus be viewed as a restricted type of decoder.

A VRP with route time constraints is addressed in this way in [94]. A true solution to the problem is obtained by sweeping the customers along the giant tour. A route ends when either the capacity or the maximum route time constraint would be exceeded by including the next customer. This customer is then used to start a new route. In [88], a polynomial-time algorithm is developed to partition the giant tour into individual routes in an optimal way. This algorithm, called SPLIT, works on an auxiliary directed acyclic graph with vertices $\{1, 2, \dots, n\}$, where vertex 1 is the depot and the other vertices are customers. An arc (i, j) is added to the graph when a route from vertex $i + 1$ to vertex j , based on the ordering in the giant tour, is feasible. The length of the arc is the total distance of the corresponding route. A solution is then obtained by solving a shortest-path problem from node 1 to node n . This can be done in polynomial time due to the acyclic nature of the graph. An example taken from [88] is illustrated in Figure 5, assuming a vehicle capacity of 10. In this figure, the demand at each vertex is shown within brackets and each link is labeled with its length. The giant tour in Figure 5(a) starts at the depot 1, visits the vertices in the order 2, 3, 4, 5, 6, and returns to the depot. The corresponding shortest path problem is illustrated in Figure 5(b). In this graph, the arc from vertex 4 to vertex 6, for example, corresponds to the route 1, 5, 6, 1 with a total distance of 90. The optimum, shown with bold arcs, is to split the tour into three feasible routes: the first route visits vertices 2 and 3, the second route visits vertex 4 only and the third route visits nodes 5 and 6, for a total distance of $55+60+90=205$. It is worth noting that a similar optimal procedure, that accounts for time windows, is reported in [128].

Before the discovery of this exact, polynomial-time tour splitting procedure, different heuristic approaches were proposed in the literature to address this problem. In [101], a local search heuristic explores different configurations of cut points along the giant tour. In [64], a cooperative co-evolutionary model, where two subpopulations evolve concurrently, is proposed. In the first subpopulation, the integer strings contain a single sequence of customers, while the strings in the second subpopulation contain a number of route sizes (in terms of number of customers). The giant tour encoded in the first string is then split according to the route sizes encoded in the second string. The two subpopulations are thus required to get a complete solution.

3.5 Ruin and Recreate and Large Neighborhoods

In this section, the genetic operators are inspired from the ruin and recreate paradigm [179], where a fraction of a solution is destroyed and then reconstructed in some alternative way. When integrated within a local search heuristic, this approach leads to large neighborhoods. For example, we can choose to remove q customers from a solution to ruin it. But each selection of q customers is likely to lead to a different neighboring solution after reconstruction. Clearly, the size of the neighborhood grows quickly with q , thus producing a large neighborhood search.

A good example of this approach is found in the work of Berger and his colleagues for the VRPTW. As the first objective of the VRPTW is usually to minimize the number of




Figure 5: Exact split procedure

vehicles and, for the same number of vehicles, to minimize the total distance, researchers in the field have soon recognized the need to pay a special attention to each objective. The line of research of Berger and his colleagues makes no exception, as explained in the following.

In [15], an insertion-based crossover *IB_X* is proposed. First, route 1 is probabilistically selected from parent 1, with a bias toward routes with large waiting times. A number of customers are then removed from this route (ruin), based on criteria indicating that a more suitable relocation of these customers into alternate routes is likely, such as a large distance to the immediate successor or a large waiting time. A subset of routes close to route 1 are then selected in parent 2, and the best candidate customers in these routes, plus any unvisited customers, are considered for insertion in route 1. This procedure is repeated for a number of routes in parent 1. At the end, the offspring solution is completed with the remaining routes of parent 1. If there are still unvisited customers, they are handled by constructing additional routes with a nearest neighbor heuristic. A second child is obtained by interchanging the role of the two parents. Three mutation operators are also used to either reduce the number of routes or to reorder customers within routes. It is worth noting that solution feasibility is always enforced. A variant of the VRPTW, where a maximum route time constraint is introduced, is also addressed with the same approach in [16].

In a follow-up work [14], the authors relax the time constraints and evolve two populations with different objectives. Population 1, which contains at least one feasible solution, minimizes the total distance, while population 2 minimizes time constraint violations, both subject to a fixed number of routes (i.e., m and $m - 1$ routes, respectively). When a feasible solution with $m - 1$ routes emerges in population 2, population 1 is replaced by population 2, the number of routes m is reduced by one, and the mutation operator *RSR_M* (see below) is applied to population 2 to obtain a new population with solutions with one fewer route. This is repeated until no feasible solution emerges in population 2. The crossover operators are *IB_X* [15] and a variant called insert-within-route neighborhood crossover

IRN_X. In this case, customers are removed from a subset of routes in parent 1. Then, the customers in the routes of parent 2 that are close enough to the subset of routes in parent 1, become candidate for insertion in this subset. A suite of five mutation operators is also proposed, among which the large neighborhood search-based mutation *LNSB_M* follows the guidelines of Shaw in [180]. Here, related customers (due to proximity or a common route assignment) are removed and reinserted with a variant of Solomon's I1 heuristic [181], within a search tree framework. The insertion order of the customers is determined through the summation of two ranks: one is based on the insertion cost and the other on the number of feasible insertion places. The smaller the insertion cost and the number of insertion places, the better the customer's overall ranking. Note that the second term is new and is aimed at alleviating the myopic behavior of the original approach. The edge exchange mutation *EE_M* is inspired from the λ -interchanges [169] and examines the insertion or exchange of a customer with the two closest routes (as evaluated by the customer's distance to the route centroid). The repair solution mutation *RS_M* is similar to *EE_M*, but focuses on solution feasibility by exploring exchanges involving infeasible customer visits. The reinsert shortest route mutation *RSR_M* is used to eliminate the route with the smallest number of customers from a solution. It first tries to insert each customer at a feasible place in the two closest routes by maximizing a so-called regret cost function. Basically, the procedure computes the best insertion place for the customer in the two routes and considers the difference between the two values. This measure indicates how much is lost by not inserting the customer at its best insertion place. This is a special case of the generalized regret measure proposed in [172], where the regret is computed over all routes instead of the two closest routes only. Finally, the reorder-customers mutation *RC_M* is applied when a new best feasible solution is found. It reorders the customers within each route by repeatedly applying Solomon's I1 heuristic [181] using different sets of randomly generated parameter values.

A parallel master-slave implementation of this GA is reported in [13]. In this implementation, the master supervises and controls the execution of the GA, while the slaves concurrently apply the genetic operators on different solutions in the two populations. In [20], a post-optimization phase is added to the GA. In this final phase, a population of routes (rather than complete solutions) is evolved, starting with the routes in the best solution produced by the original GA. The routes in the population are randomly ordered and pairs of routes are recombined based on this ordering. After a number of iterations, the best possible solution is constructed with the routes found in the final population. Four crossover operators are applied, depending on a priori probabilities. The *CE_X* operator applies CROSS exchanges [185] where two segments of consecutive customers are selected (one from each route) and exchanged. The insertion crossover *I_X* inserts one customer at a time in other routes. The rebuilding crossover *R_X* frees all customers in two parent routes and reinserts them using a least-cost insertion heuristic. The random rebuilding crossover *RR_X* is a variant of *R_X* where the customers are considered for insertion in random order. The two mutation operators considered are first-improvement local search heuristics, where a new route is accepted as soon as it provides an improvement. The permute mutation operator *P_M* is based on Or-opt exchanges [168], where a string of one, two or three consecutive customers is removed from a route and inserted elsewhere. In the reorder mutation operator *R_M*, customers are removed from a route. Then, those with a tight time window are considered first during the reinsertion.

This parallel GA was later adapted to the VRP [11, 12]. In this case, the two populations are used for diversity purposes only, as they both optimize the total distance. Since the VRP is less constrained than the VRPTW, the *IB_X* crossover and the suite of five mutation operators in [14] are also used here. An approach inspired from the granular tabu search [187] is also exploited to reduce the number of possible insertion places when applying the genetic

operators. Basically, insertion places for a given customer are restricted to surrounding routed customers located within a certain distance limit.

In the work of Potvin and Bengio [82], a genetic algorithm called GENEROUS is proposed to solve the VRPTW. One interesting feature of this work is that the encoding issue is avoided by applying the genetic operators directly on the solutions. A route-based crossover *RBX* is proposed where two routes from two different parent solutions are exchanged to form offspring. Another crossover operator, called *SBX*, is based on 2-opt* exchanges [173], where the end parts of two routes, one in each parent solution, are exchanged. In both cases, a repair operator is applied to produce feasible solutions, that is, duplicate customers are eliminated and missing customers are reinserted using a least-cost insertion heuristic. A mutation operator is specifically designed to empty small routes, by inserting customers one by one into other routes. Another mutation operator is a primitive ejection chain [144], where a customer can eject another customer from another route. The ejected customer is then feasibly inserted, if possible, at some other location. Apart from these true mutation operators, a local search heuristic based on Or-opt exchanges [168] is also applied with a certain probability to further optimize the solution. The same *SBX* crossover operator is used in [76]. The authors also generalize this idea by designing a new crossover operator, inspired from the CROSS exchanges [185], where intra-route segments from two parent solutions are exchanged to create offspring. Similar operators for the VRP and VRPTW are described in [80, 91, 110, 111, 112]. Using a multi-part solution representation called GVR, where each part represents an individual route, the authors first design a crossover operator that removes an intra-route segment from one parent and inserts it after the closest customer in a route of the second parent to create an offspring. Duplicates are then removed. A crossover operator similar to *RBX* is also proposed where a full route is taken from one parent and transferred to the other parent to create an offspring. Duplicates are then eliminated.

The GA for the VRPTW of Bräysy and Dullaert [21] emphasizes fast computation times. The algorithm is divided into two different stages that are applied a number of times within a multi-start framework. The first stage corresponds to the creation of an initial solution. The second stage corresponds to the improvement of the initial solution and is divided into two different phases. The first phase is deterministic and minimizes the number of routes using ejection chains [144]. The second phase exploits ideas found in [20] by evolving a population of routes with the aim of minimizing the total distance. The initial population is made with the routes obtained at the end of the first phase. Parent routes are paired randomly and the offspring routes are required to contain all customers found in the parent routes. The proposed *ICROSS* operator is based on CROSS exchanges [185], although exchanges that are not likely to lead to feasible routes are filtered out to limit the computational burden. The reinsert related customers crossover *RRC* removes a set of customers from two parent routes, and then tries to reinsert them, as in [180]. After reinsertion, the customers in both routes are reordered with Or-opt exchanges [168]. When no further improvement can be obtained via *ICROSS* and *RRC*, a mechanism that allows the search to explore other regions in the solution space is applied. Basically, the total waiting time along the routes is introduced in the objective and its weighting coefficient is gradually reduced until it reaches 0, in a manner reminiscent of the cooling schedule of simulated annealing [155]. At the end, each offspring route is improved with a local search heuristic based on intra-route CROSS exchanges (i.e., the two segments to be exchanged are taken in the same route).

3.6 Column generation

In the work of Alvarenga and his colleagues, a set partitioning formulation of the vehicle routing problem is exploited to partition the set of customers into a number of routes. These routes correspond to columns in the integer programming model (IP). As it is not possible

to enumerate all feasible routes, only a subset of routes is generated in a heuristic way with a GA. The scaled-down set partitioning problem is then solved exactly with an IP solver. It is worth noting that the overall approach is heuristic, because there is not guarantee that the GA will generate all routes needed to get an optimal solution.

In [4, 6], the VRPTW is addressed in this way with the unusual objective of minimizing the total distance only. Many executions of the GA are performed to produce different routes (columns). Within the GA, the crossover operator creates an offspring by alternately selecting a route from the first and second parents, while removing duplicate customers if necessary. The remaining customers are then inserted in the existing offspring routes, if possible. Otherwise, a new route is created. Thus, there is no incentive to reduce the number of routes. Many different mutation operators are proposed to reinsert a customer or exchange two customers, merge a pair of routes or split a route, with or without consideration for an improvement in the objective. The solution obtained after solving the set partitioning problem, based on the routes produced by the GA, is then divided into a number of subsets of routes. The GA is then applied again on each subset of routes to generate additional routes. By working on smaller subproblems, an intensification of the search is obtained. At the end, the final set partitioning problem produced with all columns generated from the start, is solved to get the final solution. This algorithm is shown to improve the best distances reported in the literature on many Solomon's benchmark instances [181], but at the expense of an increase in the number of vehicles. The same approach is proposed in [19]. In this work, however, the focus is on the benefits of a memory where solutions from previous iterations are preserved and used to fill a part of the current population.

