GRASP and path-relinking: Recent advances and applications

Mauricio G.C. Resende* Celso C. Ribeiro†

April 6, 2003

Abstract

This paper addresses recent advances and application of hybridizations of greedy randomized adaptive search procedures (GRASP) and path-relinking. We present a template for implementing path-relinking as an intensification procedure for GRASP. Enhancements to the procedure, recently described in the literature, are reviewed. The effectiveness of the procedure is illustrated experimentally.

1 Introduction

In this paper, we consider the minimization version of a combinatorial optimization problem defined by a cost function f and a set of feasible solutions X, where we seek an optimal solution $x^* \in X$ such that $f(x^*) \leq f(x)$, $\forall x \in X$.

A greedy randomized adaptive search procedure (GRASP) [6, 7, 12] is a multi-start metaheuristic which applies local search to starting solutions generated by a greedy randomized construction procedure. Algorithm 1 illustrates a basic GRASP heuristic for minimization.

Until recently, most implementations of GRASP assumed independence of its iterations, thus making no use of memory structures. Laguna and Martí [10] showed how path-relinking could be used in GRASP as an intensification mechanism. Path-relinking was originally proposed by Glover [9] as a way to explore trajectories between elite solutions obtained by tabu search or scatter search. Using one or more elite solutions, paths in the solution space leading to other elite solutions are explored in the search for better solutions. To generate paths, moves are selected to introduce attributes in the current solution that are present in the elite guiding solution. A number of extensions, improvements, and successful applications of GRASP with path-relinking have been reported in the literature [2, 5, 8, 11, 13, 14, 15, 16].

In this tutorial, we present recent advances and applications of GRASP with path-relinking. In Section 2, we present a basic template for implementing path-relinking

*Internet and Network Systems Research, AT&T Labs Research, 180 Park Avenue, Room C241, Florham Park, NJ 07932 USA. mgcr@research.att.com

†Department of Computer Science, Catholic University of Rio de Janeiro, Rio de Janeiro, RJ 22450, Brazil. celso@inf.puc-rio.br

```
Data      : Number of iterations i_max
Result    : Solution x* ∈ X
f* ← ∞;
for i = 1, ..., i_max do
    x ← GreedyRandomized();
    x ← LocalSearch(x);
    if f(x) < f* then
        f* ← f(x);
        x* ← x;
    end
end
```

Algorithm 1: A basic GRASP heuristic for minimization.

in a GRASP. Extensions and enhancement are discussed in Section 3. We conclude in Section 4 with some applications and computational results, showing the benefits reaped with path-relinking.

2 Path-relinking

Consider two solutions x_s and x_t on which we wish to apply path-relinking from x_s to x_t. Algorithm 2 illustrates the pseudo-code of the path-relinking procedure.

```
Data      : Starting solution x_s and target solution x_t
Result    : Best solution x* in path from x_s to x_t
Compute symmetric difference Δ(x_s, x_t);
f* ← min{f(x_s), f(x_t)};
x* ← argmin{f(x_s), f(x_t)};
x ← x_s;
while Δ(x, x_t) ≠ ∅ do
    m* ← argmin{f(x ⊕ m), ∀m ∈ Δ(x, x_t)};
    Δ(x ⊕ m*, x_t) ← Δ(x, x_t) \ {m*};
    x ← x ⊕ m*;
    if f(x) < f* then
        f* ← f(x);
        x* ← x;
    end
end
```

Algorithm 2: Path-relinking.

The procedure starts by computing the symmetric difference $\Delta(x_s, x_t)$ between the two solutions, i.e. the set of moves needed to reach x_t from x_s. A path of solutions is generated linking x_s and x_t. The best solution x^* in this path is returned by the algorithm. At each step, the procedure examines all moves $m \in \Delta(x, x_t)$ from the current solution x and selects the one which results in the least cost solution, i.e. the one which minimizes $f(x \oplus m)$, where $x \oplus m$ is the solution resulting from applying move m to solution x. The best move m^* is made, producing solution $x \oplus m^*$. The set of available moves is updated. If necessary, the best solution x^* is updated. The procedure terminates when x_t is reached, i.e. when $\Delta(x, x_t) = \emptyset$.

