

GRASP with path relinking for the two-echelon vehicle routing problem

Original

GRASP with path relinking for the two-echelon vehicle routing problem / Crainic, Teodor; Mancini, Simona; Perboli, Guido; Tadei, Roberto. - ELETTRONICO. - (2012), pp. 1-14.

Availability:

This version is available at: 11583/2502066 since:

Publisher:

Published

DOI:

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



CIRRELT

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

Interuniversity Research Centre
on Enterprise Networks, Logistics and Transportation

GRASP with Path Relinking for the Two-Echelon Vehicle Routing Problem

**Teodor Gabriel Crainic
Simona Mancini
Guido Perboli
Roberto Tadei**

August 2012

CIRRELT-2012-45

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
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

GRASP with Path Relinking for the Two-Echelon Vehicle Routing Problem

Teodor Gabriel Crainic^{1,2,*}, Simona Mancini³, Guido Perboli^{1,3}, Roberto Tadei³

¹ Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

² Department of Management and Technology, Université du Québec à Montréal, P.O. Box 8888, Station Centre-Ville, Montréal, Canada H3C 3P8

³ DAUIN, Politecnico di Torino, C.so Duca degli Abruzzi 24 10129, Torino, Italy

Abstract. We propose a meta-heuristic based on GRASP combined with Path Relinking to address the Two-Echelon Vehicle Routing Problem, an extension of the Capacitated Vehicle Routing Problem in which the delivery from a single depot to customers is achieved by routing and consolidating the freight through intermediate depots called satellites. The problem is treated by separating the depot-to-satellite transfer and the satellite-to-customer delivery, and iteratively solving the two resulting routing subproblems, while adjusting the satellite workloads that link them. The meta-heuristic scheme consists of applying a GRASP and a local search procedure in sequence. Then, the resulting solution is linked to an elite solution by means of a Path Relinking procedure. To escape from infeasible solutions, which are quite common in this kind of problem, a feasibility search procedure is applied within Path Relinking. Extensive computational results on instances with up to 50 customers and 5 satellites show that the meta-heuristic is able to improve literature results, both in efficiency and accuracy.

Keywords. Two-echelon vehicle routing, GRASP, path relinking

Acknowledgements. Partial funding for this project has been provided by the Natural Sciences and Engineering Research Council of Canada (NSERC). We also gratefully acknowledge the support of Fonds de recherche du Québec - Nature et technologies (FRQNT) through their infrastructure grants.

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: TeodorGabriel.Crainic@cirrelt.ca

1 Introduction

The aim of this paper is to present an efficient meta-heuristic to address the *Two-Echelon Vehicle Routing Problem (2E-VRP)*, named GRASP-PR. The *2E-VRP* is a variant of the Capacitated Vehicle Routing Problem (CVRP), characterized by a single depot and a given number of secondary facilities named satellites. The first level routing problem addresses depot-to-satellite delivery, while the satellite-to-customer delivery routes are built at the second level. The goal is to ensure an efficient and efficient operation of the system, where the demand is satisfied and the total cost of the traffic within the overall transportation network is minimized.

This problem is frequently faced in real-life applications, both at the strategic and tactical planning levels, and in day-to-day operations. Methods that can be applied at both levels must be accurate and fast. Thus, for planning, the *2E-VRP* is usually part of larger optimization frameworks, meaning that it must be solved many times during the optimization process, and computational times need to be limited. On the other hand, good feasible solutions are needed in a very short time when optimization problems are to be used at the operational level. Solution quality is crucial in all cases, because it directly impacts the revenues and service quality of the transportation company.

The meta-heuristic introduced in this paper is based on hybridizing GRASP and Path Relinking. More precisely, GRASP is used to generate solutions, which are post-optimized by means of a local search procedure. In order to improve the solution quality, a path between the current solution obtained by GRASP with the local search procedure and the best solution found so far is built by means of a Path Relinking procedure. The meta-heuristic is tested on medium-sized instances with 50 customers and 5 satellites, showing that the new method is able to improve existing state-of-the-art results.

The paper is organized as follows. The *2E-VRP* and the main literature results are presented in Section 2. Section 3 is dedicated to the GRASP-PR general framework, while computational tests and result analysis are reported in Section 4. Conclusions are drawn in Section 5.