In a follow-up work by Alvarenga and his colleagues [3, 5], the algorithm described above is modified to focus first on route minimization and then, on total distance. To this end, an independent GA run is performed at the start to reduce the number of vehicles. When selecting parent solutions, the hierarchical or lexicographic order of criteria proposed in [45] is applied (see Section 4), where the focus is on the route with the smallest number of customers. However, the authors note that additional criteria can be useful to distinguish between two solutions. They propose criteria of their own based on the so-called taker route, which is the route with minimum time window violations when a customer from the smallest route is inserted in it. The mutation operators from their previous work in [4] are also modified because some of them are not appropriate when the number of vehicles is considered (e.g., the mutation operator that splits a route). The new operators are aimed at emptying the smallest route, in particular through primitive forms of ejection chains [144]. The final population produced by this GA, then becomes the initial population for the GA in [4] which minimizes the total distance.

3.7 Landscape Analysis

The work reported in this section is based on an analysis of the solution space topology. As observed in [57], most instances of the VRP exhibit global convexity, which means that global optima are typically found at the bottom of big valleys (see Figure 6). Although solutions that are close from one another do not have to preserve strict convexity (and thus, are local optima), there is a general trend toward convexity. This observation implies that good solutions, like local minima, should be close to one another and should have similar features or characteristics. In [57], a GA coupled with local search is proposed to exploit global convexity. Crossover operators are designed to generate offspring that are close to their parents, using similarity measures based on common edges and customers assigned to the same vehicle. The Cluster Preserving Crossover *CPX* preserves groups or clusters of customers that are common to both parents. Clusters in the offspring are formed by intersecting pairs of routes that share the largest number of customers. Routes

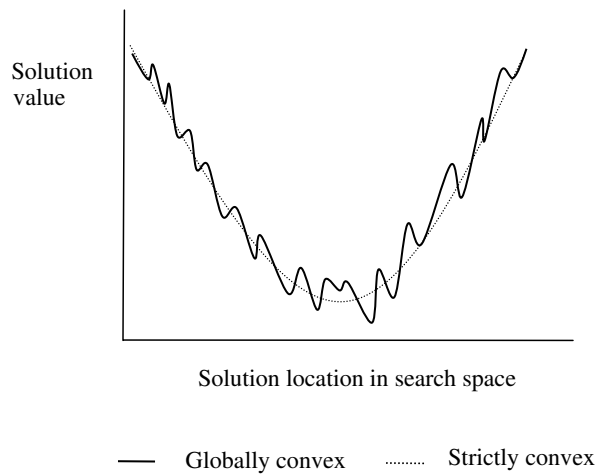


Figure 6: Global convexity

are then created from these clusters by sequencing the customers. The Common Edges Preserving Crossover *CEPX* preserves edges. It identifies paths of maximum cardinality that are common to both parents and use them to create routes in the offspring. Finally, a mixed crossover operator *CECPX* is designed to preserve both characteristics. It creates routes through paths of maximum cardinality within groups or clusters of customers that are common to both parents.

In a follow up work [58], similar ideas are explored to accelerate the local search in the GA. First, a cache is used to store the change in the objective value when a move from the current solution to a neighboring solution is performed. Since a move leaves most of the solution unchanged, many values stored in the cache are still valid for the new current solution, and can be directly fetched when evaluating the neighborhood of this solution. As opposed to the TSP, however, the capacity constraints in the VRP ask for an extension of the cache to store the feasibility status of a move. Unfortunately, due to the non negligible overhead associated with the management of the cache, the authors did not find any significant speed-up with the cache. The second technique is based on the global convexity property and emphasizes closeness when the local search is applied. In this work, after applying the edge-preserving operator *CEPX*, the local search explores a restricted neighborhood, based on 2-opt [160] and customer exchanges, where any move that would remove a common edge is forbidden.

In [48], the same ideas are exploited for a vehicle routing application in a waste management company. Starting from a central operating base, each vehicle visits a number of sectors that contains waste containers. When the vehicle is full, it goes to a dumping site to unload, before visiting the next sector. A route for an operations day is thus made of a number of route segments between two dumping sites (the dumping sites at the start and end of the day correspond to the operating base). A GA, coupled with a local search that moves sectors from one vehicle to another, is proposed to solve this problem. A crossover operator is developed to generate offspring solutions that are not too far from the parents. The features that are preserved in the parents are (1) the percentage of common arcs, (2) the percentage of common assignments of sectors to vehicles, (3) the percentage of common assignment of arcs to vehicles and (4) the percentage of common pairs of sectors assigned to a single route segment. First, each vehicle route in parent 1 is transferred to the offspring

with probability 0.5. Then, the route segments in parent 2 are inserted, whenever possible, in a route segment associated with the same vehicle in the offspring.

3.8 Cellular Approach

Typical GAs allow a solution to recombine with any other solution in the population. In cellular GAs, the population is structured according to some underlying topology (e.g., grid, mesh) and a solution can only interact through crossover with its immediate neighbors, according to the topology. As these neighborhoods overlap, good solution features can slowly diffuse in the whole population. Smaller populations can thus be used to solve problems with cellular GAs, as they are less prone to premature convergence.

In the work of Alba and Dorronsoro [1], the capacitated VRP with a maximum time- or distance-constraint is solved with a cellular GA. A 2D toroidal grid defines the connection topology. The neighborhood of a given solution is thus made of the solution itself, plus the so-called North (up), South (down), West (left) and East (right) solutions. The *ER* crossover operator [191] and reinsertion, exchange and inversion mutation operators (where a substring of customers is inverted) are used within the GA. Local optimization is also performed with 2-opt [160] and 2-interchanges [169]. In [2], this cellular GA is applied on the VRP instances of van Breedam [188], Golden et al. [149] and Taillard [184]. Nine best solutions on these instances are reported, although eight of them are obtained on van Breedam's instances, which are not as widely known as the two other sets.

3.9 Composite and Hierarchical Approaches

The term composite refers here to an approach where different, loosely integrated, metaheuristics are exploited to solve a problem. A good example is the GENSAT system which uses simulated annealing, tabu search and a GA to solve the VRPTW [115]. First, customers are clustered with the genetic sectoring heuristic GENSECT (see Section 3.3) and routes are obtained by sequencing the customers within each cluster with a least-cost insertion heuristic. Then, the solution is improved with either a local search, tabu search, simulated annealing or a hybrid of tabu search and simulated annealing. In the latter case, a probabilistic simulated annealing acceptance criterion is embedded within the tabu search. The neighborhoods of the improvement heuristics are all based on 2-interchanges [169]. A number of variants are obtained and regrouped into the GENSAT system by generating the initial solution either with GENSECT or Solomon's I1 heuristic [181] and by applying the various improvement procedures. At the time, GENSAT produced 40 new best solutions over 60 test instances, namely the 56 Solomon's benchmark instances [181] plus four instances in [178]. It also matched the best known solution on 11 instances. To be fair, GENSAT can hardly be seen as a single algorithm, but rather a collection of algorithms that are applied on a given instance, with the best solution returned at the end. In fact, the results showed that the best solutions were produced with different variants.

A parallel algorithmic scheme is proposed by LeBouthillier and Crainic in [60], where different tabu searches and GAs, explore the search space concurrently. The interaction between these various metaheuristics is realized through a central warehouse, basically an adaptive memory [175], made of elite solutions. When a metaheuristic finishes its current execution and returns its best solution, this solution is included in the warehouse if it is better than the worst solution. The diversity of solutions in the warehouse is exploited to provide new starting solutions for the metaheuristics. The proposed GAs in this system are quite standard and are based on the *ER* [191] and *OX* [167] crossovers. Through a parallel exploration of the search space, the authors have found solutions that are competitive with

the best approaches for the VRPTW on Solomon's [181] and Gehring and Homberger's [32] benchmark instances. A first refinement of this approach is proposed in [61], where a guidance mechanism is added through the identification of common patterns in over-average, average or under-average solutions in the warehouse. Frequent patterns associated with over-average solutions are fixed during a number of iterations to provide intensification. Conversely, frequent patterns associated with under-average solutions are prohibited to provide diversification. A second refinement in [59] dynamically determines when to apply the pattern-based intensification and diversification phases, based on the entropy of the solution warehouse. This entropy or uncertainty is higher when the number of different solution elements, like arcs, is larger, thus triggering intensification. Conversely, when the number of different solution elements is smaller, diversification is triggered.

In [49], a single vehicle PDPTW is addressed with a dynamic programming algorithm and a GA. If the dynamic programming algorithm cannot find an optimal solution within the allocated time, the partial solution is used to seed the initial population of the GA. This approach greatly speeds up the solution process, thus allowing applications in dynamic settings (although no such application is studied). It also allows the GA to find much better solutions when compared with a randomly generated initial population.

In [79], a hierarchical GA is used to solve a variant of the VRP where the objective is to minimize the length of the longest route based on a Manhattan distance metric. At the lower level, solutions evolve within subpopulations called niches. At the upper level, an evolutionary mechanism is applied to the niches themselves. At each upper level iteration, a number of new niches are created by merging a bad niche with a good one (where the quality of a niche is defined by its best solution). Then, the best niches obtained replace the worst ones in the original population of niches. At the lower level, solutions evolve within each niche through the application of crossover, mutation and local optimization operators. The crossover operator removes a segment of route from the first parent, remembers the immediate predecessor of the first node of the segment, finds the predecessor node in the routes of the second parent and inserts the segment just after it. Then, duplicates are removed. This is repeated as long as a feasible offspring is not obtained. The mutation operator removes a route segment and inserts it into another route. Local optimization is performed by iterating over a number of classical neighborhood structures, while always maintaining feasibility. As observed in this work, the best solutions in different niches are often the same. Accordingly, an elitist population replacement strategy, where the best solution is transferred to the next population, is applied to only one of these niches. This best solution is thus likely to disappear from the other niches. A newspaper distribution application illustrates how this two-level GA behaves.

3.10 Multi-Objective Approach

As mentioned before, vehicle routing problems often involve conflicting objectives, like minimizing the number of vehicles versus minimizing the total distance. This issue is often addressed through a weighted linear sum that aggregates the objectives or through a hierarchical approach, where the first objective is optimized followed by the second objective, while taking care not to degrade the value of the first objective. In this section, we review Multi-Objective Genetic Algorithms (MOGAs) that address vehicle routing problems using a true multi-objective approach. These algorithms are often inspired from the non dominated sorting GA reported in [137, 182]. This GA is characterized by a selection procedure based on Pareto ranks. Basically, non dominated solutions in the current population are ranked first and get the same evaluation (which can be the rank itself). After removing these solutions, the non dominated solutions in the reduced population are ranked second and get a lower evaluation, etc. Parent selection is then biased toward solutions with smaller Pareto

ranks, usually via a tournament selection procedure. To obtain a diversity of solutions along the Pareto frontier, the value of each solution is also divided by a quantity proportional to the number of solutions that are close to it (this is called *sharing* [148]). Hence, if many solutions are in the same region along the Pareto frontier, these solutions get a relatively low evaluation and are less likely to be selected as parents. Their number will thus tend to diminish. Finally, while the GA is executed, an archive that contains all non dominated solutions found during the search is maintained. This archive can also be used by the genetic operators to intensify the search around these solutions.

