

GRASP WITH PATH-RELINKING FOR THE GENERALIZED QUADRATIC ASSIGNMENT PROBLEM

GERALDO R. MATEUS, MAURICIO G.C. RESENDE, AND RICARDO M. A. SILVA

ABSTRACT. The generalized quadratic assignment problem (GQAP) is a generalization of the NP-hard quadratic assignment problem (QAP) that allows multiple facilities to be assigned to a single location as long as the capacity of the location allows. The GQAP has numerous applications, including facility design, scheduling, and network design. In this paper, we propose several GRASP with path-relinking heuristics for the GQAP using different construction, local search, and path-relinking procedures. We introduce a novel approximate local search scheme, as well as a new variant of path-relinking that deals with infeasibilities. Extensive experiments on a large set of test instances show that the best of the proposed variants is both effective and efficient.

1. INTRODUCTION

Given n facilities and m locations, the generalized quadratic assignment problem (GQAP) consists in feasibly assigning each facility to a location. Each facility demands some amount of location capacity and each location has a fixed amount of capacity to distribute among facilities. An assignment is feasible if each location has sufficient capacity to accommodate the demands of all facilities assigned to it. Given nonnegative flows between all pairs of facilities and nonnegative distances between all pairs of locations, the GQAP seeks a feasible assignment that minimizes the sum of products of flows and distances in addition to a linear assignment component. Ideally, two facilities with high inter-facility flow are either assigned to the same location or to two nearby locations. The GQAP is a generalization of the quadratic assignment problem (QAP) that allows multiple facilities to be assigned to a single location as long as the capacity of the location permits. The QAP (see, e.g., Li et al. (1994), Oliveira et al. (2004), and Pardalos et al. (1994)) is a special case of the GQAP where each facility demands one unit of capacity and each location has one unit of capacity to distribute.

Let $N = \{1, \dots, n\}$ denote the set of facilities and $M = \{1, \dots, m\}$ the set of locations. Furthermore, denote by $A_{n \times n} = (a_{ii'})$ the flow between facilities $i \in N$ and $i' \in N$, such that $a_{ii'} \in \mathfrak{R}^+$ if $i \neq i'$ and otherwise $a_{ii'} = 0$, by $B_{m \times m} = (b_{jj'})$ the distance between locations $j \in M$ and $j' \in M$, such that $b_{jj'} \in \mathfrak{R}^+$ if $j \neq j'$ and otherwise $b_{jj'} = 0$, and by $C_{n \times m} = (c_{ij})$, the cost of assigning facility $i \in N$ to location $j \in M$, such that $c_{ij} \in \mathfrak{R}^+$. Let $z \in \mathfrak{R}^+$ be a scaling factor called the unit traffic cost, $q_i \in \mathfrak{R}^+$ be the capacity demanded by facility $i \in N$, and $Q_j \in \mathfrak{R}^+$, the capacity of location $j \in M$. The GQAP consists in finding $X_{n \times m} = (x_{ij})$,

Date: December 16, 2008. Revised November 25, 2009.

Key words and phrases. Generalized quadratic assignment problem, heuristic, GRASP, path-relinking, experimental algorithm.

AT&T Labs Research Technical Report.

with $x_{ij} \in \{0, 1\}$, where facility $i \in N$ is assigned to location $j \in M$ if and only if $x_{ij} = 1$, such that the constraints

$$(1) \quad \sum_{j \in M} x_{ij} = 1, \forall i \in N,$$

$$(2) \quad \sum_{i \in N} q_i x_{ij} \leq Q_j, \forall j \in M,$$

$$(3) \quad x_{ij} \in \{0, 1\}, \forall i \in N, \forall j \in M$$

are satisfied and the objective function

$$(4) \quad \sum_{i \in N} \sum_{j \in M} c_{ij} x_{ij} + z \sum_{i \in N} \sum_{j \in M} \sum_{i' \in N, i' \neq i} \sum_{j' \in M} a_{ii'} b_{jj'} x_{ij} x_{i'j'}$$

is minimized. Constraints (1) guarantee that each facility is assigned to exactly one location, while constraints (2) ensure that location capacities are not violated.

The generalized quadratic assignment problem, introduced by Lee and Ma (2005), is a relatively new combinatorial optimization problem. Lee and Ma proposed three linear programming relaxations and a branch and bound algorithm for the GQAP, as well as suite of test problems on which they tested their algorithms. Elloumi et al. (2003) proposed three families of lower bounding techniques for the constrained module allocation problem, a special case of the GQAP. They also propose a suite of test problems. Cordeau et al. (2006) proposed a new suite of test problems and an integer linear programming formulation with which they solved small instances of their test problems (having 5 to 16 facilities and 3 to 30 locations). A heuristic based on truncating the branch and bound tree was also proposed. The authors finally describe a memetic algorithm which was shown to find optimal solutions for all small instances tested and, on larger instances (having 20 to 50 facilities and 6 to 20 locations) was able to find better solutions than those found with the truncated branch and bound. Hahn et al. (2007) use the *reformulation linearization technique* (RLT) dual ascent procedure to find optimal solutions to some problems from Cordeau et al. (2006), Lee and Ma (2005), and Elloumi et al. (2003), having up to 20 facilities and 15 locations. For problems of these dimensions, the runtimes of their algorithm were no worse than those of previous exact methods. Pessoa et al. (2008) proposed a new lower bound for the GQAP based on the Lagrangian relaxation of the RLT formulation proposed in Hahn et al. (2007). To efficiently solve the relaxed problem, they combined the Hahn et al. (2007) dual ascent procedure with the general-purpose volume algorithm of Barahona and Anbil (2000). They tested the new bounding procedure using the branch-and-bound method of Hahn et al. (2007) on a subset of instances from Cordeau et al. (2006), Lee and Ma (2005), and Elloumi et al. (2003).

GRASP, or greedy randomized adaptive search procedure, is a metaheuristic for finding approximate solutions to combinatorial optimization problems formulated as

$$\min f(x) \text{ subject to } x \in \mathcal{X},$$

where $f(\cdot)$ is an objective function to be minimized and \mathcal{X} is a discrete set of feasible solutions. It was first introduced by Feo and Resende (1989) in a paper describing a probabilistic heuristic for set covering. Since then, GRASP has experienced continued development (Feo and Resende, 1995; Resende and Ribeiro, 2002; 2010)

and has been applied in a wide range of problem areas (Festa and Resende, 2002; 2009a;b).

In this paper, we present a GRASP with path-relinking heuristic (Laguna and Martí, 1999; Resende and Ribeiro, 2005) for the generalized quadratic assignment problem. Extensive computational experiments on benchmark test problems show the effectiveness of these heuristics.

The paper is organized as follows. In Section 2, we describe the GRASP with path-relinking procedure. We consider variants of the greedy randomized construction procedure, the local search, and the path-relinking intensification procedure. Computational results are described in Section 3 and concluding remarks are made in Section 4.

2. GRASP WITH PATH-RELINKING FOR GQAP

A GRASP is a multi-start heuristic where at each iteration a greedy randomized solution is constructed to be used as a starting solution for local search. Local search repeatedly substitutes the current solution by a better solution in the neighborhood of the current solution. Each such replacement is called a *move*. If there is no better solution in the neighborhood, the current solution is declared a local minimum and the search stops. The best local minimum found over all GRASP iterations is output as the solution.

GRASP iterations are independent, i.e. solutions found in previous GRASP iterations do not influence the algorithm in the current iteration. The use of previously found solutions to influence the procedure in the current iteration can be thought of as a memory mechanism. One way to incorporate memory into GRASP is with path-relinking (Glover, 1996; Glover et al., 2000). In GRASP with path-relinking (Laguna and Martí, 1999; Resende and Ribeiro, 2005), an elite set of diverse good-quality solutions is maintained to be used during each GRASP iteration. After a solution is produced with greedy randomized construction and local search, that solution is combined with a randomly selected solution from the elite set using the path-relinking operator. The best of the combined solutions is a candidate for inclusion in the elite set and is added to the elite set if it meets quality and diversity criteria.

Algorithm 1 shows pseudo-code for a GRASP with path-relinking heuristic for the GQAP. The algorithm takes as input the set N of facilities, the set M of locations, the flow matrix A , the distance matrix B , the assignment cost matrix C , the scaling factor z , the facility demands q_i , $i \in N$, and the location capacities Q_j , $j \in M$, and outputs an assignment vector π^* specifying the location of each facility in the best solution found.

After initializing the elite set P as empty in line 1, the GRASP with path-relinking iterations are computed in lines 2 to 24 until a stopping criterion is satisfied. This criterion could be, for example, a maximum number of iterations, a target solution quality, or a maximum number of iterations without improvement. During each iteration, a greedy randomized solution π' is generated in line 3. If the elite set P does not have at least ρ elements, then if π' is feasible and sufficiently different from all other elite set solutions, π' is added to the elite set in line 22. To define the term *sufficiently different* more precisely, let $\Delta(\pi', \pi)$ be defined as the minimum number of moves needed to transform π' into π or vice-versa. For a given level of difference δ , we say π' is sufficiently different from all elite solutions

in P if $\Delta(\pi', \pi) > \delta$ for all $\pi \in P$, which we indicate with the notation $\pi' \not\approx P$. If the elite set P does have at least ρ elements, then the steps in lines 5 to 19 are computed. The greedy randomized construction procedure is not guaranteed to generate a feasible solution. It is described in detail in Subsection 2.1. If the greedy randomized procedure returns an infeasible solution, a feasible solution π' is selected uniformly at random from the elite set in line 6 to be used as a surrogate for the greedy randomized solution. An approximate local search is applied using π' as a starting point in line 8, resulting in an approximate local minimum, which we denote by π' . Since elite solutions are made up of approximate local minima, then applying an approximate local search to an elite solution will, with high probability, result in a different approximate local minimum. The approximate local search is not guaranteed to find an exact local minimum. It is described in detail in Subsection 2.2. Since π' is an approximate local minimum, the application of an approximate local search to it will, with high probability, result in a different approximate local minimum. Next, path-relinking is applied in line 10 between π' and an elite solution π^+ , randomly chosen in line 9. Solution π^+ is selected with probability proportional to $\Delta(\pi', \pi^+)$. In line 11, the approximate local search is applied to π' . If the elite set is full, then if π' is of better quality than the worst elite solution and $\pi' \not\approx P$, then it will be added to the elite set in line 14 in place of some elite solution. Among all elite solutions having cost no better than that of π' , a solution π most similar to π' , i.e. with the smallest $\Delta(\pi', \pi)$ value, is selected to be removed from the elite set. Ties are broken at random. Otherwise, if the elite set is not full, π' is simply added to the elite set in line 18 if $\pi' \not\approx P$.

2.1. Greedy randomized construction. Recall that a solution to the GQAP consists of n assignments of facilities to locations. The construction procedure builds a solution one assignment at time.

Suppose a number of assignments have already been made. To make the next assignment, the procedure needs to select a new facility and a location. Locations are made available, one at time. The procedure randomly determines whether to use a new location or a previously chosen location, favoring a new location when the previously chosen locations have insufficient or barely sufficient available capacity. Later in this section, we define this choice more precisely. If the procedure determines that a previously chosen location is to be selected, it then determines which facilities can be assigned to a previously chosen location having the maximum available capacity and randomly selects one of these facilities to be assigned. Of the locations that can accommodate this facility, one is selected at random and the assignment is made. Otherwise, if there is no previously chosen location with sufficient capacity or if the available capacity is barely sufficient, a new location is selected at random from the set of yet unchosen locations. Figure 1 illustrates this procedure. In the figure, the set N of facilities is partitioned into the set CF of assigned facilities and the set F of unassigned ones. Likewise, the set M of locations is partitioned into the set CL of previously chosen locations and the set L of unselected locations. Facilities in CF are assigned to locations in CL . The facility set T consists of all unassigned facilities with demands less than or equal to the maximum available capacity of locations in CL . After a facility from T is randomly selected, set R consists of previously selected locations that can accommodate it. A location is randomly selected from set R and the facility is assigned to it.

