

Chapter 6

Greedy Randomized Adaptive Search Procedures: Advances and Extensions



Mauricio G. C. Resende and Celso C. Ribeiro

Abstract A greedy randomized adaptive search procedure (GRASP) is a multi-start metaheuristic for combinatorial optimization problems, in which each iteration consists basically of two phases: construction and local search. The construction phase builds a feasible solution whose neighborhood is investigated until a local minimum is found during the local search phase. The best overall solution is kept as the result. In this chapter, we first describe the basic components of GRASP. Successful implementation techniques are discussed and illustrated by numerical results obtained for different applications. Enhanced or alternative solution construction mechanisms and techniques to speed up the search are also described: Alternative randomized greedy construction schemes, Reactive GRASP, cost perturbations, bias functions, memory and learning, Lagrangean constructive heuristics and Lagrangean GRASP, local search on partially constructed solutions, hashing, and filtering. We also discuss implementation strategies of memory-based intensification and post-optimization techniques using path-relinking. Restart strategies to speedup the search, hybridizations with other metaheuristics, and applications are also reviewed.

M. G. C. Resende (✉)
Amazon.com, Seattle, WA, USA
University of Washington, Seattle, WA, USA
e-mail: mgrc@uw.edu

C. C. Ribeiro
Universidade Federal Fluminense, Niterói, Rio de Janeiro, Brazil
e-mail: celso@ic.uff.br

6.1 Introduction

We consider in this chapter a combinatorial optimization problem, defined by a finite ground set $E = \{1, \dots, n\}$, a set of feasible solutions $F \subseteq 2^E$, and an objective function $f: 2^E \rightarrow \mathbb{R}$. In its minimization version, we seek an optimal solution $S^* \in F$ such that $f(S^*) \leq f(S)$, $\forall S \in F$. The ground set E , the cost function f , and the set of feasible solutions F are defined for each specific problem. For instance, in the case of the traveling salesman problem, the ground set E is that of all edges connecting the cities to be visited, $f(S)$ is the sum of the costs of all edges in S , and F is formed by all edge subsets that determine a Hamiltonian cycle.

GRASP (Greedy Randomized Adaptive Search Procedure) [90, 91] is a multi-start or iterative metaheuristic, in which each iteration consists of two phases: construction and local search. The construction phase builds a solution using a greedy randomized adaptive algorithm. If this solution is not feasible, then it is necessary to apply a repair procedure to achieve feasibility or to make a new attempt to build a feasible solution. Once a feasible solution is obtained, its neighborhood is investigated until a local minimum is found during the local search phase. The best overall solution is kept as the result.

Extensive literature surveys on greedy randomized adaptive search procedures are presented in [98–100, 212, 213, 224]. A first book on GRASP was published in 2016 by Resende and Ribeiro [215].

The pseudo-code in Fig. 6.1 illustrates the main blocks of a GRASP procedure for minimization, in which `Max_Iterations` iterations are performed and `Seed` is used as the initial seed for the pseudo-random number generator.

```

procedure GRASP(Max_Iterations, Seed)
1  Read_Input();
2  for  $k = 1, \dots, \text{Max\_Iterations}$  do
3      Solution  $\leftarrow$  Greedy_Randomized_Construction(Seed);
4      if Solution is not feasible then
5          Solution  $\leftarrow$  Repair(Solution);
6      end;
7      Solution  $\leftarrow$  Local_Search(Solution);
8      Update_Solution(Solution, Best_Solution);
9  end;
10 return Best_Solution;
end GRASP.

```

Fig. 6.1 Pseudo-code of the GRASP metaheuristic

Figure 6.2 illustrates the construction phase with its pseudo-code. At each iteration of this phase, let the set of candidate elements be formed by all elements of the ground set E that can be incorporated into the partial solution being built, without impeding the construction of a feasible solution with the remaining ground set elements. The selection of the next element for incorporation is determined by the

evaluation of all candidate elements according to a greedy evaluation function. This greedy function usually represents the incremental increase in the cost function due to the incorporation of this element into the solution under construction. The evaluation of the elements by this function leads to the creation of a restricted candidate list (RCL) formed by the best elements, i.e. those whose incorporation to the current partial solution results in the smallest incremental costs (this is the greedy aspect of the algorithm). The element to be incorporated into the partial solution is randomly selected from those in the RCL (this is the probabilistic aspect of the heuristic). Once the selected element is incorporated into the partial solution, the candidate list is updated and the incremental costs are reevaluated (this is the adaptive aspect of the heuristic). The above steps are repeated while there exists at least one candidate element. This strategy is similar to the semi-greedy heuristic proposed by Hart and Shogan [122], which is also a multi-start approach based on greedy randomized constructions, but without local search.

```

procedure Greedy_Randomized_Construction(Seed)
1  Solution  $\leftarrow$   $\emptyset$ ;
2  Initialize the set of candidate elements;
3  Evaluate the incremental costs of the candidate elements;
4  while there exists at least one candidate element do
5      Build the restricted candidate list (RCL);
6      Select an element  $s$  from the RCL at random;
7      Solution  $\leftarrow$  Solution  $\cup$   $\{s\}$ ;
8      Update the set of candidate elements;
9      Reevaluate the incremental costs;
10 end;
11 return Solution;
end Greedy_Randomized_Construction.

```

Fig. 6.2 Pseudo-code of the construction phase

A randomized greedy construction procedure is not always able to produce a feasible solution. It may be necessary to apply a repair procedure to the solution to achieve feasibility. Examples of repair procedures can be found in [80, 81, 169, 180].

The solutions generated by a greedy randomized construction are not necessarily optimal, even with respect to simple neighborhoods. The local search phase usually improves the constructed solution. A local search algorithm works in an iterative fashion by successively replacing the current solution by a better solution in its neighborhood. It terminates when no better solution is found in the neighborhood. The pseudo-code of a basic local search algorithm starting from the solution `Solution` constructed in the first phase (and possibly made feasible by the repair heuristic) and using a neighborhood N is given in Fig. 6.3.

```
procedure Local_Search(Solution)
1  while Solution is not locally optimal do
2      Find  $s' \in N(\text{Solution})$  with  $f(s') < f(\text{Solution})$ ;
3      Solution  $\leftarrow s'$ ;
4  end;
5  return Solution;
end Local_Search.
```

Fig. 6.3 Pseudo-code of the local search phase

The speed and the effectiveness of a local search procedure depend on several aspects, such as the neighborhood structure, the neighborhood search technique, the strategy used for the evaluation of the cost function value at the neighbors, and the starting solution itself. The construction phase plays a very important role with respect to this last aspect, building high-quality starting solutions for the local search. Simple neighborhoods are typically used. The neighborhood search can be implemented using either a *best-improving* or a *first-improving* strategy. In the case of the best-improving strategy, all neighbors are investigated and the current solution is replaced by the best neighbor. In the case of a first-improving strategy, the current solution moves to the first neighbor whose cost function value is smaller than that of the current solution. In practice, we observed on many applications that quite often both strategies lead to the same final solution, but in smaller computation times when the first-improving strategy is used. We also observed that premature convergence to a bad local minimum is more likely to occur with a best-improving strategy.

6.2 Construction of the Restricted Candidate List

An especially appealing characteristic of GRASP is the ease with which it can be implemented. Few parameters need to be set and tuned. Therefore, development can focus on implementing appropriate data structures for efficient construction and local search algorithms. GRASP has two main parameters: one related to the stopping criterion and the other to the quality of the elements in the restricted candidate list.

The stopping criterion used in the pseudo-code described in Fig. 6.1 is determined by the number `Max_Iterations` of iterations. Although the probability of finding a new solution improving the incumbent (current best solution) decreases with the number of iterations, the quality of the incumbent does not worsen with the number of iterations. Since the computation time does not vary much from iteration to iteration, the total computation time is predictable and increases linearly with the number of iterations. Consequently, the larger the number of iterations, the larger will be the computation time and the better will be the solution found.

For the construction of the RCL used in the first phase we consider, without loss of generality, a minimization problem as the one formulated in Sect. 6.1. We denote

by $c(e)$ the incremental cost associated with the incorporation of element $e \in E$ into the solution under construction. At any GRASP iteration, let c^{min} and c^{max} be, respectively, the smallest and the largest incremental costs.

The restricted candidate list RCL is made up of the elements $e \in E$ with the best (i.e., the smallest) incremental costs $c(e)$. This list can be limited either by the number of elements (cardinality-based) or by their quality (value-based). In the first case, it is made up of the p elements with the best incremental costs, where p is a parameter. In this chapter, the RCL is associated with a threshold parameter $\alpha \in [0, 1]$. The restricted candidate list is formed by all elements $e \in E$ which can be feasibly inserted into the partial solution under construction and whose quality is superior to the threshold value, i.e., $c(e) \in [c^{min}, c^{min} + \alpha(c^{max} - c^{min})]$. The case $\alpha = 0$ corresponds to a pure greedy algorithm, while $\alpha = 1$ is equivalent to a random construction. The pseudo-code in Fig. 6.4 is a refinement of the greedy randomized construction pseudo-code shown in Fig. 6.2. It shows that the parameter α controls the amounts of greediness and randomness in the algorithm.

```

procedure Greedy_Randomized_Construction( $\alpha$ , Seed)
1  Solution  $\leftarrow \emptyset$ ;
2  Initialize the candidate set:  $C \leftarrow E$ ;
3  Evaluate the incremental cost  $c(e)$  for all  $e \in C$ ;
4  while  $C \neq \emptyset$  do
5       $c^{min} \leftarrow \min\{c(e) \mid e \in C\}$ ;
6       $c^{max} \leftarrow \max\{c(e) \mid e \in C\}$ ;
7      RCL  $\leftarrow \{e \in C \mid c(e) \leq c^{min} + \alpha(c^{max} - c^{min})\}$ ;
8      Select an element  $s$  from the RCL at random;
9      Solution  $\leftarrow$  Solution  $\cup \{s\}$ ;
10     Update the candidate set  $C$ ;
11     Reevaluate the incremental cost  $c(e)$  for all  $e \in C$ ;
12 end;
13 return Solution;
end Greedy_Randomized_Construction.

```

Fig. 6.4 Refined pseudo-code of the construction phase

GRASP construction can be viewed as a repetitive sampling technique. Each iteration produces a sample solution from an unknown distribution, whose mean and variance are functions of the restrictive nature of the RCL. For example, if the RCL is restricted to a single element, then the same solution will be produced at all iterations. The variance of the distribution will be zero and the mean will be equal to the value of the greedy solution. If the RCL is allowed to have more elements, then many different solutions will be produced, implying a larger variance. Since greediness plays a smaller role in this case, the average solution value should be worse than that of the greedy solution. However, the value of the best solution found outperforms the average value and can be near-optimal or even optimal. It is unlikely that GRASP will find an optimal solution if the average solution value is high, even if there is a large variance in the overall solution values. On the other hand, if there is

little variance in the overall solution values, it is also unlikely that GRASP will find an optimal solution, even if the average solution is low. What often leads to good solutions are relatively low average solution values in the presence of a relatively large variance, such as is the case for $\alpha = 0.2$.

Another interesting observation is that the distances between the solutions obtained at each iteration and the best solution found increase as the construction phase moves from more greedy to more random. This causes the average time taken by the local search to increase. Very often, many GRASP solutions may be generated in the same amount of time required by the local search procedure to converge from a single random start. In these cases, the time saved by starting the local search from good initial solutions can be used to improve solution quality by performing more GRASP iterations.

These results are illustrated in Table 6.1 and Fig. 6.5, for an instance of the MAXSAT problem [219] where 1000 iterations were run. For each value of α ranging from 0 (purely random construction for maximization problems) to 1 (purely greedy construction for maximization problems), we give in Table 6.1 the average Hamming distance between each solution built during the construction phase and the corresponding local optimum obtained after local search, the average number of moves from the former to the latter, the local search time in seconds, and the total processing time in seconds. Figure 6.5 summarizes the values observed for the total processing time and the local search time. We notice that both time measures considerably decrease as α tends to 1, approaching the purely greedy choice. In particular, we observe that the average local search time taken by $\alpha = 0$ (purely random) is approximately 2.5 times longer than the time taken by $\alpha = 0.9$ (almost greedy). In this example, two to three greedily constructed solutions can be investigated in the same time needed to apply local search to one single randomly constructed solution. The appropriate choice of the value of the RCL parameter α is clearly critical and relevant to achieve a good balance between computation time and solution quality.

Table 6.1 Average number of moves and local search time as a function of the RCL parameter α for a maximization problem

α	Avg. distance	Avg. moves	Local search time (s)	Total time (s)
0.0	12.487	12.373	18.083	23.378
0.1	10.787	10.709	15.842	20.801
0.2	10.242	10.166	15.127	19.830
0.3	9.777	9.721	14.511	18.806
0.4	9.003	8.957	13.489	17.139
0.5	8.241	8.189	12.494	15.375
0.6	7.389	7.341	11.338	13.482
0.7	6.452	6.436	10.098	11.720
0.8	5.667	5.643	9.094	10.441
0.9	4.697	4.691	7.753	8.941
1.0	2.733	2.733	5.118	6.235

Prais and Ribeiro [201] show that using a single fixed value for the RCL parameter α very often hinders finding a high-quality solution, which could be found if another value is used. They propose an extension of the basic GRASP procedure, which they call *Reactive GRASP*, in which the parameter α is self-tuned and its value is periodically modified depending on the quality of the solutions obtained along the search. In particular, computational experiments on the problem of traffic assignment in communication satellites [202] show that Reactive GRASP finds better solutions than the basic algorithm for many test instances. These results motivated the study of the behavior of GRASP with different strategies for the variation of the value of the RCL parameter α :




Fig. 6.5 Total CPU time and local search CPU time as a function of the RCL parameter α for a maximization problem (1000 repetitions for each value of α)

- R: α self tuned with a Reactive GRASP procedure;
- E: α randomly chosen from a uniform discrete probability distribution;
- H: α randomly chosen from a decreasing non-uniform discrete probability distribution;
- F: α fixed, close to the purely greedy choice value.

We summarize the results obtained by the experiments reported in [200, 201]. These four strategies are incorporated into the GRASP procedures developed for four different optimization problems: (P-1) matrix decomposition for traffic assignment in communication satellites [202]; (P-2) set covering [90]; (P-3) weighted MAX-SAT [219, 221]; and (P-4) graph planarization [210, 226]. Let

$$\Psi = \{\alpha_1, \dots, \alpha_m\}$$

be the set of possible values for the parameter α for the first three strategies. The strategy for choosing and self-tuning the value of α in the case of the Reactive GRASP procedure (R) is described later in Sect. 6.3. In the case of the strategy (E) based on using the discrete uniform distribution, all choice probabilities are equal to $1/m$. The third case corresponds to the a hybrid strategy (H), in which the authors considered $p(\alpha = 0.1) = 0.5$, $p(\alpha = 0.2) = 0.25$, $p(\alpha = 0.3) = 0.125$, $p(\alpha = 0.4) = 0.03$, $p(\alpha = 0.5) = 0.03$, $p(\alpha = 0.6) = 0.03$, $p(\alpha = 0.7) = 0.01$, $p(\alpha = 0.8) = 0.01$, $p(\alpha = 0.9) = 0.01$, and $p(\alpha = 1.0) = 0.005$. Finally, in the last strategy (F), the value of α is fixed as recommended in the original references of problems P-1 to P-4 cited above, where this parameter was tuned for each problem. A subset of the literature instances was considered for each class of test problems. The results reported in [201] are summarized in Table 6.2. For each problem, we first list the number of instances considered. Next, for each strategy, we give the number of times it found the best solution (hits), as well as the average CPU time (in seconds) on an IBM 9672 model R34. The number of iterations was fixed at 10,000.

Table 6.2 Computational results for different strategies for the variation of parameter α

Problem Instances	R		E		H		F	
	Hits	Time	Hits	Time	Hits	Time	Hits	Time
P-1	36	34 579.0	35	358.2	32	612.6	24	642.8
P-2	7	7 1346.8	6	1352.0	6	668.2	5	500.7
P-3	44	22 2463.7	23	2492.6	16	1740.9	11	1625.2
P-4	37	28 6363.1	21	7292.9	24	6326.5	19	5972.0
Total	124	91	85		78		59	

Strategy (F) presented the shortest average computation times for three out the four problem types. It was also the one with the least variability in the constructed solutions and, in consequence, found the best solution the fewest times. The reactive strategy (R) is the one which most often found the best solutions, however, at the cost of computation times that are longer than those of some of the other strategies. The high number of hits observed for strategy (E) also illustrates the effectiveness of strategies based on the variation of the RCL parameter.

6.3 Alternative Construction Mechanisms

A possible shortcoming of the standard GRASP framework is the independence of its iterations, i.e., the fact that it does not learn from the search history or from solutions found in previous iterations. This is so because the basic algorithm discards information about any solution previously encountered that does not improve

the incumbent. Information gathered from good solutions can be used to implement memory-based procedures to influence the construction phase, by modifying the selection probabilities associated with each element of the RCL or by enforcing specific choices. Another possible shortcoming of the greedy randomized construction is its complexity. At each step of the construction, each yet unselected candidate element has to be evaluated by the greedy function. In cases where the difference between the number of elements in the ground set and the number of elements that appear in a solution is large, this may not be very efficient.

In this section, we consider enhancements and alternative techniques for the construction phase of GRASP. They include random plus greedy, sampled greedy, Reactive GRASP, cost perturbations, bias functions, memory and learning, local search on partially constructed solutions, and Lagrangean GRASP heuristics.

6.3.1 *Random Plus Greedy and Sampled Greedy Construction*

In Sect. 6.2, we described the semi-greedy construction scheme used to build randomized greedy solutions that serve as starting points for local search. Two other randomized greedy approaches were proposed in [216], with smaller worst-case complexities than the semi-greedy algorithm.