The authors in [75] address the VRPTW as a true bi-objective problem, where the two objectives are the number of vehicles and total distance. They used a MOGA with Pareto-ranking to this end. The crossover operator, called *BCRC* for Best Cost Route Crossover, works as follows. Routes 1 and 2 are first selected in parents 1 and 2, respectively. Then, customers selected in route 1 are removed from parent 2 and customers selected in route 2 are removed from parent 1. The removed customers are finally reinserted in each parent, in random order, with a least-cost insertion heuristic. Due to the time windows, a constrained inversion mutation operator is designed, where the length of the inverted segment is limited to three consecutive customers. The VRPTW is also addressed by Tan et al. [105, 107] with a similar approach, except that the crossover operator is *RBX* [82]. Also, three different route-based mutation operators and local search heuristics are proposed to further improve the solutions. The same authors have adapted their bi-objective GA to a more complex real-world problem [106, 108]. Here, a company must visit customers with time windows using different combinations of trucks and trailers. A truck-trailer pair moves between a port, customer warehouses, trailer exchange points (where different types of trailers are available) and container depots (where empty containers are available). Each job is made of a number of tasks, and each task is made of an origin and destination point. The *RBX* crossover operator in [107] is applied, as well as a route merge mutation. A local search heuristic that tries to merge small routes into larger routes is also used to further improve the solutions. In all cases, infeasible tasks are outsourced to other companies.

In [51, 52], a bi-objective VRP is also addressed, but with the objectives of minimizing the total distance and balancing the route lengths (more precisely, minimizing the difference between the longest and shortest routes). To obtain solutions that are dispersed along the Pareto frontier, a diversification mechanism based on elitism is proposed. As it is usually the case, an archive contains all non dominated solutions found during the search. Some solutions in this archive are included into the population at each iteration to provide a form of intensification. Diversification is obtained by maintaining additional archives of non dominated solutions where one of the objective is inverted (i.e., one objective is maximized rather than minimized). This approach is used in a context where different subpopulations evolve in parallel and are connected through a toroidal grid topology. Each subpopulation maintains the standard archive plus an additional archive where one of the original objective is inverted. When the best solutions migrate from one subpopulation to another to refresh the genetic pool, each subpopulation also sends its standard archive to all neighboring subpopulations. However, the additional archive is sent only to neighbors with the same inverted objective. The non dominated sorting GA [182] applied within each subpopulation exploits the *RBX* and *SBX* crossover operators [82]. This MOGA is hybridized with a tabu search, which is applied only at the end. The tabu search evaluates each solution in the neighborhood of the current solution based on its dominance relation with the current solution. Overall, the additional archives are shown to improve diversity along the approximate Pareto frontier generated by this GA. A related approach for a bi-objective VRPTW based on the number of vehicles and total distance, but without parallel subpopulations and diversification archives, is reported in [90].

A multi-objective pick-up and delivery problem with time windows, motivated by a

courier service application, is addressed in [62]. Here, transportation requests are received over the day and must be assigned to vehicles in real-time. Multi-criteria utility functions, aimed at evaluating the quality of a vehicle for visiting a transportation request, are evolved with a GA to approximate the (unknown) utility function used by a human dispatcher to take decisions. Basically, the GA evolves the weights of a weighted linear utility function, using real-valued strings. In a follow-up work [10], Genetic Programming (GP) [156] is used for the same purpose. GP is an evolutionary algorithm closely related to GAs, but where tree structures, rather than linear structures, are evolved. These trees encode computer programs that are made of a predefined set of functions and terminals. In the case of vehicle dispatching, the terminal set contains attributes like the detour or the delay introduced in the current vehicle route by a new request, while the functions are basic mathematical operators. This is an interesting application of GP because vehicle utility functions (programs) that are much more complex than weighted linear functions can be generated.

3.11 Parameter Optimization

Genetic algorithms have often been used for parameter optimization. As parameters are often real-valued, both bit strings and real-valued strings have been proposed for this purpose [161]. With a classical bit string representation, a lower and an upper bound on the parameter values are defined, and the bit string is then interpreted as an integer which is mapped within the parameter interval. The GA thus searches in a discrete space that corresponds to a sampling of the real interval. The sampling is more refined when the bit string is longer, but the search space then increases accordingly.

In [9], the parameter values of Solomon's I1 heuristic [181] are optimized with a GA on Solomon's VRPTW benchmark instances. Each parameter is encoded as a bit string and these strings are appended to encode groupings of parameter values. A GA based on classical one-point crossover and mutation operators is used to evolve these parameter groupings. The same approach is proposed in [83, 84] to optimize the parameter values of a parallel insertion heuristic for the VRPTW [172].

In [124], the author analyzes how the various parameters of a GA impact its performance when solving the VRP. This work thus addresses parameter optimization for the GA itself.

4 Evolution strategies

The Evolution strategies (ES) paradigm was developed in the 60's and 70's at the Technical University of Berlin by Rechenberg and Schwefel (for an introductory text on the subject, see [132]). Although mostly applied to real-valued function optimization, discrete variants have emerged for solving combinatorial optimization problems. The algorithmic framework of ES is similar to genetic algorithms as they evolve a population of solutions through mutation and recombination (crossover) operators. One distinctive feature is the concurrent evolution, with the same genetic operators, of the so-called strategy parameters which are usually properties of the mutation operator, like the mutation step size to be applied to each element of the solution. In ES, the recombination operator produces a single offspring and either involve two parents or is extended to a variable number of parents, known as a family of parents. Also, the population replacement mechanism is completely deterministic. That is, for a population of size μ , either the μ best solutions are selected out of λ offspring solutions ($\mu < \lambda$), which is known as a (μ, λ) -selection, or out of the union of the parent and offspring solutions, which is a $(\mu + \lambda)$ -selection. The latter scheme is elitist and guarantees a monotonically improving performance. A special case is the $(1 + 1)$ -ES where there is

no recombination operator, only a mutation operator. The ES search framework can be summarized as follows (note the similarities and differences with the pseudo-code of GAs):

1. Create an initial population of μ individuals, where each individual contains a solution and its associated strategy parameters.
2. Evaluate each solution.
3. Repeat for a fixed number of iterations:
 - 3.1 Repeat until λ offspring are created:
 - 3.1.1 Randomly select ρ parents in the population (with replacement).
 - 3.1.2 Apply recombination to the parent solutions to create a single offspring solution.
 - 3.1.3 Apply recombination to the parent strategy parameters to create the offspring strategy parameters.
 - 3.1.4 Append the offspring solution and strategy parameters to create a complete individual.
 - 3.1.5 Apply mutation (with a certain probability) to the offspring solution.
 - 3.1.6 Apply mutation (with a certain probability) to the offspring strategy parameters.
 - 3.2 Evaluate each offspring solution.
 - 3.3 If (μ, λ) -ES then select the μ best individuals to create the new population.
 - 3.3 If $(\mu + \lambda)$ -ES then select the μ best individuals among the union of the μ parents and λ offspring to create the new population.
4. Return the best solution found.

The first application of ES to a vehicle routing problem is found in the work of Homberger and Gehring [45], where the VRPTW is addressed with a (μ, λ) -ES. The mutation operator is a local search heuristic and the mutation step size then corresponds to the number of moves or modifications to the current solution. There is an additional strategy parameter associated with the mutation operator which is the objective to be favored (either minimization of the total distance or minimization of the number of vehicles). Two evolution strategies, called *ES1* and *ES2* are proposed. In the simple variant *ES1*, there is no recombination, only a mutation operator that modifies the solution based on the following: a random choice between reinsertion of a customer, exchange of two customers or 2-opt* [173]. This is followed by a specialized Or-opt [168] aimed at reducing the number of routes (when this objective is favored) by repeatedly inserting customers from the smallest route into alternate routes. It is worth noting that the two mutation strategy parameters are directly inherited by the offspring without any modification. This is to be opposed to the second variant *ES2*, where a uniform crossover operator [183] is applied to the mutation step size of the parents to generate new values.

The authors note that there is a bias toward the distance, when one tries to simultaneously minimize the distance and the number of vehicles. To put more emphasis on the number of vehicles, the authors propose a new approach to select the solutions that will be part of the new population. First, the λ offspring solutions are evaluated using the following hierarchical ordering of four criteria: number of vehicles, number of customers in the smallest route, minimal time delay introduced in the solution when all customers in the smallest route are inserted into other routes, total distance. The second and third criteria

measure how easily a route can be eliminated from the solution, considering the time window constraints. Based on this ordering, the best κ ($\kappa < \mu$) offspring solutions are put in the new population. Then, the λ offspring are sorted again based on the original objective which is to minimize the number of vehicles first and total distance second. The best $\mu - \kappa$ offspring are then selected to complete the new population. Depending on the ratio κ/μ , the minimization of the number of vehicles is more or less emphasized with regard to the distance.

In a follow-up paper [46], the authors describe a two-stage approach where *ES1* is first applied with the objective of minimizing the number of vehicles, followed by a tabu search heuristic to minimize the total distance. This work is an extension of a two-phase algorithm from the same authors, where a $(1, \lambda)$ -ES was used [32]. A parallel implementation of this algorithm is reported in [33, 34]. Using a coarse-grained architecture made of a number of workstations, the two-stage algorithm is executed in parallel on each machine, but with different starting points and parameter configurations. These independent search threads then communicate periodically by exchanging their best solution (to allow jumps in the search space).

In [66], an algorithm called AGES combines guided local search (GLS) [189, 190] and ES in an iterative two-stage procedure to address both the VRP and VRPTW. GLS introduces modifications into the objective function through penalties when the search gets trapped in a local optimum. Here, the penalty counter of one arc in the solution is incremented by one, each time a local optimum is reached. The arc is selected on the basis of its length (a long arc is more likely to be penalized) and current penalty counter (an arc with a high penalty counter is less likely to be penalized again). In the first stage, GLS controls the objective function of a composite local search, based on reinsertion of a customer, exchange of two customers and 2-opt* [173]. The neighborhoods considered are restricted to routes which are close to the currently penalized arc. This is called a penalty variable neighborhood. If no improvement is found for a number of iterations, the first stage is stopped and the second stage starts. In the second stage, GLS controls the objective function of a $(1+1)$ -ES. In this ES, mutation is performed on a parent solution, where the number and type of modifications considered each time are randomly chosen. When the offspring is better than the parent, it replaces the parent. The ruin and recreate principle is at the core of the proposed mutation. That is, a number of customers are removed from the current solution and are reinserted at least cost. Once again, a penalty variable neighborhood is used to restrict the search.

A refinement of this algorithm for the VRP is reported in [67]. In this implementation, the composite local search is based on some additional intra- and inter-route moves, like Or-opt [168]. The authors also explore different parameter configurations and neighborhood combinations in the composite local search, and identifies what they call a *best* AGES and a *fast* AGES. Another variant of this algorithm is proposed in [68] where the objective function is not modified with GLS. Since there are no penalized arcs, the penalty variable neighborhood cannot be used to restrict the search. Instead, an adaptive variable neighborhood search [162] is proposed based on a partition of the service area into smaller rectangular areas. The original problem is thus divided into smaller subproblems, based on the geographic proximity among customers. The rectangular areas are then processed in a given order. Afterward, they are gradually merged back together until the original service area is restored. The idea is to work on smaller problems when a lot of (time-consuming) improvements are still possible, and to gradually enlarge the neighborhood as the number of potential improving moves diminishes.

5 Particle Swarm Optimization

Particle Swarm Optimization (PSO) is a population-based metaheuristic where each solution, called a particle, is associated with an objective value and a velocity [154]. The latter defines a direction in the search space, as well as an amplitude, for modifying the solution and is determined through interaction with other solutions. The velocity basically replaces operators like crossover and mutation in classical GAs. The velocity of particle i at iteration t is obtained through a randomized weighted sum of three components: the velocity of particle i at iteration $t - 1$, the direction associated with the best solution achieved by particle i and the direction associated with the best solution achieved over all particles (or over particles that are sufficiently close to particle i , depending on the variant). Although PSO has mostly been used for solving continuous optimization problems, discrete variants also exist.

In [25], a binary version of PSO is used, where a solution or particle is a binary vector. The velocity vector is made of values between 0 and 1 and the corresponding solution is obtained by interpreting each velocity value as the probability that the corresponding element in the solution is 0 (and conversely, as 1 minus the probability that the corresponding element is 1). This approach is exploited to solve the VRP. Basically, a particle is a vector of size $n \times m$, where n is the number of customers and m the number of vehicles. That is, the vector of customers is duplicated m times, one time for each vehicle. An entry of 1 means that the customer is served by the corresponding vehicle. The PSO thus solves the assignment subproblem. As the solution obtained is not necessarily valid (i.e., exactly one position associated with a customer should be equal to 1), a repair operator is used. If the solution is valid but infeasible, the solution is discarded and the velocity is recalculated until a feasible solution is obtained. A simulated annealing heuristic is then used to sequence the customers in each route.

