

# A GRASP with restarts heuristic for the Steiner traveling salesman problem

Ruben Interian, Celso C. Ribeiro

Institute of Computing, Universidade Federal Fluminense,  
Niterói, RJ 24210-346, Brazil.  
rinterian@ic.uff.br, celso@ic.uff.br

## Abstract

Given a set of nodes and the distances between them, the traveling salesman problem (TSP) consists in finding the shortest route that visits each node exactly once and returns to the first. The Steiner traveling salesman problem (STSP) is a variant of the TSP that assumes that only a given subset of nodes must be visited by a shortest route, eventually visiting some nodes and edges more than once. In this paper, we extend some classical TSP constructive heuristics and neighborhood structures to the STSP variant. In particular, we propose a reduced 2-opt neighborhood and we show that it leads to better results in smaller computation times. Computational results with an implementation of a GRASP heuristic using path-relinking and restarts are reported. In addition, a set of test instances and best known solutions is made available for benchmarking purposes.

## 1 Introduction

The traveling salesman problem (TSP) is one of the fundamental combinatorial optimization problems [4, 12] and has numerous real-life applications in transportation, logistics, vehicle routing, genome sequencing, and other areas. Given a set of nodes and the distances between them, it consists in finding the shortest route that visits each node exactly once and returns to the first. Mathematically, the TSP can be defined as follows [2]. Given a graph  $G = (V, E)$  and a function  $w : E \rightarrow \mathbb{R}$  that associates a weight  $w(e)$  to each edge  $e \in E$ , the goal is to find a Hamiltonian cycle of minimum total weight (or cost). The traveling salesman problem is NP-hard, since its decision version is proven to be NP-complete by a simple reduction from the Hamiltonian cycle problem [4].

However, in many practical applications it is more frequent to find the following variant of the TSP. A set  $V_R \subseteq V$  of required nodes is given. Instead of searching for a Hamiltonian cycle visiting all nodes, a minimum-weight closed walk is requested that visits only the required nodes. Since only a walk is sought, nodes can be visited more than once and edges may be traversed more than once. The so-called Steiner traveling salesman problem (Steiner TSP, or STSP) was first proposed in [2, 3], where its NP-hardness is also proved. The Steiner TSP is specially suitable to model network design [1], package delivery [12, 13], and routing [7] problems. All of them are typically modeled using sparse graphs.

Most studies on the Steiner TSP focus on integer programming formulations and valid inequalities. The STSP is solved efficiently (in linear time) for series-parallel graphs in [2]. Compact, polynomial size integer programming formulations of the TSP are extended to the STSP in [7]. An extension of the Steiner TSP that adds penalties to the nodes not visited by the cycle is proposed in [11]. A network design problem consisting of multiple Steiner TSPs with order constraints is studied in [1], using an integer linear programming formulation and a branch-and-cut algorithm. An extension of the STSP in which the edge traversal costs are stochastic and correlated is studied in [6]. An online algorithm is proposed in [12, 13] to solve another extension of the STSP considering real-time edge blockages.

This paper is organized as follows. In the next section, adaptive greedy constructive heuristics for the Steiner TSP are presented. Section 3 reports on local search strategies that are explored by the GRASP with path-relinking heuristic presented in Section 4. Computational experiments are reported in Section 5 and extended in Section 6, where an improved strategy exploring periodical restarts is developed. Concluding remarks are drawn in the last section.

## 2 Greedy algorithms for the Steiner TSP

The following strategy can be applied as a heuristic for the Steiner TSP [7]. First, the instance of the STSP is reduced to a TSP instance in a complete graph defined by the set of required nodes, in which the new distances correspond to the shortest paths between every pair of required nodes in the original graph. Next, any exact or heuristic algorithm is used to solve the TSP in this new, complete graph. Finally, the solution of the TSP is converted into an STSP solution by expanding every edge by the corresponding shortest path between the two consecutive required nodes.

However, if the original STSP instance is a sparse graph, the conversion to a standard TSP instance significantly increases the number of edges, some of which may never be used. Therefore, instead of using a complete graph formed by the required nodes, we shall use the original graph for searching a minimum-weight closed walk. We use a straightforward adaptation of the nearest neighbor TSP adaptive greedy heuristic (see e.g. Chapter 3 of [10]) to the STSP described in Algorithm 1, which builds the solution greedily by choosing at each iteration the closest required node to that previously added to the walk.