Instead of combining greediness and randomness at each step of the construction procedure, the *random plus greedy* scheme applies randomness during the first p construction steps to produce a random partial solution. Next, the algorithm completes the solution with one or more pure greedy construction steps. The resulting solution is randomized greedy. One can control the balance between greediness and randomness in the construction by changing the value of the parameter p . Larger values of p are associated with solutions that are more random, while smaller values result in greedier solutions.

Similar to the random plus greedy procedure, the *sampled greedy* construction also combines randomness and greediness but in a different way. This procedure is also controlled by a parameter p . At each step of the construction process, the procedure builds a restricted candidate list by randomly sampling $\min\{p, |C|\}$ elements of the candidate set C . Each element of the RCL is evaluated by the greedy function. The element with the smallest greedy function value is added to the partial solution. This two-step process is repeated until there are no more candidate elements. The resulting solution is also randomized greedy. The balance between greediness and randomness can be controlled by changing the value of the parameter p , i.e. the number of candidate elements that are sampled. Small sample sizes lead to more random solutions, while large sample sizes lead to greedier solutions.

6.3.2 Reactive GRASP

The first strategy to incorporate a learning mechanism in the memoryless construction phase of the basic GRASP was the Reactive GRASP procedure introduced in Sect. 6.2. In this case, the value of the RCL parameter α is not fixed, but instead is randomly selected at each iteration from a discrete set of possible values. This selection is guided by the solution values found along the previous iterations. One way to accomplish this is to use the rule proposed in [202]. Let $\Psi = \{\alpha_1, \dots, \alpha_m\}$ be a set of possible values for α . The probabilities associated with the choice of each value are all initially made equal to $p_i = 1/m$, for $i = 1, \dots, m$. Furthermore, let z^* be the incumbent solution and let A_i be the average value of all solutions found using $\alpha = \alpha_i$, for $i = 1, \dots, m$. The selection probabilities are periodically reevaluated by taking $p_i = q_i / \sum_{j=1}^m q_j$, with $q_i = z^* / A_i$ for $i = 1, \dots, m$. The value of q_i will be larger for values of $\alpha = \alpha_i$ leading to the best solutions on average. Larger values of q_i correspond to more suitable values for the parameter α . The probabilities associated with the more appropriate values will then increase when they are reevaluated.

The reactive approach leads to improvements over the basic GRASP in terms of robustness and solution quality, due to greater diversification and less reliance on parameter tuning. In addition to the applications in [200–202], this approach has been used in power system transmission network planning [46], job shop scheduling [49], channel assignment in mobile phone networks [118], rural road network development [249], capacitated location [71], strip-packing [15], and a combined production-distribution problem [50].

6.3.3 Cost Perturbations

The idea of introducing some noise into the original costs is similar to that in the so-called “noising” method of Charon and Hudry [57, 58]. It adds more flexibility into the algorithm design and may be even more effective than the greedy randomized construction of the basic GRASP procedure in circumstances where the construction algorithms are not very sensitive to randomization. This is indeed the case for the shortest-path heuristic of Takahashi and Matsuyama [259], used as one of the main building blocks of the construction phase of the hybrid GRASP procedure proposed by Ribeiro et al. [233] for the Steiner problem in graphs. Another situation where cost perturbations can be very effective appears when no greedy algorithm is available for straightforward randomization. This happens to be the case of the hybrid GRASP developed by Canuto et al. [54] for the prize-collecting Steiner tree problem, which makes use of the primal-dual algorithm of Goemans and Williamson [117] to build initial solutions using perturbed costs.

In the case of the GRASP for the prize-collecting Steiner tree problem described in [54], a new solution is built at each iteration using node prizes updated by a perturbation function, based on the structure of the current solution. Two different prize

perturbation schemes were used. In *perturbation by eliminations*, the primal-dual algorithm used in the construction phase is driven to build a new solution without some of the nodes that appeared in the solution constructed in the previous iteration. In *perturbation by prize changes*, some noise is introduced into the node prizes to change the objective function, similarly to what is proposed in [57, 58].

The cost perturbation methods used in the GRASP for the minimum Steiner tree problem described in [233] incorporate learning mechanisms associated with intensification and diversification strategies. Three distinct weight randomization methods were applied. At a given GRASP iteration, the modified weight of each edge is randomly selected from a uniform distribution over an interval which depends on the selected weight randomization method applied at that iteration. The different weight randomization methods use frequency information and may be used to enforce intensification and diversification strategies. The experimental results reported in [233] show that the strategy combining these three perturbation methods is more robust than any of them used in isolation, leading to the best overall results on a quite broad mix of test instances with different characteristics. The GRASP heuristic using this cost perturbation strategy is among the most effective heuristics currently available for the Steiner problem in graphs.

6.3.4 Bias Functions

In the construction procedure of the basic GRASP, the next element to be introduced in the solution is chosen at random from the candidates in the RCL. The elements of the RCL are assigned equal probabilities of being chosen. However, any probability distribution can be used to bias the selection toward some particular candidates. Another construction mechanism was proposed by Bresina [51], where a family of such probability distributions is introduced. They are based on the rank $r(e)$ assigned to each candidate element $e \in C$, according to its greedy function value. Several bias functions were proposed, such as:

- random bias: $\text{bias}(r) = 1$;
- linear bias: $\text{bias}(r) = 1/r$;
- log bias: $\text{bias}(r) = \log^{-1}(r+1)$;
- exponential bias: $\text{bias}(r) = e^{-r}$; and
- polynomial bias of order n : $\text{bias}(r) = r^{-n}$.

Let $r(e)$ denote the rank of element $e \in C$ and let $\text{bias}(r(e))$ be one of the bias functions defined above. Once these values have been evaluated for all elements in the candidate set C , the probability $\pi(e)$ of selecting element $e \in C$ is

$$\pi(e) = \frac{\text{bias}(r(e))}{\sum_{e' \in C} \text{bias}(r(e'))}. \quad (6.1)$$

The evaluation of these bias functions may be restricted to the elements of the RCL. Bresina's selection procedure restricted to elements of the RCL was used in [49]. The standard GRASP uses a random bias function.

6.3.5 *Intelligent Construction: Memory and Learning*

Fleurent and Glover [106] observed that the basic GRASP does not use a long-term memory (information gathered in previous iterations) and proposed a long-term memory scheme to address this issue in multi-start heuristics. Long-term memory is one of the fundamentals on which tabu search relies.

Their scheme maintains a pool of elite solutions to be used in the construction phase. To become an elite solution, a solution must be either better than the best member of the pool, or better than its worst member and sufficiently different from the other solutions in the pool. For example, one can count identical solution vector components and set a threshold for rejection.

A *strongly determined variable* is one that cannot be changed without eroding the objective or changing significantly other variables. A *consistent variable* is one that receives a particular value in a large portion of the elite solution set. Let the *intensity function* $I(e)$ be a measure of the strong determination and consistency features of a solution element $e \in E$. Then, $I(e)$ becomes larger as e appears more often in the pool of elite solutions. The intensity function is used in the construction phase as follows. Recall that $c(e)$ is the greedy function, i.e. the incremental cost associated with the incorporation of element $e \in E$ into the solution under construction. Let $K(e) = F(c(e), I(e))$ be a function of the greedy and intensification functions. For example, $K(e) = \lambda c(e) + I(e)$. The intensification scheme biases selection from the RCL to those elements $e \in E$ with a high value of $K(e)$ by setting its selection probability to be $p(e) = K(e) / \sum_{s \in \text{RCL}} K(s)$.

The function $K(e)$ can vary with time by changing the value of λ . For example, λ may be set to a large value that is decreased when diversification is called for. Procedures for changing the value of λ are given by Fleurent and Glover [106] and Binato et al. [49].

6.3.6 *POP in Construction*

The Proximate Optimality Principle (POP) is based on the idea that “good solutions at one level are likely to be found ‘close to’ good solutions at an adjacent level” [115]. Fleurent and Glover [106] provided a GRASP interpretation of this principle. They suggested that imperfections introduced during steps of the GRASP construction can be “ironed-out” by applying local search during (and not only at the end of) the GRASP construction phase.

Because of efficiency considerations, a practical implementation of POP to GRASP consists in applying local search a few times during the construction phase, but not at every construction iteration. Local search was applied by Binato et al. [49] after 40% and 80% of the construction moves were performed, as well as at the end of the construction phase.

6.3.7 Lagrangean GRASP Heuristics

Lagrangean relaxation [45, 105] is a mathematical programming technique that can be used to provide lower bounds for minimization problems. Held and Karp [123, 124] were among the first to explore the use of the dual multipliers produced by Lagrangean relaxation to derive lower bounds, applying this idea in the context of the traveling salesman problem. Lagrangean heuristics further explore the use of different dual multipliers to generate feasible solutions. Beasley [43, 44] described a Lagrangean heuristic for set covering.

6.3.7.1 Lagrangean Relaxation and Subgradient Optimization

Lagrangean relaxation can be used to provide lower bounds for combinatorial optimization problems. However, the primal solutions produced by the algorithms used to solve the Lagrangean dual problem are not necessarily feasible. Lagrangean heuristics exploit dual multipliers to generate primal feasible solutions.

Given a mathematical programming problem \mathcal{P} formulated as

$$f^* = \min f(x) \tag{6.2}$$

$$g_i(x) \leq 0, \quad i = 1, \dots, m, \tag{6.3}$$

$$x \in X, \tag{6.4}$$

its Lagrangean relaxation is obtained by associating dual multipliers $\lambda_i \in \mathbb{R}_+$ with each inequality (6.3), for $i = 1, \dots, m$. This results in the following *Lagrangean relaxation problem* $LRP(\lambda)$

$$\min f'(x) = f(x) + \sum_{i=1}^m \lambda_i \cdot g_i(x) \tag{6.5}$$

$$x \in X, \tag{6.4}$$

whose optimal solution $x(\lambda)$ gives a lower bound $f'(x(\lambda))$ to the optimal value of the original problem \mathcal{P} defined by (6.2)–(6.4). The best (dual) lower bound is given by the solution of the *Lagrangean dual problem* \mathcal{D}

$$f_{\mathcal{D}} = f'(x(\lambda^*)) = \max_{\lambda \in \mathbb{R}_+^m} f'(x(\lambda)). \tag{6.6}$$

Subgradient optimization is used to solve the dual problem \mathcal{D} defined by (6.6). Subgradient algorithms start from any feasible set of dual multipliers, such as $\lambda_i = 0$, for $i = 1, \dots, m$, and iteratively generate updated multipliers.

At any iteration q , let λ^q be the current vector of multipliers and let $x(\lambda^q)$ be an optimal solution to problem $LRP(\lambda^q)$, whose optimal value is $f'(x(\lambda^q))$. Furthermore, let \bar{f} be a known upper bound to the optimal value of problem \mathcal{P} . Additionally, let $g^q \in \mathbb{R}^m$ be a subgradient of $f'(x)$ at $x = x(\lambda^q)$, with $g_i^q = g_i(x(\lambda^q))$ for $i = 1, \dots, m$. To update the Lagrangean multipliers, the algorithm makes use of a step size

$$d^q = \frac{\eta \cdot (\bar{f} - f'(x(\lambda^q)))}{\sum_{i=1}^m (g_i^q)^2}, \quad (6.7)$$

where $\eta \in (0, 2]$. Multipliers are then updated as

$$\lambda_i^{q+1} = \max\{0; \lambda_i^q - d^q \cdot g_i^q\}, \quad i = 1, \dots, m, \quad (6.8)$$

and the subgradient algorithm proceeds to iteration $q + 1$.

6.3.7.2 A Template for Lagrangean Heuristics

We describe next a template for Lagrangean heuristics that make use of the dual multipliers λ^q and of the optimal solution $x(\lambda^q)$ to each problem $LRP(\lambda^q)$ to build feasible solutions to the original problem \mathcal{P} defined by (6.2)–(6.4). In the following, we assume that the objective function and all constraints are linear functions, i.e. $f(x) = \sum_{j=1}^n c_j x_j$ and $g_i(x) = \sum_{j=1}^n d_{ij} x_j - e_i$, for $i = 1, \dots, m$.

Let \mathcal{H} be a primal heuristic that builds a feasible solution x to \mathcal{P} , starting from the initial solution $x^0 = x(\lambda^q)$ at every iteration q of the subgradient algorithm. Heuristic \mathcal{H} is first applied using the original costs c_j , i.e. using the cost function $f(x)$. In any subsequent iteration q of the subgradient algorithm, \mathcal{H} uses either the Lagrangean reduced costs $c'_j = c_j - \sum_{i=1}^m \lambda_i^q d_{ij}$ or the complementary costs $\bar{c}_j = (1 - x_j(\lambda^q)) \cdot c_j$.

Let $x^{\mathcal{H}, \gamma}$ be the solution obtained by heuristic \mathcal{H} , using a generic cost vector γ corresponding to either one of the above modified cost schemes or to the original cost vector. Its cost can be used to update the upper bound \bar{f} to the optimal value of the original problem. This upper bound can be further improved by local search and is used to adjust the step size defined by Eq. (6.7).

Figure 6.6 shows the pseudo-code of a Lagrangean heuristic. Lines 1–4 initialize the upper and lower bounds, the iteration counter, and the dual multipliers. The iterations of the subgradient algorithm are performed along the loop defined in lines 5–24. The reduced costs are computed in line 6 and the Lagrangean relaxation problem is solved in line 7. In the first iteration of the Lagrangean heuristic, the original cost vector is assigned to γ in line 9, while in subsequent iterations a modified cost vector is assigned to γ in line 11. Heuristic \mathcal{H} is applied in line 13 at the first iteration and after every H iterations thereafter (i.e., whenever the iteration counter q is a multiple of the input parameter H) to produce a feasible solution $x^{\mathcal{H}, \gamma}$ to problem \mathcal{P} . If

the cost of this solution is smaller than the current upper bound, then the best solution and its cost are updated in lines 14–18. If the lower bound $f'(x(\lambda^q))$ is greater than the current lower bound $f_{\mathcal{D}}$, then $f_{\mathcal{D}}$ is updated in line 19. Line 20 computes a subgradient at $x(\lambda^q)$ and line 21 computes the step size. The dual multipliers are updated in line 22 and the iteration counter is incremented in line 23. The best solution found and its cost are returned in line 24.

The strategy proposed by Held et al. [125] is commonly used in the implementation of Lagrangean heuristics to update the dual multipliers from one iteration to the next. Beasley [44] reported as computationally useful the adjustment of components of the subgradients to zero whenever they do not effectively contribute to the update of the multipliers, i.e., arbitrarily setting $g_i^q = 0$ whenever $g_i^q > 0$ and $\lambda_i^q = 0$, for $i = 1, \dots, m$.

```

procedure Lagrangean_Heuristic( $H$ )
1   $\bar{f} \leftarrow +\infty$ ;
2   $f_{\mathcal{D}} \leftarrow -\infty$ ;
3   $q \leftarrow 0$ ;
4   $\lambda_i^q \leftarrow 0, i = 1, \dots, m$ ;
5  repeat
6    Compute reduced costs:  $c'_j \leftarrow c_j - \sum_{i=1}^m \lambda_i^q d_{ij}, j = 1, \dots, n$ ;
7    Solve  $LRP(\lambda^q)$  to obtain a solution  $x(\lambda^q)$ ;
8    if  $q = 0$  then
9       $\gamma \leftarrow c$ ;
10   else
11     Set  $\gamma$  to the modified cost vector  $c'$  or  $\bar{c}$ ;
12   end-if;
13   if  $q$  is a multiple of  $H$  then apply heuristic  $\mathcal{H}$  with cost vector  $\gamma$  to obtain  $x^{\mathcal{H}\gamma}$ ;
14   if  $f(x^{\mathcal{H}\gamma}) < \bar{f}$ 
15     then do;
16      $x^* \leftarrow x^{\mathcal{H}\gamma}$ ;
17      $\bar{f} \leftarrow f(x^{\mathcal{H}\gamma})$ ;
18   end-if;
19   if  $f'(x(\lambda^q)) > f_{\mathcal{D}}$  then  $f_{\mathcal{D}} \leftarrow f'(x(\lambda^q))$ ;
20   Compute a subgradient:  $g_i^q \leftarrow g_i(x(\lambda^q)), i = 1, \dots, m$ ;
21   Compute the step size:  $d^q \leftarrow \eta \cdot (\bar{f} - f'(x(\lambda^q))) / \sum_{i=1}^m (g_i^q)^2$ ;
22   Update the dual multipliers:  $\lambda_i^{q+1} \leftarrow \max\{0, \lambda_i^q - d^q g_i^q\}, i = 1, \dots, m$ ;
23    $q \leftarrow q + 1$ ;
24 until stopping criterion satisfied;
25 return  $x^*, f(x^*)$ ;
end Lagrangean_Heuristic.

```

Fig. 6.6 Pseudo-code of a template for a Lagrangean heuristic

Different choices for the initial solution x^0 , for the modified costs γ , and for the primal heuristic \mathcal{H} itself lead to different variants of the above algorithm. The integer parameter H defines the frequency in which \mathcal{H} is applied. The smaller the value of H , the greater the number of times \mathcal{H} is applied. Therefore, the computa-

tion time increases as the value of H decreases. In particular, one should set $H = 1$ if the primal heuristic \mathcal{H} is to be applied at every iteration.

6.3.7.3 Lagrangean GRASP

Pessoa et al. [195, 196] proposed the hybridization of GRASP and Lagrangean relaxation leading to the Lagrangean GRASP heuristic described below. Different choices for the primal heuristic \mathcal{H} in the template of the algorithm in Fig. 6.6 lead to distinct Lagrangean heuristics. We consider two variants: the first makes use of a greedy algorithm with local search, while in the second a GRASP with path-relinking (see Sect. 6.4) is used.