Data : $N, M, A, B, C, z, q_i, Q_j$.
Result: Solution $\pi^* \in \chi$.

```

1  $P \leftarrow \emptyset$ ;
2 while stopping criterion not satisfied do
3    $\pi' \leftarrow \text{GreedyRandomized}(\cdot)$ ;
4   if elite set  $P$  has at least  $\rho$  elements then
5     if  $\pi'$  is not feasible then
6       | Randomly select a new solution  $\pi' \in P$ ;
7     end
8      $\pi' \leftarrow \text{ApproxLocalSearch}(\pi')$ ;
9     Randomly select a solution  $\pi^+ \in P$ ;
10     $\pi' \leftarrow \text{PathRelinking}(\pi', \pi^+)$ ;
11     $\pi' \leftarrow \text{ApproxLocalSearch}(\pi')$ ;
12    if elite set  $P$  is full then
13      | if  $c(\pi') \leq \max\{c(\pi) \mid \pi \in P\}$  and  $\pi' \not\approx P$  then
14        | | Replace the element most similar to  $\pi'$  among all
15        | | elements with cost worst than  $\pi'$ ;
16      | end
17    else if  $\pi' \not\approx P$  then
18      |  $P \leftarrow P \cup \{\pi'\}$ ;
19    end
20
21  else if  $\pi'$  is feasible and  $\pi' \not\approx P$  then
22    |  $P \leftarrow P \cup \{\pi'\}$ ;
23  end
24 end
25 return  $\pi^* = \min\{c(\pi) \mid \pi \in P\}$ ;

```

Algorithm 1: Pseudo-code for GRASP+PR: GRASP with path-relinking heuristic.

The above procedure is not guaranteed to produce a feasible solution. The greedy randomized construction procedure, shown in Algorithm 2, addresses this difficulty by repeatedly applying the steps described above. The main loop in lines 1 to 21 is repeated a maximum number of times or until all facilities are assigned, i.e. when $F = \emptyset$. In each iteration of the procedure, the working sets are initialized in line 2 and the threshold probability is set to 1 in line 3. The purpose of the threshold is to control whether a new location should be selected. Since it is initially set to 1, then in the first iteration of the until loop in lines 4 to 19, the procedure always selects a new (first) location. At each iteration of the until loop, the threshold is updated in line 17 such that it will be more likely that a new location is selected when there are few facilities that can be assigned to locations in the current set R . The until loop consists of two stages. With probability equal to the threshold, the first stage (lines 5 to 9) selects a new location in line 6, updates the sets L and CL in line 7, and in line 8 determines the set T of facilities that can be assigned to some selected location. In the second stage (lines 10 to 18), the procedure randomly selects a




FIGURE 1. Greedy randomized construction process. Facilities on the left are assigned to locations on the right. Facilities in set F have not yet been assigned. Locations in set L have not yet been assigned to. Facilities in set CF have already been assigned. Locations in set CL have been or can potentially be assigned to. Facilities in T have demands less than or equal to the maximum slack of facilities in set CL . After randomly selecting a facility to assign from set T , set R is defined as the set of locations that can accommodate that facility. One of these locations is selected at random to accommodate the facility.

facility from set T in line 11, updates the sets T , F , and CF in line 12, creates the location set R in line 13, randomly selects a location from that set in line 14, makes the assignment of the facility to the location in line 15, determines the set T of facilities that can be assigned to some selected location in line 16, and updates the threshold probability in line 17. The until loop is repeated until both sets T and L are empty in line 19. The while loop ends either with a valid assignment in line 25 (indicated by $F = \emptyset$) or with no solution found determined in line 23.

The construction procedure makes use of randomization to select a location in the first stage of the until loop and to select a facility and location in the second stage. To select a location in the first stage (line 6), the procedure favors locations in L with high capacity and that are close to high-capacity locations in CL . For all $k \in L$, let

$$H_k = \sum_{l \in CL} \frac{Q_k \cdot Q_l}{b_{kl}}.$$

The probability that location k is chosen is

$$\frac{H_k}{\sum_{r \in L} H_r}.$$

```

Data :  $\bar{t}$  = maximum number of tries
Result: Solution  $x \in \mathcal{X}$ 
1 while  $k < \bar{t}$  and  $F \neq \emptyset$  do
2    $F \leftarrow N$ ;  $CF \leftarrow \emptyset$ ;  $L \leftarrow M$ ;  $CL \leftarrow \emptyset$ ;  $T \leftarrow \emptyset$ ;
3   Set threshold  $\leftarrow 1$ ;
4   repeat
5     if  $L \neq \emptyset$  and  $\text{random}([0, 1]) \leq \text{threshold}$  then
6       Randomly select a location  $l \in L$ ;
7       Update sets  $L \leftarrow L \setminus \{l\}$  and  $CL \leftarrow CL \cup \{l\}$ ;
8       Set  $T \subseteq F$  to be all facilities with demands less than
          or equal to the maximum slack in  $CL$ ;
9     end
10    if  $T \neq \emptyset$  then
11      Randomly select a facility  $f \in T$ ;
12      Update sets  $T \leftarrow T \setminus \{f\}$ ;  $F \leftarrow F \setminus \{f\}$ ; and
           $CF \leftarrow CF \cup \{f\}$ ;
13      Create set  $R \subseteq CL$  to be all locations having slack
          greater than or equal to demand of facility  $f$ ;
14      Randomly select a location  $l \in R$ ;
15      Assign facility  $f$  to location  $l$ ;
16      Set  $T \subseteq F$  to be all facilities with demands less than
          or equal to the maximum slack in  $CL$ ;
17      Set threshold  $\leftarrow 1 - |T|/|F|$ ;
18    end
19    until  $T = \emptyset$  and  $L = \emptyset$ ;
20     $k \leftarrow k + 1$ ;
21 end
22 if  $F \neq \emptyset$  then
23   Solution not found;
24 else
25   return assignment  $x \in \mathcal{X}$ ;
26 end

```

Algorithm 2: Pseudo-code for GreedyRandomized: Greedy randomized construction procedure.

To select a facility in the second stage (line 11), the procedure favors facilities in T having high demand and high flows to other facilities. For all facilities $k \in T$, let

$$W_k = q_k \sum_{l \in N \setminus \{k\}} a_{kl}.$$

The probability that facility k is selected is

$$\frac{W_k}{\sum_{r \in T} W_r}.$$

To select a location in the second stage (line 14), the procedure favors locations in R that have high capacity slack and that are close to high-capacity locations in CL . Furthermore, the procedure favors locations in R for which there is a small

increase in the objective function resulting from the assignment to it of the chosen facility in T . For all locations $k \in R$, let

$$Z_k = \sum_{l \in CL \setminus \{k\}} \frac{\sigma_k \cdot Q_l}{d \cdot b_{kl}},$$

where σ_k is the available capacity of location k and

$$d = c_{tk} + z \sum_{i' \in CF} \sum_{j' \in CL} a_{ti'} b_{kj'} x_{i'j'}$$

is the increase in the objective function resulting from the assignment to it of the chosen facility in $t \in T$. The probability that location $k \in R$ is selected is

$$\frac{Z_k}{\sum_{r \in R} Z_r}.$$

2.2. Approximate local search. When the construction procedure of Subsection 2.1 produces a feasible solution π , it is not guaranteed to be locally optimal. In GRASP heuristics, a local search procedure is usually applied starting at π to find a local minimum. The heuristic described in this paper makes use of an *approximate local search* procedure. This procedure uses two neighborhood structures which we call *1-move* and *2-move*. A solution in the 1-move neighborhood of π is obtained by changing one facility-to-location assignment in π . Likewise, a solution in the 2-move neighborhood of π is obtained by simultaneously changing two facility-to-location assignments in π .

One way to carry out a local search in these neighborhoods is to evaluate moves in the 1-move neighborhood and move to the first improving solution. This is called the *first fit* local search. If no 1-move improving solution exists, 2-move neighborhood solutions are evaluated and a move is made to the first improving solution. Another way to carry out the local search is to evaluate all 1-move and 2-move neighborhood solutions and move to the best improving solution. This is called the *best fit* local search. In both variants, the search is repeated until no improving solution in the neighborhoods exists. A tradeoff approach is used here. Instead of evaluating all of the 1-move and 2-move neighborhood solutions, we sample these neighborhoods and populate a candidate list with improving solutions. One of the solutions from the candidate list is randomly selected and a move is made to that solution. The search is repeated until no improving solution is sampled. Because solutions are sampled, not all neighbors may be evaluated. Consequently, the best solution found may not be a local minimum. We call this solution an *approximate local minimum*.

Pseudo-code for the approximate local search is shown in Algorithm 3. The procedure takes as input the starting solution π and two parameters, $MaxCLS$ and $MaxItr$, which control the sampling. The repeat until loop in lines 1 to 13 is repeated until an approximate local minimum is produced. In line 2, the sampling counter $count$ and the candidate list CLS are initialized. At each iteration of the inner loop in lines 3 to 9, the 1-move and 2-move neighborhoods of π are sampled without replacement by procedure $Move(\pi)$ in line 4. If this neighbor is an improving feasible solution, it is inserted into CLS in line 6. This is repeated until either the candidate list is full or a maximum number of neighbors have been sampled. In lines 10 to 12, if the candidate list is not empty, an assignment $\pi \in CLS$ is randomly chosen. If the set CLS is empty after the sampling process,

the procedure terminates returning π as an approximate local minimum in line 14. Otherwise, the procedure moves to a solution in CLS , repeating the outer loop.

We propose two possible schemes for carrying out a move. In the first, the *greedy* scheme, the move is made to the solution in CLS having the smallest cost. In the second, the *random* scheme, the move is made at random to some solution in CLS with a bias favoring lower-cost solutions. One reason for using this approximate local search is the size of the neighborhood, which in the standard local search needs to be completely evaluated at least once.

Since there are many more 2-move neighbors than 1-move neighbors, we sample each neighbor as follows. With probability p_1 we sample from the 1-move neighborhood while with probability $p_2 = 1 - p_1$ we sample from the 2-move neighborhood. Note that by doing this, the set CLS can have both 1-move and 2-move neighbors.

```

Data :  $\pi, MaxCLS, MaxItr$ 
Result: Approximate local minimum  $\pi$ 
1 repeat
2    $count \leftarrow 0; CLS \leftarrow \emptyset;$ 
3   repeat
4      $\pi' \leftarrow \text{Move}(\pi);$ 
5     if  $\pi'$  is feasible and  $cost(\pi') < cost(\pi)$  then
6        $CLS \leftarrow CLS \cup \{\pi'\};$ 
7     end
8      $count \leftarrow count + 1;$ 
9   until  $|CLS| \geq MaxCLS$  or  $count \geq MaxItr;$ 
10  if  $CLS \neq \emptyset$  then
11    Randomly select a solution  $\pi \in CLS;$ 
12  end
13 until  $CLS = \emptyset;$ 
14 return  $\pi;$ 

```

Algorithm 3: Pseudo-code for ApproxLocalSearch: Approximate local search procedure.

The approximate local search procedure makes use of randomization to select a solution π from set CLS in line 11. The procedure favors solutions in CLS with smaller cost, according to the objective function (4). The probability that solution π is chosen is

$$\frac{G_{\pi}}{\sum_{\pi' \in CLS} G_{\pi'}},$$

where, for all $\pi' \in CLS$,

$$G_{\pi'} = \frac{1}{f(\pi')},$$

where $f(\cdot)$ is the objective function (4).

2.3. Path-relinking. Path-relinking (PR) is a search strategy that explores trajectories connecting two solutions (Glover, 1996; Glover et al., 2000). Given two solutions (one which we call the *initial* solution and the other the *target* solution), their common elements are kept fixed and the space of solutions spanned by the remaining elements is explored with a greedy algorithm. Starting from an initial solution, the scheme moves from one solution to another until the target solution is reached. The objective consists in finding a solution that is better than both the initial and target solutions. Faria Jr. et al. (2005) introduced greedy randomized adaptive path-relinking, where instead of moving between two solutions in a greedy way, the moves are done in a greedy randomized fashion: one from among a candidate list with the most promising moves in the path being investigated is randomly selected.

One way of carrying out PR is through *forward* PR, where the best of the two given solutions is used as the target solution and the other as the initial solution. Another way is through *backward* PR (Aiex et al., 2005; Resende and Ribeiro, 2003), where the best of the two given solutions is used as the initial solution and the other as the target solution. In back and forward PR, forward PR is applied and is followed by backward PR. Another way to proceed is with *mixed* PR (Glover, 1996; Ribeiro and Rosseti, 2002), where two paths are simultaneously explored, the first emanating from the initial solution and the other from the target solution, until they meet at an intermediary solution equidistant from both extremes.