2 Problem Definition and Literature Review

In the *2E-VRP*, the distribution of freight cannot be managed by direct shipping from the depot to the customers. Instead, freight must be consolidated from the depot to a satellite and then delivered from the satellite to the desired customer. This implicitly defines a two-echelon transportation system: the 1st level connecting the depot to the satellites and the 2nd one, the satellites to the customers.

We define the depot with v_0 , the set of satellites, with V_s , and the set of customers with

V_c . Let n_s be the number of satellites, and n_c the number of customers. The customers are the destinations of the freight shipments and each customer i has an associated demand d_i representing the quantity of freight that has to be delivered to it. The demand of each customer cannot be split among different vehicles at the 2nd level. For the 1st level, we consider that each satellite can be served by more than one 1st-level vehicle, therefore the aggregated freight assigned to each satellite can be split into two or more vehicles. Each 1st level vehicle can deliver the freight of one or several customers, as well as serve more than one satellite in the same route. We consider only one product, i.e., the volumes of freight belonging to different customers can be stored together and loaded into the same vehicle for both the 1st and the 2nd-level movements. We define a route made up of a 1st-level vehicle starting from the depot, serving one or more satellites, and ending up at the depot, as *1st-level route*. A *2nd-level route* is made up of a 2nd-level vehicle starting from a satellite, serving one or more customers, and ending up at the same satellite. The fleet sizes are fixed and known in advance for both levels. All vehicles belonging to the same level have the same capacity. Each vehicle may perform at most one route. Satellites have limited capacity defined as the maximum number of vehicles that can leave from it. Different satellites may have different capacities.

The literature on *2E-VRP* is still somewhat limited. A general time-dependent formulation with fleet synchronization and customer time windows has been introduced in Crainic et al. (2009) in the context of two-echelon City Logistics systems. The authors have indicated promising algorithmic directions, but no implementation has been reported. A MIP formulation for the *2E-VRP* has been presented in Perboli et al. (2011), with instances with up to 32 customers solved to optimality. In the same paper, the authors derived two math-heuristics able to address instances with up to 50 customers. Both math-heuristics are based on the MIP model presented in the paper and work on the customer-to-satellite assignment variables. The first math-heuristic, called *Diving*, considers a continuous relaxation of the model and applies a diving procedure to the customer-to-satellite assignment variables that are not integer. A restarting procedure is incorporated to recover possible unfeasibilities due to variable fixing. The second one is named *Semi-continuous*; in this method, the arc usage variables are considered continuous, while the assignment variables are still considered integer. The method solves this relaxed problem and uses the obtained values of the assignment variables to build a feasible solution for the *2E-VRP*. Several families of valid inequalities have been proposed in Perboli et al. (2010). The valid inequalities are integrated into a Branch-and-Cut scheme, which is able to drastically reduce the optimality gap. A multi-start heuristic has been presented in Crainic et al. (2011). The method is based on a clustering heuristic, which mainly works on the assignment between satellites and customers. The heuristic is used by the authors to solve large-sized instances with up to 250 customers. In Crainic et al. (2010) the authors study the effect of different spatial distributions on the total costs and a comparison with the standard CVRP solutions is given, while the impact of realistic situations in urban freight delivery where the travel costs are affected by components different from the distance, like fixed costs for using the arcs, operational costs, and environmental costs can be found in Crainic et al. (2012).

A problem from the literature quite similar to the $2E$ -VRP is the Truck and Trailer Routing Problem (TTRP), in which the use of trailers (a commonly neglected feature in the VRP) is considered where customers are served by a truck pulling a trailer. However, due to practical constraints, including government regulations, limited manoeuvring space at customer sites, road conditions, etc., some customers may only be serviced by a truck. These constraints exist in many practical situations. This problem, as the $2E$ -VRP, involves two routing levels strictly interconnected. The main difference with $2E$ -VRP is that, while in the $2E$ -VRP freight must pass through the satellites, because it must be delivered to the customer only by second level vehicles, in the TTRP the delivery of certain customers can be directly carried out by first level vehicles (truck and trailer) without passing through satellites. In Villegas et al. (2010), the Single Truck and Trailer Routing Problem with Satellite Depots (STTRPSD), a particular version of the TTRP, is introduced. In STTRPSD a vehicle composed of a truck with a detachable trailer serves the demand of a set of customers reachable only by the truck without the trailer. This accessibility constraint implies the selection of locations to park the trailer before performing the trips to the customers. This version of the problem is the most similar to the $2E$ -VRP while all deliveries must be carried out by the same kind of vehicle (truck without the trailer), even if, in this case, only one vehicle is considered, while in our problem several vehicles could be used to fulfil the customers demands.

