

# Hybrid Algorithms for Placement of Virtual Machines across Geo-Separated Data Centers

Fernando Stefanello<sup>a,\*</sup>, Vaneet Aggarwal<sup>b</sup>, Luciana S. Buriol<sup>a</sup>, Mauricio G.C. Resende<sup>c</sup>

<sup>a</sup> *Instituto de Informática, Federal University of Rio Grande do Sul, Porto Alegre, RS, Brazil.*

<sup>b</sup> *School of Industrial Engineering Purdue University, West Lafayette, IN 47907 USA.*

<sup>c</sup> *Mathematical Optimization and Planning (MOP), Amazon.com, Seattle, WA 98109 USA.*

---

## Abstract

Cloud computing has recently emerged as a new technology for hosting and supplying services over the Internet. This technology has brought many benefits, such as eliminating the need for maintaining expensive computing hardware. With an increasing demand for cloud computing, providing performance guarantees for applications that run over cloud become important. Applications can be abstracted into a set of virtual machines with certain guarantees depicting the quality of service of the application. In this paper, we consider the placement of these virtual machines across multiple data centers (VMPlacement), meeting the quality of service requirements while minimizing the bandwidth cost of the data centers. This problem is a generalization of the NP-hard Generalized Quadratic Assignment Problem (GQAP). In this paper, we present a Greedy Randomized Adaptive Search Procedure (GRASP) and a Biased Random-Key Genetic Algorithm (BRKGA), both hybridized with a Path-Relinking strategy and a local search based on Variable Neighborhood Descent (VND) for solving this problem, also tested in instances of GQAP. We show that both algorithms are effective in solving quickly small and large instances of VMPlacement problem, especially when the path-relinking is used. For GQAP, the results outperform the previous state-of-the-art.

*Keywords:* Combinatorial optimization, Cloud computing, Biased Random-Key Genetic Algorithm, GRASP, Path-Relinking

---

## 1. Introduction

Virtualization of physical servers have gained prominence in enterprise data centers. This is because virtualization offers virtually unlimited resources without any upfront capital investment and a simple pay-as-you-go charging model. Long term viability of virtualization depends, among other factors, on cost and performance. In order to attain performance

---

\*Corresponding author.

*Email addresses:* [fstefanello@inf.ufrgs.br](mailto:fstefanello@inf.ufrgs.br) (Fernando Stefanello), [vaneet@purdue.edu](mailto:vaneet@purdue.edu) (Vaneet Aggarwal), [buriol@inf.ufrgs.br](mailto:buriol@inf.ufrgs.br) (Luciana S. Buriol), [resendem@amazon.com](mailto:resendem@amazon.com) (Mauricio G.C. Resende)

guarantees, application providers can offer requirements for a number of virtual machines, bandwidth/latency requirements between virtual machines, and latency requirements between users of the service and virtual machines. Having all these performance guarantees for the application can help give an optimized service to the users. However, the service provider has to match the requirements of different applications to the placement of virtual machines with the limited bandwidth links between geographically separated data centers while minimizing its cost.

Unfortunately, today's public cloud platforms such as Amazon EC2<sup>1</sup> do not provide any performance guarantee, which in turn affects tenant cost. Specifically, the resource reservation model in today's clouds only provisions CPU and memory resources but ignores networking completely. Because of the largely oversubscribed nature of data center networks (e.g., Greenberg et al. (2009)), network bandwidth is a scarce resource shared across many tenants. In order to meet the reliability and the demand requirements, the data centers have to be placed all across the world. For instance, a teleconference call connects people from all over the world, and a data center within a reasonable distance to the end users is needed. For distributed data centers, networking cost is the major cost, which has not been accounted in the prior works on virtual machine placement to the best of our knowledge. With the limited bandwidth links between the data centers, networking intensive phases of applications collide and compete for the scarce network resources, which leads to their running times become unpredictable. The uncertainty in execution time further translates into unpredictable cost as tenants need to pay for the reserved virtual machines (VMs) for the entire duration of their jobs.