According to Ribeiro et al. (2002) and Aiex et al. (2003), better solutions are found in the backward approach than in the forward scheme, where the worst of the two solutions to be relinked is used as the initial solution. Back and forward PR usually finds better solutions than either forward or backward PR but at the expense of longer running times. Ribeiro and Rosseti (2002) showed that mixed PR outperformed forward, backward, and back and forward PR. As we will see in Section 3, for the GRASP with path-relinking heuristic proposed in this paper, the forward and backward schemes are better than the mixed scheme.

In this paper, a new variant of path-relinking is proposed, motivated by the fact that a single move guided by the target solution does not guarantee the feasibility of the newly constructed solution. For instance, suppose that among the differences between a solution π and the target solution is the location assigned to some specific facility. Performing a move in π that assigns the facility to the same location assigned to it in the target solution can result in a solution that may or not be feasible. If the capacity constraint of this location is not violated, the new solution is feasible. Otherwise, the repair procedure `makeFeasible` is applied in an attempt to make it feasible. The repair procedure is described in detail in Subsection 2.4. In a path-relinking *step*, the move described above, with or without a repair procedure, is repeated for each facility of solution π' assigned to a location that is different from the one in the target solution. Let \mathcal{G} denote the set of assignments resulting from these moves and let $\mathcal{B} \subseteq \mathcal{G}$ denote a subset of the best solutions in set \mathcal{G} .

The path-relinking process consists of a sequence of path-relinking steps. In each step, a feasible solution is either greedily or randomly selected from \mathcal{B} to become the new solution π of the next step. This process continues until the target solution is reached or when no single feasible solution is obtained from π' , i.e. when $\mathcal{B} = \emptyset$. The best solution in this path, including the start and target solutions, is returned as the final result.

Data : Starting solution π_s , target solution π_t , and candidate size factor η

Result: Best solution π^* in path from π_s to π_t

```

1  $\pi^* \leftarrow \operatorname{argmin}\{f(\pi_s), f(\pi_t)\};$ 
2  $f^* \leftarrow f(\pi^*);$ 
3  $\pi' \leftarrow \pi_s; \text{Fix} \leftarrow \emptyset; \text{nonFix} \leftarrow N;$ 
4 Compute difference  $\varphi(\pi', \pi_t)$  between solution  $\pi'$  and  $\pi_t$ ;
5 while  $\varphi(\pi', \pi_t) \neq \emptyset$  do
6    $\mathcal{B} \leftarrow \emptyset;$ 
7   for  $\forall m \in \varphi(\pi', \pi_t)$  do
8     Move the facility  $m$  in  $\pi'$  to the same location  $l$  assigned
     to  $m$  in  $\pi_t$ ;
9      $\bar{\pi} \leftarrow \text{makeFeasible}(\pi', m);$ 
10    if  $\bar{\pi}$  is feasible then
11      if  $|\mathcal{B}| \geq \eta \cdot |\varphi(\pi', \pi_t)|$  then
12        if  $c(\bar{\pi}) \leq \max\{c(\pi) \mid \pi \in \mathcal{B}\}$  and  $\bar{\pi} \notin \mathcal{B}$  then
13          replace the element most similar to  $\bar{\pi}$  among
          all elements with cost worst than  $\bar{\pi}$ ;
14        end
15      else if  $\bar{\pi} \notin \mathcal{B}$  then
16         $\mathcal{B} \leftarrow \mathcal{B} \cup \{\bar{\pi}\};$ 
17      end
18    end
19  end
20 end
21 if  $\mathcal{B} \neq \emptyset$  then
22   Randomly select a solution  $\pi \in \mathcal{B}$ ;
23   Compute difference  $\varphi(\pi, \pi_t)$  between solution  $\pi$  and  $\pi_t$ ;
24   Set  $I = \varphi(\pi', \pi_t) \setminus (\varphi(\pi', \pi_t) \cap \varphi(\pi, \pi_t))$ ;
25   Randomly select a facility  $i \in I$ ;
26    $\text{Fix} \leftarrow \text{Fix} \cup \{i\}; \text{nonFix} \leftarrow \text{nonFix} \setminus \{i\};$ 
27    $\pi' \leftarrow \pi;$ 
28   if  $f(\pi') < f^*$  then
29      $f^* \leftarrow f(\pi');$ 
30      $\pi^* \leftarrow \pi';$ 
31   end
32 else
33   return assignment  $\pi^*$ ;
34 end
35 end
36 return assignment  $\pi^*$ ;

```

Algorithm 4: Pseudo-code for PathRelinking: Path-relinking procedure.

Algorithm 4 shows pseudo-code for the path-relinking procedure. The algorithm takes as input π_s and π_t , the starting and target solutions, respectively, and outputs the best solution π^* in path from π_s to π_t . Initially, the best solution in the path is set as π^* in line 1 and its corresponding objective function is assigned to f^* in line 2. In line 3, the current solution π' is initialized with π_s , and the working sets *Fix* and *nonFix* are respectively initialized empty and with N . The while loop in lines 5 to 35 is repeated until all facilities in π' are assigned to the same locations assigned to them in π_t , i.e. the set $\varphi(\pi', \pi_t) = \{i \in N \mid \pi'(i) \neq \pi_t(i)\}$ is empty, where $\pi'(i)$ and $\pi_t(i)$ are the locations assigned to facility i in solutions π' and π_t , respectively.

After the set \mathcal{B} of best solutions is set to empty in line 6, each while loop iteration consists of two stages. The first stage in lines 7 to 20 implements the path-relinking step. It creates set \mathcal{B} with the best feasible solutions constructed from the current solution π' . In line 6, \mathcal{B} is initialized as empty. Each facility $m \in \varphi(\pi', \pi_t)$ is analyzed in lines 7 to 20 to create the set \mathcal{B} with the best feasible solutions constructed from the current solution π' . Procedure **makeFeasible** (described in detail in Subsection 2.4) is applied in line 8 to facility m to attempt to create a new solution $\bar{\pi}$ from π' . The application of **makeFeasible** to facility m can either result in a feasible or infeasible solution. This is done by assigning the location $\pi_t(m)$ to facility m in solution π' . In case **makeFeasible** returns a feasible solution $\bar{\pi} \notin \mathcal{B}$, $\bar{\pi}$ is added to \mathcal{B} if \mathcal{B} is not yet full. Otherwise, if \mathcal{B} is full and solution $\bar{\pi} \notin \mathcal{B}$ is not worse than any elite solution, then $\bar{\pi}$ is added to \mathcal{B} replacing some other elite solution.

In the second stage (lines 21 to 34), the procedure first randomly selects a solution π from set \mathcal{B} in line 22. Then, in line 25, it selects at random a facility $i \in I = \varphi(\pi', \pi_t) \setminus (\varphi(\pi', \pi_t) \cap \varphi(\pi, \pi_t))$, where I is defined in line 24 as the set containing all unfixed facilities whose locations were corrected in the previous path-relinking step. A facility is corrected when its location becomes the one assigned to it in the target solution. After fixing facility $i \in I$, sets *Fix* and *nonFix* are updated in line 26. Finally, the next path-relinking step solution π' is set as π in line 27 and, if $f(\pi') < f^*$, then the best cost f^* and best solution π^* are updated in lines 29 and 30, respectively. However, if in some path-relinking step no feasible solution is obtained from π' , the while loop is interrupted, returning the current solution π^* as the result in line 33. If the target solution is reached, then π^* is returned in line 36.

This path-relinking is different from the standard variant because given solutions π_s and π_t , their common elements are not kept fixed a priori, such that a small portion of the solution space spanned by the remaining elements is explored. The new variant fixes one facility at a time at each step.

The path-relinking procedure makes use of randomization to select a solution π from set \mathcal{B} in line 22 and to select a facility i from set I in line 25. To select a solution in \mathcal{B} , the procedure favors those solutions having smaller cost, according to the objective function $f(\cdot)$. For all $s \in \mathcal{B}$, let

$$U_s = \frac{1}{f(s)}.$$

The probability that solution π is chosen is




FIGURE 2. Single path-relinking process with repair. Facilities in set N on the left are assigned to locations in set M on the right. Facilities in set $Fix \subseteq N$ have already been fixed. Facilities in set $nonFix \subseteq N$ have not yet been fixed. Facilities in $FTL \subseteq nonFix$ are assigned to location l in solution π . Facilities in $T \subseteq FTL$ have demands less than or equal to the maximum slack of facilities in set M . After randomly selecting a facility from set T to assign, set R is defined as the set of locations that can accommodate that facility.

$$\frac{U_{\pi}}{\sum_{r \in \mathcal{B}} U_r}.$$

To select a facility i from set I , a uniform probability distribution is used.

2.4. Repair procedure. Suppose π is an intermediary solution obtained while applying path-relinking from π_s to π_t , as illustrated in Figure 2. The set N of facilities is partitioned into two sets, the set Fix of fixed facilities and the set $nonFix$ of the remaining facilities. Let $f \in nonFix$ be a facility whose location in π_t is l and in π is $u \neq l$, where $l, u \in M$.

According to Figure 2, performing a move in π that assigns facility f to location l is infeasible since the capacity Q_l of location l is insufficient to accommodate that facility. In this case, the facility set $FTL \subseteq nonFix$ is created with all unfixed facilities assigned to l in π . Next, set $T \subseteq FTL$ is constructed. It is made up of all facilities in FTL having demands at most equal to the maximum available capacity of locations in M . Let i be a facility that is randomly selected from T . Set R consists of locations in M that can accommodate facility i . A location j is selected at random from set R and facility i is assigned to it. This process is repeated until the capacity of location l has a nonnegative slack.

Algorithm 5 shows pseudo-code for the procedure `makeFeasible(\cdot)`. The algorithm takes as input a solution π , a facility $f \in \varphi(\pi, \pi_t)$, and the maximum number of tries $\bar{\tau}$, to output a feasible solution if possible. As input, facility f in π was moved to the same location l assigned to it in π_t . If the slack of location l is non-negative ($\sigma_l \geq 0$), i.e. the capacity of l is not violated, the algorithm returns the updated feasible solution π in line 2. Otherwise, in lines 3 to 24, the repair process is applied as follows. While the capacity of location l remains violated and the number k of tries is less than a maximum limit established a priori, a copy π' of solution π is made in line 6 before the repeat loop in lines 7 to 16. In this loop, a facility set $FTL \subseteq nonFix$ is created in line 8 with all yet unfixed facilities in N assigned to l in π' . Next, in line 9, set $T \subseteq FTL$ is constructed with all facilities in FTL having demands less than or equal to the maximum available capacity of locations in M . Then, if set T is not empty, a facility i is randomly selected from it in line 12 to create a set $R \subseteq M$ with all locations in M that can accommodate the facility. Finally, in line 13, a location j is randomly selected from set R and facility i is assigned to it in line 14. This loop is repeated until either the capacity of location l has a nonnegative slack or set T is empty. At the end, if the slack of l is nonnegative, the returned solution π' is feasible. Otherwise, it is infeasible.

The repair procedure makes use of randomization to select a facility i from set T in line 11 and to select a location j from set R in line 13. To select a facility in T , the procedure favors those with higher demand. Then, the probability that facility i is chosen is

$$\frac{q_i}{\sum_{r \in T} q_r},$$

where we recall that q_i is the capacity demanded by facility $i \in N$. To select a location j from set R , a uniform probability distribution is used.

3. EXPERIMENTAL RESULTS

In this section, we present results on computational experiments with the GRASP with path-relinking (GRASP-PR) heuristic introduced in this paper. First, we describe our test environment. Next, we determine an appropriated combination of values for the three most important parameters of the heuristic. Finally, we compare our implementation with other heuristics from the literature on a suite of test problems.

3.1. Test environment. All experiments with GRASP-PR were done on a Dell PE1950 computer with dual quad core 2.66 GHz Intel Xeon processors and 16 Gb of memory, running Red Hat Linux nesh version 5.1.19.6 (CentOS release 5.2, kernel 2.6.18-53.1.21.el5). The GRASP-PR heuristic was implemented in Java and compiled into bytecode with `javac` version 1.6.0_05. The random-number generator is an implementation of the Mersenne Twister algorithm (Matsumoto and Nishimura, 1998) from the COLT¹ library.