The basic scheme for implementing path-relinking with GRASP maintains an elite set P of solutions found during the optimization. The solution produced by each GRASP iteration is relinked with a solution chosen at random from the elite set. Algorithm 3 illustrates such a procedure.

Data : Number of iterations $i_{\max}$
Result : Solution $x^* \in X$
$P \leftarrow \emptyset$;
$f^* \leftarrow \infty$;
for $i = 1, \ldots, i_{\max}$ **do**
    $x \leftarrow$ `GreedyRandomized`();
    $x \leftarrow$ `LocalSearch`(x);
    if $i \geq 1$ **then**
        Choose target solution $x_t \in P$ at random;
        $x \leftarrow$ `PathRelinking`(x, x_t);
    end
    Update the elite set P with x;
    if $f(x) < f^*$ **then**
        $f^* \leftarrow f(x)$;
        $x^* \leftarrow x$;
    end
end

Algorithm 3: A basic GRASP with path-relinking heuristic for minimization.

3 Extensions and enhancements

Path-relinking maintains an elite set P of at most `MaxElite` solutions found during the search. To ensure that the elite set is sufficiently diverse to be effective, one usually implements it as follows. The first `MaxElite` distinct solutions found are inserted into the elite set. After that, a candidate solution is added to P if its cost is smaller than that of the best elite solution, or if its cost is smaller than that of the worst elite solution and it is sufficiently different from all the elite set solutions. If accepted into the elite set P, the new solution replaces the solution most similar to it from the set of elite solutions having larger cost than it [13]. The elite set can be periodically renewed [1] if no change in the elite set is observed for a specified number of GRASP iterations. One way to do this is to set the objective function values of the worse half of the elite set to infinity. This way new elite set solutions will be created.

Several alternative schemes have been proposed for path-relinking. Instead of selecting the element from the elite set to be combined with the GRASP solution uniformly at random, one can select the elite solution at random, but with probabilities proportional to the symmetric differences of the elite solutions and the GRASP solution [13]. This tends to lead to longer paths, which have higher chance of producing improved solutions. Since path-relinking can be computationally demanding, it need not be applied after every GRASP iteration, but rather periodically. Usually the paths from x_s to x_t and from x_t to x_s are different and both can be explored. Alternatively, these two paths can be explored in alternate iterations [14]. Since paths can be long, the full trajectory need not be followed. One can restrict following a truncated path starting at x_s and another starting at x_t [1]. Instead of selecting a single solution from the elite set, one can select a subset of solutions of the elite set (in one extreme a single solution and on the other all elite solutions) and perform path-relinking between all these elite solutions and the GRASP solution [2].

Path-relinking can also be done between all pairs of solutions from the elite set. This can be done intermittently during the GRASP iterations and/or at the end of the GRASP iterations, as a post-optimization phase [2, 15]. One effective way to implement this type of intensification is described in [13, 15]. It starts with the initial pool $P_0 = P$ and produces a series of pools $P_1, P_2, \ldots$, such that all pairs of solutions in pool P_k are combined and the solutions generated by this process are added to a new pool P_{k+1}, following the same constraints for elite set membership as before. This procedure is repeated until it creates a generation in which the best solution is not better than the best found in previous generations.

4 Applications and computational results

GRASP with path-relinking has been applied to a wide range of combinatorial optimization problems, e.g. [1, 2, 5, 8, 10, 11, 13, 14, 15, 16]. Different strategies for implementing GRASP with path-relinking heuristics were compared computationally with pure GRASP (without path-relinking). Figure 1 illustrates this type of comparison by considering empirical probability distributions of the random variable *time-to-target-value* on four problem types: MAX-CUT [8], three index assignment [2], private virtual circuit routing [11], and job shop scheduling [1]. The methodology to derive these plots is described in [3]. The four plots illustrate a general conclusion for many applications that GRASP indeed benefits greatly from the use of path-relinking. In all these problems, GRASP with path-relinking was much faster than pure GRASP, with speedups of up to two orders of magnitude on MAX-CUT and private virtual circuit routing.