3 GRASP with Path Relinking

GRASP is a multistart meta-heuristic for combinatorial optimization Gendreau and Potvin (2010). It consists of a constructive procedure based on a greedy randomized algorithm. In literature, this procedure is often combined with Local Search (see Feo and Resende (1995), Festa and Resende (2009a), and Festa and Resende (2009b) for a detailed survey of the method and its applications). Path Relinking is an intensification strategy that explores trajectories connecting high-quality solutions. Path Relinking was suggested as an approach to integrate intensification and diversification strategies in the context of tabu search Glover and Laguna (1997); Glover et al. (2000) and then extended to other heuristic methods Res. This approach generates new solutions by exploring trajectories that connect high-quality solutions by starting from one of these solutions, called the *starting* solution, and generating a path in the search space that leads towards the other solution, called *guiding* solution. Laguna and Martí adapted Path Relinking to the context of GRASP as a form of intensification Laguna and Martí (1999). The relinking in this context consists of finding a path between a solution found with GRASP and a chosen elite solution. Therefore, the relinking concept has a different interpretation within GRASP since the solutions found by two successive GRASP iterations are not linked by a sequence of moves. See Resende and Ribeiro (2003) for a survey and numerous examples of GRASP with Path Relinking.

3.1 GRASP with Path Relinking for the 2E-VRP

The customer-to-satellite assignment problem plays a crucial role when addressing the 2E-VRP. In fact, assuming one knows the optimal customer-to-satellite assignments, the 2E-VRP can be partitioned into at most $n_s + 1$ CVRP instances, where n_s is the number of satellites, one for the 1st-level and one for each satellite with at least one customer assigned to it. Thus, following the math-heuristics presented in Perboli et al. (2011), and in the meta-heuristics in Crainic et al. (2011), we focus on the customer-to-satellite assignment by searching for the optimal assignment, delegating state-of-the-art CVRP methods to solve the corresponding sub-problems. Both exact and heuristics methods from the literature are suitable to this purpose. Using exact methods we would obtain more precise results, while heuristics would require a more limited computational effort. After preliminary tests, we decided to use the hybrid meta-heuristic proposed in Perboli et al. (2008), which provides a good compromise between solution quality and computational time. Anyway, the effort required to evaluate the objective function, for a given assignment, is considerable. Then, heuristic methods involving large neighborhoods exploration are not suitable for the 2E-VRP, while procedures in which a rule, that allows to identify promising solutions, is applied are strongly preferable. In this work we propose a GRASP, which fits well with these requirements, combined with a Path Relinking strategy. Furthermore, intensification is applied only on promising GRASP solutions with a strong reduction of global computational time.

More precisely, The proposed method, GRASP-PR, consists of four main phases which will be described in detail:

1. A GRASP procedure;
2. A Feasibility Search (FS) phase to be applied if the solution is unfeasible;
3. A Local Search (LS) phase to improve a solution;
4. A Path Relinking phase.

The innovative aspect of this method is neither in its single components, which are well established in literature, nor in the meta-heuristic framework, but in the way the different components are combined within the framework. More in details, GRASP-PR works as follows. First of all an initial assignment is computed following the clustering constructive heuristic presented in Crainic et al. (2011) and the corresponding solution is kept as current best solution. At each iteration, a new assignment is built by means of the GRASP procedure, and the correspondent 2E-VRP solution is evaluated. If it is unfeasible, a repair procedure, named Feasibility Search, is applied. If the solution is feasible and promising, i.e. it is better or within a threshold from the current best, an intensification phase made by a local search and a path-relinking heuristic is applied, otherwise it is discarded. A pseudocode of the algorithm is reported in Algorithm 1.