¹COLT is a open source library for high performance scientific and technical computing in Java. See <http://acs.lbl.gov/~hoschek/colt/>.

```

Data : solution  $\pi$ , facility  $f$ , and
          $\bar{\tau} = \text{maximum number of tries}$ 
Result: If possible, a feasible solution
1 if  $\sigma_l \geq 0$  then
2   | return feasible solution  $\pi$ ;
3 else
4   |  $k \leftarrow 0$ ;
5   | while  $k < \bar{\tau}$  and  $\sigma_l < 0$  do
6     |  $\pi' \leftarrow \pi$ ;
7     | repeat
8       | Set  $FTL \subseteq \text{nonFix}$  to be all facilities in  $\pi'$  assigned
9       | to  $l$ ;
10      | Set  $T \subseteq FTL$  to be all facilities with demands less
11      | than or equal to the maximum slack in  $M$ ;
12      | if  $T \neq \emptyset$  then
13        | Randomly select a facility  $i \in T$ ;
14        | Set  $R \subseteq M$  to be all locations having slack
15        | greater than or equal to demand of facility  $i$ ;
16        | Randomly select a location  $j \in R$ ;
17        | Assign facility  $i$  to location  $j$ :  $\pi'(i) \leftarrow j$ ;
18      | end
19      | until  $T = \emptyset$  or  $\sigma_l \geq 0$ ;
20      |  $k \leftarrow k + 1$ ;
21   | end
22   | if  $\sigma_l < 0$  then
23     | return infeasible solution  $\pi'$ ;
24   | else
25     | return feasible solution  $\pi'$ ;
26   | end
27 end

```

Algorithm 5: Pseudo-code for `makeFeasible`: Repair procedure in path-relinking.

3.2. Parameter tuning for the GRASP-PR heuristic. Several parameters are used to describe the GRASP-PR heuristics. Some of these parameters were set to fixed values. These parameters and their fixed values are:

- size $|P|$ of elite set: 10
- minimum number ρ of elements in elite set to perform path-relinking: 2
- maximum number \bar{t} of trials in the greedy randomized construction: 10
- 1-move sample probability p_1 in the local search: 0.5
- maximum number $MaxItr$ of neighbors sampled in local search: 100
- maximum size $MaxCLS$ of the candidate list in the local search: 10
- candidate size factor η in the the non-mixed path-relinking: 0.5
- maximum number $\bar{\tau}$ of tries in `makeFeasible`: 10
- minimum difference δ with respect to elite solutions needed for trial solution to be inserted into elite set during path-relinking: 4.

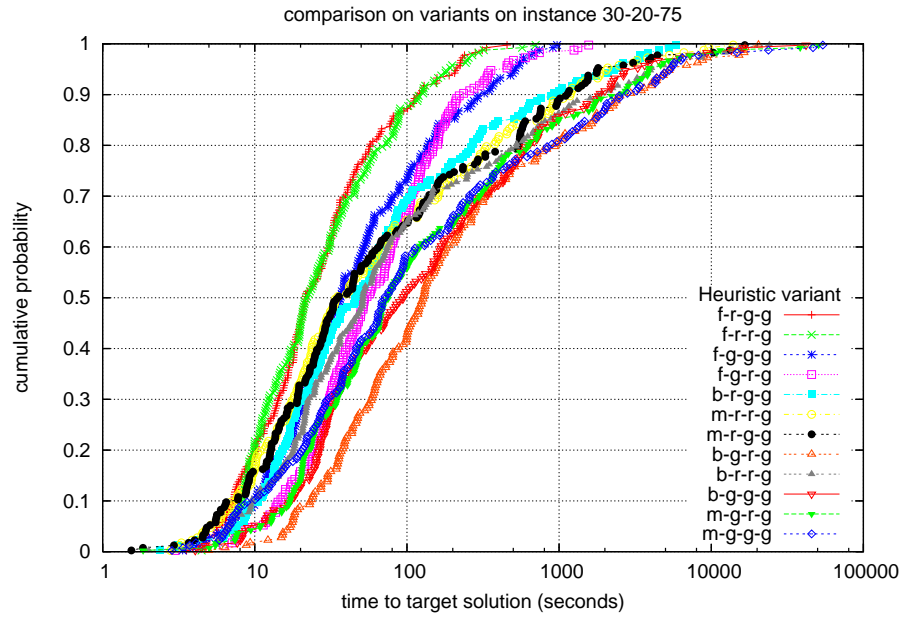
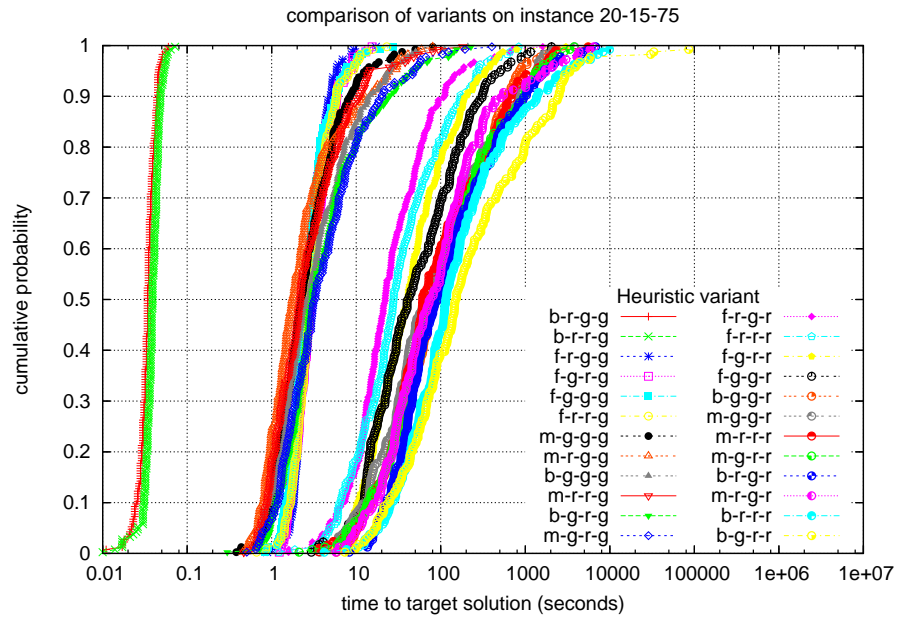


FIGURE 3. Plots of cumulative probability distributions of GRASP-PR running times for instances 20-15-75 and 30-20-75 using several variants of the GRASP with path-relinking heuristic.

To select the strategies that define the GRASP-PR variant that we use in the computational testing, we conducted the following experiment. For each combination of algorithm strategy and instance, we make 200 independent runs of the heuristic, recording the time taken to find the best known solution for the instance,

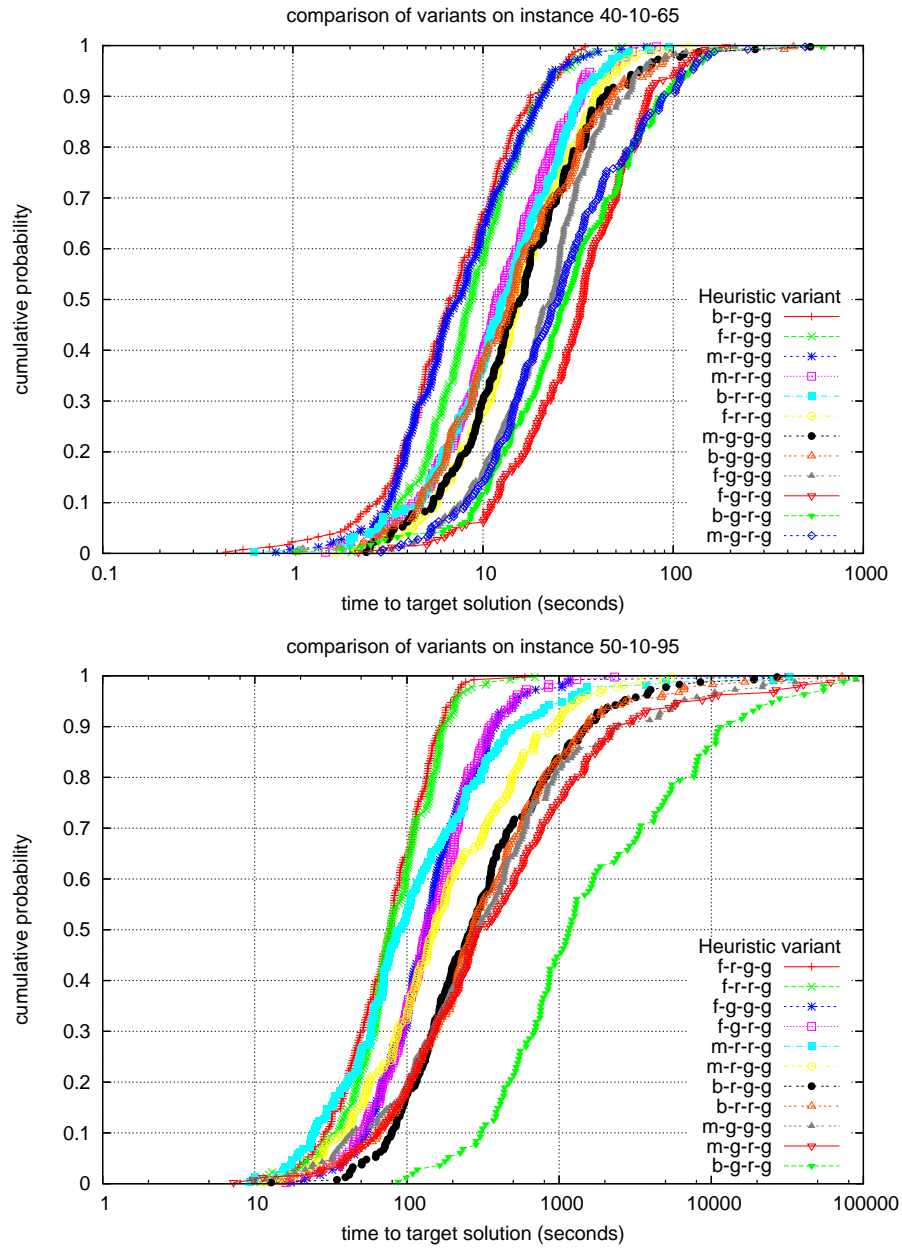


FIGURE 4. Plots of cumulative probability distributions of GRASP-PR running times for instances 40-10-65 and 50-10-95 using several variants of the GRASP with path-relinking heuristic.

and plot its runtime distribution. We consider four large instances of Cordeau et al. (2006): 20-15-75, 30-20-75, 40-10-65, and 50-10-95, and tested the GRASP-PR variants defined by the following strategies:

- Path-relinking direction: Forward (f), backward (b), mixed (m);

- Criteria to select a facility from set T in the `makeFeasible` procedure: Randomly (\mathbf{r}), such as in Algorithm 5, or greedily (\mathbf{g}), where the facility with the highest demand in T is selected;
- Criteria to select a solution from candidate list CLS in the approximate local search: Randomly (\mathbf{r}), such as in Algorithm 3, or greedily (\mathbf{g}), where the best cost solution in set CLS is selected;
- Criteria to select solution to move to in the path-relinking phase: Randomly (\mathbf{r}), as in Algorithm 4 or greedily (\mathbf{g}), where the best-valued solution in \mathcal{B} is selected.

Among the $3 \times 2^3 = 24$ possible combinations of these strategies, Figures 3 and 4 illustrate the combinations for which the GRASP-PR heuristics found the best known solution for the four test instances in less than 100,000 seconds. These figures show time-to-target plots (Aix et al., 2007) (or runtime distributions) for each algorithm variant. For each of the four instances, 200 independent runs of each variant were conducted and the time to find the target solution was recorded. The sorted running times make up the plots. Note that some of the combinations are not shown in the figures because they failed to find their target solutions within the time limit. In the figures, the strategies are represented in such a way that the first parameter corresponds to the path-relinking direction, the second parameter corresponds to the facility selection criterion, the third parameter defines the solution selection criterion, and the fourth parameter defines which path-relinking move strategy is used. For example, $\mathbf{f-g-r-g}$ represents the strategy that uses forward path-relinking, greedy facility selection, random solution selection, and greedy move in path-relinking.