The algorithm starts in line 1 by arbitrarily selecting any initial node  $i \in V_R$  to start the walk. The set of required nodes  $\mathcal{N}$  already visited by the walk is initialized in line 2. The partially built walk  $\mathcal{P}$  is initialized in line 3. The currently visited required node *current* is set in line 4. The loop in lines 5 to 11 is performed until all required nodes have been visited. At each iteration, the node *next* to be visited is set to the closest among all yet unvisited required nodes. The shortest path  $\mathcal{P}'$  from *current* to *next* is computed in line 7. The partially built walk  $\mathcal{P}$  is updated in line 8 by appending the shortest path  $\mathcal{P}'$  to it. The set of already visited required nodes  $\mathcal{N}$  and the current node are updated in lines 9 and 10, respectively. Finally, after completing the loop, the shortest path from *current* to the initial node  $i$  is appended to the walk in lines 12 and 13. The result is returned in line 14.

---

**Algorithm 1** Nearest neighbor adaptive greedy heuristic for STSP.

---

```

1: Select initial required node  $i \in V_R$ ;
2:  $\mathcal{N} \leftarrow \{i\}$ ;
3:  $\mathcal{P} \leftarrow \{i\}$ ;
4: current  $\leftarrow i$ ;
5: while  $\mathcal{N} \neq V_R$  do
6:   next  $\leftarrow$  closest node to current among all those in  $V_R \setminus \mathcal{N}$ ;
7:    $\mathcal{P}' \leftarrow$  shortest path from current to next;
8:    $\mathcal{P} \leftarrow \mathcal{P} \oplus \mathcal{P}'$ ;
9:    $\mathcal{N} \leftarrow \mathcal{N} \cup \{\textit{next}\}$ ;
10:  current  $\leftarrow \textit{next}$ ;
11: end while;
12:  $\mathcal{P}' \leftarrow$  shortest path from current to initial node  $i$ ;
13:  $\mathcal{P} \leftarrow \mathcal{P} \oplus \mathcal{P}'$ ;
14: return  $\mathcal{P}$ .
```

---

In the case of the Steiner TSP, the greedy criterion is the choice of the nearest required node to be visited.

Algorithms that add randomization to a greedy or adaptive greedy algorithm are called *semi-greedy* or *randomized greedy* algorithms. Randomization is an important feature in the implementation of effective heuristics. Semi-greedy algorithms act by replacing the deterministic greedy choice of the next element to be incorporated into the solution under construction by the random selection of an element from a restricted set of best candidate elements, called the restricted candidate list (RCL).

A simple quality-based scheme is used to define a restricted candidate list. Let  $g_{\min} = \min\{g_i : i \in V_R \setminus \mathcal{N} \text{ and } g_i \text{ is the shortest path from } \textit{current} \text{ to node } i\}$  and  $g_{\max} = \max\{g_i : i \in V_R \setminus \mathcal{N} \text{ and } g_i \text{ is the shortest path from } \textit{current} \text{ to node } i\}$ . Furthermore, let  $\alpha$  be such that  $0 \leq \alpha \leq 1$ . The RCL is formed by all yet unselected required nodes  $i \in V_R \setminus \mathcal{N}$  satisfying  $g_{\min} \leq g_i \leq g_{\min} + \alpha(g_{\max} - g_{\min})$ .

### 3 Local search

Local search procedures are used to iteratively improve the quality of an initial solution, usually obtained by a constructive heuristic. First-improving and best-improving strategies are proposed and compared in terms of their performance. Efficient objective function updates are used, without the need of recalculating the objective function values from scratch: the weight of the previous walk is used in order to find that of the walk obtained after the changes performed during each iteration.

#### 3.1 Neighborhood structure

The 2-opt neighborhood is the most commonly used neighborhood structure for the TSP problem and consists in replacing any pair of nonadjacent edges of the current solution by the unique pair of new edges that recreates a cycle.