In [129], an alternative particle representation with only $n + m - 1$ elements is used for solving the VRPTW. Each entry corresponds to the position of one of the n customers or one of the $m - 1$ copies of the depot in the solution vector. The PSO thus solves the assignment and sequencing subproblems concurrently. The objective value is the total distance plus penalties for violations of the capacity and time window constraints. The velocity is a vector of values between $-n + m - 2$ and $n + m - 2$ which is added to the solution vector to update it. Rounding is then applied to obtain integer positions.

6 Classification by Problem Types

The previous review was divided along methodological lines. The aim of this section is to group all references according to vehicle routing problem types addressed with EAs. These types have already been identified and briefly described in Section 2: the capacitated vehicle routing problem (VRP), VRP with time windows (VRPTW), VRP with time deadlines (VRPTD), Periodic VRP (PVRP), Time-dependent VRP (TVRP), VRP with Heterogeneous Fleet (VRPHF), Multi-Depot VRP (MDVRP), VRP with backhauls (VRPB) and Pickup and Delivery Problem (PDP). The catch-all type *Others* is used to refer to real-world vehicle routing applications that do not fall into any of these categories due to problem-specific features. The row headings in Table 1 correspond to the various problem types, while the column headings correspond to EAs.

Clearly, a high concentration of papers report applications of GAs. But, the few implementations of ES have been very successful (see Section 7). EAs have mostly been applied to the VRPTW, followed by the VRP, which are also the most studied problems in the ve-

Problem \ EA	GA	ES	PSO	GP
VRP	[1][2][7][8][11][12][29][41] [42][43][51][52][56][57][58] [64][70][79][80][88][91][94] [100][111][112][124]	[67][68]	[25]	
VRPTW	[3][4][5][6][9][13][14][15] [16][17][19][20][21][26][44] [47][55][60][61][63][65][69] [71][74][75][76][82][83][84] [86][90][101][102][103][104] [105][107][110][111][112] [113][114][115][116][117] [118][120][123][127][128]	[32][33][34] [45][46][66]	[129]	
VRPTD	[119][122]			
PVRP	[27][125]			
TVRP	[54]			
VRPHF	[72][73][109]			
MDVRP	[28][30][92][98][99][120] [121]			
VRPB	[31][85][109]			
PDP	[49][50][53][62][77][78][81] [95][96][97]			[10]
Others	[48][93][106][108]			

Table 1: EAs for different problem types

hicle routing literature in general. Overall, the number of papers devoted to the VRPTW is substantially larger than those devoted to the VRP. This can be explained as follows. When the first applications of EAs for vehicle routing appeared in the literature, the VRPTW was much less studied than the classical VRP. Thus, there were more opportunities for researchers to improve the best results reported in the literature on VRPTW benchmark instances. Accordingly, the first EAs were often developed for the VRPTW and only later adapted to the VRP, even if the VRP can be considered as an easier problem, due to the absence of time constraints.

7 Computational Performance

The aim of this section is to evaluate how evolutionary algorithms compare with other metaheuristic approaches for solving the VRP, VRPTW and PDPTW. These three problem classes have been chosen because of their importance in the vehicle routing literature and the availability of benchmark instances that allow a fair comparison between competing methods.

7.1 VRPTW

This section focuses on the most effective problem-solving approaches for solving the VRPTW on Solomon's [181] and Homberger and Gehring's [32] benchmark instances. It just happens that EAs, either alone or combined with other metaheuristics, are currently among the best approaches for solving the VRPTW. Another promising trend is represented by the large

neighborhood search [180], based on the ruin-and-recreate principle [179], where customer removal heuristics are applied to destroy a substantial part of the current solution before reconstructing it. All work reported here minimizes the number of vehicles first, followed by the total distance.

7.1.1 Solomon's instances

Solomon's 100-customer instances [181] are divided into six problem classes depending on the geographical distribution of the customers and the width of the time horizon. The latter is set by a time window at the depot that defines an earliest start time and latest end time for each vehicle route. The customers are randomly distributed in instances of type *R*, while they are clustered in instances of type *C*. Instances of type *RC* are a mix of types *R* and *C*. With regard to the time horizon, instances of type 1 have a short horizon and each vehicle can only service a small number of customers. Conversely, instances of type 2 have a long horizon and each vehicle can service a large number of customers. Overall, these characteristics lead to classes *R1*, *R2*, *C1*, *C2*, *RC1* and *RC2*, with 8 to 12 instances in each class. All instances are Euclidean and the time units correspond to the distance units (i.e., the travel time between two customers is the same as the Euclidean distance).

Table 2 contains the best results reported in the literature, at the time of writing, on Solomon's benchmark instances. The two rows associated with each entry correspond to the number of vehicles and total distance, respectively. At the end, the cumulative number of vehicles (CNV) and cumulative total distance (CTD) over the 56 instances are shown for each method. The computer type (CPU) and the average time in seconds T(s) used to solve each instance are also indicated. In the case of CPU, *O* stands for Opteron, *P* stands for Pentium, *X* for Xeon, *SU* for Sun Ultra, *M* for MHz and *G* for GHz. Note also that $u \times v$ in row T(s) means that the algorithm was run for v seconds on a parallel architecture with u processors.

The following algorithms are reported in Table 2:

- *B* is the reactive variable neighborhood search (VNS) of Bräysy [133]. In this algorithm, the number of vehicles is first minimized with an ejection chain-based approach [144]. Then, the distance is minimized with VNS.
- *BBB* is the GA of Berger et al. [14] where two populations are evolved (see Section 3.5). In the first population the distance is minimized based on a fixed number of routes; in the second population, time constraint violations are minimized on solutions with one fewer route.
- *BVH* is the two-stage hybrid local search of Bent and van Hentenryck [131] where simulated annealing is used in the first stage to minimize the number of vehicles and a large neighborhood search is used in the second stage to minimize the distance.
- *HG* is the algorithm of Homberger and Gehring [46] where ES is used in the first stage to minimize the number of routes and a tabu search is used in the second stage to minimize the distance (see Section 4).
- *IAL* is the problem-solving approach of Ibaraki et al. [151], where local search heuristics are embedded within different iterated and multi-start schemes.
- *LC* is the parallel search framework of Le Bouthillier and Crainic [59] where GA and tabu search threads interact through a common solution warehouse (see Section 3.9).

- *PDR* is the recent large neighborhood search of Prescott-Gagnon, Desaulniers and Rousseau [174] which uses a heuristic branch-and-price method to explore the neighborhoods.
- *PR* is the adaptive large neighborhood search of Pisinger and Ropke [170] which adaptively chooses among different removal and insertion heuristics to intensify and diversify the search.

On these small instances, the best methods all reached the same number of vehicles, both globally and for each instance class. They only differ in total distance. The results are the same, for all practical purposes, on classes *C1* and *C2* which are the easiest instances. The recent large neighborhood search *PDR* of Prescott-Gagnon, Desaulniers and Rousseau is the best on two of the four remaining classes, namely *R1* and *RC1*. The parallel algorithm of LeBouthillier and Crainic *LC* and the ES of Homberger and Gehring *HG* are the best on *R2* and *RC2*, respectively. Overall, the best cumulative total distance is obtained by *HG* (note that no CPU time can be reported in this case, because the results have been obtained over the course of many experiments involving many different parameter settings). Also noteworthy is the adaptive large neighborhood search *PR* of Pisinger and Ropke, which is very fast and still reaches a CTD which is at only .2% of the one obtained by *HG*.

7.1.2 Gehring and Homberger's instances

Gehring and Homberger [32] have designed larger VRPTW instances with 200, 400, 600 and 1 000 customers. For each size, six different classes, with 10 instances in each class, have been generated using the procedure of Solomon described in [181]. Thus, there are 60 instances for each problem size and a total of 300 instances. In Table 3, only the CNV and CTD values are reported for each problem size. They are shown in the first two rows of each entry. The computer type and average computation time in minutes are reported in the next two rows. In the case of the computer type, *A* stands for AMD Athlon.

At the time of writing, the best methods for these larger instances are (apart from *BVH*, *LC*, *PDR* and *PR* which have already been introduced in the previous section):

- *BHD* is the multi-start local search of Bräysy, Hasle and Duallert [134]. This algorithm is divided into two stages. The first stage minimizes the number of vehicles using an ejection chain-based mechanism, while the second stage minimizes the distance using CROSS-exchanges. At the end, a deterministic threshold accepting variant of simulated annealing [139] is applied to further reduce the distance.
- *GH99* and *GH01* stand for the parallel algorithms of Gehring and Homberger where a $(1, \lambda)$ - and (μ, λ) -ES, respectively, are applied in the first phase to minimize the number of routes [32, 33]. A tabu search is used in the second phase to minimize the distance (see Section 4).
- *MB* is the AGES algorithm of Mester and Bräysy [66] that combines ES, GLS and local search (see Section 4).

On these larger instances, the large neighborhood search-based *PDR* of Prescott-Gagnon, Desaulniers and Rousseau stand out with the best reported results for each problem size. The adaptive large neighborhood search of Pisinger and Ropke is still noteworthy with excellent results in very fast computation times. Admittedly, the EA-based methods trail a little bit on these instances, but the ES-based algorithm *MB* of Mester and Bräysy and the parallel algorithm *LC* of LeBouthillier and Crainic still produce competitive results in reasonable computation times.

	B	BBB	BVH	HG	IAL	LC	PDR	PR
R1	11.92	11.92	11.92	11.92	11.92	11.92	11.92	11.92
	1222.12	1221.10	1211.10	1212.73	1217.40	1213.06	1210.34	1212.39
R2	2.73	2.73	2.73	2.73	2.73	2.73	2.73	2.73
	975.12	975.43	954.27	955.03	959.11	952.73	955.74	957.72
C1	10.00	10.00	10.00	10.00	10.00	10.00	10.00	10.00
	828.38	828.48	828.38	828.38	828.38	828.38	828.38	828.38
C2	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00
	589.86	589.93	589.86	589.86	589.86	589.86	589.86	589.86
RC1	11.50	11.50	11.50	11.50	11.50	11.50	11.50	11.50
	1389.58	1389.89	1384.17	1386.44	1391.03	1385.14	1384.16	1385.78
RC2	3.25	3.25	3.25	3.25	3.25	3.25	3.25	3.25
	1128.38	1159.37	1124.46	1108.52	1122.79	1129.16	1119.44	1123.49
CNV	405	405	405	405	405	405	405	405
CTD	57710	57952	57273	57192	57444	57325	57240	57332
CPU	P-200M	P-400M	SU 10	P-400M	P3 1G	X 3.6G	O 2.3G	P4 3G
T(s)	4950	1800	7 200	n/a	15 000	5×720	1800	146

Table 2: Solomon's instances

		BHD	BVH	GH99	GH01	LC	MB	PDR	PR
200	CNV	695	697	694	696	694	694	694	694
	CTD	172 406	171 715	176 180	179 328	169 903	168 573	168 556	169 042
	CPU	A-700M	SU 10	P-200M	P-400M	X 3.6G	P4 2G	O 2.3G	P4 3G
	T(m)	2.4	n/a	4×10	4×2.1	5×10	8	53	7.7
400	CNV	1391	1393	1390	1392	1388	1389	1385	1385
	CTD	399 132	410 112	412 270	428 489	396 159	390 386	389 011	393 210
	CPU	A-700M	SU 10	P-200M	P-400M	X 3.6G	P4 2G	O 2.3G	P4 3G
	T(m)	7.9	n/a	4×20	4×7.1	5×20	17	89	15.8
600	CNV	2084	2091	2082	2079	2084	2082	2071	2071
	CTD	820 372	858 040	867 010	890 121	809 882	796 172	800 797	807 470
	CPU	A-700M	SU 10	P-200M	P-400M	X 3.6G	P4 2G	O 2.3G	P4 3G
	T(m)	16.2	n/a	4×30	4×12.9	5×30	40	105	18.3
800	CNV	2776	2778	2770	2760	2759	2765	2745	2758
	CTD	1 384 306	1 469 790	1 515 120	1 535 849	1 444 525	1 361 586	1 391 344	1 358 291
	CPU	A-700M	SU 10	P-200M	P-400M	X 3.6G	P4 2G	O 2.3G	P4 3G
	T(m)	26.2	n/a	4×40	4×23.2	5×40	145	129	22.7
1000	CNV	3465	3468	3461	3446	3439	3446	3432	3438
	CTD	2 133 376	2 266 959	2 276 390	2 290 367	2 133 673	2 078 110	2 096 823	2 110 925
	CPU	A-700M	SU 10	P-200M	P-400M	X 3.6G	P4 2G	O 2.3G	P4 3G
	T(m)	39.6	n/a	4×50	4×30.1	5×50	600	162	26.6

Table 3: Gehring and Homberger's instances

7.2 VRP

The previous section has shown that EAs are at the core of very competitive approaches for the VRPTW. This is also true for the capacitated VRP. There are many VRP benchmark instances in the literature, including those of Christofides, Mingozzi and Toth (CMT) [135], Golden et al. (GAL) [149], van Breedam [188], Augerat et al. [130] and Li, Golden and Wasil [158]. Some real-world VRP instances are also found in [142, 184]. In this section, results are reported for methods that have been applied to both CMT and GAL, the two most widely used data sets in the literature. There are 14 benchmark instances in CMT, ranging from 50 to 199 customers, and 20 larger instances with up to 483 customers in GAL.