Figures 3 and 4 show the dominance of strategies $\mathbf{f-r-g-g}$, $\mathbf{f-r-r-g}$, $\mathbf{b-r-g-g}$, and $\mathbf{b-r-r-g}$ over the other combinations. Considering the top two combinations for each instance, the frequencies of occurrence of each strategy among the winners were

$$\begin{bmatrix} \mathbf{f} : 5 \\ \mathbf{b} : 3 \\ \mathbf{m} : 0 \end{bmatrix} - \begin{bmatrix} \mathbf{r} : 8 \\ \mathbf{g} : 0 \end{bmatrix} - \begin{bmatrix} \mathbf{g} : 5 \\ \mathbf{r} : 3 \end{bmatrix} - \begin{bmatrix} \mathbf{g} : 8 \\ \mathbf{r} : 0 \end{bmatrix}.$$

We choose the most popular strategy ($\mathbf{f-r-g-g}$) as the preferred GRASP-PR heuristic in the remaining experiments. This heuristic uses forward path-relinking, randomly selects a facility from set T in procedure `makeFeasible`, greedily selects a solution from the candidate list in the approximate local search, and greedily selects a solution to move to in the path-relinking phase.

To conclude the tuning, we ran an experiment to determine the best local search strategy. We consider GRASP with path-relinking using four local search strategies:

- (1) use the approximate local search scheme during the GRASP and after path-relinking;
- (2) use the approximate local search scheme during the GRASP but do not do local search after path-relinking;
- (3) use first fit local search scheme during the GRASP and after path-relinking;
- (4) use best fit local search scheme during the GRASP and after path-relinking;

The first fit scheme searches the neighborhood of the current solution until it finds an improving solution. If none is found, then the current solution is declared a local minimum. The best fit scheme searches the entire neighborhood of the current




FIGURE 5. Plots of cumulative probability distributions of GRASP-PR running times for instances 20-15-35 and 40-10-65 using four variants for the local search strategy. The figures show that applying approximate local search is better than using best fit or first fit local search and that applying local search after path-relinking is better than not applying any local search after path-relinking.

solution looking for the best improving solution. If no improving solution is found, then the current solution is declared a local minimum.

For each variant, we made 200 independent runs on instances 20-15-35 and 40-10-65 of Cordeau et al. (2006). Each run was stopped when the best known solution for the instance was found and its running time was recorded. Runtime distribution plots for these runs are shown in Figure 5. The time to target plots clearly show that the approximate local search scheme is preferable to either first fit or best fit local search. As expected, first fit is shown to be better than best fit. Furthermore, the plots show that it is advisable to perform local search after each call to the path-relinking procedure.

Finally, to demonstrate the need for the repair procedure `makeFeasible` in the path-relinking procedure, we conducted the following experiment using instances 40-10-65 and 50-10-95 of Cordeau et al. (2006). For each instance, we repeated 200 independent runs of GRASP-PR using path-relinking with repair, path-relinking without repair (simply removing the `makeFeasible` from Algorithm 4 and stopping path-relinking when no further feasible movement could be made by the path-relinking procedure). For each of these two heuristics, we tested variants with and without the approximate local search after path-relinking (line 11 of Algorithm 1), resulting in a total of four variants: PR with repair and local search, PR with repair and no local search, PR without repair and with local search, and PR with neither repair nor local search. The results are summarized in the two plots in Figure 6.

The top plot in the figure shows runtime distributions for instance 40-10-65. For that experiment we used as target solution value the best known solution for this instance. The plot clearly shows that GRASP-PR with PR having both repair and posterior approximate local search dominates the other variants. In particular, using repair dominates not using repair.

The dominance of PR with repair can be better seen in the bottom plot, corresponding to the more difficult instance 50-10-95. In those runs, we first ran the four variants with the best known solution value of 12,845,598 as the target solution value in the stopping criterion. Both variants using repair (with and without posterior local search) found the solution on all 200 runs in less than 600 seconds (about 70% of the runs terminated in less than 100 seconds). On the other hand, neither variant without repair was able to find a single solution in less than 15,000 seconds. To show the relative performance of those variants we used the weak target value of 16,004,409 for the variant without repair but with local search and the even weaker value 17,7461,596 for the variant without either repair or local search. Even with these weaker target values, the two variants only found the target solution in less than 15,000 seconds in fewer than 30% of the runs, as can be seen in the lower plot in Figure 6. Using those weaker targets, the variants using repair (with and without local search) found the target solutions in less than one second on all 200 independent runs each.

3.3. Comparison of the GRASP-PR heuristic with other algorithms. In the experiments to follow, we make use of the test problems used by Lee and Ma (2005), Cordeau et al. (2006), and Hahn et al. (2007).

Among the largest instances are those presented by Cordeau et al. (2006) with 20 to 50 facilities and 6 to 20 locations. Cordeau et al. developed a memetic algorithm which found promising results for these larger instances. Table 1 compares these results with ones obtained with the GRASP-PR heuristic. The first column of the

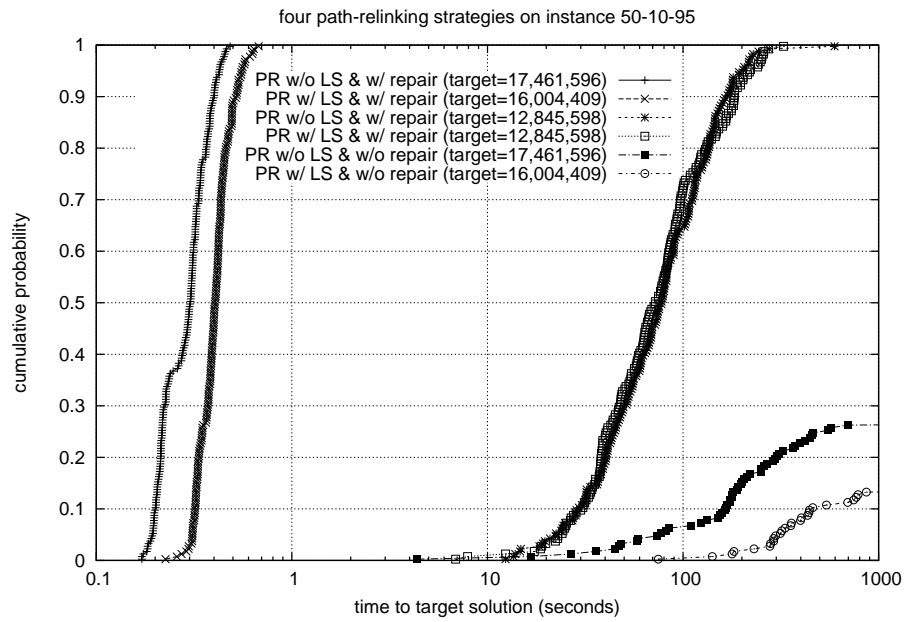
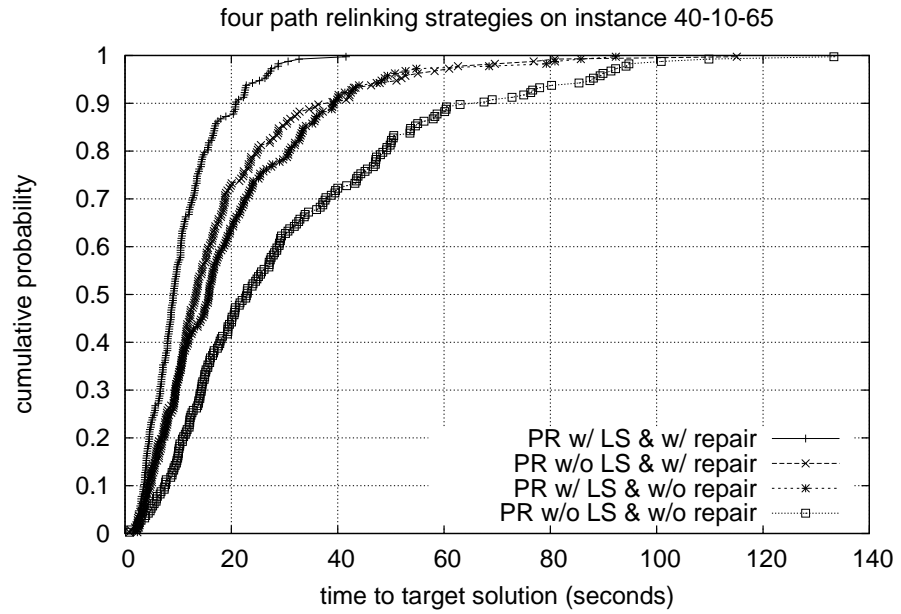


FIGURE 6. Plots of cumulative probability distributions of GRASP-PR running times for instances 40-10-65 and 50-10-95 using four variants for the path-relinking strategy. 200 independent runs of the algorithms were made and each run was stopped after 15,000 seconds. The bottom plot for instance 50-10-95 shows all measured times for six combinations of path-relinking strategy and target value. The figures show that, independent of using approximate local search after path-relinking, it is always better to use the repair procedure in path-relinking.

TABLE 1. Summary of results for memetic and GRASP-PR heuristics on the Cordeau et al. (2006) instances. Times are given in 1.2 GHz Sun workstation seconds running C for the memetic algorithm and 2.66 GHz Intel Xeon processor running Java for GRASP-PR.

instance	Memetic algorithm		GRASP-PR times to target solution				
	solution	time	min	max	avg	s.d.	0.95
20-15-35	1471896	96.0	0.16	38.87	7.09	6.48	21.04
20-15-55	1723638	102.0	0.24	14.42	2.88	2.18	7.69
20-15-75	1953188	102.0	0.26	12.82	2.02	1.72	5.25
30-06-95	5160920	114.0	0.55	23.81	2.60	2.22	6.44
30-07-75	4383923	156.0	0.50	38.47	7.86	5.47	18.18
30-08-55	3501695	96.0	0.18	4.89	1.62	0.95	3.60
30-10-65	3620959	210.0	2.75	1032.80	122.55	146.54	514.82
30-20-35	3379359	564.0	1.08	4441.40	79.43	314.14	166.21
30-20-55	3593105	462.0	1.28	150.11	25.28	21.22	66.82
30-20-75	4050938	522.0	2.11	759.81	41.64	68.66	148.43
30-20-95	5710645	5232.0	833.99	2533608.00	543019.01	747962.39	2186440.80
35-15-35	4456670	456.0	8.41	1717.94	307.65	242.74	775.25
35-15-55	4639128	384.0	4.33	75.69	21.24	11.92	42.47
35-15-75	6301723	396.0	5.18	621.83	68.57	74.39	183.19
35-15-95	6670264	864.0	6.61	19171.48	1454.00	3057.43	6949.08
40-07-75	7405793	180.0	4.53	377.06	59.66	51.29	159.00
40-09-95	7667719	1140.0	6.18	5017.56	419.08	612.62	1490.31
40-10-65	7265559	240.0	0.84	115.06	17.96	15.90	52.73
50-10-65	10513029	504.0	2.52	84.64	24.68	16.33	64.04
50-10-75	11217503	606.0	22.79	24507.34	1352.41	3085.42	4404.50
50-10-95	12845598	1254.0	9.97	1059.59	89.36	91.95	200.20

table shows the names of the instances. These instances are labeled with three parameters: $n = |N|$, $m = |M|$, and $s \in \{1, \dots, 100\}$. Parameter s controls the tightness of the capacity constraints. The higher the value of s , the higher is the tightness of the capacity constraints.

The third column lists the time (in seconds) reported by Cordeau et al. (2006) for their memetic algorithm to find the solutions whose values are given in the second column. It should be noted that Cordeau et al. implemented their memetic algorithm in C and ran their experiment on a 1.2 GHz Sun workstation. The solution values in the second column are the best values found by the memetic algorithm.

