
Crossover and Diversity: A Study about GVR

Jorge Tavares
CISUC
Pinhal de Marrocos
3030 Coimbra, Portugal
jast@dei.uc.pt

Francisco B. Pereira
ISEC
Quinta da Nora
3030 Coimbra, Portugal
xico@dei.uc.pt

Penousal Machado
ISEC
Quinta da Nora
3030 Coimbra, Portugal
machado@dei.uc.pt

Ernesto Costa
CISUC
Pinhal de Marrocos
3030 Coimbra, Portugal
ernesto@dei.uc.pt

Abstract

Genetic Vehicle Representation (GVR) is a two level representational scheme, designed to deal in an effective way with all the information needed by candidate solutions, for the Vehicle Routing Problem (VRP). In this paper, we present an analysis on the influence of two crossover operators in the algorithm performance. A first study on diversity is also presented, regarding the issues of diversity measurement and possible relations to the algorithm performance. Results show that for GVR one type of crossover is more suited for solving VRP instances, and both operators may not avoid the loss of diversity. Nevertheless, solutions discovered by GVR are competitive and are the best ones found by an evolutionary algorithm.

1 Introduction

The Vehicle Routing Problem (VRP) is a complex combinatorial optimization problem, which can be seen as a merge of two well-known problems: the Travelling Salesperson (TSP) and the Bin Packing (BPP). We can describe it as follows: given a fleet of vehicles with uniform capacity, a common depot, and several customer demands, find the set of routes with overall minimum route cost which service all the demands. All the itineraries start and end at the depot, and they must be designed in such a way that each customer is served only once and just by one vehicle. The VRP is NP-hard. The fact that VRP is both of theoretical and practical interest (e.g. distribution, which is a major part of logistics and a substantial cost for many companies), explains the amount of research made in the past years.

Due to the nature of the problem it is not viable to use exact methods for large instances of the VRP. For instances with few nodes, the branch and bound technique is well suited and gives the best possible solution [Kohl et al., 1999]. For larger instances most approaches rely on heuristics that provide approximate solutions [Cordeau et al., 2000]. Another alternative is to apply optimization techniques, such as tabu search [Duncan, 1995], simulated annealing [Bent and Hentenryck, 2001], or ant colony optimization [Gambardella et al., 1999].

Applications of evolutionary computation (EC) techniques to the VRP are also used (see [Bräysy, 2001] for a brief overview). Most researchers rely on hybrid approaches that combine the power of an EC algorithm with the use of specific heuristics [Thangiah, 1995], [Tan et al., 2001], [Jung and Moon, 2002], or use simplified versions of the problem. One common simplification is to pre-set the number of vehicles that is going to be used in the solution [Zhu, 2000]. The first attempts to apply standard EC algorithms to the most generic version of VRP attained a limited success [Duncan, 1995], [Machado et al., 2002].

Our previous research shows that the representation is a key issue in the application of EC techniques to the VRP. This lead to the proposal of a new representational scheme, Genetic Vehicle Representation (GVR) [Pereira et al., 2002], [Tavares et al., 2002], which deals efficiently with the two levels of information that a candidate solution must encode: clustering of the demands (i.e., allocation of all the demands to different vehicles) and specification of the delivery ordering for each one of the routes. GVR also enables an easy adjustment of the number of vehicles required. It is important to notice that our methodology does not use any specific heuristic. We performed a comprehensive set of experiments with some well-known benchmarks and confirmed that GVR enabled the discovery of several new best solu-

tions. The efficiency of the GVR approach was verified in a systematic study, where GVR performance is compared with standard EC models, based on Path Representation [Tavares et al., 2003].

In this paper, we present a study involving the analysis of the performance of two crossover operators developed for GVR. An important issue that is part of this study is the maintenance of population diversity. Diversity measurements are elaborated and tested to access its relations with the operators used.

The paper has the following structure: in section 2 we give a formal definition of the problem variants used. In section 3 we present a description of GVR and the diversity measures used. Section 4 comprises the presentation of the experimental results with a brief analysis. Finally, in section 5, we draw some overall conclusions.