The best methods for minimizing the total distance on these benchmark instances are:

- *DK* is the solution attribute-based tabu search of Derigs and Kaiser [138]. Results are shown for the variant *DKP* where the initial routes are constructed in parallel with the savings heuristic of Clarke and Wright [136] and the variant *DKS* where the routes are constructed one by one with the sequential savings heuristic.
- *LGW* is the record-to-record travel heuristic of Li, Golden and Wasil [158]. This heuristic is a deterministic variant of simulated annealing, where a new solution is accepted if its value is within an acceptable margin of the best known solution value (the record) [140].
- *MB* is the AGES algorithm of Mester and Bräysy [67] (see Section 4). Both the best results *MB Best* and those obtained with a fast version of this algorithm *MB Fast* are shown.
- *N* is the GA of Nagata [70], based on the *EAX* crossover operator (see Section 3.1).
- *P* is the GA of Prins [88], where a giant tour is split into feasible routes (see Section 3.4).
- *PR* is the adaptive large neighborhood search of Pisinger and Ropke [170].
- *T* is the tabu search with adaptive memory of Tarantilis [186].

Table 4 reports the best results obtained with these methods on the CMT and GAL instances. The row % refers to the gap in percentage with the best known solutions at the time of writing (this way of reporting results is the norm in the VRP literature) As in the previous tables, row $T(m)$ refers to the average computation time in minutes for solving an instance and row *CPU* is the type of processor used. Note also that *C* stands for Celeron.

The GA of Nagata reaches the optimum on the CMT instances, but was applied on only 7 instances out of 14. Also, its performances degrades on the largest GAL instances. The ES-based algorithm *MB* of Mester and Bräysy looks particularly attractive here. The *Best* variant produces the best solution values on both data sets, if we consider only methods that have been tested on all benchmark instances. Furthermore, the *Fast* variant still produces good solutions, but in a fraction of the computation time required by the other methods (except *LGW* which is also very fast).

7.3 PDPTW

The PDPTW has been included in this study, because this is a problem where EAs have not yet convincingly demonstrated their effectiveness. In fact, not so much work is reported

		DKP	DKS	LGW	MB Best	MB Fast	N	P	PR	T
CMT	%	0.28	0.21	0.41 ^a	0.03	0.08	0.00 ^a	0.24	0.11	0.18
	T(m)	4.54	5.84	0.32	7.72	0.27	2.37	5.19	17.50	6.60
GAL	%	0.79	0.87	1.05	0.11	1.01	2.83 ^b	1.03	0.60	0.71
	T(m)	113.34	112.14	0.97	72.94	0.63	41.37	66.90	107.67	45.48
CPU		C 2.4G	C 2.4G	A 1G	P4 2G	P4 2G	X 3.2G	P3 1G	P4 3G	P2 400M

Table 4: CMT and GAL instances

^aAverage on 7 out of the 14 instances

^bAverage on 12 out of the 20 instances

on the PDPTW in general. Benchmark instances for the PDPTW are typically derived from VRPTW instances by pairing two customers (to get a pickup and a delivery node for each customer request). For example, Nanry and Barnes [166] and Lau and Liang [157] have generated PDPTW instances from Solomon's instances [181] in this way. An extended data set with 100, 200, 400, 600, 800 and 1 000 nodes is also reported by Li and Lim [159]. Recently, a new set of instances with randomly generated pickup and delivery nodes has been produced by Ropke, Cordeau and Laporte [176].

The data set of Li and Lim has emerged as a standard in the last few years. Unfortunately, a single paper reports results on a subset of these instances, thus precluding a true assessment of EAs with the best known methods. In a number of cases, home-made instances have been used to experiment with EAs. In other cases, dynamic variants where new customer requests occur over time have been addressed. Furthermore, some authors minimize the total distance only, while others minimize the number of vehicles first (as opposed to the VRPTW, where most researchers have adopted the same objective, there is no consensus yet). The only application of a GA on benchmark instances comes from the work of Pankratz in [77] (see Section 3.3), where results are reported on the nine 100-node instances of Nanry and Barnes and the 56 100-node instances of Li and Lim, with the objective of minimizing the total distance.

On the instances of Li and Lim, the CNV and CTD of the GA of Pankratz are 410 and 58 191, respectively. On the same instances, Ropke and Pisinger [177] report 402 and 58 060, respectively, with an adaptive large neighborhood search with fixed parameter settings, and with the objective of minimizing the number of vehicles first. It means that they improved upon the CTD of the GA, even if it was not their primary objective. Furthermore, the best results reported by Ropke and Pisinger in [177] over a number of experiments with different parameter settings, is much lower at 56 060.

8 Conclusion

This survey has shown that EAs have been applied in many different ways to address difficult vehicle routing problems. Among EAs, genetic algorithms have been the most widely used, although the few implementations of evolution strategies reported in the literature proved to be very successful. The VRP and VRPTW have been the topic of most of these studies, but recent developments are now observed for less known vehicle routing variants, like the PDPTW. Also, a few real world applications have been reported lately.

Although this survey has focused on pure vehicle routing problems, EAs are now used to address these problems in relation with other logistics problems. For example, the work in [89] describes a GA for a location-routing problem, where depots for the vehicles are located in a transportation network in addition to the optimization of their routes. EAs are thus very well and alive and used to address an ever widening range of vehicle routing applications.

Acknowledgments. Financial support for this work was provided by the Canadian Natural Sciences and Engineering Research Council (NSERC). This support is gratefully acknowledged.

References

- [1] E. Alba and B. Dorronsoro. Solving the vehicle routing problem by using cellular genetic algorithms. *Lecture Notes in Computer Science*, J. Gottlieb and G.R. Raidl, eds., Springer, Berlin, 3004:11–20, 2004.
- [2] E. Alba and B. Dorronsoro. Computing nine new best-so-far solutions for capacitated VRP with a cellular genetic algorithm. *Information Processing Letters*, 98:225–230, 2006.
- [3] G.B. Alvarenga, R.M. de Abreu Silva and R.M. Sampaio. A hybrid algorithm for the vehicle routing problem with time window. *INFOCOMP Journal of Computer Science*, 4:9–16, 2005.
- [4] G.B. Alvarenga and G.R. Mateus. A two-phase genetic and set partitioning approach for the vehicle routing problem with time windows. In *Proceedings of the Fourth International Conference on Hybrid Intelligent Systems*, M. Ishikawa et al., eds., IEEE Computer Society Press, Los Alamitos, CA, 428–433, 2004.
- [5] G.B. Alvarenga and G.R. Mateus. Hierarchical tournament selection genetic algorithm for the vehicle routing problem with time windows. In *Proceedings of the Fourth International Conference on Hybrid Intelligent Systems*, M. Ishikawa et al., eds., IEEE Computer Society Press, Los Alamitos, CA, 410–415, 2004.
- [6] G.B. Alvarenga, G.R. Mateus and G. de Tomi. A genetic and set partitioning two-phase approach for the vehicle routing problem with time windows. *Computers & Operations Research*, 34:1561–1584, 2007.
- [7] B.M. Baker and M.A. Ayechev. A genetic algorithm for the vehicle routing problem. *Computers & Operations Research*, 30:787–800, 2003.
- [8] J.C. Bean. Genetic algorithms and random keys for sequencing and optimization. *ORSA Journal on Computing*, 6:154–160, 1994.
- [9] I. Benyahia and J.-Y. Potvin. Generalization and refinement of route construction heuristics using genetic algorithms. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, IEEE Press, Piscataway, NJ, volume I, 39–43, 1995.
- [10] I. Benyahia and J.-Y. Potvin. Decision support for vehicle dispatching using genetic programming. *IEEE Transactions on Systems, Man and Cybernetics*, 28:306–314, 1998.
- [11] J. Berger and M. Barkaoui. A hybrid genetic algorithm for the capacitated vehicle routing problem. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Lecture Notes in Computer Science 2723, E. Cantú-Paz et al., eds., Chicago, IL, 646–656, 2003.
- [12] J. Berger and M. Barkaoui. A new hybrid genetic algorithm for the capacitated vehicle routing problem. *Journal of the Operational Research Society*, 54:1254–1262, 2003.
- [13] J. Berger and M. Barkaoui. A parallel hybrid genetic algorithm for the vehicle routing problem with time windows. *Computers & Operations Research*, 31:2037–2053, 2004.
- [14] J. Berger, M. Barkaoui and O. Bräysy. A route-directed hybrid genetic approach for the vehicle routing problem with time windows. *INFOR*, 41:179–194, 2003.

- [15] J. Berger, M. Salois and R. Bégin. A hybrid genetic algorithm for the vehicle routing problem with time windows. In *Lecture Notes in Computer Science*, R.E. Mercer and E. Neufeld, eds., Springer, London, UK, 1418:114–127, 1998.
- [16] J. Berger, M. Sassi and M. Salois. A hybrid genetic algorithm for the vehicle routing problem with time windows and itinerary constraints. In *Proceedings of the Genetic and Evolutionary Computation Conference*, W. Banhaf et al., eds., Morgan Kaufmann, San Francisco, CA, 44–51, 1999.
- [17] J.L. Blanton and R.L. Wainwright. Multiple vehicle routing with time and capacity constraints using genetic algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, S. Forrest, ed., Morgan Kaufmann, San Mateo, CA, 451–459, 1993.
- [18] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*, 35:268–308, 2003.
- [19] H.C Brandao de Oliveira, J.L. Alexandrino and M. Moreira de Souza. Memetic and genetic algorithms: A comparison among different approaches to solve vehicle routing problem with time windows. In *Proceedings of the Sixth International Conference on Hybrid Intelligent Systems*, IEEE Computer Society Press, Los Alamitos, CA, 55, 2006.
- [20] O. Bräysy, J. Berger and M. Barkaoui. A new hybrid evolutionary algorithm for the vehicle routing problem with time windows. *Working Paper*. Presented at Route 2000 Workshop, Skodsborg, Denmark, 2000.
- [21] O. Bräysy and W. Dullaert. A fast evolutionary metaheuristic for the vehicle routing problem with time windows. *International Journal on Artificial Intelligence Tools*, 12:153–172, 2003.
- [22] O. Bräysy, W. Dullaert and M. Gendreau. Evolutionary algorithms for the vehicle routing problem with time windows. *Journal of Heuristics*, 10:587–611, 2004.
- [23] O. Bräysy and M. Gendreau. Vehicle routing problem with time windows, Part II: Metaheuristics. *Transportation Science*, 39:119–139, 2005.
- [24] P. Calégari, G. Coray, A. Hertz, D. Kobler and P. Kluonen. A taxonomy of evolutionary algorithms in combinatorial optimization. *Journal of Heuristics*, 5:145–158, 1999.
- [25] A.-L. Chen, G.K. Yang and Z.M. Wu. Hybrid discrete particle swarm optimization algorithm for capacitated vehicle routing problem. *Journal of Zhejiang University SCIENCE A*, 7:607–614, 2006.
- [26] A.J. Chin, H.W. Kit and A. Lim. A new GA approach for the vehicle routing problem. In *Proceedings of the 11th IEEE International Conference on Tools with Artificial Intelligence*, IEEE Press, Piscataway, NJ, 307–310, 1999.
- [27] L.M.A. Drummond, L.S. Ochi and D.S. Vianna. An asynchronous parallel metaheuristic for the period vehicle routing problem. *Future Generation Computer Systems*, 17:379–386, 2001.
- [28] M. Filipec, D. Skrlec and S. Krajcar. Darwin meets computers: New approach to multiple depot capacitated vehicle routing problem. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, IEEE Press, Piscataway, NJ, Vol. 1, 421–426, 1997.