For each instance, we made 200 independent runs of GRASP-PR, with the exception of instance 30-20-95 for which we made only 18 runs. Each run stopped when a solution value as good as the one in column 2 was found. Columns 4 to 7 give the minimum, maximum, and average times, as well as the standard deviation of these runs to find solutions with values equal to those listed in column 2. Finally,

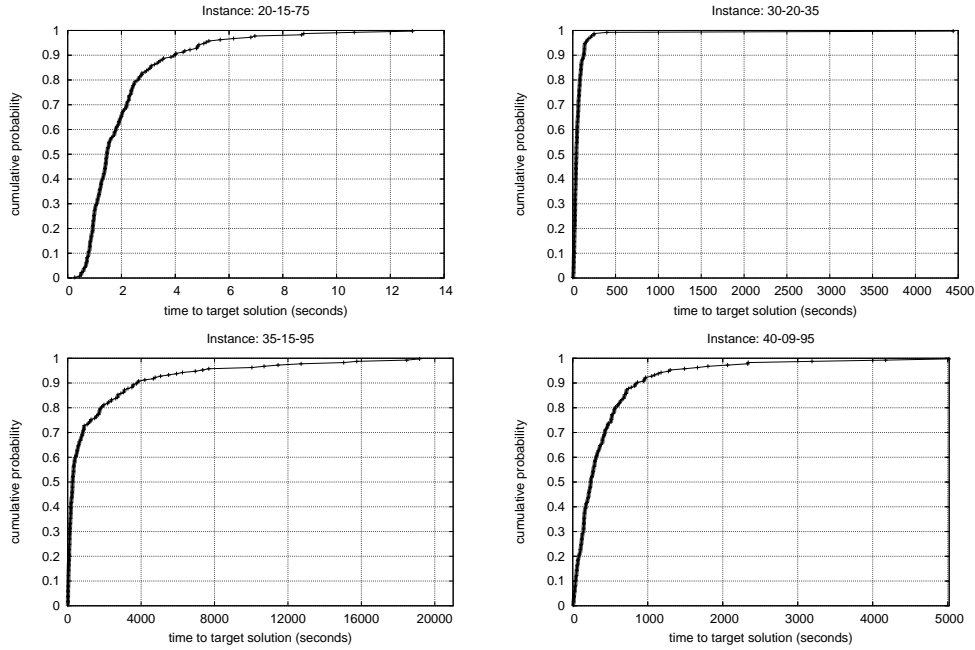


FIGURE 7. Time-to-target plots. Plots of cumulative probability distributions of GRASP-PR running times for four instances: 20-15-75, 30-20-35, 35-15-95, and 40-09-95.

column 8 lists the time for 95% of the GRASP-PR runs to find solutions having those same target values.

On all runs for each instance in Table 1, the GRASP-PR heuristic found solutions with the target values shown in column 2. The average performance improvement with respect to the memetic algorithm varied between a factor of 1.482 and 59.186, except for instances 30-20-95, 35-15-95, and 50-10-75, for which the times reported by Cordeau et al. were smaller than the average times for the GRASP-PR heuristic. Note, however, that on all instances, including 30-20-95, 35-15-95, and 50-10-75, the smallest GRASP-PR running times were far less than those reported for the memetic algorithm. Figure 7 shows time-to-target-plots (Aiex et al., 2007) for GRASP-PR on four instances from Table 1. For each group with identical number of facilities, we chose to plot the running-time distribution for the instance on which the memetic algorithm took longest to solve.

To optimally solve the GQAP, Lee and Ma (2005) proposed a branch and bound (B&B) algorithm having single assignment as the branching rule, as done by Burkard (1991) for the quadratic assignment problem (QAP). Furthermore, they proposed three linearization methods, F-Y, K-B, and L3, based on the Frieze and Yadegar (1983), Kaufman and Broeckx (1978), and Padberg and Rijal (1996) linearizations for the QAP, respectively. Lee and Ma created a suite of test problems with 10 to 16 facilities and 3 to 8 locations on which they tested their algorithms. Table 2 compares these results with ones obtained with the GRASP-PR heuristic.

The first and second columns of Table 2 show the instance names and their corresponding optimum values. These instances are labeled in three parts: the

TABLE 2. Summary of results for the four algorithms of Lee and Ma (2005) and the GRASP-PR heuristic on the Lee and Ma (2005) instances. Times in seconds are measured on a 677 MHz Intel Pentium III running AMPL/CPLEX 8.1 for the Lee and Ma (2005) algorithms and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	solution	Lee and Ma (2005) times				GRASP-PR time to target solution				
		F-Y	K-B	L3	B&B	min	max	avg	s.d.	0.95
$10 \times 5EG$	1021050	39.2	12.2	11.5	161.5	0.031	0.663	0.208	0.118	0.466
$10 \times 5EH$	1188500	49.8	21.4	20.4	361.1	0.038	0.459	0.186	0.096	0.386
$10 \times 5FG$	1433500	51.6	5.2	21.3	196.2	0.040	0.785	0.282	0.148	0.543
$10 \times 5FH$	1471950	96.2	26.5	17.1	723.9	0.039	0.559	0.244	0.124	0.467
$12 \times 3EG$	1323900	5.4	2.0	2.1	72.1	0.039	0.417	0.147	0.070	0.263
$12 \times 3EH$	1771450	17.6	8.4	5.7	418.3	0.038	0.622	0.218	0.112	0.453
$12 \times 3FG$	1788750	19.6	6.3	6.0	251.4	0.032	0.474	0.137	0.078	0.304
$12 \times 3FH$	1970500	17.3	14.8	6.0	359.2	0.039	0.433	0.149	0.090	0.345
$12 \times 4EG$	1500150	31.8	11.2	11.3	202.6	0.038	0.995	0.353	0.181	0.689
$12 \times 4EH$	2081370	102.9	29.6	30.0	626.5	0.045	0.717	0.329	0.134	0.564
$12 \times 4FG$	1770500	88.5	55.3	20.5	925.3	0.046	0.790	0.327	0.164	0.614
$12 \times 4FH$	2213500	82.0	5.9	17.8	251.4	0.072	1.117	0.527	0.206	0.858
$12 \times 5EG$	1519100	119.1	65.4	32.3	565.3	0.054	1.264	0.285	0.153	0.517
$12 \times 5EH$	2086100	567.9	20.6	59.3	510.6	0.023	1.399	0.608	0.217	0.951
12×6	1960750	651.5	18.2	229.7	641.9	0.086	1.110	0.529	0.205	0.845
12×8	1244300	1091.0	1333.9	278.2	1445.8	0.079	1.060	0.405	0.217	0.804
14×6	2351000	408.6	52.5	78.3	1030.0	0.068	1.295	0.636	0.249	1.083
14×7	2365650	1670.2	172.1	315.0	2459.1	0.092	1.563	0.560	0.191	0.811
14×8	2622470	2717.6	256.4	275.3	2034.5	0.073	0.936	0.398	0.182	0.746
14×9	2326370	3759.1	7881.7	483.2	3278.9	0.057	0.919	0.481	0.157	0.748
15×6	2707850	1614.1	109.4	284.8	2160.1	0.109	7.769	1.015	0.838	2.084
15×7	2720070	1971.3	658.7	296.2	3119.2	0.201	2.090	0.802	0.372	1.653
15×8	2856850	4056.1	397.5	1242.1	5007.9	0.414	21.727	2.348	2.582	8.032
16×6	2709300	1815.7	630.3	338.4	4884.4	0.121	0.934	0.460	0.146	0.730
16×7	2809870	22423.7	988.7	7208.4	5921.9	0.472	13.962	1.761	1.580	4.431

TABLE 3. Summary of results for RLT1_B&B algorithm and GRASP-PR heuristic. Times in seconds are measured on a 733 MHz Intel Itanium processor (Intel Fortran 8.0 compiler) for the RLT1_B&B algorithm and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	solution	RLT1_B&B algorithm		GRASP-PR times to target solution				
		tot. time	time to bks	min	max	avg	s.d.	0.95
2005De	2302	17640.0	14390.0	0.091	6.304	1.465	0.896	3.302
2005Aa	3059	136.0	128.0	0.829	25.745	7.051	4.756	19.268
2408Ca	1028	6.3	3.3	0.149	2.307	0.714	0.447	1.594
2408Aa	5643	719862.0	185726.0	1.604	50.445	10.288	8.369	29.735
20-15-35	1471896	735.0	528.8	0.160	38.870	7.090	6.480	21.040
$14 \times 9EG$	2326370	169.0	90.7	0.057	0.919	0.481	0.157	0.748
15×8	2856850	130.0	39.5	0.414	21.727	2.348	2.582	8.032
$16 \times 7-1115$	2809870	548.0	512.6	0.472	13.962	1.761	1.580	4.431

number before the \times is $n = |N|$, the number after the \times is $m = |M|$, and the optional last two letters indicate the following additional characteristics: facilities having a wide range of demands (E), facilities having roughly the same demand (F), locations having a wide range of capacities (G), locations having roughly the same capacity (H). Columns 3 to 6 list the times (in seconds) reported by Lee and Ma (2005) for their four algorithms (F-Y, K-B, L3, and B&B) to find solutions whose values are given in the second column. Times in boldface are the best among the Lee and Ma algorithms.

For each instance, we made 200 independent runs of GRASP-PR. Each run stopped when a solution value as good as the one in column 2 was found. Columns 7 to 10 denote, respectively, the minimum, maximum, and average times, and standard deviation of these runs to find solutions with values equal to those of column 2. Finally, column 11 lists the time for 95% of the runs to find solutions having those same target values. GRASP-PR found the target value on all 200 runs for each of the instances in Table 2, with an average performance improvement varying between a factor of 11.2 and 1004.6 in relation to the best time of the algorithms of Lee and Ma.

Hahn et al. (2007) proposed a level-1 reformulation-linearization technique (RLT) dual ascent procedure in a branch-and-bound scheme (RLT1_B&B) for the GQAP. They solved four instances (c2005De, 2005Aa, 2408Ca, and 2408Aa) from the website² of S. Elloumi, one instance (20-15-35) from Cordeau et al. (2006), and three instances ($14 \times 9EG$, 15×8 , and $16 \times 7-1115$) from Lee and Ma (2005), as shown in the first column of Table 3. That table compares the results of the RLT1_B&B algorithm with ones obtained with the GRASP-PR heuristic. Column 3 lists the time (in seconds on a 733 MHz Intel Itanium processor) reported by Hahn et al. (2007) for RLT1_B&B to find the optimum solutions with the values given in column 2.

²http://cedric.cnam.fr/oc/TAP/all_CTAP_instances/

For each instance, we made 200 independent runs of the GRASP-PR heuristic. Each of these runs stopped when the optimal solution value was found. Columns 5 to 8 denote, respectively, the minimum, maximum, and average times, and standard deviation of the time to find optimal solutions. Finally, column 9 lists the times needed for 95% of GRASP-PR runs to find optimally-valued solutions.

On all runs, GRASP-PR achieved the target values (optimal solutions) on all instances in Table 3. The average performance improvement, with respect to the RLT1_B&B algorithm, varied between a factor of 8.8 and over 69 thousand. With respect to the time taken by RLT1_B&B to first find an optimally-valued solution, the improvement of GRASP-PR varied from 4.6 and over 18 thousand.

Pessoa et al. (2008) proposed a new lower bound for the GQAP based on the Lagrangian relaxation of the RLT formulation proposed in Hahn et al. (2007). As shown in the first column of Table 4, Pessoa et al. tested their approach on four instances (c2005De, 2005Aa, 2408Ca, and 2408Aa) from the website of S. Elloumi, 12 instances (20-15-35, 20-15-55, 20-15-75, 30-08-55, 30-20-35, 30-20-55, 30-20-75, 30-20-95, 30-15-35, 35-15-55, 35-15-75, and 35-15-95) from Cordeau et al. (2006), and three instances ($14 \times 9EG$, 15×8 , and $16 \times 17-1115$) from Lee and Ma (2005). Table 4 compares these results with ones obtained with the GRASP-PR heuristic. The column 3 lists the time (in seconds on a 1.6 GHz Ultra 45 Sun workstation) reported by Pessoa et al. (2008) for their hybrid B&B algorithm to find the optimal solutions and prove their optimality, using the optimum values given in column 2 as upper bounds.

For each instance, we made 200 independent runs of the GRASP-PR heuristic. Each of these runs stopped when an optimally-valued solution (column 2) was found. Columns 4 to 6 denote the minimum, maximum, and average times, and standard deviation of the time needed to find the optimally-valued solutions. Finally, column 8 gives the time needed for 95% of these runs to find solutions having those same target values. The GRASP-PR heuristic achieved the target values on all instances in Table 4, with an average performance improvement, with respect to the hybrid B&B algorithm, varying between a factor of 132.7 and over 100 thousand, with the exception of instance 30-20-95. The smallest GRASP-PR running time for instance 30-20-95, however, was less than the running time reported by Pessoa et al. (2008) for the hybrid B&B algorithm.