The following property holds: Let  $W = (v_1, \dots, v_i, \dots, v_j, \dots, v_m)$  be any optimal solution of the Steiner TSP. Then, the subpath  $(v_i, \dots, v_j)$  is also a shortest path between the required nodes  $v_i$  and  $v_j$ . This is true because, if this subpath was not the shortest, then  $W$  would not be optimal. Therefore, it is not necessary to investigate moves that involve changes in the order in which the non-required nodes are visited. Then, the problem amounts to determining the order in which the required nodes should be visited and then finding the shortest path between any pair of consecutive required nodes in the walk.

In consequence, we explore a 2-opt neighborhood for the STSP that is formed by all moves that replace the paths between two pairs of consecutive required nodes in the walk by the two unique pairs of shortest paths that reconnect a closed walk.

#### 3.2 Reduced 2-opt neighborhood

A reduced 2-opt neighborhood can be defined in order to take advantage of the problem structure. In fact, convergence can be faster if only a few, promising moves in the neighborhood are considered.

We implement this idea in the following way. For each required node  $v$ , let  $I(v)$  be the set formed by all required nodes that are reachable from  $v$  by a shortest path that does not visit any other required node. In other words,  $I(v)$  represent the set of required nodes that are closer to  $v$ , in the sense that they necessarily belong to paths to farther nodes.

Using this auxiliary data structure, we restrict the 2-opt moves to pairs of consecutive required nodes  $(v_1, v_2)$  and  $(w_1, w_2)$  satisfying the condition that  $w_1 \in I(v_1)$ .

### 4 GRASP with path-relinking heuristic

GRASP (which stands for *greedy randomized adaptive search procedures*) is a multi-start metaheuristic, in which each iteration consists of two main phases: construction and local search. The first phase is the construction of a feasible solution, usually by a greedy randomized algorithm. Once a feasible solution is obtained, its neighborhood is investigated until a local minimum is found during the second phase of local search. The best overall solution is kept as the result. The reader is referred to Resende and Ribeiro [10] for a complete account of GRASP.

We used the adaptive greedy randomized heuristic presented in Section 2 and the local search strategies described in Section 3 to customize a GRASP with path-relinking heuristic for the Steiner TSP.

Path-relinking is an intensification strategy that explores trajectories connecting elite solutions produced by metaheuristics. Path-relinking is usually carried out between two solutions: one is the initial solution  $S_i$ , while the other is the guiding solution  $S_g$ . A path that connects these solutions is constructed in the search for better solutions. Local search may be applied to the best solution in the path, since there is no guarantee that this solution is locally optimal. In the context of GRASP, path-relinking may be used to connect solutions obtained after the local search step with elite solutions produced during previous iterations, providing a sort of memory mechanism.

More specifically, in the context of the STSP, path-relinking attempts to preserve common characteristics of good walks, i.e. common subpaths. As explained below, path-relinking matches the positions of the largest common subpath to the initial and guiding solutions and then swaps the positions of nodes that do not belong to this common subpath.

We first observe that any solution of the STSP has no unique representation as a sequence of the visited required nodes, since any closed walk can start from different nodes and can be traversed in two directions (forward and backward). Therefore, the representation of the initial and guiding solutions must be adjusted to facilitate the operation of relinking them. With this purpose, before applying path-relinking, we adjust the representations of the initial and guiding solutions by detecting the largest common subpath  $w_l = (v_i \dots v_j)$  between them.

In our implementation, we choose to detect the largest (or longest) common subpath, instead of the longest common subsequence, in order to prioritize consecutive sequences of nodes in both solutions. This problem is known as the largest (or longest) common substring (LCS) problem and can be solved in  $O(n)$  time and space [5].

The guiding solution  $S_g$  and the initial solution  $S_i$  are oriented in the same direction according with  $w_l$ . Next, the initial nodes of the walks associated with  $S_g$  and  $S_i$  are made to coincide with the initial node  $v_i$  of  $w_l$ .

To move from the initial to the guiding solution, path-relinking considers a restricted neighborhood. Each move in this restricted neighborhood involves the swap of two required nodes in the walk corresponding to the current solution that are not in the same positions as they are visited in the guiding solution. In addition, each move should place at least one of the two involved nodes in the appropriate position corresponding to the order in which it will be visited in the guiding solution. After two required nodes are swapped, the shortest paths from their predecessors and to their successors are updated. Since at least one node is placed in the appropriate position of the guiding solution at each iteration, path-relinking will take at most as many iterations as the number of required nodes that were misplaced in the initial solution with respect to the guiding solution.