Greedy heuristic: This heuristic greedily repairs the solution $x(\lambda^q)$ produced in line 7 of the Lagrangean heuristic described in Fig. 6.6 to make it feasible for problem \mathcal{P} . It makes use of the modified costs (c' or \bar{c}). Local search can be applied to the resulting solution, using the original cost vector c . We refer to this approach as a *greedy Lagrangean heuristic* (GLH).

GRASP heuristic: Instead of simply performing one construction step followed by local search, as GLH does, this variant applies a GRASP heuristic to repair the solution $x(\lambda^q)$ produced in line 7 of the Lagrangean heuristic to make it feasible for problem \mathcal{P} .

Although the GRASP heuristic produces better solutions than the greedy heuristic, the greedy heuristic is much faster. To appropriately address this trade-off, we adapt line 10 of Fig. 6.6 to use the GRASP heuristic with probability β and the greedy heuristic with probability $1 - \beta$, where β is a parameter of the algorithm.

We note that this strategy involves three main parameters: the number H of iterations after which the basic heuristic is always applied, the number Q of iterations performed by the GRASP heuristic when it is chosen as the primal heuristic, and the probability β of choosing the GRASP heuristic as \mathcal{H} . We shall refer to the Lagrangean heuristic that uses this hybrid strategy as $\text{LAGRASP}(\beta, H, Q)$.

We next summarize computational results obtained for 135 instances of the set k -covering problem. These instances have up to 400 constraints and 4000 binary variables. The set k -covering, or set multi-covering, problem is an extension of the classical set covering problem, in which each element is required to be covered at least k times. The problem finds applications in the design of communication networks and in computational biology.

The first experiment with the GRASP Lagrangean heuristic established the relationship between running times and solution quality for different parameter settings. Parameter β , the probability of GRASP being applied as the heuristic \mathcal{H} , was set to

0, 0.25, 0.50, 0.75, and 1. Parameter H , the number of iterations between successive calls to the heuristic \mathcal{H} , was set to 1, 5, 10, and 50. Parameter Q , the number of iterations carried out by the GRASP heuristic, was set to 1, 5, 10, and 50. By combining some of these parameter values, 68 variants of the hybrid LAGRASP(β, H, Q) heuristic were created. Each variant was applied eight times to a subset of 21 instances, with different initial seeds being given to the random number generator.

The plot in Fig. 6.7 summarizes the results for all variants evaluated, displaying points whose coordinates are the values of the average deviation from the best known solution value and the total time in seconds for processing the eight runs on all instances, for each combination of parameter values. Eight variants of special interest are identified and labeled with the corresponding parameters β, H , and Q , in this order. These variants correspond to selected Pareto points in the plot in Fig. 6.7. Setting $\beta = 0$ and $H = 1$ corresponds to the greedy Lagrangean heuristic (GLH) or, equivalently, to LAGRASP(0,1,-), whose average deviation (in percentage) from the best value amounts to 0.12% in 4859.16 s of total running time. Table 6.3 shows the average deviation from the best known solution value and the total time for each of the eight selected variants.

In another experiment, all 135 test instances were considered for the comparison of the above selected eight variants of LAGRASP. Table 6.4 summarizes the results




Fig. 6.7 Average deviation from the best value and total running time for 68 different variants of LAGRASP on a reduced set of 21 instances of the set k -covering problem: each point represents a unique combination of parameters β, H , and Q

Table 6.3 Summary of the numerical results obtained with the selected variants of the GRASP Lagrangean heuristic on a reduced set of 21 instances of the set k -covering problem

Heuristic	Average deviation	Total time (s)
LAGRASP(1,1,50)	0.09%	399,101.14
LAGRASP(0.50,1,1)	0.11%	6198.46
LAGRASP(0,1,-)	0.12%	4859.16
LAGRASP(0.25,5,10)	0.24%	4373.56
LAGRASP(0.25,5,5)	0.25%	2589.79
LAGRASP(0.25,5,1)	0.26%	1101.64
LAGRASP(0.25,50,5)	0.47%	292.95
LAGRASP(0,50,-)	0.51%	124.26

These values correspond to the coordinates of the selected variants in Fig. 6.7. The total time is given in seconds

obtained by the eight selected variants. It shows that LAGRASP(1,1,50) found the best solutions, with an average deviation from the best values amounting to 0.079%. It also found the best known solutions in 365 runs (each variant was run eight times on each instance), again with the best performance when the eight variants are evaluated side by side, although its running times are the largest. On the other hand, the smallest running times were observed for LAGRASP(0,50,-), which was over 3000 times faster than LAGRASP(1,1,50) but found the worst-quality solutions among the eight variants considered.

Table 6.4 Summary of the numerical results obtained with the selected variants of the GRASP Lagrangean heuristic on the full set of 135 instances of the set k -covering problem

Heuristic	Average deviation	Hits	Total time (s)
LAGRASP(1,1,50)	0.079%	365	1,803,283.64
LAGRASP(0.50,1,1)	0.134%	242	30,489.17
LAGRASP(0,1,-)	0.135%	238	24,274.72
LAGRASP(0.25,5,10)	0.235%	168	22,475.54
LAGRASP(0.25,5,5)	0.247%	163	11,263.80
LAGRASP(0.25,5,1)	0.249%	164	5347.78
LAGRASP(0.25,50,5)	0.442%	100	1553.35
LAGRASP(0,50,-)	0.439%	97	569.30

The total time is given in seconds

Figure 6.8 illustrates the merit of the proposed approach for one of the test instances. We first observe that all variants reach the same lower bounds, as expected, since they depend exclusively on the common subgradient algorithm. However, as the lower bound appears to stabilize, the upper bound obtained by GLH (LAGRASP(0,1,-)) also seems to freeze. On the other hand, the other variants continue to make improvements by discovering better upper bounds, since the randomized GRASP construction helps them to escape from locally optimal solutions and find new, improved upper bounds.




Fig. 6.8 Evolution of lower and upper bounds over the iterations for different variants of LAGRASP. The number of iterations taken by each LAGRASP variant depends on the step-size, which in turn depends on the upper bounds produced by each heuristic

Finally, we provide a comparison between GRASP with backward path-relinking and the LAGRASP variants on all 135 test instances when the same time limits are used to stop all heuristics. Eight runs were performed for each heuristic and each instance, using different initial seeds for the random number generator. Each heuristic was run a total of $(8 \times 135 =)$ 1080 times. The results in Table 6.5 show that all variants of LAGRASP outperformed GRASP with backward path-relinking and were able to find solutions whose costs are very close to or as good as the best known solution values, while GRASP with backward path-relinking found solutions whose costs are on average 4.05% larger than the best known solution values.

Table 6.5 Summary of results for the best variants of LAGRASP and GRASP

Heuristic	Average deviation	Hits
LAGRASP(1,1,50)	3.30%	0
LAGRASP(0.50,1,1)	0.35%	171
LAGRASP(0,1,-)	0.35%	173
LAGRASP(0.25,5,10)	0.45%	138
LAGRASP(0.25,5,5)	0.45%	143
LAGRASP(0.25,5,1)	0.46%	137
LAGRASP(0.25,50,5)	0.65%	97
LAGRASP(0,50,-)	0.65%	93
GRASP with backward path-relinking	4.05%	0

Figure 6.9 displays for one test instance the typical behavior of these heuristics. As opposed to the GRASP with path-relinking, the Lagrangean heuristics are able to escape from local optima for longer and keep on improving the solutions to obtain the best results.

We note that an important feature of Lagrangean heuristics is that they provide not only a feasible solution (which gives an upper bound, in the case of a minimization problem), but also a lower bound that may be used to give an estimate of the optimality gap that may be considered as a stopping criterion.

6.4 Path-Relinking

The LAGRASP heuristics presented in Sect. 6.3.7.3 made use of path-relinking. Path-relinking is another enhancement to the basic GRASP procedure, leading to significant improvements in both solution quality and running times. This technique was originally proposed by Glover [112] as an intensification strategy to explore trajectories connecting elite solutions obtained by tabu search or scatter search [113, 115, 116].

We consider the undirected graph associated with the solution space $G = (S, M)$, where the nodes in S correspond to feasible solutions and the edges in M correspond to moves in the neighborhood structure, i.e. $(i, j) \in M$ if and only if $i \in S$, $j \in S$, $j \in N(i)$ and $i \in N(j)$, where $N(s)$ denotes the neighborhood of a node $s \in S$. Path-relinking is usually carried out between two solutions: one is called the *initial*




Fig. 6.9 Evolution of solution costs with time for the best variants of LAGRASP and GRASP with backward path-relinking (GPRb)

solution, while the other is the *guiding solution*. One or more paths in the solution space graph connecting these solutions are explored in the search for better solutions. Local search is applied to the best solution in each of these paths, since there is no guarantee that the best solution is locally optimal.

Let $s \in S$ be a node on the path between an initial solution and a guiding solution $g \in S$. Not all solutions in the neighborhood $N(s)$ are candidates to follow s on the path from s to g . We restrict the choice only to those solutions that are more similar to g than s . This is accomplished by selecting moves from s that introduce attributes contained in the guiding solution g . Therefore, path-relinking may be viewed as a strategy that seeks to incorporate attributes of high quality solutions (i.e. the guiding elite solutions), by favoring these attributes in the selected moves.

The use of path-relinking within a GRASP procedure, as an intensification strategy applied to each locally optimal solution, was first proposed by Laguna and Martí [147]. It was followed by several extensions, improvements, and successful applications [8, 9, 22, 54, 104, 183, 206, 212, 216, 217, 227, 233, 249]. A survey of GRASP with path-relinking can be found in [213].

Enhancing GRASP with path-relinking almost always improves the performance of the heuristic. As an illustration, Fig. 6.10 shows time-to-target plots [7, 10, 228, 235, 236] for GRASP and GRASP with path-relinking implementations for four different applications. These time-to-target plots show the empirical cumulative probability distributions of the *time-to-target* random variable when using pure GRASP and GRASP with path-relinking, i.e., the time needed to find a solution at least as good as a prespecified target value. For all problems, the plots show that GRASP with path-relinking is able to find target solutions faster than GRASP.

GRASP with path-relinking makes use of an *elite set* to collect a diverse pool of high-quality solutions found during the search. This pool is limited in size, i.e. it can have at most `Max_Elite` solutions. Several schemes have been proposed for the implementation of path-relinking, which may be applied as:

- an intensification strategy, between each local optimum obtained after the local search phase and one or more elite solutions;
- a post-optimization step, between every pair of elite solutions;
- an intensification strategy, periodically (after a fixed number of GRASP iterations since the last intensification phase) submitting the pool of elite solutions to an evolutionary process (see Sect. 6.4.7);
- a post-optimization phase, submitting the pool of elite solutions to an evolutionary process; or
- any other combination of the above schemes.

The pool of elite solutions is initially empty. Each locally optimal solution obtained by local search and each solution resulting from path-relinking is considered as a candidate to be inserted into the pool. If the pool is not yet full, the candidate is simply added to the pool. Otherwise, if the candidate is better than the incumbent (best solution found so far), it replaces an element of the pool. In case the candidate is better than the worst element of the pool but not better than the best element, then it replaces some element of the pool if it is sufficiently different from every other solution currently in the pool. To balance the impact on pool quality and diversity, the element selected to be replaced is the one that is most similar to the entering solution among those elite solutions of quality no better than the entering solution [216].




Fig. 6.10 Time to target plots comparing running times of pure GRASP and GRASP with path-relinking on four instances of distinct problem types: three index assignment, maximum satisfiability, bandwidth packing, and quadratic assignment

Given a local optimum s_1 produced at the end of a GRASP iteration, we need to select at random a solution s_2 from the pool to apply path-relinking between s_1 and s_2 . In principle, any pool solution could be selected. However, we may want to avoid pool solutions that are too similar to s_1 , because relinking two solutions that are similar limits the scope of the path-relinking search. If the solutions are represented by 0–1 indicator vectors, we should favor pairs of solutions that are far from each other, based on their Hamming distance (i.e., the number of components that take on different values in each solution). A strategy introduced in Resende and Werneck [216] is to select a pool element s_2 at random with a probability proportional to the Hamming distance between the pool element and the local optimum s_1 . Since the number of paths between two solutions grows exponentially with their Hamming distance, this strategy favors pool elements with a large number of paths connecting them to and from s_1 .

After determining which solution (s_1 or s_2) will be designated the initial solution i and which will be the guiding solution g , the algorithm starts by computing the set $\Delta(i, g)$ of components in which i and g differ. This set corresponds to the moves which should be applied to i to reach g . Starting from the initial solution, the best move in $\Delta(i, g)$ still not performed is applied to the current solution, until the guiding solution is reached. By best move, we mean the one that results in the highest quality solution in the restricted neighborhood. The best solution found along this trajectory is submitted to local search and returned as the solution produced by the path-relinking algorithm.

```

procedure GRASP+PR(Seed);
1   Set pool of elite solutions  $\mathcal{E} \leftarrow \emptyset$ ;
2   Set best solution value  $f^* \leftarrow \infty$ ;
3   while stopping criterion not satisfied do
4     Solution  $\leftarrow$  Greedy_Randomized_Construction(Seed);
5     if Solution is not feasible then
6       Solution  $\leftarrow$  Repair(Solution);
7     end-if;
8     Solution  $\leftarrow$  Local_Search(Solution);
9     if  $|\mathcal{E}| > 0$  then
10      Select an elite solution Solution' at random from  $\mathcal{E}$ ;
11      Solution  $\leftarrow$  PR(Solution, Solution');
12    end-if;
13    if  $f(\text{Solution}) < f^*$  then
14      Best_Solution  $\leftarrow$  Solution;
15       $f^* \leftarrow f(S)$ ;
16    end-if;
17    Update the pool of elite solutions  $\mathcal{E}$  with Solution;
18  end-while;
19  return Best_Solution;
end GRASP+PR.

```

Fig. 6.11 Pseudo-code of a template of a GRASP with path-relinking for a minimization problem

The pseudo-code shown in Fig. 6.11 summarizes the steps of a GRASP with path-relinking for a minimization problem. The pseudo-code follows the structure of the basic GRASP algorithm in Fig. 6.1. Lines 1 and 2 initialize the pool of elite solutions and the best solution value, respectively. Path-relinking is performed in line 11 between the solution `Solution` obtained at the end of the local search phase (line 8) and a solution `Solution'` randomly selected from the pool of elite solutions \mathcal{E} (line 10). Procedure `PR(Solution, Solution')` could make use, for example, of any variant of a pure or combined path-relinking strategy. The best overall solution found `Best_Solution` is returned in line 19 after the stopping criterion is satisfied.

Several alternatives have been considered and combined in recent implementations of path-relinking. These include forward, backward, back and forward, mixed, truncated, greedy randomized adaptive, evolutionary, and external path-relinking. All these alternatives, which are described in the following, involve trade-offs between computation time and solution quality.

6.4.1 *Forward Path-Relinking*

In *forward* path-relinking, the GRASP local optimum is designated as the initial solution and the pool solution is made the guiding solution. This is the original scheme proposed by Laguna and Martí [147].

6.4.2 *Backward Path-Relinking*

In *backward* path-relinking, the pool solution is designated as the initial solution and the GRASP local optimum is made the guiding one. This scheme was originally proposed in Aiex et al. [9] and Resende and Ribeiro [212]. The main advantage of this approach over forward path-relinking comes from the fact that, in general, there are more high-quality solutions near pool elements than near GRASP local optima. Backward path-relinking explores more thoroughly the neighborhood around the pool solution, whereas forward path-relinking explores more the neighborhood around the GRASP local optimum. Experiments in [9, 212] have shown that backward path-relinking usually outperforms forward path-relinking.

6.4.3 *Back and Forward Path-Relinking*

Back and forward path-relinking combines forward and backward path-relinking. As shown in [9, 212], it finds solutions at least as good as forward path-relinking or backward path-relinking, but at the expense of taking about twice as long to

run. The reason that back and forward path-relinking often finds solutions of better quality than simple backward or forward path-relinking stems from the fact that it thoroughly explores the neighborhoods of both solutions s_1 and s_2 .

6.4.4 Mixed Path-Relinking

Mixed path-relinking shares the benefits of back and forward path-relinking, i.e. it thoroughly explores both neighborhoods, but does so in about the same time as forward or backward path-relinking alone. This is achieved by interchanging the roles of the initial and guiding solutions at each step of the path-relinking procedure. Therefore, two paths are generated, one starting at s_1 and the other at s_2 . The paths evolve and eventually meet at some solution about half way between s_1 and s_2 . The joined path relinks these two solutions. Mixed path-relinking was suggested by Glover [112] and was first implemented and tested by Ribeiro and Rosseti [227], where it was shown to outperform forward, backward, and back and forward path-relinking. Figure 6.12 shows a comparison of pure GRASP and four variants of path-relinking: forward, backward, back and forward, and mixed. The time-to-target plots show that GRASP with mixed path-relinking has the best running time profile among the variants compared.

6.4.5 Truncated Path-Relinking

Since good-quality solutions tend to be near other good-quality solutions, one would expect to find the best solutions with path-relinking near the initial or guiding solution. Indeed, Resende et al. [222] showed that this is the case for instances of the max-min diversity problem, as shown in Fig. 6.13. In that experiment, a back and forward path-relinking scheme was tested. The figure shows the average number of best solutions found by path-relinking taken over several instances and several applications of path-relinking. The 0–10% range in this figure corresponds to subpaths near the initial solutions for the forward path-relinking phase as well as the backward phase, while the 90–100% range are subpaths near the guiding solutions. As the figure indicates, exploring the subpaths near the extremities may produce solutions about as good as those found by exploring the entire path. There is a higher concentration of better solutions close to the initial solutions explored by path-relinking.