- [29] M. Filipec, D. Skrlec and S. Krajcar. An efficient implementation of genetic algorithms for constrained vehicle routing problem. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, IEEE Press, Piscataway, NJ, Vol. 3, 2231–2236, 1998.
- [30] M. Filipec, D. Skrlec and S. Krajcar. Genetic algorithm approach for multiple depot capacitated vehicle routing problem solving with heuristic improvements. *International Journal of Modelling & Simulation*, 20:320–328, 2000.
- [31] K. Ganesh and T.T. Narendran. CLOVES: A cluster-and-search heuristic to solve the vehicle routing problem with delivery and pick-up. *European Journal of Operational Research*, 178:699–717, 2007.
- [32] H. Gehring and J. Homberger. A parallel hybrid evolutionary metaheuristic for the vehicle routing problem with time windows. In *Proceedings of EUROGEN99 - Short Course on Evolutionary Algorithms in Engineering and Computer Science*, K. Miettinen, M.M. Mäkelä and J. Toivanen, eds., Reports of the Department of Mathematical Information Technology, No. A 2/1999, University of Jyväskylä, 57–64, 1999.
- [33] H. Gehring and J. Homberger. A parallel two-phase metaheuristic for routing problems with time windows. *Asia-Pacific Journal of Operational Research*, 18:35–47, 2001.
- [34] H. Gehring and J. Homberger. Parallelization of a two-phase metaheuristic for routing problems with time windows. *Journal of Heuristics*, 8:251–276, 2002.
- [35] M. Gendreau, G. Laporte and J.-Y. Potvin. Vehicle routing: Modern heuristics. In *Local Search in Combinatorial Optimization*, J.K. Lenstra and E.H.L. Aarts, eds., Wiley, Chichester, 311–336, 1997.
- [36] M. Gendreau, G. Laporte and J.-Y. Potvin. Metaheuristics for the capacitated VRP. In *The Vehicle Routing Problem*, P. Toth and D. Vigo, eds., SIAM Series on Discrete Mathematics and Applications, Philadelphia, PA, 129–154, 2002.
- [37] M. Gendreau, G. Laporte, J.-Y. Potvin and F. Semet. Classical and modern heuristics for the vehicle routing problem. *International Transactions in Operational Research*, 7:285–300, 2000.
- [38] M. Gendreau, G. Laporte, J.-Y. Potvin and F. Semet. Métaheuristiques pour le problème des tournées de véhicules. In *Résolution de Problèmes de RO par les Métaheuristiques*, M. Pirlot and J. Teghem, eds., Hermès, Paris, 49–70, 2003.
- [39] M. Gendreau, J.-Y. Potvin. Metaheuristics in combinatorial optimization. *Annals of Operations Research*, 140:189–213, 2005.
- [40] M. Gendreau, J.-Y. Potvin, O. Bräysy, G. Hasle and A. Løkketangen. Metaheuristics for the vehicle routing problem and extensions: A categorized bibliography. Forthcoming in *The Vehicle Routing Problem: Latest Advances and New Challenges*, B.L. Golden, S. Raghavan and E. Wasil, eds., Springer.
- [41] S. Han. A hybrid meta-heuristic algorithm for vehicle scheduling problem: Genetic algorithm and tabu search. *NUCB Journal of Economics and Information Science*, 48:343–358, 2004.
- [42] S. Han and Y. Tabata, A hybrid genetic algorithm for the vehicle routing problem with controlling lethal gene. *International Journal of Asian Pacific Management Review*, 7:287–298, 2002.

- [43] F.T. Hanshar and B.M Ombuki-Berman, Dynamic vehicle routing using genetic algorithms. *Applied Intelligence*, 27:89-99, 2007.
- [44] W.-K. Ho, J. Chin and A. Lim, A hybrid search algorithm for the vehicle routing problem with time windows. *International Journal on Artificial Intelligence Tools*, 10:431-449, 2001.
- [45] J. Homberger and H. Gehring. Two evolutionary metaheuristics for the vehicle routing problem with time windows. *INFOR*, 37:297-318, 1999.
- [46] J. Homberger and H. Gehring. A two-phase hybrid metaheuristics for the vehicle routing problem with time windows. *European Journal of Operational Research*, 162:220-238, 2005.
- [47] H.-S. Hwang. An improved model for vehicle routing with time constraint based on a genetic algorithm. *Computers & Industrial Engineering*, 42:361-369, 2002.
- [48] A. Jaszkievicz and P. Kominek. Genetic local search with distance preserving recombination operator for a vehicle routing problem. *European Journal of Operational Research*, 151:352-364, 2003.
- [49] W.-R. Jih and J. Yung-Jen Hsu. Dynamic vehicle routing using hybrid genetic algorithms. In *Proceedings of the IEEE International Conference on Robotics and Automation*, IEEE Press, Piscataway, NJ, 1:453-458, 1999.
- [50] R.M. Jorgensen, J. Larsen and K.B. Bergvinsdottir. Solving the dial-a-ride problem using genetic algorithms. *Journal of the Operational Research Society*, 58:1321-1331, 2007.
- [51] N. Jozefowicz, F. Semet and E.G. Talbi. Parallel and hybrid models for multi-objective optimization: Application to the vehicle routing problem. *Lecture Notes in Computer Science*, J. Guervós et al., eds., Springer, Berlin, 2439:271-280, 2002.
- [52] N. Jozefowicz, F. Semet and E.G. Talbi. Enhancements of NSGA-II and its application to the vehicle routing problem with route balancing. *Lecture Notes in Computer Science*, E.G. Talbi et al., eds., Springer, Berlin, 3871:131-142, 2006.
- [53] S. Jung and A. Haghani. Genetic algorithm for a pickup and delivery problem with time windows. In *Transportation Research Record*, 1733:1-7, 2000.
- [54] S. Jung and A. Haghani. Genetic algorithm for the time-dependent vehicle routing problem. In *Transportation Research Record*, 1771:164-171, 2001.
- [55] S. Jung and B.-R. Moon. A hybrid genetic algorithm for the vehicle routing problem with time windows. In *Proceedings of the Genetic and Evolutionary Computation Conference*, W.B. Langdon et al., eds., Morgan Kaufmann, San Francisco, CA, 1309-1316, 2002.
- [56] S. Krajcar, D. Skrllec, B. Pribicevic and S. Blagajac. GA Approach to Solving Multiple Vehicle Routing Problem. *Lecture Notes in Computer Science*, C.A. Pinto-Ferreira and N.J. Mamede, eds., Springer, Berlin, 990:473-481, 1995.
- [57] M. Kubiak. Systematic construction of recombination operators for the vehicle routing problem. *Foundations of Computing and Decision Sciences*, 29:205-226, 2004.
- [58] M. Kubiak and P. Wesolek. Accelerating local search in a memetic algorithm for the capacitated vehicle routing problem. *Lecture Notes in Computer Science*, C. Cotta and J. van Hemert, eds., Springer, Berlin, 4446:96-107, 2007.

- [59] A. Le Bouthillier. Recherches coopératives pour la résolution de problèmes d'optimisation combinatoire. Thèse de Doctorat, Département d'informatique et de recherche opérationnelle, Université de Montréal, Canada, 2007.
- [60] A. Le Bouthillier and T.G. Crainic. A cooperative parallel meta-heuristic for the vehicle routing problem with time windows. *Computers & Operations Research*, 32:1685–1708, 2005.
- [61] A. Le Bouthillier, T.G. Crainic and P. Kropf. A guided cooperative search for the vehicle routing problem with time windows. *IEEE Intelligent Systems*, 20:36–42, 2005.
- [62] F. Leclerc and J.-Y. Potvin. Genetic algorithms for vehicle dispatching. *International Transactions in Operational Research*, 4:391–400, 1997.
- [63] S.J. Louis, X. Yin and Z.Y. Yuan. Multiple vehicle routing with time windows using genetic algorithms. In *Proceedings of the 1999 Congress on Evolutionary Computation*, P.J. Angeline et al., eds., IEEE Press, Piscataway, NJ, 1804–1808, 1999.
- [64] P. Machado, J. Tavares, F.B. Pereira, and E. Costa. Vehicle routing problem: Doing it the evolutionary way. In *Proceedings of the Genetic and Evolutionary Computation Conference*, W.B. Langdon et al., eds., Morgan Kaufmann, San Francisco, CA, 690, 2002.
- [65] O. Maeda, M. Nakamura, B. Ombuki and K. Onaga. A genetic algorithm approach to vehicle routing problem with time deadlines in geographical information systems. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, IEEE Press, Piscataway, NJ, 595–600, 1999.
- [66] D. Mester and O. Bräysy. Active guided evolution strategies for large-scale vehicle routing problems with time windows. *Computers & Operations Research*, 32:1593–1614, 2005.
- [67] D. Mester and O. Bräysy. Active guided evolution strategies for large-scale capacitated vehicle routing problems. *Computers & Operations Research*, 34:2964–2975, 2007.
- [68] D. Mester, O. Bräysy and W. Dullaert. A multi-parametric evolution strategies algorithm for vehicle routing problems. *Experts Systems with Applications*, 32:508–517, 2007.
- [69] N.H. Moin. Hybrid genetic algorithms for vehicle routing problems with time windows. *International Journal of the Computer, the Internet and Management*, 10:51–67, 2002.
- [70] Y. Nagata. Edge assembly crossover for the capacitated vehicle routing problem. *Lecture Notes in Computer Science*, C. Cotta and J. van Hemert, eds., Springer, Berlin, 4446:142–153, 2007.
- [71] Y. Nagata. Effective memetic algorithm for the vehicle routing problem with time windows: Edge assembly crossover for the VRPTW. In *Proceedings of the Seventh Metaheuristics International Conference*, Montreal, Canada, 2007 (CD-ROM).
- [72] L.S. Ochi, D.S. Vianna, L.M.A. Drummond and A.O. Victor. An evolutionary hybrid metaheuristic for solving the vehicle routing problem with heterogeneous fleet. *Lecture Notes in Computer Science*, 1391:187–195, 1998.
- [73] L.S. Ochi, D.S. Vianna, L.M.A. Drummond and A.O. Victor. A parallel evolutionary algorithm for the vehicle routing problem with heterogeneous fleet. *Future Generation Computer Systems*, 14:285–292, 1998.