Algorithm 2 presents the pseudo-code of path-relinking from the initial solution  $S_i$  to the guiding solution  $S_g$ . The current solution  $S$  and the best solutions  $S^*$  are initialized in line 1. The cost  $f^*$  of the best solution found by path-relinking is initialized in line 2. The loop in lines 3 to 10 is performed until the current solution reaches the guiding solution.  $S'$  is set to the best solution in the restricted neighborhood of the current solution  $S$  in line 4. The best solution  $S^*$  found by path-relinking and its cost  $f^*$  are updated in lines 6 and 7, respectively, if the new solution  $S'$  improved the previous best. The current solution is updated in line 9 and a new path-relinking iteration resumes. The best solution found by path-relinking is returned in line 11.

---

**Algorithm 2** Path-relinking algorithm for STSP.

---

```

1:  $S, S^* \leftarrow S_i$ ;
2:  $f^* \leftarrow cost(S_i)$ ;
3: while  $S \neq S_g$  do
4:    $S' \leftarrow$  best solution in the restricted neighborhood of  $S$ ;
5:   if  $cost(S') < f^*$  then
6:      $S^* \leftarrow S'$ ;
7:      $f^* \leftarrow cost(S')$ ;
8:   end if;
9:    $S \leftarrow S'$ ;
10: end while;
11: return  $S^*$ .

```

---

The pseudo-code in Algorithm 3 summarizes the main steps of the proposed GRASP with path-relinking (GRASP+PR) heuristic, following the same structure proposed in Section 9.3 of [10]. The set of elite solutions is initialized in line 1. The loop in lines 2 to 12 is performed until some stopping criterion is satisfied. An initial solution is built in line 3 by the greedy randomized constructive heuristic

described in Section 2. A local search procedure is used in line 4 to improve the solution obtained at the end of the construction phase. Except for the first iteration, when the elite set is still empty, lines 6 to 9 amount to the application of path-relinking. Line 6 randomly selects an elite solution  $S'$  from the elite set  $\mathcal{E}$ . The representation of solution  $S$  is adjusted considering the selected elite solution  $S'$ . Backward path-relinking is applied from the initial solution  $S_i = S'$  to the guiding solution  $S_g = S$ . Local search is applied in line 9 to the solution obtained by path-relinking. The elite set  $\mathcal{E}$  is updated with the new solution  $S$  in line 11. The best elite solution is returned in line 13.

---

**Algorithm 3** GRASP+PR algorithm for STSP
 

---

```

1:  $\mathcal{E} \leftarrow \emptyset$ ;
2: while stopping criterion not satisfied do
3:    $S \leftarrow \text{RandomizedGreedy}$ ;
4:    $S \leftarrow \text{LocalSearch}(S)$ ;
5:   if  $|\mathcal{E}| > 0$  then
6:     Select solution  $S'$  at random from  $\mathcal{E}$ ;
7:      $S \leftarrow \text{AdjustRepresentation}(S, S')$ 
8:      $S \leftarrow \text{PathReLinking}(S, S')$ ;
9:      $S \leftarrow \text{LocalSearch}(S)$ ;
10:  end if;
11:   $\text{UpdateEliteSet}(S, \mathcal{E})$ ;
12: end while;
13: Return the best solution  $S$  in  $\mathcal{E}$ .

```

---

## 5 Computational experiments

Several experiments were performed to assess the performance of the algorithms presented above and their variants. The algorithms were implemented in C# programming language and compiled by Roslyn, a reference C# compiler, in a Intel Core i5 machine with a 2.9 GHz processor and 8 GB of random-access memory, running under the Windows 10 operating system.