Truncated path-relinking can be applied to either forward, backward, backward and forward, or mixed path-relinking. Instead of exploring the entire path, truncated path-relinking only explores a fraction of the path and, consequently, takes a fraction of the time to run. Truncated path-relinking has been applied in [22, 222].

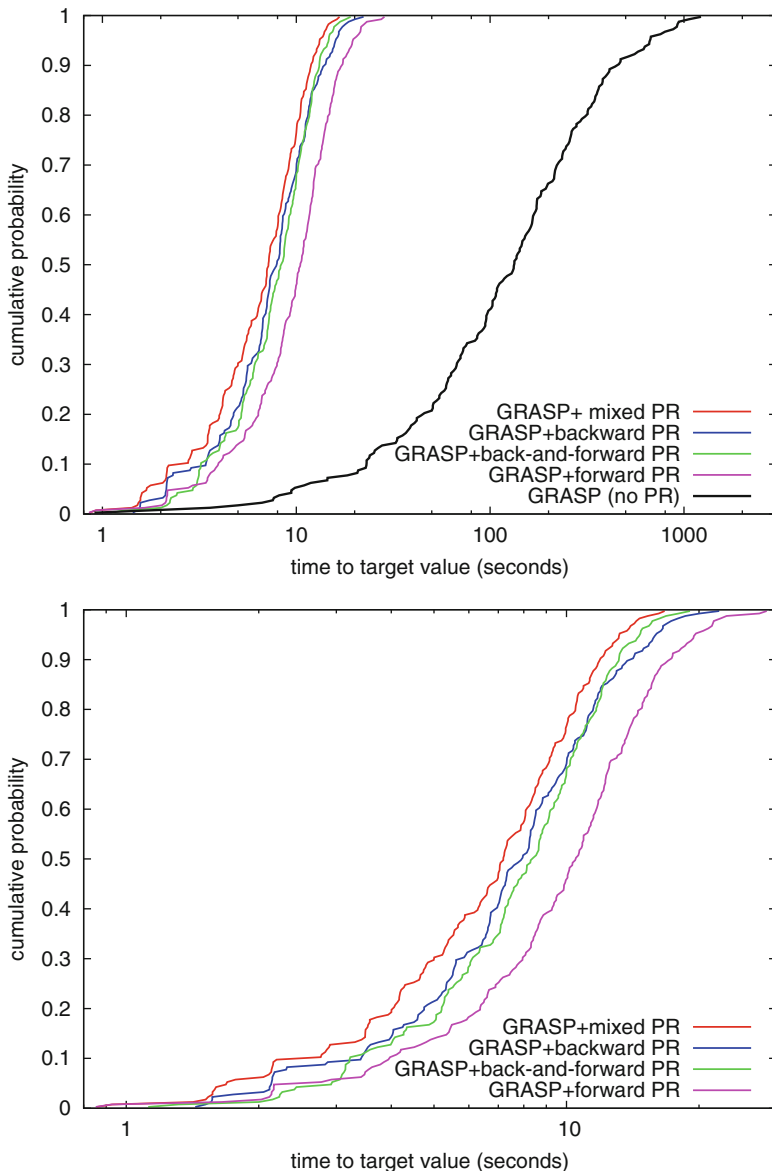


Fig. 6.12 Time-to-target plots for pure GRASP and four variants of GRASP with path-relinking (forward, backward, back and forward, and mixed) on an instance of the 2-path network design problem

6.4.6 Greedy Randomized Adaptive Path-Relinking

In path-relinking, the best not yet performed move in set $\Delta(i, g)$ is applied to the current solution, until the guiding solution is reached. If ties are broken deterministically, this strategy will always produce the same path between the initial and guiding solutions. Since the number of paths connecting i and g is exponential in $|\Delta(i, g)|$, exploring a single path can be somewhat limiting.

Greedy randomized adaptive path-relinking, introduced by Binato et al. [47], is a semi-greedy version of path-relinking. Instead of taking the best move in $\Delta(i, g)$ still not performed, a restricted candidate list of good moves still not performed is set up and a randomly selected move from the latter is applied. By applying this strategy several times between the initial and guiding solutions, several paths can be explored. Greedy randomized adaptive path-relinking has been applied in [22, 86, 222].




Fig. 6.13 Average number of best solutions found at different depths of the path from the initial solution to the guiding solution on instances of the max-min diversity problem

6.4.7 Evolutionary Path-Relinking

GRASP with path-relinking maintains a pool of elite solutions. Applying path-relinking between pairs of pool solutions may result in an even better pool of solutions. Aiex et al. [9] applied path-relinking between all pairs of elite solutions as an intensification scheme to improve the quality of the pool and as a post-optimization step. The application of path-relinking was repeated until no further improvement was possible.

Resende and Werneck [216, 217] described an *evolutionary* path-relinking scheme applied to pairs of elite solutions and used as a post-optimization step. The pool resulting from the GRASP with path-relinking iterations is referred to as population P_0 . At step k , all pairs of elite set solutions of population P_k are relinked and the resulting solutions are made candidates for inclusion in population P_{k+1} of the next generation. The same rules for acceptance into the pool during GRASP with path-relinking are used for acceptance into P_{k+1} . If the best solution in P_{k+1} is better than the best in P_k , then k is incremented by one and the process is repeated. Resende et al. [222] describe another way to implement evolutionary path-relinking, where a single population is maintained. Each pair of elite solutions is relinked and the resulting solution is a candidate to enter the elite set. If accepted, it replaces an existing elite solution. The process is continued while there are still pairs of elite solutions that have not yet been relinked.




Fig. 6.14 An internal path (red arcs, red nodes) from solution S to solution T and two external (blue arcs, blue nodes) paths, one emanating from solution S and the other from solution T . These paths are produced by internal and external path-relinking

Andrade and Resende [21] used this evolutionary scheme as an intensification process every 100 GRASP iterations. During the intensification phase, every solution in the pool is relinked with the two best ones. Since two elite solutions may be relinked more than once in different calls to the intensification process, greedy randomized adaptive path-relinking was used.

Resende et al. [222] showed that a variant of GRASP with evolutionary path-relinking outperformed several other heuristics using GRASP with path-relinking, simulated annealing, and tabu search for the max-min diversity problem.

6.4.8 External Path-Relinking and Diversification

So far in this section, we have considered variants of path-relinking in which a path in the search space graph connects two feasible solutions by progressively introducing in one of them (the initial solution) attributes of the other (the guiding solution). Since attributes common to both solutions are not changed and all solutions visited belong to a path between the two solutions, we may also refer to this type of path-relinking as *internal path-relinking*.

External path-relinking extends any path connecting two feasible solutions S and T beyond its extremities. To extend such a path beyond S , attributes not present in either S or T are introduced in S . Symmetrically, to extend it beyond T , attributes not present in either S or T are introduced in T . In its greedy variant, all moves are evaluated and the solution chosen to be next in the path is one with best cost or, in case they are all infeasible, the one with least infeasibility. In either direction, the procedure stops when all attributes that do not appear in either S or T have been tested for extending the path. Once both paths are complete, local search may be applied to the best solution in each of them. The best of the two local minima is returned as the solution produced by the external path-relinking procedure.

Figure 6.14 illustrates internal and external path-relinking. The path with red nodes and edges is the one resulting from internal path-relinking applied with S as the initial solution and T as the guiding solution. We observe that the orientation introduced by the arcs in this path is due only to the choice of the initial and guiding solutions. If the roles of solutions S and T were interchanged, it could have been computed and generated in the reverse direction. The same figure also illustrates two paths obtained by external path-relinking, one emanating from S and the other from T , both represented with blue nodes and edges. The orientations of the arcs in each of these paths indicate that they necessarily emanate from either solution S or T .

To conclude, we establish a parallel between internal and external path-relinking. Since internal path-relinking works by fixing all attributes common to the initial and guiding solutions and searches for paths between them satisfying this property, it is clearly an intensification strategy. Contrarily, external path-relinking progressively removes common attributes and replaces them by others that do not appear in either one of the initial or guiding solution. Therefore, it can be seen as a diversification strategy which produces solutions increasingly farther from both the initial and the guiding solutions. External path-relinking becomes therefore a tool for search diversification.

External path-relinking was introduced by Glover [114] and first applied by Duarte et al. [84] in a heuristic for differential dispersion minimization.

6.5 Restart Strategies

Figure 6.15 shows a typical iteration count distribution for a GRASP with path-relinking. Observe in this example that for most of the independent runs whose iteration counts make up the plot, the algorithm finds a target solution in relatively few iterations: about 25% of the runs take at most 101 iterations; about 50% take at most 192 iterations; and about 75% take at most 345. However, some runs take much longer: 10% take over 1000 iterations; 5% over 2000; and 2% over 9715 iterations. The longest run took 11,607 iterations to find a solution at least as good as the target. These long tails contribute to a large average iteration count as well as to a high standard deviation. This section proposes strategies to reduce the tail of the distribution, consequently reducing the average iteration count and its standard deviation.




Fig. 6.15 Typical iteration count distribution of GRASP with path-relinking

Consider again the distribution in Fig. 6.15. The distribution shows that each run will take over 345 iterations with a probability of about 25%. Therefore, any time the algorithm is restarted, the probability that the new run will take over 345 iterations is also about 25%. By restarting the algorithm after 345 iterations, the new run will take more than 345 iterations with probability of also about 25%. Therefore, the probability that the algorithm will be still running after $345 + 345 = 690$ iterations is the probability that it takes more than 345 iterations multiplied by the probability that it takes more than 690 iterations given that it took more than 345 iterations, i.e., about $(1/4) \times (1/4) = (1/4)^2$. It follows by induction that the probability that

the algorithm will still be running after k periods of 345 iterations is $1/(4^k)$. In this example, the probability that the algorithm will be running after 1725 iterations will be about 0.1%, i.e., much less than the 5% probability that the algorithm will take over 2000 iterations without restart.

A *restart strategy* is defined as an infinite sequence of time intervals $\tau_1, \tau_2, \tau_3, \dots$ which define epochs $\tau_1, \tau_1 + \tau_2, \tau_1 + \tau_2 + \tau_3, \dots$ when the algorithm is restarted from scratch. It can be shown that the optimal restart strategy uses $\tau_1 = \tau_2 = \dots = \tau^*$, where τ^* is some (unknown) constant. Strategies for speeding up stochastic local search algorithms using restarts were first proposed by Luby et al. [156], where they proved the existence of an optimal restart strategy. Restart strategies in metaheuristics have been addressed in [67, 139, 182, 187, 250]. Further work on restart strategies can be found in [251, 252].

Implementing the optimal strategy may be difficult in practice because it requires the constant value τ^* . Runtimes can vary greatly for different combinations of algorithm, instance, and solution quality sought. Since usually one has no prior information about the runtime distribution of the stochastic search algorithm for the optimization problem under consideration, one runs the risk of choosing a value of τ^* that is either too small or too large. On the one hand, a value that is too small can cause the restart variant of the algorithm to take much longer to converge than a no-restart variant. On the other hand, a value that is too large may never lead to a restart, causing the restart-variant of the algorithm to take as long to converge as the no-restart variant. Figure 6.16 illustrates the restart strategies with time-to-target




Fig. 6.16 Time-to-target plot for target solution value of 554 for a GRASP with path-linking with restart on the maximum cut instance *G12* using different values of τ

plots for the maximum cut instance *G12* [126] on an 800-node graph with edge density of 0.63% with target solution value 554 for $\tau = 6, 9, 12, 18, 24, 30,$ and 42 s. For each value of τ , 100 independent runs of a GRASP with path-relinking with restarts were performed. The variant with $\tau = \infty$ corresponds to the heuristic without restart. The figure shows that, for some values of τ , the resulting heuristic outperformed its counterpart with no restart by a large margin.

In GRASP with path-relinking, the number of iterations between improvements of the incumbent solution tends to vary less than the runtimes for different combinations of instance and solution quality sought. If one takes this into account, a simple and effective restart strategy for GRASP with path-relinking is to keep track of the last iteration when the incumbent solution was improved and restart the GRASP with path-relinking if κ iterations have gone by without improvement. We shall call such a strategy $\text{restart}(\kappa)$. A restart consists in saving the incumbent and emptying out the elite set.

```

procedure GRASP+PR+Restarts(Seed);
1   Set pool of elite solutions  $\mathcal{E} \leftarrow \emptyset$ ;
2   Set best solution value  $f^* \leftarrow \infty$ ;
3   LastImprov  $\leftarrow 0$ ;
4   CurrentIter  $\leftarrow 0$ ;
5   while stopping criterion not satisfied do
6     CurrentIter  $\leftarrow$  CurrentIter + 1;
7     Solution  $\leftarrow$  Greedy-Randomized.Construction(Seed);
8     if Solution is not feasible then
9       Solution  $\leftarrow$  Repair(Solution);
10    end-if;
11    Solution  $\leftarrow$  Local-Search(Solution);
12    if  $|\mathcal{E}| > 0$  then
13      Select an elite solution Solution' at random from  $\mathcal{E}$ ;
14      Solution  $\leftarrow$  forward-PR(Solution, Solution');
15    end-if;
16    if  $f(\text{Solution}) < f^*$  then
17      Best_Solution  $\leftarrow$  Solution;
18       $f^* \leftarrow f(S)$ ;
19      LastImprov  $\leftarrow$  CurrentIter;
20    end-if;
21    if CurrentIter - LastImprov  $> \kappa$  then
22       $\mathcal{E} \leftarrow \emptyset$ ;
23      LastImprov  $\leftarrow$  CurrentIter;
24    else
25      Update the pool of elite solutions  $\mathcal{E}$  with Solution;
26    end-if;
27  end-while;
28  return Best_Solution;
end GRASP+PR+Restarts.

```

Fig. 6.17 Pseudo-code of a template of a GRASP with path-relinking with restarts for a minimization problem

The pseudo-code shown in Fig. 6.17 summarizes the steps of a GRASP with path-relinking using the restart(κ) strategy for a minimization problem. The algorithm keeps track of the current iteration (`CurrentIter`), as well as of the last iteration when an improving solution was found (`LastImprov`). If an improving solution is detected in line 16, then this solution and its cost are saved in lines 17 and 18, respectively, and the iteration of the last improvement is set to the current iteration in line 19. If, in line 21, it is determined that more than κ iterations have gone by since the last improvement of the incumbent, then a restart is triggered, emptying out the elite set in line 22 and resetting the iteration of the last improvement to the current iteration in line 23. If restart is not triggered, then in line 25 the current solution is tested for inclusion in the elite set and the set is updated if it is accepted. The best overall solution found `Best_Solution` is returned in line 28 after the stopping criterion is satisfied.

As an illustration of the use of the restart(κ) strategy within a GRASP with path-relinking, consider the maximum cut instance *G12*. For the values $\kappa = 50, 100, 200, 300, 500, 1000, 2000,$ and 5000 , the heuristic was run independently 100 times, and was stopped when a cut of weight 554 or higher was found. A strategy without restarts was also implemented. Figures 6.18 and 6.19, as well as Table 6.6, summarize these runs, showing the average time to target solution as a function of the value of κ and the time-to-target plots for different values of κ . These figures illustrate well the effect on running time of selecting a value of κ that is either too small ($\kappa = 50, 100$) or too large ($\kappa = 2000, 5000$). They further show that there is a wide range of κ values ($\kappa = 200, 300, 500, 1000$) that result in lower runtimes when compared to the strategy without restarts.

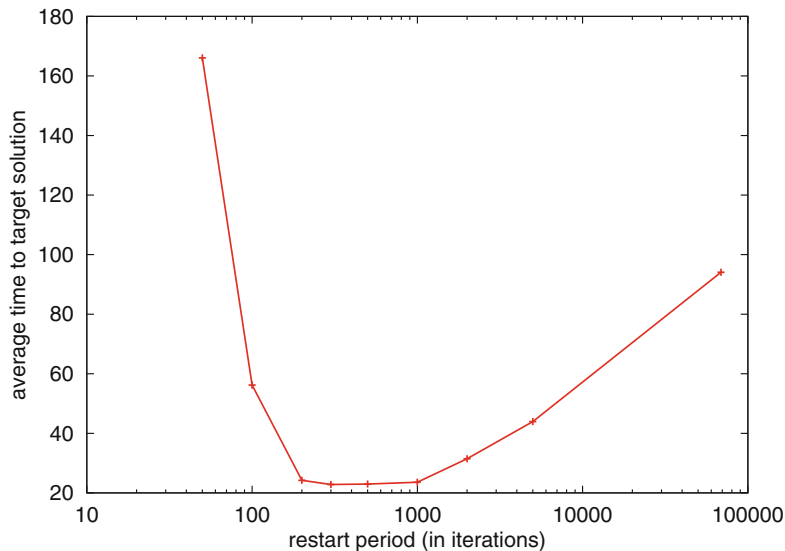


Fig. 6.18 Average time to target solution for maximum cut instance *G12* using different values of κ . All runs of all strategies have found a solution at least as good as the target value of 554

Figure 6.20 further illustrates the behavior of the restart(100), restart(500), and restart(1000) strategies for the previous example, when compared with the strategy without restarts on the same maximum cut instance *G12*. However, in this figure, for each strategy, we plot the number of iterations to the target solution value. It is interesting to note that, as expected, each strategy restart(κ) behaves exactly like the strategy without restarts for the κ first iterations, for $\kappa = 100, 500, 1000$. After this point, each trajectory deviates from that of the strategy without restarts. Among these strategies, restart(500) is the one with the best performance.