Algorithm 1 GRASP with Path Relinking

Compute an initial solution by means of the clustering constructive heuristic presented in Crainic et al. (2011);

while a maximum number of iterations $ITERMAX$ is not reached **do**

 Compute a solution CS by means of the GRASP procedure

if CS is unfeasible **then**

 Try to repair CS by means of the Feasibility Search (FS)

if CS is still unfeasible **then**

 Discard CS

end if

end if

if CS is feasible or it has been successfully repaired by FS **then**

if CS cost is within a given threshold t from the best solution cost **then**

 Apply the Local Search (LS)

 Apply Path Relinking

else

 Discard CS

end if

else

 Discard CS

end if

end while

Note that the initial solution is computed in order to have a current feasible solution. It is considered as the current best solution at the first iteration of GRASP-PR and used to determine if the solutions obtained by the GRASP procedure are promising or not, i.e. if Local Search and Path Relinking should be applied or not. In the following subsections each component of the meta-heuristic framework is described in detail.

3.2 GRASP

The GRASP procedure assigns customers to satellites. The core of the GRASP procedure is the clustering constructive heuristic presented in Crainic et al. (2011), where the customers are assigned according to a less-distance-based rule. The assignment of customer i to satellite l is made with probability p_{il}

$$p_{il} = \frac{1 - \frac{D_{il}}{\sum_{l \in V_s} D_{il}}}{n_s - 1}, \quad (1)$$

where D_{ij} is the distance between customer i and satellite j . The rationale is to assign customer i to satellite j with a probability inversely proportional to the distance between them. The k assignments with the highest probability are considered and one of them is randomly selected.

The algorithm is adaptive. Thus, if by assigning customer i to satellite j the satellite capacity is exceeded, i.e., the number of vehicles required at the satellite is larger than the number of vehicles available for that satellite, or the global required number of vehicles is larger than the fleet size, the correspondent assignment becomes forbidden by setting its probability to zero. In this way, at each step of the greedy algorithm, previous choices are taken into account. When all customers are assigned to satellites, the original problem can be split into several CVRP subproblems, which are solved by means of the hybrid meta-heuristic in Perboli et al. (2008). This procedure differs from the greedy algorithm used in Crainic et al. (2011) to find an initial solution because, in GRASP, the customer-to-satellite assignment is probabilistic, while in the greedy algorithm it follows a deterministic rule assigning each customer to its nearest available satellite. Notice that, the GRASP procedure does not include a local search phase to improve the routes, as this optimization is delegated to the meta-heuristic used to solve the CVRP subproblems.

3.3 Feasibility Search

The GRASP procedure does not guarantee the feasibility of the obtained solution, because, even when the satellite capacity is satisfied, the global fleet size constraint may be violated. When this happens, we try to rebuild a feasible solution by means of the Feasibility Search (FS) procedure.

The FS does not imply a neighborhood exploration, it rather proceeds in a straightforward customer-moving procedure, aiming to empty vehicles that are in excess at some satellites. More in detail, customers are selected based on a distance criterion maximizing their distance from the satellite. A selected customer is then moved, from its assigned satellite (the one with the less filled vehicle) to another randomly chosen satellite, in order to free the exceeding vehicle. These moves are repeated until the global fleet size constraint is satisfied. If no move allows the feasibility of the obtained solution, this solution is discarded.

3.4 Local Search

The Local Search phase is performed only if the solution obtained by GRASP or GRASP and Feasibility Search is both feasible and promising, i.e., its cost is better (or within a given threshold) than the cost of the best solution found so far. The procedure adopted for the local search is based on the Clustering Improvement algorithm presented in Crainic et al. (2008). The neighborhood explored by the local search is composed by all the solutions differing by exactly one customer-to-satellite assignment from the ones of the current solution.

The order according to which the solutions in the neighborhood are analyzed is given by an assignment list. Customers are sorted in non decreasing order of the difference between

the distance to the satellite to which they are assigned in the solution, and the distance to the nearest not-assigned satellite. The choice of this sorting order is based on the observation that customers displaying smaller such differences lead to improved solutions with a much higher frequency than the others. The exploration follows a First Improvement criterion and terminates after LS_{max} iterations or when the entire neighborhood has been explored without finding any improvement to the best solution. Preliminary computational experiments shown that this neighborhood exploration strategy is much more efficient than a standard random one, and that following a First Improvement criterion obtained better results than a Best Improvement one. A pseudocode of this procedure is reported in Algorithm 2.