We considered the same test problems used by Letchford et al. [6, 7] and Zhang et al. [12, 13], created by a random generator described in [7]. This generator was designed to create graphs that resemble real-life road networks. It creates connected sparse graphs and a fraction of required nodes is specified for each instance. In addition to the graphs from [7], we considered some larger instances, with up to 300 nodes. Altogether, ten sparse weighted graphs with 50 to 300 nodes were used to assess the performance of the heuristics. Each graph generated two instances: one with  $\lfloor \frac{N}{3} \rfloor$  required nodes and another with  $\lfloor \frac{2 \cdot N}{3} \rfloor$  required nodes, where  $N$  is the total number of nodes, corresponding to 20 different instances. We observe that individual optimal values for each of these instances have not been previously reported in [7]. Due to space limitations, we report here only numerical results for the instances with  $\lfloor \frac{2 \cdot N}{3} \rfloor$  required nodes. Additional computational results will be reported in the final, extended version of this work.

### 5.1 Selecting the quality measure for the RCL

In order to compare the effect of the value of  $\alpha$  in the quality-based scheme used to define a restricted candidate list, we ran the randomized nearest neighbor constructive heuristic with  $\alpha = 0.1$  and  $\alpha = 0.05$ . The greedy randomized algorithm was applied to all instances. Average and best values over 100 runs are presented in Table 1. As for all tables that follow, the best solution values found for each instance are depicted in boldface. The randomized heuristic with  $\alpha = 0.05$  found significantly more better solutions. We observe that the use of a better constructive method for building the initial solutions is likely to improve the quality of the solutions produced by the GRASP heuristic.

Table 1: Greedy randomized heuristic: average and best value over 100 runs.

| Nodes | $\alpha = 0.10$ |             | $\alpha = 0.05$ |             |
|-------|-----------------|-------------|-----------------|-------------|
|       | average         | best        | average         | best        |
| 50    | 1207.51         | <b>979</b>  | 1213.34         | 1031        |
| 75    | 1308.21         | 1125        | 1309.49         | <b>1094</b> |
| 100   | 1607.99         | 1375        | 1599.44         | <b>1299</b> |
| 125   | 1929.31         | 1627        | 1911.41         | <b>1623</b> |
| 150   | 2110.56         | <b>1875</b> | 2102.40         | 1899        |
| 175   | 2248.23         | 1970        | 2265.40         | <b>1913</b> |
| 200   | 2615.95         | 2369        | 2614.78         | <b>2354</b> |
| 225   | 2817.25         | 2524        | 2844.12         | <b>2522</b> |
| 250   | 3022.70         | 2748        | 2988.79         | <b>2723</b> |
| 300   | 3242.32         | <b>2952</b> | 3264.87         | 2977        |

## 5.2 Reduced 2-opt neighborhood

We now address the benefits of using the reduced 2-opt neighborhood, designed specifically for this problem. Preliminary computational experiments have shown that the use of this reduced neighborhood led to some target objective function values much faster than the use of the entire 2-opt neighborhood. These empirical observations were explored by the implementation of an alternative VND (variable neighborhood descent) local search procedure. First, only moves in the reduced 2-opt neighborhood are applied. The full neighborhood is explored only after a local minimum is obtained in the reduced 2-opt neighborhood.

Table 2 illustrates the efficiency of the VND approach when compared with the classic use of the full 2-opt neighborhood. It presents the solution values and the computation times in seconds for 100 iterations of the pure GRASP (without path-relinking) heuristic using both the pure 2-opt neighborhood and the VND approach for local search. In both cases, local search has been implemented following a best-improvement strategy.

Table 2: 2-opt vs reduced 2-opt neighborhoods, 100 GRASP iterations.

| Nodes | 2-opt neighborhood |         | reduced 2-opt |         |
|-------|--------------------|---------|---------------|---------|
|       | value              | seconds | value         | seconds |
| 50    | <b>978</b>         | 0.484   | <b>978</b>    | 0.36    |
| 75    | <b>1035</b>        | 1.078   | 1045          | 0.985   |
| 100   | 1239               | 2.359   | <b>1208</b>   | 1.610   |
| 125   | 1496               | 4.454   | <b>1469</b>   | 2.921   |
| 150   | 1643               | 7.468   | <b>1615</b>   | 4.703   |
| 175   | 1743               | 11.313  | <b>1719</b>   | 6.687   |
| 200   | 1976               | 17.469  | <b>1925</b>   | 9.187   |
| 225   | 2094               | 24.891  | <b>2047</b>   | 12.828  |
| 250   | 2224               | 32.546  | <b>2170</b>   | 15.297  |
| 300   | 2409               | 55.718  | <b>2285</b>   | 26.578  |