2 The Vehicle Routing Problem

2.1 Capacitated Vehicle Routing Problem

The simplest version of the VRP is the Capacitated Vehicle Routing Problem (CVRP), which can be formally described in the following way: there is one central depot 0 , which uses k independent delivery vehicles, with identical delivery capacity C , to service demands d_i from n customers, $i = 1, \dots, n$. The vehicles must accomplish the delivery with a minimum total length cost, where the cost c_{ij} is the distance from customer i to customer j , with $i, j \in [1, n]$. The distance between customers is symmetric, i.e., $c_{ij} = c_{ji}$ and also $c_{ii} = 0$. A solution for the CVRP would be a partition $\{R_1, \dots, R_q\}$ of the n demands into k routes, each route R_q satisfying

$$\sum_{p \in R_q} d_p \leq C. \quad (1)$$

Associated with each partition is a permutation of the demands belonging to it, specifying the delivery order of the vehicles. In figure 1 we present an illustration of the problem, viewed as a graph, where the nodes represent the customers.

2.2 Vehicle Routing Problem with Time Windows

An important extension to the problem is the Vehicle Routing Problem with Time Windows (VRPTW) which adds several time constraints to the previous definition. Associated with each customer i there is a time window $[e_i, l_i]$ during which it has to be served.




Figure 1: Vehicle Routing Problem.

Route 1	3	2	7	
Route 2	4	9	1	10
Route 3	8	6	5	

Figure 2: An example of a GVR chromosome.

This way, a vehicle must arrive to i no sooner than e_i (the earliest arrival time) and no later than l_i (the latest arrival time). Additionally, there is a service time f_i for each customer and a limit T , defining the maximum travel time permitted for any vehicle. The travel time from customer i to customer j is $t_{ij} = c_{ij}$.

Like in the CVRP, a candidate solution would be a partition R_1, \dots, R_k of the n demands into k routes, each route R_q satisfying both the delivery capacity and the time constraints.

3 Evolutionary Model

3.1 Genetic Vehicle Representation

A candidate solution to an instance of the CVRP or VRPTW must specify the number of vehicles, the partition of the demands through all these vehicles and also the delivery order for each route. In GVR, the genetic material of an individual contains several routes, each of them is composed by an ordered subset of the customers. All demands belonging to the problem being solved must be present in one of the routes. As an example, the chromosome from figure 2 represents the solution presented in figure 1.

The information encoded in the chromosome must be interpreted in such a way that it yields a legal solution. When a vehicle exceeds its capacity, the according route is split in several ones. An example illustrates this adjustment: assume that the original route

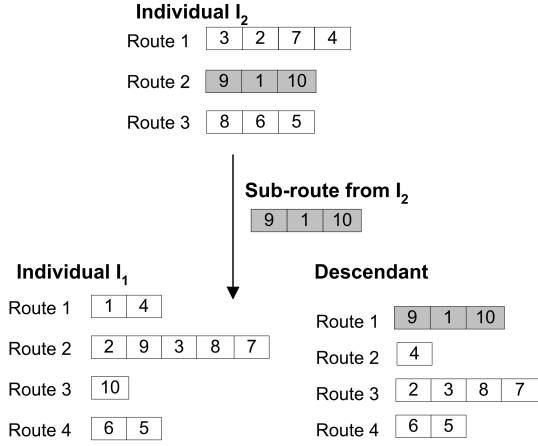


Figure 3: Example of the generic crossover.

$\{a, b, c, d, e, f\}$ causes the vehicle to exceed its capacity at node d . When this situation occurs, the itinerary is divided in two sections: $\{a, b, c\}$ and $\{d, e, f\}$, and a new vehicle is added to the solution. If necessary, further divisions can be made in the second section.

For the VRPTW, another specific situation arises, when time constraints are not satisfied. In this case, three different types of violation may occur: early arrival at a customer, late arrival at a customer or late arrival at the depot. The first situation is easily solved, since it's only required that the vehicle waits until it meets the earliest arrival time of the window. To resolve the other two cases, a new section is created on the itinerary, providing a valid route by adding a new vehicle to the solution. Notice that all these changes only occur at the interpretation level and, therefore, the information codified in the chromosome is not altered.

3.2 Genetic Operators

The EC algorithm processes the individuals in a straightforward way. Assuming that the population size is N , in each generation, N parents are chosen and N descendants are obtained through the application of genetic operators to the elements of the selected set.