Fig. 6.19 Time-to-target plots for maximum cut instance *G12* using different values of κ . The figure also shows the time-to-target plot for the strategy without restarts. All runs of all strategies found a solution at least as good as the target value of 554

We make some final observations about these experiments. The effect of the restart strategies can be mainly observed in the column corresponding to the fourth quartile of Table 6.6. Entries in this quartile correspond to those in the heavy tails of the distributions. The restart strategies in general did not affect the other quartiles of the distributions, which is a desirable characteristic. Compared to the no-restart strategy, restart strategies restart(500) and restart(1000) were able to reduce the maximum number of iterations, as well as the average and the standard deviation. Strategy restart(100) did so, too, but not as much as restart(500) and restart(1000). Restart strategies restart(500) and restart(1000) were clearly the best strategies of those tested.

Table 6.6 Summary of computational results on maximum cut instance *G12* with four strategies

Strategy	Iterations in quartile				Average	st.dev.
	1st	2nd	3rd	4th		
No restarts	326	550	1596	68,813	4525.1	11,927.0
restart(1000)	326	550	1423	5014	953.2	942.1
restart(500)	326	550	1152	4178	835.0	746.1
restart(100)	509	1243	3247	8382	2055.0	2005.9

For each strategy, 100 independent runs were executed, each stopped when a solution as good as the target solution value 554 was found. For each strategy, the table shows the distribution of the number of iterations by quartile. For each quartile, the table gives the maximum number of iterations taken by all runs in that quartile, i.e., the slowest of the fastest 25% (1st), 50% (2nd), 75% (3rd), and 100% (4th) of the runs. The average number of iterations over the 100 runs and the standard deviation (st.dev.) are also given for each strategy

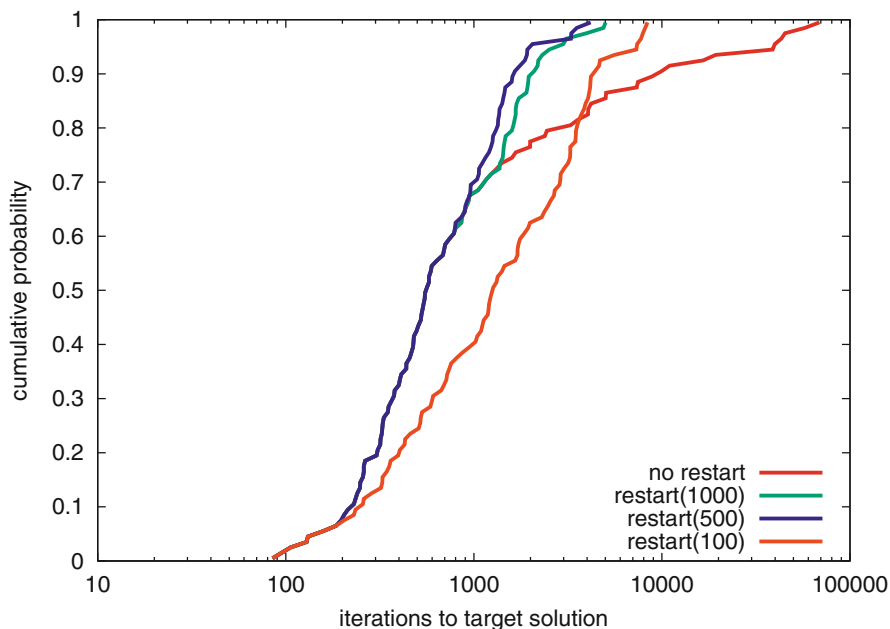


Fig. 6.20 Comparison of the iterations-to-target plots for maximum cut instance *G12* using strategies restart(100), restart(500), and restart(1000). The figure also shows the iterations-to-target plot for the strategy without restarts. All runs of all strategies found a solution at least as good as the target value of 554

The restart(κ) strategy for GRASP with path-relinking discussed in this section was originally proposed by Resende and Ribeiro [214]. Besides the experiments presented in this chapter for the maximum cut instance *G12*, that paper also considered five other instances of maximum cut, maximum weighted satisfiability, and bandwidth packing. Interian and Ribeiro [136] implemented restart strategies for GRASP with path-relinking for the Steiner traveling salesman problem.

6.6 Extensions

In this section, we comment on some extensions, implementation strategies, and hybridizations of GRASP.

The use of hash tables to avoid cycling in conjunction with tabu search was proposed by Woodruff and Zemel [266]. A similar approach was later explored by Ribeiro et al. [232] in their tabu search algorithm for query optimization in relational databases. In the context of GRASP implementations, hash tables were first used by Martins et al. [168] in their multi-neighborhood heuristic for the Steiner problem in graphs, to avoid the application of local search to solutions already visited in previous iterations.

Filtering strategies are used to speed up the iterations of GRASP, see e.g. [93, 168, 202]. With filtering, local search is not applied to all solutions obtained at the end of the construction phase, but only to some more promising unvisited solutions, defined by a threshold with respect to the incumbent.

Almost all randomization effort in the basic GRASP algorithm involves the construction phase. Local search stops at the first local optimum. On the other hand, strategies such as VNS (Variable Neighborhood Search), proposed by Hansen and Mladenović [121, 172], rely almost entirely on the randomization of the local search to escape from local optima. With respect to randomization, GRASP and variable neighborhood strategies can be considered complementary and potentially capable of leading to effective hybrid methods. A first attempt in this direction was made by Martins et al. [168] where the construction phase of a hybrid heuristic for the Steiner problem in graphs follows the greedy randomized strategy of GRASP, while the local search phase makes use of two different neighborhood structures, like the VND (variable neighborhood descent) procedure [121, 172]. That heuristic was later improved by Ribeiro et al. [233], where one of the key components of the new algorithm was another strategy for the exploration of different neighborhoods. Ribeiro and Souza [229] also combined GRASP with VND in a hybrid heuristic for the degree-constrained minimum spanning tree problem. Festa et al. [102] studied different variants and combinations of GRASP and VNS for the maximum cut problem, finding and improving the best known solutions for some open instances from the literature.

GRASP has also been used in conjunction with genetic algorithms. The greedy randomized strategy used in the construction phase of a GRASP heuristic is applied to generate the initial population for a genetic algorithm. As an example, consider the genetic algorithm of Ahuja et al. [4] for the quadratic assignment problem. It makes use of the GRASP heuristic proposed by Li et al. [150] to create the initial population of solutions. A similar approach was used by Armony et al. [31], with the initial population made up of both randomly generated solutions and those built by a GRASP heuristic.

The hybridization of GRASP with tabu search was first studied by Laguna and González-Velarde [146]. Delmaire et al. [71] considered two approaches. In the first, GRASP is applied as a powerful diversification strategy in the context of a tabu search procedure. The second approach is an implementation of the Reac-

tive GRASP algorithm presented in Sect. 6.3.2, in which the local search phase is strengthened by tabu search. Results reported for the capacitated location problem show that the hybrid approaches perform better than the isolated methods previously used. Two two-stage heuristics are proposed in [1] for solving the multi-floor facility layout problem. GRASP/TS applies a GRASP to find the initial layout and tabu search to refine it.

Iterated Local Search (ILS) iteratively builds a sequence of solutions generated by the repeated application of local search and perturbation of the local optima found by local search [42]. Lourenço et al. [155] point out that ILS has been re-discovered many times and is also known as iterated descent [40, 41], large step Markov chains [165], iterated Lin-Kernighan [137], and chained local optimization [164]. A GRASP/ILS hybrid can be obtained by replacing the standard local search of GRASP by ILS. The GRASP construction produces a solution which is passed to the ILS procedure. Ribeiro and Urrutia [230] presented a hybrid GRASP with ILS for the mirrored traveling tournament problem, in which perturbations are achieved by randomly generating solutions in the game rotation ejection chain [110, 111] neighborhood.

6.7 Applications

The first application of GRASP was described in the literature in 1989 [90]. In that paper, GRASP was applied to difficult set covering problems. Since then, GRASP has been applied to a wide range of problems. The main applications areas are summarized below with links to specific references:

- Assignment problems [4, 9, 89, 106, 150, 153, 154, 169, 170, 177, 178, 183, 188, 190, 198, 202, 204, 218, 241]
- Biology [23, 64, 68, 76, 97, 108, 231]
- Computer vision [53, 132, 246, 247]
- Covering, packing, and partitioning [13, 15, 16, 28, 29, 72, 77, 90, 109, 119, 192, 195, 196, 223, 239, 244, 245]
- Diversity and dispersion [79, 84, 163, 222]
- Finance [19, 127]
- Graph and map drawing [66, 96, 147, 159, 160, 162, 184, 210, 226]
- Location and layout [1, 60, 66, 71, 120, 133, 141, 171, 181, 185, 253, 255, 260, 261]
- Logic [75, 104, 189, 208, 219, 221]
- Minimum Steiner tree [54, 166–168, 233]
- Optimization in graphs [2, 3, 5, 12, 32, 56, 73, 82, 83, 93, 101, 103, 134, 148, 149, 157, 160, 161, 168, 179, 191, 193, 207, 210, 220, 226, 233, 248, 257]
- Power systems [25, 46, 48, 86, 203, 263]
- Robotics [144, 240]
- Routing [30, 33, 38, 52, 55, 63, 136, 143, 145, 151, 176, 181, 206, 262, 264, 265]
- Software engineering [158]

- Sports [26, 140, 225, 230]
- Telecommunications [2, 17, 18, 20, 22, 31, 62, 107, 141, 153, 174, 175, 194, 197, 199, 202, 207, 209, 212, 234, 258]
- Timetabling, scheduling, and manufacturing [8, 11, 14, 20, 22, 24, 35–37, 39, 49, 50, 59, 61, 65, 69, 70, 74, 78, 85, 87, 88, 92, 94, 95, 138, 142, 146, 152, 173, 180, 186, 205, 230, 237, 238, 242, 243, 267, 268]
- Transportation [30, 34, 87, 89, 256]
- VLSI design [27, 28]

The reader is referred to Festa and Resende [100] and the book by Resende and Ribeiro [215] for extended annotated bibliographies of GRASP applications.

6.8 Concluding Remarks

The results described in this chapter reflect successful applications of GRASP to a large number of classical combinatorial optimization problems, as well as to problems that arise in real-world situations in different areas of business, science, and technology.

We underscore the simplicity of implementation of GRASP, which makes use of simple building blocks (solution construction procedures and local search methods) that are often readily available. Contrary to what occurs with most other metaheuristics, such as tabu search or genetic algorithms, that make use of a large number of parameters in their implementations, the basic variant of GRASP requires the adjustment of a single parameter, i.e. the restricted candidate list (RCL) parameter α .

Recent developments, presented in this chapter, show that different extensions of the basic procedure allow further improvements in the solutions found by GRASP. Among these, we highlight reactive GRASP, which automates the adjustment of the restricted candidate list parameter; variable neighborhoods, which permit accelerated and intensified local search; path-relinking, which beyond allowing the implementation of intensification strategies based on the memory of elite solutions, opens the way for the development of very effective cooperative parallel strategies [6, 8, 9, 227]; and restart strategies to speedup the search.

These and other extensions make up a set of tools that can be added to simpler heuristics to find better-quality solutions. To illustrate the effect of additional extensions on solution quality, Fig. 6.21 shows some results obtained for the prize-collecting Steiner tree problem (PCSTP), as discussed by Canuto et al. in [54]. The figure shows results for 11 different levels of solution accuracy (varying from optimal to 10% from optimal) on 40 PCSTP instances. For each level of solution accuracy, the figure shows the number of instances for which each component found solutions within the accuracy level. The components are the primal-dual constructive algorithm (GW) of Goemans and Williamson [117], GW followed by local search (GW+LS), corresponding to the first GRASP iteration, 500 iterations of GRASP with path-relinking (GRASP+PR), and the complete algorithm, using variable neighborhood search as a post-optimization procedure (GRASP+PR+VNS).

We observe that the number of optimal solutions found goes from six, using only the constructive algorithm, to a total of 36, using the complete algorithm described in [54]. The largest relative deviation with respect to the optimal value decreases from 36.4% in the first case, to only 1.1% for the complete algorithm. It is easy to notice the contribution made by each additional extension.




Fig. 6.21 Performance of GW approximation algorithm, a single GRASP iteration (GW followed by local search), 500 iterations of GRASP with path-relinking, and 500 iterations of GRASP with path-relinking followed by VNS for series C prize-collecting Steiner tree problems

The structure of GRASP makes it very amenable to straightforward, efficient parallel implementations that benefit from the computer architecture. Parallel implementations of GRASP [6, 8, 9, 227] are quite robust and lead to linear speedups both in independent and cooperative strategies. Cooperative strategies are based on the collaboration between processors through path-relinking and a global pool of elite solutions. This allows the use of more processors to find better solutions in less computation time. Many parallel implementations of GRASP have been reported in the literature, see e.g. [166, 168, 178, 188, 189]. In many of these papers, a common observation was made: the speedups in the measured running times were proportional to the number of processors. This observation can be explained if the random variable *time-to-target-solution-value* is exponentially distributed. Aiex et al. [7] developed a graphical methodology based on runtime distributions to empirically show that the running times of GRASP heuristics fit exponential distributions, as summarized below.

Runtime distributions or time-to-target plots display on the ordinate axis the probability that an algorithm will find a solution at least as good as a given target value within a given running time, shown on the abscissa axis. They provide a very useful tool to characterize the running times of stochastic algorithms for combinatorial optimization problems and to compare different algorithms or strategies for solving a given problem. Time-to-target plots were first used by Feo et al. [93] and have been widely used as a tool for algorithm design and comparison. Runtime distributions have also been advocated by Hoos and Stützle [135] as a way to characterize the running times of stochastic local search algorithms for combinatorial optimization. In particular, they have been largely applied to evaluate and compare the efficiency of different strategies of sequential and parallel implementations of GRASP with (and without) path-relinking heuristics. Aiex et al. [7] used time-to-target plots to show experimentally that the running times of GRASP heuristics fit shifted (or two-parameter) exponential distributions, reporting computational results for 2400 runs of GRASP heuristics for each of five different problems: maximum stable set, quadratic assignment, graph planarization [210, 211, 226], maximum weighted satisfiability, and maximum covering. Aiex et al. [10] developed a Perl program to create time-to-target plots for measured times that are assumed to fit a shifted exponential distribution, following closely the work in [7]. Ribeiro et al. [235] developed a closed form result to compare two exponential algorithms and an iterative procedure to compare two algorithms following generic runtime distributions. This work was extended by Ribeiro et al. [236] and was also applied in the comparison of parallel heuristics. Ribeiro and Rosseti [228] developed a code to compare runtime distributions of randomized algorithms.

To conclude, this chapter provides the reader with the tools to build a basic GRASP to find optimal or near-optimal solutions to a combinatorial optimization problem. The chapter also provides the means to add more advanced features to this basic GRASP, like path-relinking and restart strategies, that enable better performance, both with respect to solution quality and solution run time. Left out of this chapter is the use of GRASP for solving continuous optimization problems. The interested reader is pointed to [128–131, 215, 254] for an introduction to C-GRASP, or Continuous GRASP, as well as to some software and applications of C-GRASP.