- [74] B. Ombuki, M. Nakamura and O. Maeda. A hybrid search based on genetic algorithms and tabu search for vehicle routing. In *Proceedings of the 6th IASTED International Conference on Artificial Intelligence and Soft Computing*, ACTA Press, Calgary, Canada, 176–181, 2002.
- [75] B. Ombuki, B.J. Ross and F. Hanshar. Multi-objective genetic algorithms for vehicle routing problem with time windows. *Applied Intelligence*, 24:17–30, 2006.
- [76] H.T. Ozdemir and C.K. Mohan. Evolving schedule graphs for the vehicle routing problem with time windows. In *Proceedings of the 2000 Congress on Evolutionary Computation*, IEEE Press, Piscataway, NJ, vol. 2, 888–895, 2000.
- [77] G. Pankratz. A grouping genetic algorithm for the pick up and delivery problem with time windows. *OR Spectrum*, 27:21–41, 2005.
- [78] G. Pankratz. Dynamic vehicle routing by means of a genetic algorithm. *International Journal of Physical Distribution & Logistics*, 35:362–383, 2005.
- [79] J.P. Pedroso. Niche search: An application in vehicle routing. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, IEEE Press, Piscataway, NJ, 177–182, 1998.
- [80] F.B. Pereira, J. Tavares, P. Machado and E. Costa. GVR: A new genetic representation for the vehicle routing problem. In *Lecture Notes in Computer Science*, M. O’Neill et al. eds., Springer, Berlin, 2464:95–102, 2002.
- [81] T. Potter and T. Bossomaier. Solving vehicle routing problems with genetic algorithms. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, IEEE Press, Piscataway, NJ, 2:788–793, 1995.
- [82] J.-Y. Potvin and S. Bengio, The vehicle routing problem with time windows - Part II: Genetic search”, *INFORMS Journal on Computing*, 8:165-172, 1996.
- [83] J.-Y. Potvin and D. Dubé. Improving a vehicle routing heuristic through genetic search. In *Proceedings of the IEEE International Conference on Evolutionary Computation*, IEEE Press, Piscataway, NJ, 194–199, 1994.
- [84] J.-Y. Potvin, D. Dubé and C. Robillard. A hybrid approach to vehicle routing using neural networks and genetic algorithms”. *Applied Intelligence*, 6:241–252, 1996.
- [85] J.-Y. Potvin, C. Duhamel and F. Guertin. A genetic algorithm for vehicle routing with backhauling. *Applied Intelligence*, 6:345-355, 1996.
- [86] J.-Y. Potvin and F. Guertin. Coupling a greedy route construction heuristic with a genetic algorithm for the vehicle routing problem with time windows. In *Advances in Metaheuristics, Optimization and Stochastic Modeling Technologies*, R.S. Barr, R.V. Helgason and J.L. Kennington, eds., Kluwer, Boston, MA, 423–442, 1997.
- [87] J.-Y. Potvin and S.R. Thangiah. Vehicle routing through simulation of natural processes. In *Fusion of Neural Networks, Fuzzy Sets and Genetic Algorithms: Industrial Applications*, L.C. Jain and N.M. Martin, eds., CRC Press, Boca Raton, FL, 141–168, 1999.
- [88] C. Prins. A simple and effective evolutionary algorithm for the vehicle routing problem. *Computers & Operations Research*, 31:1985–2002, 2004.

- [89] C. Prins, C. Prodhon and R.W. Calvo. A memetic algorithm with population management (MA/PM) for the capacitated location-routing problem. In *Lecture Notes in Computer Science*, J. Gottlieb and G.R. Raidl, eds., Springer, Berlin, 3906:183–194, 2006.
- [90] M. Rahoual, B. Kitoun, M.-H. Mabed, V. Bachelet and F. Benameur. Multicriteria genetic algorithms for the vehicle routing problem with time windows. In *Proceedings of the 4th Metaheuristics International Conference*, Porto, Portugal, 527–532, 2001.
- [91] M.A. Russell and G.B. Lamont. A genetic algorithm for unmanned aerial vehicle routing. In *Proceedings of the Genetic and Evolutionary Computation Conference*, H.-G. Beyer et al., eds., ACM Press, New York, NY, 1523–1530, 2005.
- [92] S. Salhi, S.R. Thangiah and F. Rahman. A genetic clustering method for the multi-depot vehicle routing problem. In *Artificial Neural Networks and Genetic Algorithms*, G.D. Smith, N.C. Steele and R.F. Albrecht, eds., Springer Verlag, Austria, 234–237, 1998.
- [93] H.G. Santos, L.S. Ochi, E.H. Marinho and L.M.A. Drummond. Combining an evolutionary algorithm with data mining to solve a single-vehicle routing problem. *Neuro-computing*, 70:70–77, 2006.
- [94] L.J. Schmitt. An empirical computational study of genetic algorithms to solve order based problems: An emphasis on the TSP and VRPTC. Ph.D. dissertation, Fogelman College of Business and Economics, University of Memphis, TN, 1994.
- [95] J. Schönberger and H. Kopfer. A memetic algorithm for a pickup and delivery selection problem. In *Proceedings of the International Conference on Computational Intelligence for Modeling, Control and Automation*, M. Mohammadian, ed., Vienna, Austria, 2003 (CD-ROM).
- [96] J. Schönberger and H. Kopfer. A combined approach to solve the pickup and delivery selection problem. In *Operations Research Proceedings 2002*, U. Leopold-Wildburger, F. Rendl and G. Wäscher, eds., Springer, Berlin, 150–155, 2003.
- [97] J. Schönberger and H. Kopfer. Planning the incorporation of logistics service providers to fulfill precedence and time window-constrained transport requests in a most profitable way. In *Distribution Logistics: Advanced Solutions to Practical Problems*, B. Fleischmann and A. Klose, eds., Springer, Berlin, 141–158, 2004.
- [98] M. Skok, D. Skrlec and S. Krajcar. The non-fixed destination multiple depot capacitated vehicle routing problem and genetic algorithms. In *Proceedings of the 22nd International Conference on Information Technology Interfaces*, SRCE University Computing Centre, University of Zagreb, Zagreb, Croatia, 403–408, 2000.
- [99] M. Skok, D. Skrlec and S. Krajcar. The genetic algorithm method for multiple depot capacitated vehicle routing problem solving. In *Proceedings of the Fourth International Conference on Knowledge-Based Intelligent Engineering Systems & Allied Technologies*, IEEE Press, Piscataway, NJ, 520–526, 2000.
- [100] D. Skrlec, M. Filipec and S. Krajcar. A heuristic modification of genetic algorithm used for solving the single depot capacitated vehicle routing problem. In *Proceedings of Intelligent Information Systems*, IEEE Computer Society, Los Alamitos, CA, 184–188, 1997.
- [101] K.C. Tan, L.H. Lee and K. Ou. Hybrid genetic algorithms in solving vehicle routing problems with time window constraints. *Asia-Pacific Journal of Operational Research*, 18:121–130, 2001.

- [102] K.C. Tan, L.H. Lee and K. Ou. Artificial intelligence heuristics in solving vehicle routing problems with time window constraints. *Engineering Applications of Artificial Intelligence*, 14:825–837, 2001.
- [103] K.C. Tan, T.H. Lee, K. Ou and L.H. Lee. A messy genetic algorithm for the vehicle routing problem with time window constraints. In *Proceedings of the 2001 Congress on Evolutionary Computation*, IEEE Press, Piscataway, NJ, Vol. 1, 679–686, 2001.
- [104] K.C. Tan, L.H. Lee, K.Q. Zhu and K. Ou. Heuristic methods for vehicle routing problem with time windows. *Artificial Intelligence in Engineering*, 15:281–295, 2001.
- [105] K.C. Tan, T.H. Lee, Y.H. Chew and L.H. Lee. A multiobjective evolutionary algorithm for solving vehicle routing problem with time windows. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, IEEE Press, Piscataway, NJ, Vol. 1, 361–366, 2003.
- [106] K.C. Tan, T.H. Lee, Y.H. Chew and L.H. Lee. A hybrid multiobjective evolutionary algorithm for solving truck and trailer vehicle routing problems. In *Proceedings of the 2003 Congress on Evolutionary Computation*, IEEE Press, Piscataway, NJ, Vol. 3, 2134–2141, 2003.
- [107] K.C. Tan, Y.H. Chew and L.H. Lee. A hybrid multiobjective evolutionary algorithm for solving vehicle routing problem with time windows. *Computational Optimization and Applications*, 34:115–151, 2006.
- [108] K.C. Tan, Y.H. Chew and L.H. Lee. A hybrid multi-objective evolutionary algorithm for solving truck and trailer vehicle routing problems. *European Journal of Operational Research*, 172:855–885, 2006.
- [109] R. Tavakkoli-Moghaddam, A.R. Saremi and M.S. Ziaee. A memetic algorithm for the vehicle routing problem with backhauls. *Applied Mathematics and Computation*, 181:1049–1060, 2006.
- [110] J. Tavares, F.B. Pereira, P. Machado and E. Costa. GVR delivers it on time. In *Proceedings of the 4th Asia-Pacific Conference on Simulated Evolution and Learning*, Singapore, 745–749, 2002.
- [111] J. Tavares, F.B. Pereira, P. Machado and E. Costa. On the influence of GVR in vehicle routing. In *Proceedings of the ACM Symposium on Applied Computing*, Melbourne, FL, 753–758, 2003.
- [112] J. Tavares, F.B. Pereira, P. Machado and E. Costa. Crossover and diversity: A study about GVR. In *Proceedings on the Analysis and Design of Representations and Operators: A Bird-of-a-Feather Workshop at the 2003 Genetic and Evolutionary Computation Conference*, Chicago, IL, 27–33, 2003.
- [113] S.R. Thangiah. An adaptive clustering method using a geometric shape for vehicle routing problems with time windows. In *Proceedings of the 6th International Conference on Genetic Algorithms*, L.J. Eshelman, ed., Morgan Kaufmann, San Francisco, CA, 536–543, 1995.
- [114] S.R. Thangiah. Vehicle routing with time windows using genetic algorithms. In *Practical Handbook of Genetic Algorithms: New Frontiers*, L. Chambers, ed., CRC Press, Boca Raton, FL, vol. II: 253–277, 1995.

- [115] S.R. Thangiah. A Hybrid genetic algorithm, simulated annealing and tabu search heuristic for vehicle routing problems with time windows. In *Practical Handbook of Genetic Algorithms: Complex Coding Systems*, L. Chambers, ed., CRC Press, Boca Raton, FL, vol. III: 347–384, 1999.
- [116] S.R. Thangiah and A.V. Gubbi. Effect of genetic sectoring on vehicle routing problems with time windows. In *Proceedings of the IEEE International Conference on Developing and Managing Intelligent System Projects*, IEEE Computer Society Press, Los Alamitos, CA, 146–153, 1993.
- [117] S.R. Thangiah and K.E. Nygard. School bus routing using genetic algorithms. In *Applications of Artificial Intelligence X: Knowledge-Based Systems*, G. Biswas, ed., SPIE, 387–398, 1992.
- [118] S.R. Thangiah, K.E. Nygard and P.L. Juell. GIDEON: A genetic algorithm system for vehicle routing with time windows. In *Proceedings of Seventh IEEE Conference on Artificial Intelligence Applications*, IEEE Computer Society Press, Los Alamitos, CA, 322–328, 1991.
- [119] S.R. Thangiah, I.H. Osman, R. Vinayagamoorthy and T. Sun. Algorithms for the vehicle routing problems with time deadlines. *American Journal of Mathematical and Management Sciences*, 13:323–355, 1995.
- [120] S.R. Thangiah and P. Petrovic. Introduction to genetic heuristics and vehicle routing problems with complex constraints. In *Advances in Computational and Stochastic Optimization, Logic Programming, and Heuristic Search*, D.L. Woodruff, ed., Kluwer, Boston, MA, 253–286, 1998.
- [121] S.R. Thangiah and S. Salhi. Genetic clustering: An adaptive heuristic for the multi-depot vehicle routing problem. *Applied Artificial Intelligence*, 15:361–383, 2001.
- [122] S.R. Thangiah, R. Vinayagamoorthy and A.V. Gubbi. Vehicle routing with time deadlines using genetic and local algorithms. In *Proceedings of the 5th International Conference on Genetic Algorithms*, S. Forrest, ed., Morgan Kaufmann, San Mateo, CA, 506–513, 1993.
- [123] Z. Tong, L. Ning and S. Debao. Genetic algorithm for vehicle routing problem with time window with uncertain vehicle number. In *Proceedings of the 5th World Congress on Intelligent Control and Automation*, Hangzhou, China, 2846–2849, 2004.
- [124] A. Van Breedam. An analysis of the effect of local improvement operators in genetic algorithm and simulated annealing for the vehicle routing problem. RUCA Working Paper 96/14, University of Antwerp, Belgium, 1996.
- [125] D.S. Vianna, L.S. Ochi and L.M.A. Drummond. A parallel evolutionary metaheuristic for the period vehicle routing problem. *Lecture Notes in Computer Science*, J.P.D. Rolim et al., eds., Springer, Berlin, 1586:183–191, 1999.
- [126] M. Yagiura and T. Ibaraki. On metaheuristic algorithms for combinatorial optimization problems. *Systems and Computers in Japan*, 32:33–55, 2001.
- [127] K.Q. Zhu. A new genetic algorithm for VRPTW. In *Proceedings of the International Conference on Artificial Intelligence*, H.R. Arabnia, ed., CSREA Press, 1–10, 2000.
- [128] K.Q. Zhu. A diversity-controlling adaptive genetic algorithm for the vehicle routing problem with time windows. In *Proceedings of the 15th International Conference on Tools with Artificial Intelligence*, IEEE Computer Society, Los Alamitos, CA, 176–183, 2003.