The VND local search strategy starting by the reduced 2-opt neighborhood led to the best solutions for nine out of the ten test problems. In addition, its computation times have been significantly smaller for all instances. As an example, in the case of the largest instance with 300 nodes, the time taken by 100 GRASP iterations using the VND local search strategy amounted to only 47.8% of the time taken when exclusively the complete 2-opt neighborhood is used. The GRASP heuristic using the VND local search strategy performs better both in terms of solution quality and computation times.

### 5.3 Probabilistic choice of $\alpha$

We have already shown in Section 5.1 that, although choosing  $\alpha = 0.05$  most often leads to better results than  $\alpha = 0.10$  for the greedy randomized heuristic proposed for the STSP, for some instances the latter was a better choice than the former. Different probabilistic strategies were considered in [8] for the choice of the RCL parameter  $\alpha$ , in contrast with the commonly used choice of fixing its value (see also [10]). It was shown that a randomly chosen  $\alpha$  from a decreasing nonuniform discrete probability distribution offers a good compromise between the running time of the algorithm and the quality of the solutions produced by the randomized heuristic. We relied on this work to sustain that it can be a good choice, in addition to the previously used “good” value  $\alpha = 0.05$ , to consider also other higher values for this parameter with smaller probabilities of being chosen at each iteration. Therefore, in the following experiments, we used  $\alpha = 0.05$  with a probability 70% and the values  $\alpha = 0.10$  and  $\alpha = 0.20$  with probabilities 20% and 10%, respectively.

### 5.4 Path-relinking

In this section, we address the impact of path-relinking in the search process. Table 3 shows the lengths of the solutions produced by the nearest neighbor adaptive greedy heuristic, by the pure GRASP heuristic running for 200 iterations (together with its running time in seconds), and by the GRASP with backward path-relinking running by the same time taken by 200 pure GRASP iterations. The randomized heuristic used in the construction phase of both the pure GRASP and GRASP with path-relinking algorithms makes use of the probabilistic criterion for the choice of  $\alpha$  discussed in Section 5.3. The local search phase of both the pure GRASP and the GRASP with path-relinking algorithms was implemented using the best improvement VND strategy starting by the reduced 2-opt neighborhood. Path-relinking made use of elite sets formed by at most ten elements.

Table 3: GRASP and GRASP with path-relinking, 200 pure GRASP iterations.

| Nodes | Greedy ( $\alpha = 0$ ) |            | GRASP (200 iterations) |             | GRASP+PR (same running time) |            |
|-------|-------------------------|------------|------------------------|-------------|------------------------------|------------|
|       | value                   | value      | seconds                | value       | seconds                      | iterations |
| 50    | 1356                    | <b>978</b> | 0.750                  | <b>978</b>  | 0.750                        | 179        |
| 75    | 1204                    | 1031       | 1.781                  | <b>1029</b> | 1.782                        | 183        |
| 100   | 1290                    | 1211       | 3.484                  | <b>1193</b> | 3.485                        | 182        |
| 125   | 1915                    | 1450       | 6.297                  | <b>1427</b> | 6.313                        | 196        |
| 150   | 1847                    | 1615       | 9.812                  | <b>1567</b> | 9.844                        | 184        |
| 175   | 2085                    | 1704       | 14.64                  | <b>1667</b> | 14.641                       | 191        |
| 200   | 2484                    | 1952       | 19.828                 | <b>1895</b> | 19.828                       | 177        |
| 225   | 2549                    | 2024       | 27.266                 | <b>1973</b> | 27.281                       | 191        |
| 250   | 2523                    | 2165       | 35.078                 | <b>2080</b> | 35.172                       | 182        |
| 300   | 2759                    | 2308       | 62.328                 | <b>2219</b> | 62.609                       | 193        |

Path-relinking considerably improved GRASP performance, leading to better solutions for all instances in the same running time and fewer iterations than the pure GRASP. Time-to-target plots for pure GRASP and GRASP with path-relinking algorithms for a 200-node instance are shown in Figure 1. The target value is 2000. Each algorithm was run 200 times. The plots in this figure provide empirical evidence that algorithm GRASP+PR outperforms pure GRASP for this instance and target value.