Two categories of operators are considered: crossover and mutation. They must be able to deal with the two levels of the representation. Thus, they should be capable to change the delivery order within a specific route and to modify the allocation of demands to vehicles. In this last situation, they cannot only switch customers from one route to another, but also modify the number of vehicles belonging to a solution

(adding and removing routes). Another important requirement is that the genetic operators must always generate legal solutions.

The crossover operator used in our approach does not promote a mutual exchange of genetic material between two parents. The crossover operates in the following way: when an individual from the selected set is submitted to this kind of operation, it receives a fragment of genetic material (more precisely, a route) from another parent and inserts it as its first route. After insertion, a repair process checks the original routes from the receiving individual and removes all customers that also appear in the inherited itinerary. This ensures that the new chromosome is legal, since there will be no repeated points. The donor is not modified. The example from figure 3 illustrates how crossover acts.

As an alternative to the *generic crossover*, we developed a more specific operator, sensitive to the geographical locations of customers. When accepting the fragment $\{a_1, \dots, a_n\}$ from the parent, the receiving individual determines which customer c is geographically closer to a_1 . Then it inserts $\{a_1, \dots, a_n\}$ immediately after c . The example from figure 4 helps to illustrate how this kind of crossover acts.

Descendants resulting from crossover can be subject to mutation. We consider four operators, based on proposals usually applied to order based representations:

- **Swap:** selects two customers and swaps them. Selected points can belong to the same or to different routes.
- **Inversion:** selects a sub-route and reverses the visiting order of the customers belonging to it.
- **Insertion:** selects a customer and inserts it in another place. The route where it is inserted is selected randomly. It is possible to create a new itinerary with this single customer, with probability $\frac{1}{2 \times V}$ (V represents the number of vehicles of the current solution).
- **Displacement:** selects a sub-route and inserts it in another place. This operator can perform intra or inter displacement. Just like in the previous operator, it is also possible to create a new route with the subsequence.

All genetic operators described have a specific probability of application to a single individual.

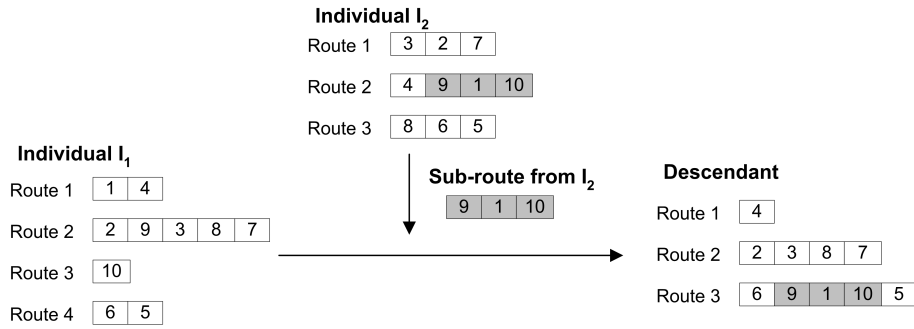


Figure 4: Example of the specific crossover.

3.3 Diversity Measures

An important factor on evolution is the diversity of the population. Prior empirical analysis of the results achieved by GVR has lead us to believe that this approach is prone to premature convergence. In the beginning of the simulation, GVR quickly identifies some local optimal solutions. Then there is a high evolutionary pressure towards these areas and it is almost impossible to escape premature convergence. This loss of genetic diversity could partially explain the shortcomings of GVR in dealing with the VRPTW variant [Tavares et al., 2002]. In order to gather evidence to give further support to these empirical conclusions, several measures were developed, so that a better comprehension of the population diversity evolution could be observed.

Since GVR isn't a binary representation, standard diversity measures cannot be used [Louis and Rawlins, 1993]. In fact, for individuals in GVR representation it is difficult to determine which factors are relevant to assess similarity. The number of clusters for the demands, the sequence order of the customers or just the fitness value? [Burke et al., 2002a] and [Burke et al., 2002b] provide a good overview and discussion on the topic of measuring genetic diversity. Since some of these measures are related to structural factors of the solution – and thus the phenotype – while others are more concerned with the genotype, we decided to develop a set of measures concerning these issues:

- **Routes Similarity:** it gives us a weighted percentage of individuals that have the same number of routes. It allows us to verify structural similarity.
- **Fitness Similarity:** measures the percentage of individuals whose fitness is within a 1% range of the best individual.

- **Lost Links:** a link is established between a pair of consecutive genes belonging to one of the routes. A lost link is a link formed in a generation and lost in the next one. This measure is the percentage of lost links.