Algorithm 2 Local search

Given the current solution, the customers are sorted by non-decreasing order of the reassignment cost, defined as $RC_i = c_{ij} - c_{ik}$, where i is a customer, j is the satellite to which i is assigned in the current solution, and $k \neq j$ is the satellite such that, moving i from satellite j to satellite k , the capacity constraints on the global second-level vehicle fleet and the satellite k are satisfied and the cost c_{ik} is minimum among the satellites $k \neq j$. This is equivalent to order the customers according to non-decreasing order of the estimation of the change in the solution quality due to the assignment of one customer from the present satellite to its second-best choice. Let be CL the ordered list of the customers.

repeat

 Consider the first customer i in CL ;

if k exists **then**

 remove i from CL ;

else

 terminate the LS algorithm and return the best solution;

end if

 Solve the CVRPs of satellites j and k ;

 Update the demand of each satellite according to the new assignment and solve the first-level CVRP;

 Compute the objective function of the new solution and compare it to the cost of the current solution;

if the new solution is better **then**

 Keep it as new current solution and exit from the neighborhood;

else

if the new solution has an objective function which is worse than a fixed percentage threshold γ from the objective function of the current solution **then**

 Terminate the LS algorithm and return the best solution;

else

 Consider the next customer in the list

end if

end if

until CL is empty

Even if the neighborhood size is not so large, $O(n_c)$, where n_c is the number of customer,

the computational time could grow up due to the need of recompute the CVRPs after a change in the customer-satellite assignments. This is the rationale of adding the additional heuristic stopping criterion when the reassignment has an objective function which is significantly worse than the current solution, i.e. it is larger of more than a given percentage threshold, γ , with respect to the current solution. In fact, being the customers ordered by non-decreasing order of RC_i and being RC_i related to the change in the objective function when we assign the customer to another satellite, if the objective function of a neighbor is deteriorating too much, it is unlikely that the following neighbors may bring us an improving solution.

3.5 Path Relinking

The Path Relinking phase consists of starting from the local optimum S_{LS} obtained by the Local Search procedure and “relinking” it to best solution S_b . The relinking is performed in a backward way, from S_b towards S_{LS} , inserting an element of S_{LS} into S_b at each step. More precisely, the Path Relinking procedure considers a customer assigned to satellite s_1 in S_{LS} and to satellite s_2 in S_b . It then assigns this customer to satellite s_2 in S_{LS} , without changing the other assignments. If the new solution is unfeasible, then the Feasibility Search is applied. If it is still unfeasible it is discarded. The procedure terminates when S_{LS} becomes equal to S_b . The order according to which the customers are selected is given by a list in which customers are ordered in non decreasing order of the difference of distances between the customer and the satellites to which it is assigned in S_b and the customer and the satellites to which it has been assigned in S_{LS} . In this way we first analyze most promising moves, i.e., solutions characterized by a customer-satellite change minimizing its perturbation to the solution. Preliminary tests have shown the effectiveness of this Path Relinking strategy Mancini (2011).

4 Computational Results

In this section, we present computational results and analyze the performance of the method we propose. An analysis of the impact of each component of the algorithm is reported and the results of GRASP-PR are compared with the literature.

Computational tests were effectuated on instances with 50 customers and 5 satellites, introduced in Crainic et al. (2010). The instances present different combinations of customer distributions and types of satellite locations. Three customer distributions are considered representing a regional area, a large city, and a small town. Three types of satellite locations are considered as well, namely, random around the customer area, sliced around the customer area, and within part of the ring around the customer area, the latter representing city settings with limited accessibility due to geographical constraints (e.g., near to natural barriers such as the sea, a lake, or a mountain). Two instances were randomly generated for each combination of

customer and satellite distributions. Table 1 displays the list of instances and their layout characteristics. Computational tests were performed on a computer with a Core 2 Duo processor at 2.5 GHz. The number of iterations, ITERMAX, and the maximum number of local search iterations, LS_{max} , were fixed to 25 and 250, respectively. These values come out of a tuning phase. In Table 2 we report best and average (over 10 trials) results obtained by the different steps of the methods. More in details the table is organized as follows:

- Column 1: instance name;
- Column 2: results of GRASP and Feasibility Search (phases 1 and 2 of the algorithm);
- Column 3: results of GRASP, Feasibility Search, and Local Search (phases 1, 2, and 3);
- Column 4: results of the overall GRASP-PR, in which all phases are applied (GRASP, Feasibility Search, Local Search, and Path Relinking);

A comparison with the literature is reported in Tables 3 and 4. More in detail, objective function values are reported in Table 3 while the corresponding computational times, expressed in CPU seconds, are reported in Table 4. Both tables are organized as follows:

- Column 1: instance name;
- Column 2: results of GRASP and Feasibility Search (phases 1 and 2 of the algorithm);
- Column 3: results of GRASP, Feasibility Search, and Local Search (phases 1, 2, and 3);
- Column 4: results of the overall GRASP-PR, in which all phases are applied (GRASP, Feasibility Search, Local Search, and Path Relinking);
- Columns 5-7: results of Multi-Start heuristic (MS) proposed in Crainic et al. (2011), the math-heuristics (MH) presented in Perboli et al. (2011), and the Branch-and-Cut (BC) proposed in Perboli et al. (2010) with a time-limit of 10000 seconds (BC also yields, on average, the overall best solutions in the literature).

Each row reports the values of a single instance, while the last two rows give the mean values and the gaps with respect to the BC (in Table 3 only), respectively.

When compared to the results of the Multi-Start heuristic (Column MS), the pure GRASP obtains comparable results, while the introduction of Local Search yields a gain of more than 6%. The complete meta-heuristic GRASP-PR shows an improvement of 8.7% with respect to MS and outperforms the best heuristic results in the literature (MH) by 1.5%. This behavior confirms the trend reported in the literature, which encourages the use of Path Relinking as an enhancement of GRASP.

Instances	Customer Distribution	Satellite Location
Instance50-s5-37.dat	regional area	random
Instance50-s5-38.dat	regional area	random
Instance50-s5-39.dat	regional area	sliced
Instance50-s5-40.dat	regional area	sliced
Instance50-s5-41.dat	regional area	limited accessibility
Instance50-s5-42.dat	regional area	limited accessibility
Instance50-s5-43.dat	large city	random
Instance50-s5-44.dat	large city	random
Instance50-s5-45.dat	large city	sliced
Instance50-s5-46.dat	large city	sliced
Instance50-s5-47.dat	large city	limited accessibility
Instance50-s5-48.dat	large city	limited accessibility
Instance50-s5-49.dat	small town	random
Instance50-s5-50.dat	small town	random
Instance50-s5-51.dat	small town	sliced
Instance50-s5-52.dat	small town	sliced
Instance50-s5-53.dat	small town	limited accessibility
Instance50-s5-54.dat	small town	limited accessibility

Table 1: Instance layout characteristics

GRASP-PR achieves slightly worst results with respect to the exact solution method BC, but this small loss in accuracy (1.89%) is highly compensated by a significant reduction in computational effort of more than 2 orders of magnitude, as shown in Table 4.

5 Conclusion

We presented GRASP-PR, a GRASP with Path Relinking meta-heuristic for the *2E-VRP*, an extension of the classical vehicle routing problem, in which the delivery from a single depot to customers is managed by routing and consolidating the freight through intermediate facilities. Computational tests show that the method we propose outperforms the methods from the literature. Due to the good quality of the obtained solutions and the limited computational effort, GRASP-PR could be adopted both for long term planning and on-demand optimization.

Acknowledgments

Partial funding for this project has been provided by the Natural Sciences and Engineering Council of Canada (NSERC). We also gratefully acknowledge the support of Fonds de