Path-relinking is also a very effective strategy to introduce cooperation in parallel implementations of GRASP heuristics. In this context, a number of processors perform GRASP iterations, while a single processor maintains a centralized set of elite solutions. Results obtained for the 2-path network design problem [14] are illustrated in Figure 2 by the same type of plot, showing the speedup obtained by the cooperative strategy with respect to the independent one on a cluster of eight processors. Much larger improvements can be obtained if a bigger number of processors is available.

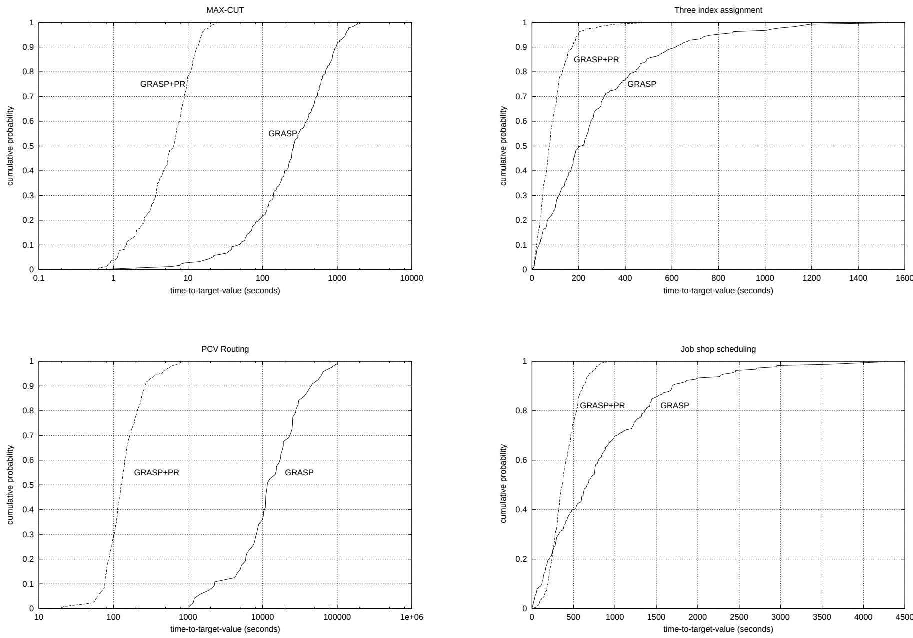

Figure 1: Probability distributions of time-to-target-value on instances of the MAX-CUT, three index assignment, private virtual circuit routing, and job shop scheduling problems, for GRASP and GRASP with path-relinking.

Finally, we notice that path-relinking can also be successfully used in conjunction with implementations of other metaheuristics such as VNS and ant colonies, as recently reported respectively by Festa et al. [8] and Aloise et al. [4].

References

[1] R.M. Aiex, S. Binato, and M.G.C. Resende. Parallel GRASP with path-relinking for job shop scheduling. *Parallel Computing*, 29:393–430, 2003.

[2] R.M. Aiex, M.G.C. Resende, P.M. Pardalos, and G. Toraldo. GRASP with path-relinking for the three-index assignment problem. *INFORMS Journal on Computing*, to appear.

[3] R.M. Aiex, M.G.C. Resende, and C.C. Ribeiro. Probability distribution of solution time in GRASP: An experimental investigation. *Journal of Heuristics*, 8:343–373, 2002.

[4] D.J. Aloise, D. Aloise, C.T.M. Rocha, J.D. Melo, and C.C. Ribeiro. Scheduling terrestrial oil production drills. Technical report, Catholic University of Rio de Janeiro, 2003.