4 Experimental Results

4.1 Settings

To evaluate our approach we performed a collection of tests with instances from some well-known benchmarks. For the CVRP we used 3 instances from Augerat Set A (instances A54k7 and A80k10) and Augerat Set B (instance B78k10). For each instance of the datasets, the number of customers is given by the first number on the instance name. The main difference between these sets of problems is their tightness (the ratio between demand and capacity) and the location of customers. For the VRPTW we used 4 instances from Solomon's benchmarks, from the R collection, random customer distribution, usually considered the most difficult. All instances consist of 100 customers plus the depot and are characterized by a short scheduling horizon, allowing only a few customers per route.

The settings of the EC algorithm are the following: Number of generations: 50000; Population size: 200; Tournament selection with tourney size: 5; Elitist strategy; Crossover rate: {0.6, 0.75}; Mutation rates: swap: 0.05; inversion: 0.15; insertion: 0.05; displacement: 0.2. We performed an extended set of tests, which are not shown in this paper due to lack of space, with different probabilities of application for the genetic operators. For every set of parameters we performed 30 runs with the same initial conditions. All initial populations were randomly generated. Statistical analysis was performed with a level of significance $\alpha = 0.05$.

Table 1: Summary of GVR results.

Instances	Previous	Generic CX				Specific CX			
		CX = 0.6		CX = 0.75		CX = 0.6		CX = 0.75	
		Best	Avg.	Best	Avg.	Best	Avg.	Best	Avg.
a54k7	1167.0	1178.0	1258.53	1197.0	1267.83	1167.0	1191.27	1167.0	1191.33
a80k10	1763.0	1868.0	1964.63	1892.0	1971.03	1785.0	1831.10	1789.0	1812.63
b78k10	1221.0	1283.0	1319.47	1274.0	1318.03	1221.0	1266.60	1224.0	1247.67
R101	1637.7	1675.8	1726.01	1681.6	1722.08	1684.0	1724.40	1775.0	1878.22
R102	1466.6	1512.6	1561.57	1504.4	1547.00	1502.5	1549.07	1580.3	1649.24
R103	1208.7	1252.3	1314.84	1259.0	1323.43	1246.6	1295.45	1297.2	1341.84
R104	982.0	1008.5	1088.30	1041.9	1089.70	1013.2	1070.39	1022.8	1074.25

4.2 Performance

In table 1, we present, for all instances, the results achieved by GVR. The table shows the best solutions found by the two types of crossover, as well as the averages of the best solution found in each of the 30 runs. The Column “Previous” indicates the best solutions known in the literature when our research started. A brief perusal of the results reveals that GVR was able to find, with both crossovers, good solutions. A closer inspection of the best results column, indicates that the *specific crossover* performed better, achieving not only overall better solutions, but was also capable of finding some of the previous best known solutions (instances a54k7 and b78k10).

Examining the column with the averages, the values for *specific crossover* are also consistently better than the averages of the generic one. The distances between the average of the best solutions to the best known solution range between 2% and 14% with an average of 6%, whilst for the *generic crossover*, these distances are ranging between 5% and 12% with an average of 9%. These differences are statistical significant for all instances with only three exceptions: instances R101, R102 (both with crossover probability of 0.6) and R104 (with crossover probability of 0.75).

By looking into table 2, it can also be confirmed that statistical significant differences exist between the application of crossover probabilities. For the *specific crossover* these differences are all significant with the exception of instances a54k7 and R104. For the *generic crossover*, the opposite situation occurs : only instances b78k10 and R102 have significant differences. This analysis strengthens the observations found in table 1, where a lower crossover probability has performed better than a higher one. The best results were all discovered with a probability of 0.6, regardless of the type of operator chosen. This brief analysis shows that the usage of a different crossover, even if similar in

concept, introduces differences on the results attained by the evolutionary approach to the VRP.

Table 2: Summary of the statistical analysis ($\alpha = 0.05$, G = Generic CX, S = Specific CX, 1 means there is a statistical significance difference).