INST	GRASP		GRASP+LS		GRASP-PR	
	BEST	AVG	BEST	AVG	BEST	AVG
Instance50-s5-37.dat	1599.86	1615.34	1586.23	1586.23	1545.99	1545.99
Instance50-s5-38.dat	1335.22	1335.22	1222.27	1222.27	1172.83	1172.83
Instance50-s5-39.dat	1657.27	1657.27	1580.19	1580.19	1535.28	1535.28
Instance50-s5-40.dat	1260.14	1409.40	1197.00	1197.00	1197.00	1197.00
Instance50-s5-41.dat	1817.17	1817.17	1687.96	1687.96	1687.96	1687.96
Instance50-s5-42.dat	1509.39	1509.39	1191.46	1191.46	1191.46	1191.46
Instance50-s5-43.dat	1607.28	1607.99	1603.56	1603.56	1593.06	1593.06
Instance50-s5-44.dat	1111.28	1111.28	1063.25	1063.25	1047.96	1047.96
Instance50-s5-45.dat	1801.99	1801.99	1480.32	1480.32	1480.32	1480.32
Instance50-s5-46.dat	1248.41	1248.41	1074.88	1074.88	1074.88	1074.88
Instance50-s5-47.dat	1807.40	1807.40	1786.17	1786.17	1683.13	1683.13
Instance50-s5-48.dat	1178.88	1188.8	1178.88	1178.88	1078.28	1078.28
Instance50-s5-49.dat	1697.96	1705.9	1539.89	1546.77	1500.39	1510.98
Instance50-s5-50.dat	1201.11	1201.11	1201.11	1201.11	1072.42	1072.42
Instance50-s5-51.dat	1590.00	1590.00	1535.18	1535.18	1435.83	1435.83
Instance50-s5-52.dat	1132.20	1132.20	1132.20	1132.20	1132.20	1132.20
Instance50-s5-53.dat	1599.09	1599.09	1598.66	1598.66	1598.66	1598.66
Instance50-s5-54.dat	1206.97	1304.97	1201.90	1201.90	1201.90	1201.90
MEAN	1464.53	1479.83	1381.17	1381.56	1346.09	1346.67

Table 2: Best and average results

INST	GRASP	GRASP+LS	GRASP-PR	MS	MH	BC
Instance50-s5-37.dat	1599.86	1586.23	1545.99	1586.23	1587.95	1528.73
Instance50-s5-38.dat	1335.22	1222.27	1172.83	1340.49	1186.02	1187.39
Instance50-s5-39.dat	1657.27	1580.19	1535.28	1604.32	1525.24	1528.25
Instance50-s5-40.dat	1260.14	1197.00	1197.00	1387.28	1226.79	1179.64
Instance50-s5-41.dat	1817.17	1687.96	1687.96	1762.62	1726.04	1681.04
Instance50-s5-42.dat	1509.39	1191.46	1191.46	1559.39	1324.38	1232.87
Instance50-s5-43.dat	1607.28	1603.56	1593.06	1687.28	1453.11	1422.29
Instance50-s5-44.dat	1111.28	1063.25	1047.96	1227.26	1063.64	1061.25
Instance50-s5-45.dat	1801.99	1480.32	1480.32	1756.60	1497.91	1444.82
Instance50-s5-46.dat	1248.41	1074.88	1074.88	1148.31	1173.12	1068.50
Instance50-s5-47.dat	1807.40	1786.17	1683.13	1683.13	1620.7	1581.57
Instance50-s5-48.dat	1178.88	1178.88	1078.28	1319.96	1122.18	1092.32
Instance50-s5-49.dat	1697.96	1539.89	1500.39	1500.39	1508.87	1441.64
Instance50-s5-50.dat	1201.11	1201.11	1072.42	1131.65	1170.89	1089.67
Instance50-s5-51.dat	1590.00	1535.18	1435.83	1600.83	1456.12	1440.64
Instance50-s5-52.dat	1132.20	1132.20	1132.20	1145.54	1185.05	1109.52
Instance50-s5-53.dat	1599.09	1598.66	1598.66	1647.67	1569.59	1554.58
Instance50-s5-54.dat	1206.97	1201.90	1201.90	1201.90	1189.14	1135.39
MEAN	1464.53	1381.17	1346.09	1460.60	1365.93	1321.12
GAP	10.86%	4.55%	1.89%	10.56%	3.39%	