References

1. S. Abdinnour-Helm, S.W. Hadley, Tabu search based heuristics for multi-floor facility layout. *Int. J. Prod. Res.* **38**, 365–383 (2000)
2. J. Abello, P.M. Pardalos, M.G.C. Resende, On maximum clique problems in very large graphs, in *External Memory Algorithms and Visualization*, ed. by J. Abello, J. Vitter. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 50 (American Mathematical Society, Providence, 1999), pp. 199–130
3. J. Abello, M.G.C. Resende, S. Sudarsky, Massive quasi-clique detection, in *LATIN 2002: Theoretical Informatics*, ed. by S. Rajsbaum. Lecture Notes in Computer Science, vol. 2286 (Springer, Berlin, 2002), pp. 598–612

4. R.K. Ahuja, J.B. Orlin, A. Tiwari, A greedy genetic algorithm for the quadratic assignment problem. *Comput. Oper. Res.* **27**, 917–934 (2000)
5. R.K. Ahuja, J.B. Orlin, D. Sharma, Multi-exchange neighborhood structures for the capacitated minimum spanning tree problem. *Math. Program.* **91**, 71–97 (2001)
6. R.M. Aiex, M.G.C. Resende, Parallel strategies for GRASP with path-relinking, in *Metaheuristics: Progress as Real Problem Solvers*, ed. by T. Ibaraki, K. Nonobe, M. Yagiura (Springer, New York, 2005), pp. 301–331
7. R.M. Aiex, M.G.C. Resende, C.C. Ribeiro, Probability distribution of solution time in GRASP: an experimental investigation. *J. Heuristics* **8**, 343–373 (2002)
8. R.M. Aiex, S. Binato, M.G.C. Resende, Parallel GRASP with path-relinking for job shop scheduling. *Parallel Comput.* **29**, 393–430 (2003)
9. R.M. Aiex, P.M. Pardalos, M.G.C. Resende, G. Toraldo, GRASP with path-relinking for three-index assignment. *INFORMS J. Comput.* **17**, 224–247 (2005)
10. R.M. Aiex, M.G.C. Resende, C.C. Ribeiro, TTTPLOTS: a perl program to create time-to-target plots. *Optim Lett.* **1**, 355–366 (2007)
11. E. Alekseeva, M. Mezma, D. Tuyttens, N. Melab, Parallel multi-core hyper-heuristic GRASP to solve permutation flow-shop problem. *Concurrency Comput. Pract. Exp.* **29**, e3835 (2017)
12. D. Aloise, C.C. Ribeiro, Adaptive memory in multistart heuristics for multicommodity network design. *J. Heuristics* **17**, 153–179 (2011)
13. R. Álvarez-Valdés, F. Parreno, J.M. Tamarit, A GRASP algorithm for constrained two-dimensional non-guillotine cutting problems. *J. Oper. Res. Soc.* **56**, 414–425 (2005)
14. R. Álvarez-Valdés, E. Crespo, J.M. Tamarit, F. Villa, GRASP and path relinking for project scheduling under partially renewable resources. *Eur. J. Oper. Res.* **189**, 1153–1170 (2008)
15. R. Alvarez-Valdesa, F. Parreno, J.M. Tamarit, Reactive GRASP for the strip-packing problem. *Comput. Oper. Res.* **35**, 1065–1083 (2008)
16. R. Alvarez-Valdes, F. Parreño, J.M. Tamarit, A GRASP/path relinking algorithm for two- and three-dimensional multiple bin-size bin packing problems. *Comput. Oper. Res.* **40**, 3081–3090 (2013)
17. E. Amaldi, A. Capone, F. Malucelli, Planning UMTS base station location: optimization models with power control and algorithms. *IEEE Trans. Wirel. Commun.* **2**, 939–952 (2003)
18. E. Amaldi, A. Capone, F. Malucelli, F. Signori, Optimization models and algorithms for downlink UMTS radio planning, in *Proceedings of Wireless Communications and Networking*, vol. 2 (2003), pp. 827–831
19. K.P. Anagnostopoulos, P.D. Chatzoglou, S. Katsavounis, A reactive greedy randomized adaptive search procedure for a mixed integer portfolio optimization problem. *Manag. Financ.* **36**, 1057–1065 (2010)
20. D.V. Andrade, M.G.C. Resende, A GRASP for PBX telephone migration scheduling, in *Proceedings of the Eighth INFORMS Telecommunications Conference* (2006)
21. D.V. Andrade, M.G.C. Resende, GRASP with evolutionary path-relinking. Technical Report TD-6XPTS7, AT&T Labs Research, Florham Park, 2007
22. D.V. Andrade, M.G.C. Resende, GRASP with path-relinking for network migration scheduling, in *Proceedings of the International Network Optimization Conference* (2007)
23. A.A. Andreatta, C.C. Ribeiro, Heuristics for the phylogeny problem. *J. Heuristics* **8**, 429–447 (2002)
24. C. Andrés, C. Miralles, R. Pastor, Balancing and scheduling tasks in assembly lines with sequence-dependent setup times. *Eur. J. Oper. Res.* **187**, 1212–1223 (2008)
25. C.H. Antunes, E. Oliveira, P. Lima, A multi-objective GRASP procedure for reactive power compensation planning. *Optim. Eng.* **15**, 199–215 (2014)
26. A.P.F. Araújo, C. Boeres, V.E.F. Rebello, C.C. Ribeiro, S. Urrutia, Exploring grid implementations of parallel cooperative metaheuristics: a case study for the mirrored traveling tournament problem, in *Metaheuristics: Progress in Complex Systems Optimization*, ed. by K.F. Doerner, M. Gendreau, P. Greistorfer, W. Gutjahr, R.F. Hartl, M. Reimann (Springer, New York, 2007), pp. 297–322

27. S.M. Areibi, GRASP: an effective constructive technique for VLSI circuit partitioning, in *Proceedings of the IEEE Canadian Conference on Electrical and Computer Engineering*, Edmonton, pp. 462–467 (1999)
28. S. Areibi, A. Vannelli, A GRASP clustering technique for circuit partitioning, in *Satisfiability Problems*, ed. by J. Gu, P.M. Pardalos. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 35 (American Mathematical Society, Providence, 1997), pp. 711–724
29. M.F. Argüello, T.A. Feo, O. Goldschmidt, Randomized methods for the number partitioning problem. *Comput. Oper. Res.* **23**, 103–111 (1996)
30. M.F. Argüello, J.F. Bard, G. Yu, A GRASP for aircraft routing in response to groundings and delays. *J. Comb. Optim.* **1**, 211–228 (1997)
31. M. Armony, J.C. Kliniewicz, H. Luss, M.B. Rosenwein, Design of stacked self-healing rings using a genetic algorithm. *J. Heuristics* **6**, 85–105 (2000)
32. J.E.C. Arroyo, P.S. Vieira, D.S. Vianna, A GRASP algorithm for the multi-criteria minimum spanning tree problem. *Ann. Oper. Res.* **159**, 125–133 (2008)
33. J.B. Atkinson, A greedy randomised search heuristic for time-constrained vehicle scheduling and the incorporation of a learning strategy. *J. Oper. Res. Soc.* **49**, 700–708 (1998)
34. J.F. Bard, An analysis of a rail car unloading area for a consumer products manufacturer. *J. Oper. Res. Soc.* **48**, 873–883 (1997)
35. J.F. Bard, T.A. Feo, Operations sequencing in discrete parts manufacturing. *Manage. Sci.* **35**, 249–255 (1989)
36. J.F. Bard, T.A. Feo, An algorithm for the manufacturing equipment selection problem. *IIE Trans.* **23**, 83–92 (1991)
37. J.F. Bard, T.A. Feo, S. Holland, A GRASP for scheduling printed wiring board assembly. *IIE Trans.* **28**, 155–165 (1996)
38. J.F. Bard, L. Huang, P. Jaillet, M. Dror, A decomposition approach to the inventory routing problem with satellite facilities. *Transp. Sci.* **32**, 189–203 (1998)
39. J.F. Bard, Y. Shao, A.I. Jarrah, A sequential GRASP for the therapist routing and scheduling problem. *J. Scheduling* **17**, 109–133 (2014)
40. E.B. Baum, Iterated descent: a better algorithm for local search in combinatorial optimization problems. Technical Report, California Institute of Technology, 1986
41. E.B. Baum, Towards practical ‘neural’ computation for combinatorial optimization problems, in *AIP Conference Proceedings 151 on Neural Networks for Computing* (American Institute of Physics Inc., Woodbury, 1987), pp. 53–58
42. J. Baxter, Local optima avoidance in depot location. *J. Oper. Res. Soc.* **32**, 815–819 (1981)
43. J.E. Beasley, An algorithm for set-covering problems. *Eur. J. Oper. Res.* **31**, 85–93 (1987)
44. J.E. Beasley, A Lagrangian heuristic for set-covering problems. *Nav. Res. Logist.* **37**, 151–164 (1990)
45. J.E. Beasley, Lagrangean relaxation, in *Modern Heuristic Techniques for Combinatorial Problems*, ed. by C.R. Reeves (Blackwell Scientific Publications, Oxford, 1993), pp. 243–303
46. S. Binato, G.C. Oliveira, A reactive GRASP for transmission network expansion planning, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 81–100
47. S. Binato, H. Faria Jr., M.G.C. Resende, Greedy randomized adaptive path relinking, in *Proceedings of the IV Metaheuristics International Conference*, ed. by J.P. Sousa, pp. 393–397 (2001)
48. S. Binato, G.C. Oliveira, J.L. Araújo, A greedy randomized adaptive search procedure for transmission expansion planning. *IEEE Trans. Power Syst.* **16**, 247–253 (2001)
49. S. Binato, W.J. Hery, D. Loewenstern, M.G.C. Resende, A GRASP for job shop scheduling, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 59–79
50. M. Boudia, M.A.O. Louly, C. Prins, A reactive GRASP and path relinking for a combined production-distribution problem. *Comput. Oper. Res.* **34**, 3402–3419 (2007)

51. J.L. Bresina, Heuristic-biased stochastic sampling, in *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, pp. 271–278 (1996)
52. A.M. Campbell, B.W. Thomas, Probabilistic traveling salesman problem with deadlines. *Transp. Sci.* **42**, 1–21 (2008)
53. R.G. Cano, G. Kunifuji, C.C. de Souza, P.J. de Rezende, A hybrid GRASP heuristic to construct effective drawings of proportional symbol maps. *Comput. Oper. Res.* **40**, 1435–1447 (2013)
54. S.A. Canuto, M.G.C. Resende, C.C. Ribeiro, Local search with perturbations for the prize-collecting Steiner tree problem in graphs. *Networks* **38**, 50–58 (2001)
55. C. Carreto, B. Baker, A GRASP interactive approach to the vehicle routing problem with backhauls, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 185–199
56. W.A. Chaovalitwongse, C.A.S. Oliveira, B. Chiarini, P.M. Pardalos, M.G.C. Resende, Revised GRASP with path-relinking for the linear ordering problem. *J. Comb. Optim.* **22**, 572–593 (2011)
57. I. Charon, O. Hudry, The noising method: a new method for combinatorial optimization. *Oper. Res. Lett.* **14**, 133–137 (1993)
58. I. Charon, O. Hudry, The noising methods: a survey, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 245–261
59. M. Chica, O. Cerdón, S. Damas, J. Bautista, A multiobjective GRASP for the 1/3 variant of the time and space assembly line balancing problem, in *Trends in Applied Intelligent Systems*, ed. by N. García-Pedrajas, F. Herrera, C. Fyfe, J. Benítez, M. Ali. *Lecture Notes in Computer Science*, vol. 6098 (Springer, Berlin, 2010), pp. 656–665
60. R. Colomé, D. Serra, Consumer choice in competitive location models: formulations and heuristics. *Pap. Reg. Sci.* **80**, 439–464 (2001)
61. C.W. Commander, S.I. Butenko, P.M. Pardalos, C.A.S. Oliveira, Reactive GRASP with path relinking for the broadcast scheduling problem, in *Proceedings of the 40th Annual International Telemetry Conference*, pp. 792–800 (2004)
62. C. Commander, C.A.S. Oliveira, P.M. Pardalos, M.G.C. Resende, A GRASP heuristic for the cooperative communication problem in ad hoc networks, in *Proceedings of the VI Metaheuristics International Conference*, pp. 225–330 (2005)
63. A. Corberán, R. Martí, J.M. Sanchis, A GRASP heuristic for the mixed Chinese postman problem. *Eur. J. Oper. Res.* **142**, 70–80 (2002)
64. R. Cordone, G. Lulli, A GRASP metaheuristic for microarray data analysis. *Comput. Oper. Res.* **40**, 3108–3120 (2013)
65. J.F. Corcher, M.T. Alonso, F. Parre no, R. Alvarez-Valdes, Solving a large multicontainer loading problem in the car manufacturing industry. *Comput. Oper. Res.* **82**, 139–152 (2017)
66. G.L. Cravo, G.M. Ribeiro, L.A. Nogueira Lorena, A greedy randomized adaptive search procedure for the point-feature cartographic label placement. *Comput. Geosci.* **34**, 373–386 (2008)
67. M.M. D’Apuzzo, A. Migdalas, P.M. Pardalos, G. Toraldo, Parallel computing in global optimization, in *Handbook of Parallel Computing and Statistics*, ed. by E. Kontogiorghes (Chapman & Hall/CRC, Boca Raton, 2006)
68. S. Das, S.M. Idicula, Application of reactive GRASP to the biclustering of gene expression data, in *Proceedings of the International Symposium on Biocomputing (ACM, Calicut, 2010)*, p. 14
69. P. De, J.B. Ghosj, C.E. Wells, Solving a generalized model for con due date assignment and sequencing. *Int. J. Prod. Econ.* **34**, 179–185 (1994)
70. R. De Leone, P. Festa, E. Marchitto, Solving a bus driver scheduling problem with randomized multistart heuristics. *Int. Trans. Oper. Res.* **18**, 707–727 (2011)
71. H. Delmaire, J.A. Díaz, E. Fernández, M. Ortega, Reactive GRASP and Tabu Search based heuristics for the single source capacitated plant location problem. *INFOR* **37**, 194–225 (1999)
72. X. Delorme, X. Gandibleux, F. Degoutin, Evolutionary, constructive and hybrid procedures for the bi-objective set packing problem. *Eur. J. Oper. Res.* **204**, 206–217 (2010)

73. Y. Deng, J.F. Bard, A reactive GRASP with path relinking for capacitated clustering. *J. Heuristics* **17**, 119–152 (2011)
74. Y. Deng, J.F. Bard, G.R. Chacon, J. Stuber, Scheduling back-end operations in semiconductor manufacturing. *IEEE Trans. Semicond. Manuf.* **23**, 210–220 (2010)
75. A.S. Deshpande, E. Triantaphyllou, A greedy randomized adaptive search procedure (GRASP) for inferring logical clauses from examples in polynomial time and some extensions. *Math. Comput. Model.* **27**, 75–99 (1998)
76. S. Dharan, A.S. Nair, Biclustering of gene expression data using reactive greedy randomized adaptive search procedure. *BMC Bioinf.* **10**(Suppl 1), S27 (2009)
77. J.A. Díaz, D.E. Luna, J.-F. Camacho-Vallejo, M.-S. Casas-Ramírez, GRASP and hybrid GRASP-Tabu heuristics to solve a maximal covering location problem with customer preference ordering. *Expert Syst. Appl.* **82**, 67–76 (2017)
78. A. Drexler, F. Salewski, Distribution requirements and compactness constraints in school timetabling. *Eur. J. Oper. Res.* **102**, 193–214 (1997)
79. A. Duarte, R. Martí, Tabu search and GRASP for the maximum diversity problem. *Eur. J. Oper. Res.* **178**, 71–84 (2007)
80. A. Duarte, C.C. Ribeiro, S. Urrutia, A hybrid ILS heuristic to the referee assignment problem with an embedded MIP strategy. *Lect. Notes Comput. Sci.* **4771**, 82–95 (2007)
81. A.R. Duarte, C.C. Ribeiro, S. Urrutia, E.H. Haeusler, Referee assignment in sports leagues. *Lect. Notes Comput. Sci.* **3867**, 158–173 (2007)
82. A. Duarte, R. Martí, M.G.C. Resende, R.M.A. Silva, GRASP with path relinking heuristics for the antibandwidth problem. *Networks* **58**, 171–189 (2011)
83. A. Duarte, R. Martí, A. Álvarez, F. Ángel-Bello, Metaheuristics for the linear ordering problem with cumulative costs. *Eur. J. Oper. Res.* **216**, 270–277 (2012)
84. A. Duarte, J. Sánchez-Oro, M.G.C. Resende, F. Glover, R. Martí, GRASP with exterior path relinking for differential dispersion minimization. *Inform. Sci.* **296**, 46–60 (2015)
85. M. Essafi, X. Delorme, A. Dolgui, Balancing lines with CNC machines: a multi-start and based heuristic. *CIRP J. Manuf. Sci. Technol.* **2**, 176–182 (2010)
86. H. Faria Jr., S. Binato, M.G.C. Resende, D.J. Falcão, Transmission network design by a greedy randomized adaptive path relinking approach. *IEEE Trans. Power Syst.* **20**, 43–49 (2005)
87. T.A. Feo, J.F. Bard, Flight scheduling and maintenance base planning. *Manag. Sci.* **35**, 1415–1432 (1989)
88. T.A. Feo, J.F. Bard, The cutting path and tool selection problem in computer-aided process planning. *J. Manufact. Syst.* **8**, 17–26 (1989)
89. T.A. Feo, J.L. González-Velarde, The intermodal trailer assignment problem: Models, algorithms, and heuristics. *Transp. Sci.* **29**, 330–341 (1995)
90. T.A. Feo, M.G.C. Resende, A probabilistic heuristic for a computationally difficult set covering problem. *Oper. Res. Lett.* **8**, 67–71 (1989)
91. T.A. Feo, M.G.C. Resende, Greedy randomized adaptive search procedures. *J. Glob. Optim.* **6**, 109–133 (1995)
92. T.A. Feo, K. Venkatraman, J.F. Bard, A GRASP for a difficult single machine scheduling problem. *Comput. Oper. Res.* **18**, 635–643 (1991)
93. T.A. Feo, M.G.C. Resende, S.H. Smith, A greedy randomized adaptive search procedure for maximum independent set. *Oper. Res.* **42**, 860–878 (1994)
94. T.A. Feo, J.F. Bard, S. Holland, Facility-wide planning and scheduling of printed wiring board assembly. *Oper. Res.* **43**, 219–230 (1995)
95. T.A. Feo, K. Sarathy, J. McGahan, A GRASP for single machine scheduling with sequence dependent setup costs and linear delay penalties. *Comput. Oper. Res.* **23**, 881–895 (1996)
96. E. Fernández, R. Martí, GRASP for seam drawing in mosaicking of aerial photographic maps. *J. Heuristics* **5**, 181–197 (1999)
97. P. Festa, On some optimization problems in molecular biology. *Math. Biosci.* **207**, 219–234 (2007)