Instances	0.6	0.75	G	S
	G/S	G/S	0.6/0.75	0.6/0.75
a54k7	1	1	0	0
a80k10	1	1	0	1
b78k10	1	1	1	1
R101	0	1	0	1
R102	0	1	1	1
R103	1	1	0	1
R104	1	0	0	0

4.3 Diversity

So far we have only observed the practical effects of both crossover operators in the process of finding good candidate solutions. Our previous work shows that the GVR approach is well-suited in solving the CVRP, consistently finding the best known solutions to several instances of the problem, and even discovering new lower bounds. However, when applying GVR to the time windows variant, there was a significant decrease of performance.

The evolution of the fitness of the best individual seems to indicate that it reaches local optima very fast and is not able to escape it. The reasons for this kind of behavior are not clear, but empirical data has shown that it might be related to the diversity of the population. In this regard, crossover operators may have an influence in maintaining or promoting diversity among the populations.

During the evolution process, we collected data that enable us to measure the diversity of the populations.

Table 3: Generic CX diversity measures.

Instances	CX = 0.6			CX = 0.75		
	Route	Fitness	Links	Route	Fitness	Links
a54k7	62.90	56.73	82.40	62.45	57.46	82.27
a80k10	61.30	59.11	84.88	61.97	60.33	84.99
b78k10	61.82	56.70	83.86	62.28	57.82	85.13
R101	44.25	62.60	67.31	44.57	62.64	69.42
R102	45.84	60.86	77.71	43.88	60.43	73.96
R103	47.05	58.81	80.09	47.40	59.41	79.70
R104	47.08	53.83	80.77	49.18	55.70	82.50

Table 4: Specific CX diversity measures.

Instances	CX = 0.6			CX = 0.75		
	Route	Fitness	Links	Route	Fitness	Links
a54k7	46.26	26.30	74.47	46.81	22.83	73.15
a80k10	46.08	29.62	75.47	46.83	29.23	76.09
b78k10	47.14	29.55	77.62	46.65	27.42	75.10
R101	8.75	33.64	69.42	7.32	36.48	70.27
R102	8.73	29.18	73.57	7.12	30.43	69.14
R103	9.73	26.50	73.37	7.04	23.71	68.65
R104	11.82	24.77	73.76	8.90	21.88	73.12

In table 3 and table 4, we show the summary of the three different measures used. As we can observe in the columns “Routes” in table 3, the structural similarity of the individuals is extremely high. For both versions of the problem, approximately half of the population has the same number of routes. In fact, for the CVRP these values are always above 60%, while for the VRPTW, they range between 44% and 50%. *Fitness similarity* shows a similar behavior: for all instances, regardless of probability, *fitness similarity* possesses values between 53% and 62%. These values indicate a high loss of population diversity. This lack of diversity indicates, that the *generic crossover* is unable to explore the search space, producing new genetic material with good quality. This can be explained by the third diversity measure. The percentage of new links lost from the previous generation is very high, around 80%. This loss is possibly a result of crossover breaking good links, making it more difficult to individuals with lower fitness to transmit quality links.

For the *specific crossover*, the situation is not entirely different when considering the premature convergence of the algorithm, but the magnitude of values for the diversity measures is different. From table 4 it is possible to observe that the loss of links follows the same previously described pattern. The percentage is a bit lower, around 70% in average, but still a high value. But when we take a closer examination, we can

verify that route similarity is lower, specially for the VRPTW instances. For this variant, the number of individuals with equal number of routes is below 10%, while for CVRP, it is higher (45%). The fitness measure has an average of 30% for all instances. These values allow to verify the previous arguments in favor of the *specific crossover* against *generic crossover* regarding the algorithm performance. Nevertheless, even with lower structural similarity, *specific crossover* is not able to discover solutions with a significant higher quality than those produced by *generic crossover* and is also unable to keep a lower destructive effect on individuals. Results attained with both crossover operators show that GVR is an effective representation for this problem. Nevertheless, a better understanding of its operators is still required. Moreover, the design of a new crossover operator may be necessary in order to achieve the best known solutions for the VRPTW problem variant.

5 Conclusion

In this paper we presented a study which compares the performance of two crossover operators designed for GVR. Results show that the choice of operators is important when using GVR for the CVRP or time windows variant. In CVRP, *specific crossover* is an efficient operator, able to discover the

best known solutions and, finding new lower bounds [Pereira et al., 2002]. In what concerns the VRPTW, both operators are unable to achieve the previous performance [Tavares et al., 2002].