Elloumi et al. (2003) proposed three linearization methods (L1, L2, and L3), three semidefinite programming formulations (S0, S1, and S2), and a Lagrangian decomposition (D0) algorithm for the *constrained module allocation problem (CMAP)*, or *constrained task assignment problem (CTAP)*. The CTAP is a special case of the GQAP, where a slightly different cost function is used. For these experiments, we adopted this cost function to be able to compare the GRASP-PR heuristic with their algorithms. Elloumi et al. used the instances introduced in Elloumi (1991) and also used in Roupin (2004). In these instances, four configurations (A, B, C, and D) are considered. For each configuration, two classes of instances are generated: a class with a complete communication graph and a second class where the density of the communication graph is 50%, resulting in a total of eight types of instances. For each type, five instances (a-e), with 10 facilities and three locations, and five instances (a-e), with 20 facilities and five locations, are generated. These instances were downloaded from a website of S. Elloumi. Only a subset of these instances have their optimal values available in the website. Therefore, we decided

TABLE 4. Summary of results for Hybrid_B&B algorithm and GRASP-PR heuristic. Times are given in seconds on a 1.6 GHz Ultra 45 Sun workstation (C++ compiler) for the Hybrid_B&B algorithm and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	Hybrid_B&B algorithm		GRASP-PR times				
	opt soln	time	min	max	avg	s.d.	0.95
c2005De	5435	10085.0	0.083	16.764	2.140	1.747	4.965
2005Aa	3059	3018.0	0.829	25.745	7.051	4.756	19.268
2408Ca	1028	117.0	0.149	2.307	0.714	0.447	1.594
2408Aa	5643	1031637.0	1.604	50.445	10.288	8.369	29.735
20-15-35	1471896	2729.0	0.16	38.87	7.09	6.48	21.04
20-15-55	1723638	4370.0	0.24	14.42	2.88	2.18	7.69
20-15-75	1953188	1340.0	0.26	12.82	2.02	1.72	5.25
30-08-55	3501695	2826.0	0.18	4.89	1.62	0.95	3.60
30-20-35	3379359	228177.0	1.08	4441.40	79.43	314.14	166.21
30-20-55	3593105	707027.0	1.28	150.11	25.28	21.22	66.82
30-20-75	4050938	18803.0	2.11	759.81	41.64	68.66	148.43
30-20-95	5710645	2612.0	833.99	2533608.00	543019.01	747962.39	2186440.80
35-15-35	4456670	978664.0	8.41	1717.94	307.65	242.74	775.25
35-15-55	4639128	413241.0	4.33	75.69	21.24	11.92	42.47
35-15-95	6670264	205931.0	14.40	18470.45	1551.78	3533.59	7356.88
14 × 9	2326370	139.0	0.057	0.919	0.481	0.157	0.748
15 × 8	2856850	359.0	0.414	21.727	2.348	2.582	8.032
16 × 7	2809870	311.0	0.472	13.962	1.761	1.580	4.431

to divide the instances into two tables. Table 5 compares the running times of the algorithms of Elloumi et al. with those of GRASP-PR on the CTAP instances with known optimal values listed in the website of S. Elloumi. Tables 6–8 list the other instances, jointly with the best solution obtained with the GRASP-PR heuristic. We computed these values because they could not be found in the website. The tables list the average CPU times, and in parentheses, their corresponding average errors with the best known solutions for the seven bounds of Elloumi et al. (2003).

In Tables 5–8, the first and second columns show the instance names and their corresponding optimum or best known values. These instances are labeled in five parts: a letter “c”, if the communication graph is complete, $n = |N|$, $m = |M|$, the configuration (A, B, C, or D), and the instance identifier (a, b, c, d, or e).

Columns 3 to 9 list the average time (in seconds) reported by Elloumi et al. (2003) per instance category for their seven methods (L1, L2, L3, S0, S1, S2, and D0) to find solutions whose values are given in the second column. Least mean error values are indicated by boldface with ties broken by running time.

For each instance, we made 200 independent runs of the GRASP-PR heuristic. Each run stopped when a solution value equal to one in column 2 was found. Columns 10 to 13 denote, respectively, the minimum, maximum, and average times, and standard deviation of these times to find solutions with values equal to those of column 2. Finally, column 14 lists the times taken by the faster 95% of these 200 runs.

The GRASP-PR heuristic found the target values on all 200 runs for each instance in Tables 5–8, with an improvement varying between a factor of 7.3 and over five thousand with respect to the best method of Elloumi et al. (2003). For each problem group, the improvement was computed by taking the mean value of the average GRASP-PR running times for that group with respect to the running times in boldface.

4. CONCLUDING REMARKS

In this paper, we propose a new GRASP with path-relinking heuristic for the generalized quadratic assignment problem. To better explore the solution space, this heuristic incorporates randomization in three novel components: a construction procedure, a local search procedure, and a path-relinking procedure with repair.

With the exception of the construction procedure, randomization is not common in the components of a GRASP heuristic. Randomization in construction is usually confined to the selection of a solution element from the restricted candidate list. Similarly, path-relinking procedures do not usually incorporate randomization. A well-known exception is the greedy randomized adaptive path-relinking variant of Faria Jr. et al. (2005).

In addition to enabling better exploration of the solution space, randomization is useful to help deal with infeasibilities. The randomized construction steps are repeated until either a feasible solution is constructed or a maximum number of repetitions is reached. The main role of randomization in the path-relinking phase is to enable the repair of infeasible solutions that can be produced when an attribute of the guiding solution is incorporated in the current solution.

Randomization in the local search is used to sample moves from those leading to an improving solution. Since the entire neighborhood is not evaluated, this sampling speeds up the search without compromising its effectiveness. Since not all

TABLE 5. Summary (Part I) of results for the seven algorithms of Elloumi et al. (2003) and the GRASP-PR heuristic on the instances used in Elloumi et al. (2003). Times in seconds are measured on a 400 MHz Pentium II for the algorithms of Elloumi et al. (2003) and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	bks	Algs. of Elloumi et al.: avg time in secs (avg error = (bks - lower)/bks)							GRASP-PR times to bks				
		L1	L2	L3	S0	S1	S2	D0	min	max	avg	s.d.	0.95
c1003Aa	1616								0.038	0.927	0.192	0.140	0.518
c1003Ab	1390								0.035	0.991	0.211	0.154	0.505
c1003Ac	1730	1(68%)	1(12%)	1(23%)	6(69%)	948(9%)	884(3%)	30(15%)	0.037	0.609	0.205	0.133	0.482
c1003Ad	1289								0.031	0.535	0.170	0.114	0.419
c1003Ae	1048								0.040	0.996	0.323	0.181	0.670
c1003Ba	1299								0.023	0.815	0.210	0.127	0.505
c1003Bb	865								0.041	0.838	0.333	0.194	0.684
c1003Bc	1154	1(94%)	1(29%)	1(37%)	144(94%)	998(26%)	890(12%)	25(33%)	0.032	1.142	0.337	0.222	0.880
c1003Bd	834								0.034	0.768	0.250	0.151	0.546
c1003Be	812								0.031	0.668	0.189	0.126	0.462
c1003Ca	455								0.033	0.516	0.186	0.117	0.426
c1003Cb	467								0.033	1.089	0.115	0.099	0.301
c1003Cc	475	1(4%)	1(1%)	1(4%)	3(11%)	96(1%)	274(1%)	19(1%)	0.034	0.558	0.166	0.110	0.408
c1003Cd	472								0.038	0.340	0.105	0.075	0.285
c1003Ce	350								0.033	0.349	0.113	0.091	0.298
c1003Da	843								0.022	0.756	0.186	0.134	0.470
c1003Db	879								0.032	0.662	0.151	0.108	0.417
c1003Dc	1230	1(100%)	1(25%)	1(39%)	2(100%)	798(22%)	568(8%)	27(32%)	0.032	0.606	0.197	0.122	0.477
c1003Dd	956								0.029	0.494	0.186	0.114	0.435
c1003De	848								0.036	0.618	0.200	0.125	0.478

TABLE 6. Summary (Part II) of results for the seven algorithms of Elloumi et al. (2003) and the GRASP-PR heuristic on the instances used in Elloumi et al. (2003). Times in seconds are measured on a 400 MHz Pentium II for the algorithms of Elloumi et al. (2003) and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	bks	Algs. of Elloumi et al.: avg time in secs (avg error = (bks - lower)/bks)							GRASP-PR times to bks				
		L1	L2	L3	S0	S1	S2	D0	min	max	avg	s.d.	0.95
1003Aa	731								0.037	0.657	0.169	0.112	0.441
1003Ab	713								0.020	0.522	0.115	0.084	0.318
1003Ac	645	1(37%)	1(21%)	1(27%)	4(41%)	575(12%)	458(5%)	26(25%)	0.034	0.570	0.158	0.113	0.393
1003Ad	688								0.030	0.268	0.071	0.031	0.128
1003Ae	715								0.018	0.158	0.063	0.022	0.109
1003Ba	306								0.021	0.921	0.220	0.146	0.513
1003Bb	528								0.043	0.710	0.213	0.141	0.487
1003Bc	326	1(85%)	1(48%)	1(57%)	12(85%)	860(29%)	898(11%)	25(57%)	0.029	0.682	0.204	0.143	0.517
1003Bd	364								0.022	0.466	0.147	0.083	0.295
1003Be	324								0.018	0.245	0.090	0.041	0.165
1003Ca	346								0.031	0.498	0.110	0.086	0.330
1003Cb	424								0.040	0.861	0.181	0.134	0.465
1003Cc	347	1(2%)	1(1%)	1(3%)	4(4%)	189(2%)	156(1%)	19(1%)	0.033	0.409	0.075	0.056	0.230
1003Cd	434								0.029	0.114	0.061	0.018	0.097
1003Ce	285								0.027	0.052	0.040	0.005	0.048
1003Da	219								0.018	1.009	0.266	0.188	0.629
1003Db	402								0.035	1.495	0.267	0.193	0.619
1003Dc	297	1(100%)	1(43%)	1(52%)	3(100%)	1631(27%)	855(13%)	25(55%)	0.031	0.734	0.180	0.130	0.431
1003Dd	445								0.022	0.737	0.163	0.119	0.398
1003De	358								0.028	0.811	0.076	0.062	0.147

TABLE 7. Summary (Part III) of results for the seven algorithms of Elloumi et al. (2003) and the GRASP-PR heuristic on the instances used in Elloumi et al. (2003). Times in seconds are measured on a 400 MHz Pentium II for the algorithms of Elloumi et al. (2003) and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	bks	Algs. of Elloumi et al.: avg time in secs (avg error = (bks - lower)/bks)							GRASP-PR times to bks				
		L1	L2	L3	S0	S1	S2	D0	min	max	avg	s.d.	0.95
2005Aa	3059								0.829	25.745	7.051	4.756	19.268
2005Ab	2954								0.247	6.902	2.139	1.179	4.585
2005Ac	3012	1(85%)	82(12%)	4(27%)	59(85%)	1800(10%)	1800(2%)	238(19%)	0.441	6.160	2.142	0.978	3.961
2005Ad	3174								0.232	4.648	1.423	0.637	2.510
2005Ae	3054								0.162	9.210	1.774	1.318	4.570
2005Ba	2442								1.042	58.466	19.982	13.163	49.609
2005Bb	2088								0.371	44.036	4.916	5.536	17.237
2005Bc	1986	1(98%)	75(17%)	4(32%)	600(98%)	1800(16%)	1800(3%)	223 (26%)	0.455	13.507	3.705	2.512	9.428
2005Bd	2449								0.422	8.548	2.562	1.657	6.331
2005Be	2453								0.358	8.042	1.931	1.116	4.177
2005Ca	783								0.082	1.062	0.313	0.169	0.686
2005Cb	636								0.058	1.320	0.397	0.223	0.843
2005Cc	772	1(19%)	78(1%)	4(10%)	71(31%)	1800(1%)	1800(0%)	239.8 (1%)	0.100	2.054	0.539	0.371	1.387
2005Cd	682								0.056	1.041	0.143	0.079	0.239
2005Ce	732								0.058	0.403	0.164	0.064	0.300
2005Da	2413								0.586	48.446	8.282	7.880	25.535
2005Db	2316								0.569	15.026	3.771	2.440	8.646
2005Dc	1965	1(100%)	66(18%)	4(30%)	34(100%)	1800(14%)	1800(3%)	206(24%)	0.067	5.591	1.583	0.951	3.343
2005Dd	2211								0.084	2.822	0.800	0.369	1.402
2005De	2302								0.091	6.304	1.465	0.896	3.302