98. P. Festa, M.G.C. Resende, GRASP: An annotated bibliography, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 325–367
99. P. Festa, M.G.C. Resende, An annotated bibliography of GRASP, part I: algorithms. *Int. Trans. Oper. Res.* **16**, 1–24 (2009)
100. P. Festa, M.G.C. Resende, An annotated bibliography of GRASP, part II: applications. *Int. Trans. Oper. Res.* **16**, 131–172 (2009)
101. P. Festa, P.M. Pardalos, M.G.C. Resende, Algorithm 815: FORTRAN subroutines for computing approximate solution to feedback set problems using GRASP. *ACM Trans. Math. Softw.* **27**, 456–464 (2001)
102. P. Festa, M.G.C. Resende, P. Pardalos, C.C. Ribeiro, GRASP and VNS for Max-Cut, in *Extended Abstracts of the Fourth Metaheuristics International Conference*, Porto, pp. 371–376 (2001)
103. P. Festa, P.M. Pardalos, M.G.C. Resende, C.C. Ribeiro, Randomized heuristics for the MAX-CUT problem. *Optim. Methods Softw.* **7**, 1033–1058 (2002)
104. P. Festa, P.M. Pardalos, L.S. Pitsoulis, M.G.C. Resende, GRASP with path-relinking for the weighted MAXSAT problem. *ACM J. Exp. Algorithmics* **11**, 1–16 (2006)
105. M.L. Fisher, The Lagrangian relaxation method for solving integer programming problems. *Manag. Sci.* **50**, 1861–1871 (2004)
106. C. Fleurent, F. Glover, Improved constructive multistart strategies for the quadratic assignment problem using adaptive memory. *INFORMS J. Comput.* **11**, 198–204 (1999)
107. E. Fonseca, R. Fuchsuber, L.F.M. Santos, A. Plastino, S.L. Martins, Exploring the hybrid metaheuristic DM-GRASP for efficient server replication for reliable multicast, in *International Conference on Metaheuristics and Nature Inspired Computing*, Hammamet (2008)
108. R.D. Frinhani, R.M. Silva, G.R. Mateus, P. Festa, M.G.C. Resende, GRASP with path-relinking for data clustering: a case study for biological data, in *Experimental Algorithms*, ed. by P.M. Pardalos, S. Rebennack. *Lecture Notes in Computer Science*, vol. 6630 (Springer, Berlin, 2011), pp. 410–420
109. J.B. Ghosh, Computational aspects of the maximum diversity problem. *Oper. Res. Lett.* **19**, 175–181 (1996)
110. F. Glover, New ejection chain and alternating path methods for traveling salesman problems, in *Computer Science and Operations Research: New Developments in Their Interfaces*, ed. by O. Balci, R. Sharda, S. Zenios (Elsevier, Amsterdam, 1992), pp. 449–509
111. F. Glover, Ejection chains, reference structures and alternating path methods for traveling salesman problems. *Discret. Appl. Math.* **65**, 223–254 (1996)
112. F. Glover, Tabu search and adaptive memory programming – advances, applications and challenges, in *Interfaces in Computer Science and Operations Research*, ed. by R.S. Barr, R.V. Helgason, J.L. Kennington (Kluwer Academic Publishers, Boston, 1996), pp. 1–75
113. F. Glover, Multi-start and strategic oscillation methods – principles to exploit adaptive memory, in *Computing Tools for Modeling, Optimization and Simulation: Interfaces in Computer Science and Operations Research*, ed. by M. Laguna, J.L. González-Velarde (Kluwer Academic Publishers, Boston, 2000), pp. 1–24
114. F. Glover, Exterior path relinking for zero-one optimization. *Int. J. Appl. Metaheuristic Comput.* **5**, 1–8 (2014)
115. F. Glover, M. Laguna, *Tabu Search* (Kluwer Academic Publishers, Boston, 1997)
116. F. Glover, M. Laguna, R. Martí, Fundamentals of scatter search and path relinking. *Control Cybern.* **39**, 653–684 (2000)
117. M.X. Goemans, D.P. Williamson, The primal dual method for approximation algorithms and its application to network design problems, in *Approximation Algorithms for NP-Hard Problems*, ed. by D. Hochbaum (PWS Publishing Co., Boston, 1996), pp. 144–191
118. F.C. Gomes, C.S. Oliveira, P.M. Pardalos, M.G.C. Resende, Reactive GRASP with path-relinking for channel assignment in mobile phone networks, in *Proceedings of the 5th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications* (ACM Press, New York, 2001), pp. 60–67

119. P.L. Hammer, D.J. Rader Jr., Maximally disjoint solutions of the set covering problem. *J. Heuristics* **7**, 131–144 (2001)
120. B.T. Han, V.T. Raja, A GRASP heuristic for solving an extended capacitated concentrator location problem. *Int. J. Inf. Technol. Decis. Mak.* **2**, 597–617 (2003)
121. P. Hansen, N. Mladenović, Developments of variable neighborhood search, in *Essays and Surveys in Metaheuristics*, ed. by C.C. Ribeiro, P. Hansen (Kluwer Academic Publishers, Boston, 2002), pp. 415–439
122. J.P. Hart, A.W. Shogan, Semi-greedy heuristics: an empirical study. *Oper. Res. Lett.* **6**, 107–114 (1987)
123. M. Held, R.M. Karp, The traveling-salesman problem and minimum spanning trees. *Oper. Res.* **18**, 1138–1162 (1970)
124. M. Held, R.M. Karp, The traveling-salesman problem and minimum spanning trees: part II. *Math. Program.* **1**, 6–25 (1971)
125. M. Held, P. Wolfe, H.P. Crowder, Validation of subgradient optimization. *Math. Program.* **6**, 62–88 (1974)
126. C. Helmberg, F. Rendl, A spectral bundle method for semidefinite programming. *SIAM J. Optim.* **10**, 673–696 (2000)
127. A.J. Higgins, S. Hajkowicz, E. Bui, A multi-objective model for environmental investment decision making. *Comput. Oper. Res.* **35**, 253–266 (2008)
128. M.J. Hirsch, GRASP-based heuristics for continuous global optimization problems. Ph.D. thesis, Department of Industrial and Systems Engineering, University of Florida, Gainesville, 2006
129. M.J. Hirsch, C.N. Meneses, P.M. Pardalos, M.G.C. Resende, Global optimization by continuous GRASP. *Optim. Lett.* **1**, 201–212 (2007)
130. M.J. Hirsch, P.M. Pardalos, M.G.C. Resende, Solving systems of nonlinear equations with continuous GRASP. *Nonlinear Anal. Real World Appl.* **10**, 2000–2006 (2009)
131. M.J. Hirsch, P.M. Pardalos, M.G.C. Resende, Speeding up continuous GRASP. *Eur. J. Oper. Res.* **205**, 507–521 (2010)
132. M.J. Hirsch, P.M. Pardalos, M.G.C. Resende, Correspondence of projected 3D points and lines using a continuous GRASP. *Int. Trans. Oper. Res.* **18**, 493–511 (2011)
133. K. Holmqvist, A. Migdalas, P.M. Pardalos, Greedy randomized adaptive search for a location problem with economies of scale, in *Developments in Global Optimization*, ed. by I.M. Bomze et al. (Kluwer Academic Publishers, Dordrecht, 1997), pp. 301–313
134. K. Holmqvist, A. Migdalas, P.M. Pardalos, A GRASP algorithm for the single source uncapacitated minimum concave-cost network flow problem, in *Network Design: Connectivity and Facilities Location*, ed. by P.M. Pardalos, D.-Z. Du. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 40 (American Mathematical Society, Providence, 1998), pp. 131–142
135. H.H. Hoos, T. Stützle, Evaluation of Las Vegas algorithms - Pitfalls and remedies, in *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, ed. by G. Cooper, S. Moral (Morgan Kaufmann, Madison, 1998), pp. 238–245
136. R. Interian, C.C. Ribeiro, A GRASP heuristic using path-relinking and restarts for the Steiner traveling salesman problem. *Int. Trans. Oper. Res.* **24**, 1307–1323 (2017)
137. D.S. Johnson, Local optimization and the traveling salesman problem, in *Proceedings of the 17th Colloquium on Automata*. LNCS, vol. 443 (Springer, Berlin, 1990), pp. 446–461
138. E.H. Kampke, J.E.C. Arroyo, A.G. Santos, Reactive GRASP with path relinking for solving parallel machines scheduling problem with resource-assignable sequence dependent setup times, in *Proceedings of the World Congress on Nature and Biologically Inspired Computing*, Coimbatore (IEEE, New York, 2009), pp. 924–929
139. H. Kautz, E. Horvitz, Y. Ruan, C. Gomes, B. Selman, Dynamic restart policies, in *Proceedings of the Eighteenth National Conference on Artificial Intelligence* (American Association for Artificial Intelligence, Edmonton, 2002), pp. 674–681
140. G. Kendall, S. Knust, C.C. Ribeiro, S. Urrutia, Scheduling in sports: an annotated bibliography. *Comput. Oper. Res.* **37**, 1–19 (2010)

141. J.G. Klincewicz, Avoiding local optima in the p -hub location problem using tabu search and GRASP. *Ann. Oper. Res.* **40**, 283–302 (1992)
142. J.G. Klincewicz, A. Rajan, Using GRASP to solve the component grouping problem. *Nav. Res. Log.* **41**, 893–912 (1994)
143. G. Kontoravdis, J.F. Bard, A GRASP for the vehicle routing problem with time windows. *ORSA J. Comput.* **7**, 10–23 (1995)
144. M. Kulich, J.J. Miranda-Bront, L. Preucil, A meta-heuristic based goal-selection strategy for mobile robot search in an unknown environment. *Comput. Oper. Res.* **84**, 178–187 (2017)
145. N. Labadi, C. Prins, M. Reghioi, GRASP with path relinking for the capacitated arc routing problem with time windows, in *Advances in Computational Intelligence in Transport, Logistics, and Supply Chain Management*, ed. by A. Fink, F. Rothlauf (Springer, Berlin, 2008), pp. 111–135
146. M. Laguna, J.L. González-Velarde, A search heuristic for just-in-time scheduling in parallel machines. *J. Intell. Manuf.* **2**, 253–260 (1991)
147. M. Laguna, R. Martí, GRASP and path relinking for 2-layer straight line crossing minimization. *INFORMS J. Comput.* **11**, 44–52 (1999)
148. M. Laguna, R. Martí, A GRASP for coloring sparse graphs. *Comput. Optim. Appl.* **19**, 165–178 (2001)
149. M. Laguna, T.A. Feo, H.C. Elrod, A greedy randomized adaptive search procedure for the two-partition problem. *Oper. Res.* **42**, 677–687 (1994)
150. Y. Li, P.M. Pardalos, M.G.C. Resende, A greedy randomized adaptive search procedure for the quadratic assignment problem, in *Quadratic Assignment and Related Problems*, ed. by P.M. Pardalos, H. Wolkowicz. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 16 (American Mathematical Society, Providence, 1994), pp. 237–261
151. A. Lim, F. Wang, A smoothed dynamic tabu search embedded GRASP for m -VRPTW, in *Proceedings of ICTAI 2004*, pp. 704–708 (2004)
152. A. Lim, B. Rodrigues, C. Wang, Two-machine flow shop problems with a single server. *J. Sched.* **9**, 515–543 (2006)
153. X. Liu, P.M. Pardalos, S. Rajasekaran, M.G.C. Resende, A GRASP for frequency assignment in mobile radio networks, in *Mobile Networks and Computing*, ed. by B.R. Badrinath, F. Hsu, P.M. Pardalos, S. Rajasekaran. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 52 (American Mathematical Society, Providence, 2000), pp. 195–201
154. H.R. Lourenço, D. Serra, Adaptive approach heuristics for the generalized assignment problem. *Mathw. Soft Comput.* **9**, 209–234 (2002)
155. H.R. Lourenço, O.C. Martin, T. Stützle, Iterated local search, in *Handbook of Metaheuristics*, ed. by F. Glover, G. Kochenberger (Kluwer Academic Publishers, Boston, 2003), pp. 321–353
156. M. Luby, A. Sinclair, D. Zuckerman, Optimal speedup of Las Vegas algorithms. *Inf. Process. Lett.* **47**, 173–180 (1993)
157. M. Luis, S. Salhi, G. Nagy, A guided reactive GRASP for the capacitated multi-source Weber problem. *Comput. Oper. Res.* **38**, 1014–1024 (2011)
158. C.L.B. Maia, R.A.F. Carmo, F.G. Freitas, G.A.L. Campos, J.T. Souza, Automated test case prioritization with reactive GRASP. *Adv. Softw. Eng.* **2010**, Article ID 428521 (2010)
159. R. Martí, Arc crossing minimization in graphs with GRASP. *IEEE Trans.* **33**, 913–919 (2001)
160. R. Martí, Arc crossing minimization in graphs with GRASP. *IEEE Trans.* **33**, 913–919 (2002)
161. R. Martí, V. Estruch, Incremental bipartite drawing problem. *Comput. Oper. Res.* **28**, 1287–1298 (2001)
162. R. Martí, M. Laguna, Heuristics and meta-heuristics for 2-layer straight line crossing minimization. *Discret. Appl. Math.* **127**, 665–678 (2003)
163. R. Martí, F. Sandoya, GRASP and path relinking for the equitable dispersion problem. *Comput. Oper. Res.* **40**, 3091–3099 (2013)
164. O. Martin, S.W. Otto, Combining simulated annealing with local search heuristics. *Ann. Oper. Res.* **63**, 57–75 (1996)

165. O. Martin, S.W. Otto, E.W. Felten, Large-step Markov chains for the traveling salesman problem. *Complex Syst.* **5**, 299–326 (1991)
166. S.L. Martins, C.C. Ribeiro, M.C. Souza, A parallel GRASP for the Steiner problem in graphs, in *Proceedings of IRREGULAR'98 – 5th International Symposium on Solving Irregularly Structured Problems in Parallel*, ed. by A. Ferreira, J. Rolim. Lecture Notes in Computer Science, vol. 1457 (Springer, Berlin, 1998), pp. 285–297
167. S.L. Martins, P.M. Pardalos, M.G.C. Resende, C.C. Ribeiro, Greedy randomized adaptive search procedures for the steiner problem in graphs, in *Randomization Methods in Algorithmic Design*, P.M. Pardalos, S. Rajasejaran, J. Rolim. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 43 (American Mathematical Society, Providence, 1999), pp. 133–145
168. S.L. Martins, M.G.C. Resende, C.C. Ribeiro, P.M. Pardalos, A parallel GRASP for the Steiner tree problem in graphs using a hybrid local search strategy. *J. Glob. Optim.* **17**, 267–283 (2000)
169. G.R. Mateus, M.G.C. Resende, R.M.A. Silva, GRASP with path-relinking for the generalized quadratic assignment problem. *J. Heuristics* **17**, 527–565 (2011)
170. T. Mavridou, P.M. Pardalos, L.S. Pitsoulis, M.G.C. Resende, A GRASP for the biquadratic assignment problem. *Eur. J. Oper. Res.* **105**, 613–621 (1998)
171. M. Mestria, L.S. Ochi, S.L. Martins, GRASP with path relinking for the symmetric Euclidean clustered traveling salesman problem. *Comput. Oper. Res.* **40**, 3218–3229 (2013)
172. N. Mladenović, P. Hansen, Variable neighborhood search. *Comput. Oper. Res.* **24**, 1097–1100 (1997)
173. S.K. Monkman, D.J. Morrice, J.F. Bard, A production scheduling heuristic for an electronics manufacturer with sequence-dependent setup costs. *Eur. J. Oper. Res.* **187**, 1100–1114 (2008)
174. R.E.N. Moraes, C.C. Ribeiro, Power optimization in ad hoc wireless network topology control with biconnectivity requirements. *Comput. Oper. Res.* **40**, 3188–3196 (2013)
175. L.F. Morán-Mirabal, J.L. González-Velarde, M.G.C. Resende, R.M.A. Silva, Randomized heuristics for handover minimization in mobility networks. *J. Heuristics* **19**, 845–880 (2013)
176. L.F. Morán-Mirabal, J.L. González-Velarde, M.G.C. Resende, Randomized heuristics for the family traveling salesperson problem. *Int. Trans. Oper. Res.* **21**, 41–57 (2014)
177. R.A. Murphey, P.M. Pardalos, L.S. Pitsoulis, A greedy randomized adaptive search procedure for the multitarget multisensor tracking problem, in *Network Design: Connectivity and Facilities Location*, ed. by P.M. Pardalos, D.-Z. Du. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 40 (American Mathematical Society, Providence, 1998), pp. 277–301
178. R.A. Murphey, P.M. Pardalos, L.S. Pitsoulis, A parallel GRASP for the data association multidimensional assignment problem, in *Parallel Processing of Discrete Problems*, ed. by P.M. Pardalos. The IMA Volumes in Mathematics and Its Applications, vol. 106 (Springer, New York, 1998), pp. 159–180
179. M.C.V. Nascimento, L. Pitsoulis, Community detection by modularity maximization using GRASP with path relinking. *Comput. Oper. Res.* **40**, 3121–3131 (2013)
180. M.C.V. Nascimento, M.G.C. Resende, F.M.B. Toledo, GRASP heuristic with path-relinking for the multi-plant capacitated lot sizing problem. *Eur. J. Oper. Res.* **200**, 747–754 (2010)
181. V.-P. Nguyen, C. Prins, C. Prodhon, Solving the two-echelon location routing problem by a GRASP reinforced by a learning process and path relinking. *Eur. J. Oper. Res.* **216**, 113–126 (2012)
182. E. Nowicki, C. Smutnicki, An advanced tabu search algorithm for the job shop problem. *J. Sched.* **8**, 145–159 (2005)
183. C.A. Oliveira, P.M. Pardalos, M.G.C. Resende, GRASP with path-relinking for the quadratic assignment problem, in *Proceedings of III Workshop on Efficient and Experimental Algorithms*, vol. 3059, ed. by C.C. Ribeiro, S.L. Martins (Springer, New York, 2004), pp. 356–368
184. I.H. Osman, B. Al-Ayoubi, M. Barake, A greedy random adaptive search procedure for the weighted maximal planar graph problem. *Comput. Ind. Eng.* **45**, 635–651 (2003)