Table 3: Comparison of objective function values

INST	GRASP	GRASP+LS	GRASP-PR	MS	MH	BC
Instance50-s5-37.dat	5	25	33	2	71	10000
Instance50-s5-38.dat	5	27	38	2	68	10000
Instance50-s5-39.dat	5	30	37	2	66	10000
Instance50-s5-40.dat	2	64	69	5	73	10000
Instance50-s5-41.dat	7	55	59	32	97	10000
Instance50-s5-42.dat	8	49	54	2	67	10000
Instance50-s5-43.dat	7	36	48	1	66	10000
Instance50-s5-44.dat	12	12	44	14	66	10000
Instance50-s5-45.dat	3	13	13	1	73	10000
Instance50-s5-46.dat	2	22	22	7	69	10000
Instance50-s5-47.dat	28	80	90	53	76	10000
Instance50-s5-48.dat	15	15	45	13	74	10000
Instance50-s5-49.dat	4	8	34	21	86	10000
Instance50-s5-50.dat	5	54	67	1	98	10000
Instance50-s5-51.dat	5	53	65	1	82	10000
Instance50-s5-52.dat	5	55	55	12	67	10000
Instance50-s5-53.dat	5	38	38	1	45	10000
Instance50-s5-54.dat	17	77	96	32	30	10000
MEAN	8	40	50	11	71	10000

Table 4: Comparison of computational times in CPU seconds

recherche du Québec through their infrastructure grants.

References

- T.G. Crainic, S. Mancini, G. Perboli, and R. Tadei. Clustering-based heuristics for the two-echelon capacitated vehicle routing problem. Technical report, CIRRELT 2008-46, CIRRELT, Montreal, Canada, 2008.
- T.G. Crainic, N. Ricciardi, and G. Storchi. Models for evaluating and planning city logistics systems. *Transportation Science*, 43:432–454, 2009.
- T.G. Crainic, S. Mancini, G. Perboli, and R. Tadei. Two-echelon vehicle routing problem: A satellite location analysis. *Procedia Social and Behavioral Sciences*, 2:5944–5955, 2010.
- T.G. Crainic, S. Mancini, G. Perboli, and R. Tadei. Multi-start heuristics for the two-echelon vehicle routing problem. *Lecture Notes in Computer Science*, 6622:179–190, 2011.
- T.G. Crainic, S. Mancini, G. Perboli, and R. Tadei. Impact of generalized travel costs on satellite location in two-echelon vehicle routing problem. *Procedia Social and Behavioral Sciences*, 39:195–204, 2012.
- T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133, 1995.
- P. Festa and M.G.C. Resende. An annotated bibliography of GRASP - Part I: Algorithms. *International Transactions on Operational Research*, 16:1–24, 2009a.
- P. Festa and M.G.C. Resende. An annotated bibliography of GRASP - Part II: Applications. *International Transactions on Operational Research*, 16:131–172, 2009b.
- M. Gendreau and J.-Y. Potvin. *Handbook of Metaheuristics - Second Edition*. Springer, 2010.
- F. Glover and M. Laguna. *Tabu Search*. Kluwer, Boston, MA, 1997.
- F. Glover, M. Laguna, and R. Martí. Fundamentals of scatter search and path relinking. *Control and Cybernetics*, 29(3):653–684, 2000.
- M. Laguna and R. Martí. Grasp and path relinking for 2-layer straight line crossing minimization. *INFORMS Journal on Computing*, 11:44–52, 1999.
- S. Mancini. *The Two-Echelon Vehicle Routing Problem*. PhD thesis, Politecnico di Torino, 2011.
- G. Perboli, F. Pezzella, and R. Tadei. EVE-OPT: an hybrid algorithm for the capability vehicle routing problem. *Mathematical Methods of Operations Research*, 68:361–382, 2008.
- G. Perboli, R. Tadei, and F. Masoero. New families of valid inequalities for the two-echelon vehicle routing problem. *Electronic Notes in Discrete Mathematics*, 36:639–646, 2010.

- G. Perboli, R. Tadei, and D. Vigo. The two-echelon capacitated vehicle routing problem: models and math-based heuristics. *Transportation Science*, 45:364–380, 2011.
- M.G.C. Resende and C.C. Ribeiro. GRASP and path-relinking: Recent advances and applications. In *Proceedings of the Fifth Metaheuristics International Conference (MIC 2003)*, 2003.
- J. Villegas, C. Prins, C. Prodhon, A. Medaglia, and N. Velasco. GRASP/VND and multi-start evolutionary local search for the single truck and trailer routing problem with satellite depots. *Engineering Application of Artificial Intelligence*, 23:780–794, 2010.