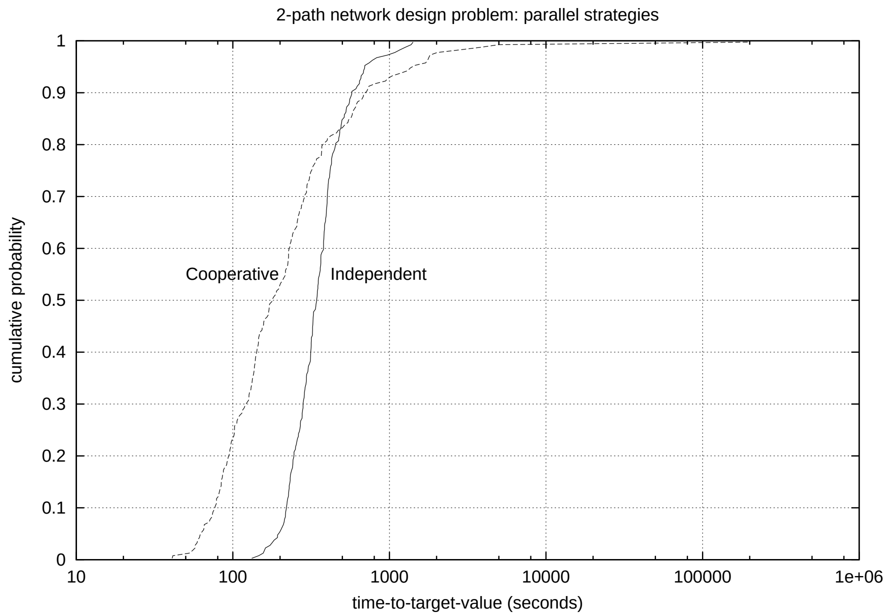

Figure 2: Probability distributions of time-to-target-value on an instance of the 2-path network design problem for cooperative and independent parallel implementations of GRASP with path-relinking.

[5] S.A. Canuto, M.G.C. Resende, and C.C. Ribeiro. Local search with perturbations for the prize-collecting Steiner tree problem in graphs. *Networks*, 38:50–58, 2001.

[6] T.A. Feo and M.G.C. Resende. A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters*, 8:67–71, 1989.

[7] T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133, 1995.

[8] P. Festa, P.M. Pardalos, M.G.C. Resende, and C.C. Ribeiro. Randomized heuristics for the MAX-CUT problem. *Optimization Methods & Software*, 7:1033–1058, 2002.

[9] F. Glover. Tabu search and adaptive memory programing – Advances, applications and challenges. In R.S. Barr, R.V. Helgason, and J.L. Kennington, editors, *Interfaces in Computer Science and Operations Research*, pages 1–75. Kluwer, 1996.

[10] M. Laguna and R. Martí. GRASP and path relinking for 2-layer straight line crossing minimization. *INFORMS Journal on Computing*, 11:44–52, 1999.

[11] M.G.C. Resende and C.C. Ribeiro. A GRASP with path-relinking for private virtual circuit routing. *Networks*, 41:104–114, 2003.

[12] M.G.C. Resende and C.C. Ribeiro. Greedy randomized adaptive search procedures. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, pages 219–249. Kluwer Academic Publishers, 2003.

[13] M.G.C. Resende and R.F. Werneck. A GRASP with path-relinking for the p-median problem. Technical report, Internet and Network Systems Research Center, AT&T Labs Research, Florham Park, NJ, 2002.

[14] C.C. Ribeiro and I. Rosseti. A parallel GRASP heuristic for the 2-path network design problem. *Lecture Notes in Computer Science*, 2400:922–926, 2002.

[15] C.C. Ribeiro, E. Uchoa, and R.F. Werneck. A hybrid GRASP with perturbations for the Steiner problem in graphs. *INFORMS Journal on Computing*, 14:228–246, 2002.

[16] M.C. Souza, C. Duhamel, and C.C. Ribeiro. A GRASP with path-relinking heuristic for the capacitated minimum spanning tree problem. In M.G.C. Resende and J.P. Sousa, editors, *Metaheuristics: Computer Decision-Making*. Kluwer Academic Publishers, to appear.