185. J.A. Pacheco, S. Casado, Solving two location models with few facilities by using a hybrid heuristic: a real health resources case. *Comput. Oper. Res.* **32**, 3075–3091 (2005)
186. A.V.F. Pacheco, G.M. Ribeiro, G.R. Mauri, A GRASP with path-relinking for the workover rig scheduling problem. *Int. J. Nat. Comput. Res.* **1**, 1–14 (2010)
187. G. Palubeckis, Multistart tabu search strategies for the unconstrained binary quadratic optimization problem. *Ann. Oper. Res.* **131**, 259–282 (2004)
188. P.M. Pardalos, L.S. Pitsoulis, M.G.C. Resende, A parallel GRASP implementation for the quadratic assignment problem, in *Parallel Algorithms for Irregularly Structured Problems – Irregular’94*, ed. by A. Ferreira, J. Rolim (Kluwer Academic Publishers, Dordrecht, 1995), pp. 115–133
189. P.M. Pardalos, L.S. Pitsoulis, M.G.C. Resende, A parallel GRASP for MAX-SAT problems. *Lect. Notes Comput. Sci.* **1184**, 575–585 (1996)
190. P.M. Pardalos, L.S. Pitsoulis, M.G.C. Resende, Algorithm 769: Fortran subroutines for approximate solution of sparse quadratic assignment problems using GRASP. *ACM Trans. Math. Softw.* **23**, 196–208 (1997)
191. P.M. Pardalos, T. Qian, M.G.C. Resende, A greedy randomized adaptive search procedure for the feedback vertex set problem. *J. Comb. Optim.* **2**, 399–412 (1999)
192. F. Parreño, R. Alvarez-Valdes, J.M. Tamarit, J.F. Oliveira, A maximal-space algorithm for the container loading problem. *INFORMS J. Comput.* **20**, 412–422 (2008)
193. R.A. Patterson, H. Pirkul, E. Rolland, A memory adaptive reasoning technique for solving the capacitated minimum spanning tree problem. *J. Heuristics* **5**, 159–180 (1999)
194. O. Pedrola, M. Ruiz, L. Velasco, D. Careglio, O. González de Dios, J. Comellas, A GRASP with path-relinking heuristic for the survivable IP/MPLS-over-WSON multi-layer network optimization problem. *Comput. Oper. Res.* **40**, 3174–3187 (2013)
195. L.S. Pessoa, M.G.C. Resende, C.C. Ribeiro, Experiments with the LAGRASP heuristic for set k -covering. *Optim. Lett.* **5**, 407–419 (2011)
196. L.S. Pessoa, M.G.C. Resende, C.C. Ribeiro, A hybrid Lagrangean heuristic with GRASP and path-relinking for set k -covering. *Comput. Oper. Res.* **40**, 3132–3146 (2013)
197. E. Pinana, I. Plana, V. Campos, R. Martí, GRASP and path relinking for the matrix bandwidth minimization. *Eur. J. Oper. Res.* **153**, 200–210 (2004)
198. L.S. Pitsoulis, P.M. Pardalos, D.W. Hearn, Approximate solutions to the turbine balancing problem. *Eur. J. Oper. Res.* **130**, 147–155 (2001)
199. F. Poppe, M. Pickavet, P. Arijs, P. Demeester, Design techniques for SDH mesh-restorable networks, in *Proceedings of the European Conference on Networks and Optical Communications, Volume 2: Core and ATM Networks*, pp. 94–101, (1997)
200. M. Prais, C.C. Ribeiro, Parameter variation in GRASP implementations, in *Extended Abstracts of the Third Metaheuristics International Conference*, Angra dos Reis, pp. 375–380 (1999)
201. M. Prais, C.C. Ribeiro, Parameter variation in GRASP procedures. *Investigación Operativa* **9**, 1–20 (2000)
202. M. Prais, C.C. Ribeiro, Reactive GRASP: an application to a matrix decomposition problem in TDMA traffic assignment. *INFORMS J. Comput.* **12**, 164–176 (2000)
203. M. Rahmani, M. Rashidinejad, E.M. Carreno, R.A. Romero, Evolutionary multi-move path-relinking for transmission network expansion planning, in *2010 IEEE Power and Energy Society General Meeting*, Minneapolis (IEEE, New York, 2010), pp. 1–6
204. M.C. Rangel, N.M.M. Abreu, P.O. Boaventura Netto, GRASP in the QAP: an acceptance bound for initial solutions. *Pesquisa Operacional* **20**, 45–58 (2000)
205. M.G. Ravetti, F.G. Nakamura, C.N. Meneses, M.G.C. Resende, G.R. Mateus, P.M. Pardalos, Hybrid heuristics for the permutation flow shop problem. Technical Report, AT&T Labs Research Technical Report, Florham Park, 2006
206. M. Reghioui, C. Prins, N. Labadi, GRASP with path relinking for the capacitated arc routing problem with time windows, in *Applications of Evolutionary Computing*, ed. by M. Giacobini et al. Lecture Notes in Computer Science, vol. 4448 (Springer, Berlin, 2007), pp. 722–731
207. M.G.C. Resende, Computing approximate solutions of the maximum covering problem using GRASP. *J. Heuristics* **4**, 161–171 (1998)

208. M.G.C. Resende, T.A. Feo, A GRASP for satisfiability, in *Cliques, Coloring, and Satisfiability: The Second DIMACS Implementation Challenge*, ed. by D.S. Johnson, M.A. Trick. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 26 (American Mathematical Society, Providence, 1996), pp. 499–520
209. L.I.P. Resende, M.G.C. Resende, A GRASP for frame relay permanent virtual circuit routing, in *Extended Abstracts of the III Metaheuristics International Conference*, ed. by C.C. Ribeiro, P. Hansen, Angra dos Reis, pp. 397–401 (1999)
210. M.G.C. Resende, C.C. Ribeiro, A GRASP for graph planarization. *Networks* **29**, 173–189 (1997)
211. M.G.C. Resende, C.C. Ribeiro, Graph planarization, in *Encyclopedia of Optimization*, vol. 2, ed. by C. Floudas, P.M. Pardalos (Kluwer Academic Publishers, Boston, 2001), pp. 368–373
212. M.G.C. Resende, C.C. Ribeiro, A GRASP with path-relinking for private virtual circuit routing. *Networks* **41**, 104–114 (2003)
213. M.G.C. Resende, C.C. Ribeiro, GRASP with path-relinking: recent advances and applications, in *Metaheuristics: Progress as Real Problem Solvers*, ed. by T. Ibaraki, K. Nonobe, M. Yagiura (Springer, Boston, 2005), pp. 29–63
214. M.G.C. Resende, C.C. Ribeiro, Restart strategies for GRASP with path-relinking heuristics. *Optim. Lett.* **5**, 467–478 (2011)
215. M.G.C. Resende, C.C. Ribeiro, *Optimization by GRASP: Greedy Randomized Adaptive Search Procedures* (Springer, New York, 2016)
216. M.G.C. Resende, R.F. Werneck, A hybrid heuristic for the p -median problem. *J. Heuristics* **10**, 59–88 (2004)
217. M.G.C. Resende, R.F. Werneck, A hybrid multistart heuristic for the uncapacitated facility location problem. *Eur. J. Oper. Res.* **174**, 54–68 (2006)
218. M.G.C. Resende, P.M. Pardalos, Y. Li, Algorithm 754: Fortran subroutines for approximate solution of dense quadratic assignment problems using GRASP. *ACM Trans. Math. Softw.* **22**, 104–118 (1996)
219. M.G.C. Resende, L.S. Pitsoulis, P.M. Pardalos, Approximate solution of weighted MAX-SAT problems using GRASP, in *Satisfiability Problems*, ed. by J. Gu, P.M. Pardalos. DIMACS Series on Discrete Mathematics and Theoretical Computer Science, vol. 35 (American Mathematical Society, Providence, 1997), pp. 393–405
220. M.G.C. Resende, T.A. Feo, S.H. Smith, Algorithm 787: Fortran subroutines for approximate solution of maximum independent set problems using GRASP. *ACM Trans. Math. Softw.* **24**, 386–394 (1998)
221. M.G.C. Resende, L.S. Pitsoulis, P.M. Pardalos, Fortran subroutines for computing approximate solutions of MAX-SAT problems using GRASP. *Discret. Appl. Math.* **100**, 95–113 (2000)
222. M.G.C. Resende, R. Martí, M. Gallego, A. Duarte, GRASP and path relinking for the maximum diversity problem. *Comput. Oper. Res.* **37**, 498–508 (2010)
223. A.P. Reynolds, B. de la Iglesia, A multi-objective GRASP for partial classification. *Soft Comput.* **13**, 227–243 (2009)
224. C.C. Ribeiro, GRASP: Une métaheuristique gloutonne et probabiliste, in *Optimisation Approchée en Recherche Opérationnelle*, ed. by J. Teghem, M. Pirlot (Hermès, Paris, 2002), pp. 153–176
225. C.C. Ribeiro, Sports scheduling: problems and applications. *Int. Trans. Oper. Res.* **19**, 201–226 (2012)
226. C.C. Ribeiro, M.G.C. Resende, Algorithm 797: Fortran subroutines for approximate solution of graph planarization problems using GRASP. *ACM Trans. Math. Softw.* **25**, 342–352 (1999)
227. C.C. Ribeiro, I. Rosseti, Efficient parallel cooperative implementations of GRASP heuristics. *Parallel Comput.* **33**, 21–35 (2007)
228. C.C. Ribeiro, I. Rosseti, tttplots-compare: A perl program to compare time-to-target plots or general runtime distributions of randomized algorithms. *Optim. Lett.* **9**, 601–614 (2015)
229. C.C. Ribeiro, M.C. Souza, Variable neighborhood search for the degree constrained minimum spanning tree problem. *Discret. Appl. Math.* **118**, 43–54 (2002)

230. C.C. Ribeiro, S. Urrutia, Heuristics for the mirrored traveling tournament problem. *Eur. J. Oper. Res.* **179**, 775–787 (2007)
231. C.C. Ribeiro, D.S. Vianna, A GRASP/VND heuristic for the phylogeny problem using a new neighborhood structure. *Int. Trans. Oper. Res.* **12**, 325–338 (2005)
232. C.C. Ribeiro, C.D. Ribeiro, R.S. Lanzelotte, Query optimization in distributed relational databases. *J. Heuristics* **3**, 5–23 (1997)
233. C.C. Ribeiro, E. Uchoa, R.F. Werneck, A hybrid GRASP with perturbations for the Steiner problem in graphs. *INFORMS J. Comput.* **14**, 228–246 (2002)
234. C.C. Ribeiro, S.L. Martins, I. Rosseti, Metaheuristics for optimization problems in computer communications. *Comput. Comun.* **30**, 656–669 (2007)
235. C.C. Ribeiro, I. Rosseti, R. Vallejos, On the use of run time distributions to evaluate and compare stochastic local search algorithms, in *Engineering Stochastic Local Search Algorithms*, ed. by T. Sttzle, M. Biratari, and H.H. Hoos. Lecture Notes in Computer Science, vol. 5752 (Springer, Berlin, 2009), pp. 16–30
236. C.C. Ribeiro, I. Rosseti, R. Vallejos, Exploiting run time distributions to compare sequential and parallel stochastic local search algorithms. *J. Glob. Optim.* **54**, 405–429 (2012)
237. R.Z. Ríos-Mercado, J.F. Bard, Heuristics for the flow line problem with setup costs. *Eur. J. Oper. Res.* **110**, 76–98 (1998)
238. R.Z. Ríos-Mercado, J.F. Bard, An enhanced TSP-based heuristic for makespan minimization in a flow shop with setup costs. *J. Heuristics* **5**, 57–74 (1999)
239. R.Z. Ríos-Mercado, E. Fernández. A reactive GRASP for a commercial territory design problem with multiple balancing requirements. *Comput. Oper. Res.* **36**, 755–776 (2009)
240. A. Riva, F. Amigoni, A GRASP metaheuristic for the coverage of grid environments with limited-footprint tools, in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS '17*, Richland, SC, pp. 484–491. International Foundation for Autonomous Agents and Multiagent Systems (2017)
241. A.J. Robertson, A set of greedy randomized adaptive local search procedure (GRASP) implementations for the multidimensional assignment problem. *Comput. Optim. Appl.* **19**, 145–164 (2001)
242. P.L. Rocha, M.G. Ravetti, G.R. Mateus, The metaheuristic GRASP as an upper bound for a branch and bound algorithm in a scheduling problem with non-related parallel machines and sequence-dependent setup times, in *Proceedings of the 4th EU/ME Workshop: Design and Evaluation of Advanced Hybrid Meta-Heuristics*, vol. 1 (2004), pp. 62–67
243. F.J. Rodríguez, C. Blum, C. García-Martínez, M. Lozano, GRASP with path-relinking for the non-identical parallel machine scheduling problem with minimising total weighted completion times. *Ann. Oper. Res.* **201**, 383–401 (2012)
244. F.J. Rodríguez, F. Glover, C. García-Martínez, R. Martí, M. Lozano, Grasp with exterior path-relinking and restricted local search for the multidimensional two-way number partitioning problem. *Comput. Oper. Res.* **78**, 243–254 (2017)
245. M.A. Salazar-Aguilar, R.Z. Ríos-Mercado, J.L. González-Velarde, GRASP strategies for a bi-objective commercial territory design problem. *J. Heuristics* **19**, 179–200 (2013)
246. J. Santamaría, O. Cordón, S. Damas, R. Martí, R.J. Palma, GRASP & evolutionary path relinking for medical image registration based on point matching, in *2010 IEEE Congress on Evolutionary Computation* (IEEE, New York, 2010), pp. 1–8
247. J. Santamaría, O. Cordón, S. Damas, R. Martí, R.J. Palma, GRASP and path relinking hybridizations for the point matching-based image registration problem. *J. Heuristics* **18**, 169–192 (2012)
248. D. Santos, A. de Sousa, F. Alvelos, A hybrid column generation with GRASP and path relinking for the network load balancing problem. *Comput. Oper. Res.* **40**, 3147–3158 (2013)
249. M. Scaparra, R. Church, A GRASP and path relinking heuristic for rural road network development. *J. Heuristics* **11**, 89–108 (2005)
250. I.V. Sergienko, V.P. Shilo, V.A. Roshchin, Optimization parallelizing for discrete programming problems. *Cybern. Syst. Anal.* **40**, 184–189 (2004)
251. O.V. Shylo, T. Middelkoop, P.M. Pardalos, Restart strategies in optimization: parallel and serial cases. *Parallel Comput.* **37**, 60–68 (2011)
252. O.V. Shylo, O.A. Prokopyev, J. Rajgopal, On algorithm portfolios and restart strategies. *Oper. Res. Lett.* **39**, 49–52 (2011)

253. F. Silva, D. Serra, Locating emergency services with different priorities: the priority queuing covering location problem. *J. Oper. Res. Soc.* **59**, 1229–1238 (2007)
254. R.M.A. Silva, M.G.C. Resende, P.M. Pardalos, M.J. Hirsch, A Python/C library for bound-constrained global optimization with continuous GRASP. *Optim. Lett.* **7**, 967–984 (2013)
255. R.M.A. Silva, M.G.C. Resende, P.M. Pardalos, G.R. Mateus, G. de Tomi, GRASP with path-relinking for facility layout, in *Models, Algorithms, and Technologies for Network Analysis*, ed. by B.I. Goldenhorin, V.A. Kalyagin, P.M. Pardalos. *Springer Proceedings in Mathematics and Statistics*, vol. 59 (Springer, Berlin, 2013), pp. 175–190
256. D. Sosnowska, Optimization of a simplified fleet assignment problem with metaheuristics: simulated annealing and GRASP, in *Approximation and Complexity in Numerical Optimization*, ed. by P.M. Pardalos (Kluwer Academic Publishers, Dordrecht, 2000)
257. M.C. Souza, C. Duhamel, C.C. Ribeiro, A GRASP heuristic for the capacitated minimum spanning tree problem using a memory-based local search strategy, in *Metaheuristics: Computer Decision-Making*, ed. by M.G.C. Resende, J. Souza (Kluwer Academic Publisher, Dordrecht, 2004), pp. 627–658
258. A. Srinivasan, K.G. Ramakrishnan, K. Kumaram, M. Aravamudam, S. Naqvi, Optimal design of signaling networks for Internet telephony, in *IEEE INFOCOM 2000*, vol. 2 (2000), pp. 707–716
259. H. Takahashi, A. Matsuyama, An approximate solution for the Steiner problem in graphs. *Math. Jpn.* **24**, 573–577 (1980)
260. T.L. Urban, Solution procedures for the dynamic facility layout problem. *Ann. Oper. Res.* **76**, 323–342 (1998)
261. T.L. Urban, W.-C. Chiang, R.A. Russel, The integrated machine allocation and layout problem. *Int. J. Prod. Res.* **38**, 2913–2930 (2000)
262. F.L. Usberti, P.M. França, A.L.M. França, GRASP with evolutionary path-relinking for the capacitated arc routing problem. *Comput. Oper. Res.* **40**, 3206–3217 (2013)
263. J.X. Vianna Neto, D.L.A. Bernert, L.S. Coelho, Continuous GRASP algorithm applied to economic dispatch problem of thermal units, in *Proceedings of the 13th Brazilian Congress of Thermal Sciences and Engineering*, Uberlandia (2010)
264. J.G. Villegas, C. Prins, C. Prodhon, A.L. Medaglia, N. Velasco, GRASP/VND and multi-start evolutionary local search for the single truck and trailer routing problem with satellite depots. *Eng. Appl. Artif. Intelli.* **23**, 780–794 (2010)
265. J.G. Villegas, C. Prins, C. Prodhon, A.L. Medaglia, N. Velasco, A GRASP with evolutionary path relinking for the truck and trailer routing problem. *Comput. Oper. Res.* **38**, 1319–1334 (2011)
266. D.L. Woodruff, E. Zemel, Hashing vectors for tabu search. *Ann. Oper. Res.* **41**, 123–137 (1993)
267. J.Y. Xu, S.Y. Chiu, Effective heuristic procedure for a field technician scheduling problem. *J. Heuristics* **7**, 495–509 (2001)
268. J. Yen, M. Carlsson, M. Chang, J.M. Garcia, H. Nguyen, Constraint solving for inkjet print mask design. *J. Imaging Sci. Technol.* **44**, 391–397 (2000)