We propose several diversity measures for individuals in GVR representation. These measures try to assess genotypic and phenotypic similarity factors. The analysis of the experimental results indicates that, crossover operators capable of preserving diversity and linkage in GVR individuals, might produce better results in the VRPTW variant. This is possibly due to the fact, that a less destructive operator may be more well-suited, because of the additional constraints.

Adopting a good representational scheme and genetic operators is essential. The choice of a representation that is sensitive to the structure of the solutions provides an important advantage. For the VRP this is an important and essential factor [Tavares et al., 2003].

References

- [Bent and Hentenryck, 2001] Bent, R. and Hentenryck, P. V. (2001). A two-stage hybrid local search for the vehicle routing problem with time windows. Technical Report CS-01-06, Brown University.
- [Bräysy, 2001] Bräysy, O. (2001). Genetic algorithm for the vehicle routing problem with time windows. *Arpakannus, Special Issue on Bioinformatics and Genetic Algorithms*, pages 33–38.
- [Burke et al., 2002a] Burke, E., Gustafson, S., and Kendall, G. (2002a). A survey and analysis of diversity measures in genetic programming. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 716–723.
- [Burke et al., 2002b] Burke, E., Gustafson, S., Kendall, G., and Krasnogor, N. (2002b). Advanced population diversity measures in genetic programming. In *Parallel Problem Solving from Nature - PPSN VII*, page 341 ff.
- [Cordeau et al., 2000] Cordeau, J.-F., Desautniers, G., Desrosiers, J., Solomon, M. M., and Soumis, F. (2000). The VRP with Time Windows. In *The Vehicle Routing Problem*, chapter 7. Monographs on Discrete Mathematics and Applications.
- [Duncan, 1995] Duncan, T. (1995). Experiments in the use of neighbourhood search techniques for vehicle routing. Technical Report AIAI-TR-176, Artificial Intelligence Applications Institute, University of Edinburgh.
- [Gambardella et al., 1999] Gambardella, L. M., Taillard, É., and Agazzi, G. (1999). MACS-VRPTW: A multiple ant colony system for vehicle routing problems with time windows. In *New Ideas in Optimization*, pages 63–76. McGraw-Hill.
- [Jung and Moon, 2002] Jung, S. and Moon, B.-R. (2002). A hybrid genetic algorithm for the vehicle routing problem with time windows. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1309–1316.
- [Kohl et al., 1999] Kohl, N., Desrosiers, J., Madsen, O. B. G., Solomon, M. M., and Soumis, F. (1999). 2-path cuts for the vehicle routing problem with time windows. *Transportation Science*, 33(1):101–116.
- [Louis and Rawlins, 1993] Louis, S. J. and Rawlins, G. J. E. (1993). Syntactic analysis of convergence in genetic algorithms. In *Foundations of genetic algorithms 2*, pages 141–151.
- [Machado et al., 2002] Machado, P., Tavares, J., Pereira, F. B., and Costa, E. (2002). Vehicle routing problem: Doing it the evolutionary way. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, page 690.
- [Pereira et al., 2002] Pereira, F. B., Machado, J. T. P., and Costa, E. (2002). Gvr: a new representation for the vehicle routing problem. In *Artificial Intelligence and Cognitive Science: 13th Irish Conference Proceedings*, pages 95–102.
- [Tan et al., 2001] Tan, K. C., Lee, L., and Ou, K. (2001). Hybrid genetic algorithms in solving vehicle routing problems with time window constraints. *Asia-Pacific Journal of Operational Research*, 18(1):121–130.
- [Tavares et al., 2002] Tavares, J., Pereira, F. B., Machado, P., and Costa, E. (2002). GVR Delivers It On Time. In *SEAL02 4th Asia-Pacific Conference on Simulated Evolution And Learning*, pages 745–749, Singapore.
- [Tavares et al., 2003] Tavares, J., Pereira, F. B., Machado, P., and Costa, E. (2003). On the Influence of GVR in Vehicle Routing. In *Symposium of Applied Computing (SAC2003)*, Florida, USA.
- [Thangiah, 1995] Thangiah, S. (1995). Vehicle routing with time windows using genetic algorithms. In *Application Handbook of Genetic Algorithms: New Frontiers, Volume II*, pages 253–277. CRC Press.
- [Zhu, 2000] Zhu, K. Q. (2000). A new genetic algorithm for vrptw. *Proceedings of the International Conference on Artificial Intelligence*.