GRASP WITH PATH-RELINKING FOR THE GQAP

TABLE 8. Summary (Part IV) of results for the seven algorithms of Elloumi et al. (2003) and the GRASP-PR heuristic on the instances used in Elloumi et al. (2003). Times in seconds are measured on a 400 MHz Pentium II for the algorithms of Elloumi et al. (2003) and on a 2.66 GHz Intel Xeon running Java for GRASP-PR.

instance	bks	Algs. of Elloumi et al.: avg time in secs (avg error = (bks - lower)/bks)							GRASP-PR times to bks				
		L1	L2	L3	S0	S1	S2	D0	min	max	avg	s.d.	0.95
2408Aa	5643	Instances in http://cedric.cnam.fr/oc/TAP/all_CTAP_instances/ but not in Elloumi et al. (2003).							1.604	50.445	10.288	8.369	29.735
2408Ca	1028								0.149	2.307	0.714	0.447	1.594
c2005Aa	6412	1(69%) 72(29%) 4(40%) 85(70%) 1800(12%) 1800(3%) 214(35%)							1.209	18.117	5.307	3.276	13.211
c2005Ab	6260								0.253	15.489	3.141	2.046	7.442
c2005Ac	6491								0.346	20.114	4.488	3.146	11.287
c2005Ad	6267								0.237	8.763	1.944	1.085	3.982
c2005Ae	6194								0.185	11.353	2.938	2.133	7.264
c2005Ba	5420								1(98%) 81(17%) 4(32%) 199(98%) 1800(16%) 1800(3%) 156 (26%)				
c2005Bb	5370	0.665	24.880	4.455	3.293	11.138							
c2005Bc	5645	1.163	40.923	9.494	6.926	25.171							
c2005Bd	5420	0.260	22.465	3.362	3.032	9.235							
c2005Be	5836	0.725	25.527	3.920	3.164	10.125							
c2005Ca	1181	1(2%) 69(0%) 4(7%) 58(12%) 1800(1%) 1800(0%) 228(1%)											
c2005Cb	1017								0.071	3.385	0.379	0.287	0.746
c2005Cc	1197								0.110	3.558	1.001	0.811	3.040
c2005Cd	1038								0.061	0.402	0.172	0.060	0.295
c2005Ce	1166								0.057	0.564	0.196	0.084	0.387
c2005Da	5139								1(100%) 69(56%) 4(63%) 34(100%) 1800(20%) 1800(9%) 127 (66%)				
c2005Db	5519	0.548	29.663	4.848	3.940	12.613							
c2005Dc	5907	0.834	55.999	9.995	9.810	36.641							
c2005Dd	5494	0.505	3.977	1.502	0.695	2.878							
c2005De	5435	0.083	16.764	2.140	1.747	4.965							

neighborhood solutions are evaluated and the procedure terminates when no improving solution is sampled, the resulting solution may not be a true local minimum. We call such solutions approximate local minima.

The algorithm was implemented in Java and tested extensively on a set of 128 instances available in the literature (Cordeau et al., 2006; Lee and Ma, 2005; Elloumi et al., 2003). Previous papers on this subject only considered small subsets of these instances. For each instance (except 30-20-95 of Cordeau et al. (2006)), 200 independent runs of the heuristic were done.³ Each run terminated when the best-known solution for that instance was found. For all instances, the heuristic found the best-known solution on all runs, demonstrating its robustness. For all instances, except 30-20-95 of Cordeau et al. (2006), the average running time of the GRASP with path-relinking heuristic was never greater than the running time reported for the other heuristics in the literature. In fact, our average running times were often orders of magnitude smaller than those reported in the literature. Even on instance 30-20-95, where we measured the longest running times, the fastest of the 18 solution times of the GRASP with path-relinking heuristic implemented in Java was 834 seconds on a 2.66 GHz processor while the time reported for the C implementation by Cordeau et al. (2006) was 5232 seconds on a 1.2 GHz processor. Taking into account the effects of processor speed and programming languages, we believe that the heuristic proposed in this paper is competitive with the algorithms proposed in the literature for solving the generalized quadratic assignment problem.

REFERENCES

- R.M. Aiex, S. Binato, and M.G.C. Resende. Parallel GRASP with path-relinking for job shop scheduling. *Parallel Computing*, 29:393–430, 2003.
- R.M. Aiex, P.M. Pardalos, M.G.C. Resende, and G. Toraldo. GRASP with path-relinking for three-index assignment. *INFORMS J. on Computing*, 17:224–247, 2005.
- R.M. Aiex, M.G.C. Resende, and C.C. Ribeiro. TTTPLOTS: A perl program to create time-to-target plots. *Optimization Letters*, 1:201–212, 2007.
- F. Barahona and R. Anbil. The volume algorithm: producing primal solutions with a subgradient method. *Mathematical Programming*, 87:385–399, 2000.
- R. E. Burkard. Locations with Spatial Interactions: The Quadratic Assignment Problem. In *Discrete Location Theory*, pages 89–98. Wiley, New York, 1991.
- J.-F. Cordeau, M. Gaudioso, G. Laporte, and L. Moccia. A memetic heuristic for the generalized quadratic assignment problem. *INFORMS Journal on Computing*, 18:433–443, 2006.
- S. Elloumi. Contribution à la résolution des programmes non linéaires en variables 0-1, application aux problèmes de placement de tâches dans les systèmes distribués (thèse de doctorat in informatique). Technical report, Conservatoire National des Arts et Métiers, Paris, 1991.
- S. Elloumi, F. Roupin, and E. Soutif. Comparison of different lower bounds for the constrained module allocation problem. Technical Report 473, CNAM-Laboratoire Cédric, 292 Rue St Martin, FR-75141 Paris Cedex 03 France, 2003.
- H. Faria Jr., S. Binato, M. G. C. Resende, and D. J. Falção. Transmission network design by a greedy randomized adaptive path relinking approach. *IEEE Transactions on Power Systems*, 20(1):43–49, 2005.

³Because of long running times, on instance 30-20-95, only 18 independent runs were done.

- T.A. Feo and M.G.C. Resende. A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters*, 8:67–71, 1989.
- T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. *J. of Global Optimization*, 6:109–133, 1995.
- P. Festa and M.G.C. Resende. GRASP: An annotated bibliography. In C.C. Ribeiro and P. Hansen, editors, *Essays and Surveys on Metaheuristics*, pages 325–367. Kluwer Academic Publishers, 2002.
- P. Festa and M.G.C. Resende. An annotated bibliography of GRASP – Part I: Algorithms. *International Transactions on Operational Research*, 16:1–24, 2009a.
- P. Festa and M.G.C. Resende. An annotated bibliography of GRASP – Part II: Applications. *International Transactions on Operational Research*, 2009b. In press.
- A. Frieze and J. Yadegar. On the quadratic assignment problem. *Discrete Applied Mathematics*, 5:89–98, 1983.
- F. Glover. Tabu search and adaptive memory programming – Advances, applications and challenges. In R.S. Barr, R.V. Helgason, and J.L. Kennington, editors, *Interfaces in Computer Science and Operations Research*, pages 1–75. Kluwer, 1996.
- F. Glover, M. Laguna, and R. Martí. Fundamentals of scatter search and path relinking. *Control and Cybernetics*, 39:653–684, 2000.
- P.M. Hahn, B.-J. Kim, M. Guignard, J. MacGregor Smith, and Y.-R. Zhu. An algorithm for the generalized quadratic assignment problem. *Computational Optimization and Applications*, 2007. Published online November 14, 2007, DOI: 10.1007/s10589-007-9093-1.
- L. Kaufman and F. Broeckx. An algorithm for the quadratic assignment problem using benders decomposition. *European Journal of Operational Research*, 2:204–211, 1978.
- M. Laguna and R. Martí. GRASP and path relinking for 2-layer straight line crossing minimization. *INFORMS Journal on Computing*, 11:44–52, 1999.
- C.-G. Lee and Z. Ma. The generalized quadratic assignment problem. Technical Report MIEOR TR2005-01, Department of Mechanical and Industrial Engineering at the University of Toronto, 2005.
- Y. Li, P.M. Pardalos, and M.G.C. Resende. A greedy randomized adaptive search procedure for the quadratic assignment problem. In P.M. Pardalos and H. Wolkowicz, editors, *Quadratic assignment and related problems*, volume 16 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 237–261. American Mathematical Society, New Providence, RI, 1994.
- M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation*, 8:3–30, 1998.
- C.A.S. Oliveira, P.M. Pardalos, and M.G.C. Resende. GRASP with path-relinking for the quadratic assignment problem. In C.C. Ribeiro and S.L. Martins, editors, *Efficient and Experimental Algorithms*, volume 3059 of *Lecture Notes in Computer Science*, pages 356–368. Springer-Verlag, 2004.
- M. Padberg and M. Rijal. *Location, Scheduling, Design and Integer Programming*. Kluwer Academic Publishers, 1996.
- P.M. Pardalos, F. Rendl, and H. Wolkowicz. The quadratic assignment problem: A survey and recent development. In P.M. Pardalos and H. Wolkowicz, editors, *The*

- Quadratic Assignment and Related Problems*, volume 16 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 1–42. American Mathematical Society, New Providence, RI, 1994.
- A.A. Pessoa, P.M. Hahn, M. Guignard, and Y.-R. Zhu. An improved algorithm for the generalized quadratic assignment problem. Technical Report MIEOR TR2005-01, Electrical and Systems Engineering at the University of Pennsylvania, 2008.
- M.G.C. Resende and C.C. Ribeiro. Greedy randomized adaptive search procedures. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, pages 219–249. Kluwer Academic Publishers, 2002.
- M.G.C. Resende and C.C. Ribeiro. A GRASP with path-relinking for private virtual circuit routing. *Networks*, 41(1):104–114, 2003.
- M.G.C. Resende and C.C. Ribeiro. GRASP with path-relinking: Recent advances and applications. In T. Ibaraki, K. Nonobe, and M. Yagiura, editors, *Metaheuristics: Progress as Real Problem Solvers*, pages 29–63. Springer, 2005.
- M.G.C. Resende and C.C. Ribeiro. Greedy randomized adaptive search procedures: Advances and applications. In M. Gendreau and J.-Y. Potvin, editors, *Handbook of Metaheuristics*. Springer Science+Business Media, 2nd edition, 2010.
- C.C. Ribeiro and I. Rosseti. A parallel GRASP for the 2-path network design problem. *Lecture Notes in Computer Science*, 2004:922–926, 2002.
- C.C. Ribeiro, E. Uchoa, and R.F. Werneck. A hybrid grasp with perturbations for the steiner problem in graphs. *INFORMS Journal on Computing*, 14:228–246, 2002.
- F. Roupin. From linear to semidefinite programming: an algorithm to obtain semidefinite relaxations for bivalent quadratic problems. *Journal of Combinatorial Optimization*, 8:469–493, 2004.
- (G.R. Mateus) DEPT. OF COMPUTER SCIENCE, FEDERAL UNIVERSITY OF MINAS GERAIS, CEP 31270-010, BELO HORIZONTE, MG, BRAZIL
E-mail address, G.R. Mateus: `mateus@dcc.ufmg.br`
- (M.G.C. Resende) ALGORITHMS AND OPTIMIZATION RESEARCH DEPARTMENT, AT&T LABS RESEARCH, 180 PARK AVENUE, ROOM C241, FLORHAM PARK, NJ 07932 USA.
E-mail address, M.G.C. Resende: `mgcr@research.att.com`
- (R.M.A. Silva) COMPUTATIONAL INTELLIGENCE AND OPTIMIZATION GROUP, DEPT. OF COMPUTER SCIENCE, FEDERAL UNIVERSITY OF LAVRAS, C.P. 3037, CEP 37200-000, LAVRAS, MG, BRAZIL
E-mail address, R.M.A. Silva: `rmas@dcc.ufla.br`