## 6 Restart strategies for GRASP with path-relinking

Resende and Ribeiro [9] have shown that restart strategies are able to reduce the running time to reach a target solution value for many problems. We apply the same type of restart( $\kappa$ ) strategy, in which the elite set is emptied and the heuristic restarted from scratch after  $\kappa$  consecutive iterations have been performed

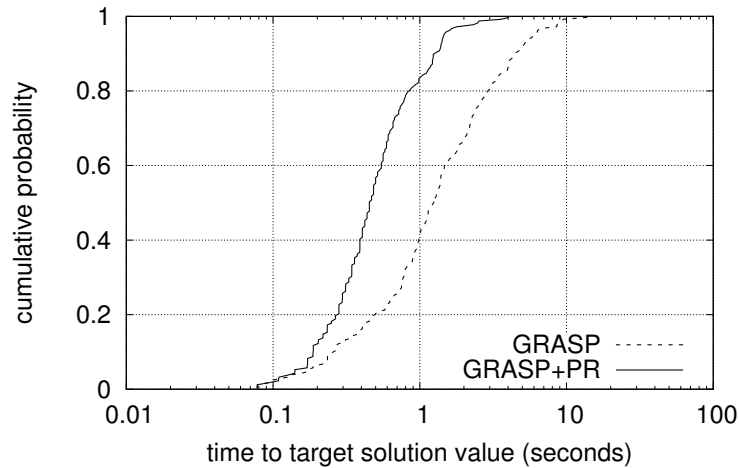


Figure 1: Time-to-target plot for 200-node instance and target value set to 2000.

without improvement in the best solution found. Computational results for restart strategies for STSP are displayed in Table 4, showing that they contribute to find better solutions in the same number of iterations, mainly when the problem size increases.

Table 4: Restart strategies for 1000 iterations.

| Nodes | no restarts |         | restart(100) |         | restart(200) |         |
|-------|-------------|---------|--------------|---------|--------------|---------|
|       | value       | seconds | value        | seconds | value        | seconds |
| 50    | <b>978</b>  | 4.03    | <b>978</b>   | 3.96    | <b>978</b>   | 3.95    |
| 75    | <b>1029</b> | 9.46    | <b>1029</b>  | 9.45    | <b>1029</b>  | 9.26    |
| 100   | <b>1193</b> | 18.95   | <b>1193</b>  | 18.73   | <b>1193</b>  | 18.43   |
| 125   | 1421        | 33.04   | <b>1417</b>  | 33.93   | 1420         | 33.09   |
| 150   | 1565        | 53.39   | 1564         | 53.04   | <b>1562</b>  | 52.23   |
| 175   | 1657        | 77.25   | 1665         | 76.84   | <b>1652</b>  | 76.23   |
| 200   | 1883        | 105.53  | 1885         | 105.68  | <b>1867</b>  | 105.20  |
| 225   | 1941        | 148.96  | 1953         | 144.68  | <b>1928</b>  | 142.26  |
| 250   | 2054        | 173.68  | 2073         | 175.01  | <b>2035</b>  | 176.04  |
| 300   | 2203        | 304.40  | <b>2192</b>  | 310.03  | 2205         | 296.76  |

As previously observed in [9, 10], the effect of restart strategies can be mainly noticed in the longest runs. Considering the 200 runs for the 200-node instance with the target value set to 1900, they are associated with the column corresponding to the fourth quartile of Table 5. Entries in this quartile correspond to those in the heavy tails of the runtime distributions. The restart strategies in general do not affect too much the other quartiles of the distributions, which is a desirable characteristic. Compared to the norestart strategy, the restart(200) strategy was able to reduce not only the average running time in the fourth quartile, but also in the third and second quartiles. Consequently, strategy restart(200) performed the best among those tested, with the smallest average running times over the 200 runs.

## 7 Concluding remarks

In the Steiner TSP, one seeks a minimum-weight closed walk that visits a subset of required nodes. Since only a walk is sought, nodes can be visited more than once and edges may be traversed more than once.