- [129] Q. Zhu, L. Qian, Y. Li and S. Zhu. An improved particle swarm optimization algorithm for vehicle routing problem with time windows. In *Proceedings of the IEEE Congress on Evolutionary Computation International*, IEEE Press, Piscataway, NJ, 1386–1390, 2006.

Auxiliary references

- [130] P. Augerat, J.M. Belenguer, E. Benavent, A. Corberán, D. Naddef and G. Rinaldi. Computational results with a branch and cut code for the capacitated vehicle routing problem. *Research Report 949-M*, Université Joseph Fourier, Grenoble, France.
- [131] R. Bent and P. Van Hentenryck. A two-stage hybrid local search for the vehicle routing problem with time windows. *Transportation Science*, 38:515–530, 2004.
- [132] H.-G. Beyer and H.-P. Schwefel. Evolution strategies: A comprehensive introduction. *Natural Computing*, 1:3–52, 2002.
- [133] O. Bräysy. A reactive variable neighborhood search for the vehicle routing problem with time windows. *INFORMS Journal on Computing*, 15:347–368, 2003.
- [134] O. Bräysy, G. Hasle and W. Dullaert. A multi-start local search algorithm for the vehicle routing problem with time windows. *European Journal of Operational Research*, 159:586–605, 2004.
- [135] N. Christofides, A. Mingozzi and P. Toth. The vehicle routing problem. In *Combinatorial Optimization*, N. Christofides, A. Mingozzi, P. Toth and C. Sandi, eds., Wiley, Chichester, UK, 315–338, 1979.
- [136] G. Clarke and J. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12:568–581, 1964.
- [137] K. Deb, S. Agrawa, A. Pratap and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. *Lecture Notes in Computer Science*, M. Schoenauer et al., eds., Springer, Berlin, 1917:849–858, 2000.
- [138] U. Derigs and R. Kaiser. Applying the attribute based hill climber heuristic to the vehicle routing problem. *European Journal of Operational Research*, 177:719–732, 2007.
- [139] G. Dueck and T. Scheurer. Threshold accepting: a general purpose optimization algorithm. *Journal of Computational Physics*, 90:161–175, 1990.
- [140] G. Dueck. New optimization heuristics: the great deluge algorithm and the record-to-record travel. *Journal of Computational Physics*, 104:86–92, 1993.
- [141] D. Feillet, P. Dejax and M. Gendreau. Traveling salesman problem with profits. *Transportation Science*, 39:188–205, 2005.
- [142] M.L. Fisher. Optimal solution of vehicle routing problems using minimum k-trees. *Operations Research*, 42:626–642, 1994.
- [143] B.R. Fox and M.B. McMahon. Genetic operators for sequencing problems. In *Foundations of Genetic Algorithms*, G.J.E. Rawlins, ed., Morgan Kaufmann, San Mateo, CA, 284–300, 1991.
- [144] F. Glover. Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discrete Applied Mathematics*, 65:223–253, 1996.

- [145] F. Glover and M. Laguna. *Tabu Search*. Kluwer, Boston, MA, 1997.
- [146] D.E. Goldberg, B. Korb and K. Deb. Messy genetic algorithms: Motivation, analysis and first results. *Complex Systems*, 3:493-530, 1989.
- [147] D.E. Goldberg and J. Richardson. Alleles, loci, and the traveling salesman problem. In *Proceedings of the First International Conference on Genetic Algorithms and Their Applications*, J.J. Grefenstette, ed., Lawrence Erlbaum Associates, Hillsdale, NJ, 154–159, 1988.
- [148] D.E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Applications*, J.J. Grefenstette, ed., Lawrence Erlbaum Associates, Hillsdale, NJ, 41–49, 1987.
- [149] B.L. Golden, E.A. Wasil, J.P. Kelly and I.-M. Chao. The impact of metaheuristics on solving the vehicle routing problem: Algorithms, problem sets and computational results. In *Fleet Management and Logistics*, T.G. Crainic and G. Laporte, eds., Kluwer, Norwell, MA, 33–56, 1998.
- [150] J.H. Holland. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, MI, 1975 (reprinted in 1992 by The MIT Press, Cambridge, MA).
- [151] T. Ibaraki, S. Imahori, M. Kubo, T. Masuda, T. Uno and M. Yaigura. Effective local search algorithms for routing and scheduling problems with general time window constraints. *Transportation Science*, 39:206–232, 2005.
- [152] D.R. Jones and M.A. Beltramo. Solving partitioning problems with genetic algorithms. In *Proceedings of the 4th International Conference on Genetic Algorithms*, R.K. Belew and L.B. Booker, eds., Morgan Kaufmann, San Mateo, CA, 442–449, 1991.
- [153] S. Jung and B.-R. Moon. The natural crossover for the 2D Euclidean TSP. In *Proceedings of the Genetic and Evolutionary Computation Conference*, D. Whitley et al., eds., Morgan Kaufmann, San Francisco, CA, 1003–1010, 2000.
- [154] J. Kennedy and R.C. Eberhart. *Swarm Intelligence*. Morgan Kaufmann, San Francisco, CA, 2001.
- [155] S. Kirkpatrick, C.D. Gelatt Jr. and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
- [156] J.R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press, Cambridge, MA, 1992.
- [157] H.C. Lau and Z. Liang. Pickup and delivery with time windows: algorithms and test cases. In *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence*, IEEE Computer Society, Los Alamitos, CA, 333–340, 2001.
- [158] F. Li, B.L. Golden and E.A. Wasil. Very large-scale vehicle routing: new test problems, algorithms and results. *Computers & Operations Research*, 32:1165–1179, 2005.
- [159] H. Li and A. Lim. A metaheuristic for solving the pickup and delivery problem with time windows. In *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence*, IEEE Computer Society, Los Alamitos, CA, 160–167, 2001.
- [160] S. Lin. Computer solutions of the traveling salesman problem. *Bell System Technical Journal*, 44:2245–2269, 1965.

- [161] Z. Michalewicz. Genetic Algorithms + Data Structures = Evolution Programs. Third Edition, Springer, Berlin, 1996.
- [162] N. Mladenović and P. Hansen. Variable neighborhood search, *Computers & Operations Research*, 24:1097–1100, 1997.
- [163] P. Moscato. On evolution, search, optimization, genetic algorithms and martial arts: Towards memetic algorithms. *Technical Report Caltech Concurrent Computation Program*, 826, California Institute of Technology, Pasadena, CA, 1989.
- [164] P. Moscato and C. Cotta. A gentle introduction to memetic algorithms. In *Handbook of Metaheuristics*, F. Glover and G.A. Kochenberger, eds., Kluwer, Boston, 105–144, 2003.
- [165] Y. Nagata and S. Kobayashi. Edge assembly crossover: A high-power genetic algorithm for the traveling salesman problem. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, T. Bäck, ed., Morgan Kaufmann, San Francisco, CA, 450–457, 1997.
- [166] W.P. Nanry and J.W. Barnes. Solving the pickup and delivery problem with time windows using reactive tabu search. *Transportation Research B*, 34:107–121, 2000.
- [167] I.M. Oliver, D.J. Smith and J.R.C. Holland. A study of permutation crossover operators on the traveling salesman problem. In *Proceedings of the Second International Conference on Genetic Algorithms and Their Applications*, J.J. Grefenstette, ed., Lawrence Erlbaum Associates, Hillsdale, NJ, 224–230, 1987.
- [168] I. Or. Traveling salesman-type combinatorial problems and their relation to the logistics of blood banking. Ph.D. Thesis, Dept. of Industrial Engineering and Management Sciences, Northwestern University, 1976.
- [169] I.H. Osman. Metastrategy simulated annealing and tabu search algorithms for the vehicle routing problem. *Annals of Operations Research*, 41:421–451, 1993.
- [170] D. Pisinger and S. Ropke. A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34:2403–2435, 2007.
- [171] J.-Y. Potvin. Genetic algorithms for the traveling salesman problem. *Annals of Operations Research*, 63:339–370, 1996.
- [172] J.-Y. Potvin and J.-M. Rousseau. A parallel route building algorithm for the vehicle routing and scheduling problem with time windows. *European Journal of Operational Research*, 66:331–340, 1993.
- [173] J.-Y. Potvin and J.-M. Rousseau. An exchange heuristic for routing problems with time windows. *Journal of the Operational Research Society*, 46:1433–1446, 1995.
- [174] E. Prescott-Gagnon, G. Desaulniers and L.-M. Rousseau. A branch-and-price-based large neighborhood search algorithm for the vehicle routing problem with time windows. Technical Report G-2007-67, GERAD, Montreal, Canada.
- [175] Y. Rochat and É.D. Taillard. Probabilistic diversification and intensification in local search for vehicle routing. *Journal of Heuristics*, 1:147–167, 1995.
- [176] S. Ropke, J.-F. Cordeau and G. Laporte. Models and branch-and-cut algorithms for pickup and delivery problems with time windows. *Networks*, 49:258–272, 2007.

- [177] S. Ropke and D. Pisinger. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40:455–472, 2006.
- [178] R.A. Russell. Hybrid heuristics for the vehicle routing problem with time windows. *Transportation Science*, 29:156–166, 1995.
- [179] G. Schrimpf, J. Schneider, H. Stamm-Wilbrandt and G. Duek. Record breaking optimization results using the ruin and recreate principle. *Journal of Computational Physics*, 159:139–171, 2000.
- [180] P. Shaw. Using constraint programming and local search methods to solve vehicle routing problems. *Lecture Notes in Computer Science*, M. Maher and J.-F. Puget, eds., Springer, New-York, 1520:417–431, 1998.
- [181] M.M. Solomon. Time window constrained routing and scheduling problems. *Operations Research*, 35:254–265, 1987.
- [182] N. Srinivas and K. Deb. Multi-objective function optimization using non-dominated sorting genetic algorithms. *Evolutionary Computation Journal*, 2:221–248, 1994.
- [183] G. Syswerda. Uniform crossover in genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms*, J.D. Schaffer, ed., Morgan Kaufmann, San Mateo, CA, 2–9, 1989.
- [184] É.D. Taillard. Parallel iterative search methods for vehicle routing problems. *Networks*, 23:661–673, 1993.
- [185] É.D. Taillard, P. Badeau, M. Gendreau, F. Guertin and J.-Y. Potvin. A tabu search heuristic for the vehicle routing problem with soft time windows. *Transportation Science*, 31:170–186, 1997.
- [186] C.D. Tarantilis. Solving the vehicle routing problem with adaptive memory programming methodology. *Computers & Operations Research*, 32:2309–2327, 2005.
- [187] P. Toth and D. Vigo. The granular tabu search and its application to the vehicle-routing problem. *INFORMS Journal on Computing*, 15:333–346, 2003.
- [188] A. Van Breedam. 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. Dissertation, Faculty of Applied Economics, University of Antwerp, Antwerp, Belgium, 1994.
- [189] C. Voudouris and E.P.K. Tsang. Guided local search. Technical Report CSM-247, Department of Computer Science, University of Essex, Colchester, UK, 1995.
- [190] C. Voudouris and E.P.K. Tsang. Guided local search. In *Handbook of Metaheuristics*, F. Glover and G.A. Kochenberger, eds., Kluwer, Boston, MA, 185–218, 2003.
- [191] D. Whitley. The GENITOR algorithm and selection pressure: why rank-based allocation of reproductive trials is best. In *Proceedings of the Third International Conference on Genetic Algorithms*, J.D. Schaffer, ed., Morgan Kaufmann, San Mateo, CA, 116–121, 1989.