We developed a GRASP with path-relinking and restarts for solving the Steiner Traveling Salesman Problem. The algorithm used in the construction phase is a randomized extension of the nearest



Table 5: Summary of computational results for each restart strategy for the 200-node instance: 200 independent runs were executed for each strategy. Each run was made to stop when a solution as good as the target solution value 1900 was found. For each strategy, the table shows the distribution of the running times by quartile. For each quartile, the table gives the average running times in seconds over all runs in that quartile. The average running times over the 200 runs are also given for each strategy.

| Strategy         | Average running times in quartile (seconds) |              |               |               |               |
|------------------|---|--------------|---------------|---------------|---------------|
|                  | 1st   | 2nd          | 3rd           | 4th           | average       |
| Without restarts | 3.648                                       | 9.915        | 17.952        | 37.355        | 17.218        |
| restart(100)     | 2.933                                       | 8.466        | 17.067        | 37.509        | 16.494        |
| restart(200)     | 2.955                                       | <b>8.093</b> | <b>15.410</b> | <b>34.878</b> | <b>15.334</b> |

neighbor heuristic for the Traveling Salesman Problem. A variable neighborhood descent (VND) strategy exploring a reduced 2-opt neighborhood is used to optimize a best improving local search scheme. Path-relinking and restart strategies are used to improve the efficiency of the GRASP algorithm.

Extensive computational results for a set of instances previously used in the literature are reported. Since neither optimal values nor even upper bounds have been previously reported for these instances, the solutions obtained by the GRASP with path-relinking and restarts heuristic proposed in this work cannot be directly be compared with other solutions.

As a step towards avoiding this difficulty and facilitating the research on this problem, we made all test instances considered in this paper (together with their best known solutions and their costs) available at <http://www2.ic.uff.br/~rinterian/instances/allinstances.html>. This website will be continuously updated with information provided by other researchers working on this problem with optimal values, upper bounds and best known feasible solutions for these and other benchmarking instances.

## References

- [1] S. Borne, A. R. Mahjoub, and R. Taktak. A branch-and-cut algorithm for the multiple Steiner TSP with order constraints. *Electronic Notes in Discrete Mathematics*, 41:487–494, 2013.
- [2] G. Cornuéjols, J. Fonlupt, and D. Naddef. The traveling salesman problem on a graph and some related integer polyhedra. *Mathematical Programming*, 33:1–27, 1985.
- [3] B. Fleischmann. A cutting plane procedure for the travelling salesman problem on road networks. *European Journal of Operational Research*, 21(3):307–317, 1985.
- [4] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
- [5] L. C. K. Hui. Color set size problem with application to string matching. In *Proceedings of the Third Annual Symposium on Combinatorial Pattern Matching, CPM '92*, pages 230–243. Springer-Verlag, 1992.
- [6] A. N. Letchford and S. D. Nasiri. The Steiner travelling salesman problem with correlated costs. *European Journal of Operational Research*, 245:62–69, 2015.
- [7] A. N. Letchford, S. D. Nasiri, and D. Oliver Theis. Compact formulations of the Steiner traveling salesman problem and related problems. *European Journal of Operational Research*, 228:83–92, 2013.
- [8] M. Prais and C. C. Ribeiro. Parameter variation in GRASP procedures. *Investigación Operativa*, 9:1–20, 2000.

- 
- [9] M. G. C. Resende and C. C. Ribeiro. Restart strategies for GRASP with path-relinking heuristics. *Optimization Letters*, 5:467–478, 2011.
- [10] M. G. C. Resende and C. C. Ribeiro. *Optimization by GRASP*. Springer, New York, 2016.
- [11] J.-J. Salazar-González. The Steiner cycle polytope. *European Journal of Operational Research*, 147:671–679, 2003.
- [12] H. Zhang, W. Tong, Y. Xu, and G. Lin. The Steiner traveling salesman problem with online edge blockages. *European Journal of Operational Research*, 243:30–40, 2015.
- [13] H. Zhang, W. Tong, Y. Xu, and G. Lin. The Steiner traveling salesman problem with online advanced edge blockages. *Computers & Operations Research*, 70:26–38, 2016.