Placement of virtual machines within a data center have been widely explored (Guo et al., 2010; Ballani et al., 2011; Xie and Hu, 2012). These papers account for the networking needs in addition to the CPU and memory needs within a data center. For example, Guo et al. (2010) propose bandwidth reservation between every pair of VMs. Ballani et al. (2011) propose a simpler virtual cluster (VC) model where all virtual machines are connected to a virtual switch with links of bandwidth  $B$ . Xie and Hu (2012) extend these approaches to consider time-varying network requirement. However, all these works account for a single data center where the bandwidths are much larger as compared to the bandwidths across data centers. Instead, this paper deals with the placement of virtual machines across geographically separated data centers.

In this paper, we consider multiple data centers that are connected with limited bandwidth links. The latency between every pair of data centers is known. In order to meet the application's quality of service guarantees, there is a required minimum bandwidth and maximum latency between each pair of virtual machines. We assume that there are multiple users who would use these services, and users are connected to some data center. In order to meet the overall application performance, there is an additional requirement of maximum latency between users and the virtual machines. Intuitively, if there is a set of VMs needed by a user and the set does not have any requirement with any other user or VM, it can be placed in a single data center. However, a VM interacts with multiple VMs which may be

---

<sup>1</sup><http://aws.amazon.com/ec2/>

needed by other users, thus increasing the set of options for placement. There is a cost of transferring data between data centers and the placement minimizes this cost thus preferring placement of all VMs in a single data center which may not be feasible due to the quality of service requirements.

This problem has similarity with the problem of *Virtual Network Embedding* (VNE), since the set of virtual machines and the relation between them can be considered as the *Virtual Network* (VN), while the set of data centers and the links can be mapped as the *Substrate Network* (SN). [Fischer et al. \(2013\)](#) describe a general framework for VNE. The VMPlacement problem can be described using this framework, but the problem has some differences from other works since we consider a complete graph connecting the data center, a quadratic cost function for the communication cost, and latency constraints that usually are ignored from other works.

This problem is a generalization of the NP-hard Generalized Quadratic Placement Problem (GQAP) given in [Lee and Ma \(2004\)](#). Solving this problem optimally is possible only in very small instances, which may not represent the size found in real-world applications. Thus, we propose two metaheuristic approaches to solve the Virtual Machine Placement Problem (VMPlacement). We extend the BRKGA proposed in [Stefanello et al. \(2015a\)](#) and proposed a new GRASP algorithm, both coupled with a path-relinking and an extensive local search procedure. We test the performance of both algorithms in a dataset comprised of instances of sizes ranging from small to large. Both algorithms have similar performance, although a slight advantage for BRKGA was observed in small instances while GRASP has a slight advantage in larger instances. We show that the algorithms are able to quickly find feasible solutions and find high-quality final solutions for the VMPlacement problem, especially when the path-relinking procedure is used. Also, the result obtained with the proposed algorithms outperforms the previous state-of-the-art results for GQAP.

This paper is an extension of the papers [Stefanello et al. \(2015a\)](#) and [Stefanello et al. \(2015b\)](#) and has several contributions from the previous papers: i) We proposed a path-relinking strategy as an intensification method, and a local search method based in a large exploration of the neighborhood; ii) We extend the BRKGA proposed in [Stefanello et al. \(2015a\)](#) and proposed a new GRASP algorithm; iii) We integrate a path-relinking strategy to BRKGA; iv) We provide experimental results for CPLEX, BRKGA and GRASP applied in a set of instances for VMPlacement and GQAP.

The rest of the paper is organized as follows. In [Section 2](#), we present mathematical models for the Virtual Machine Placement Problem in multiple data centers. Two metaheuristic approaches are described in [Section 3](#). Computational results are presented in [Section 4](#). Finally, conclusions are drawn in [Section 5](#).

## 2. Virtual Machine Placement Problem

In the Virtual Machine Placement Problem (VMPlacement), the objective is to place a set  $K$  of virtual machines (VM) in a set  $N$  of data centers (DC) in order to minimize the communication cost among virtual machines over the network.

In this problem, each data center has a capacity  $a_i$ , which represents the number of virtual machines that can be placed in DC  $i$ . Also, between two data centers  $i$  and  $j$ , there are a bandwidth capacity ( $B_{ij}$ ), a latency ( $L_{ij}$ ), and a cost  $C_{ij}$  to transfer a data unit between the pair of data centers. In order to meet the reliability and demand requirements of the applications, certain bandwidth and latency requirements can be imposed on the different VMs that are placed on the data centers. Each pair of virtual machines  $v$  and  $w$  has a required bandwidth ( $b_{vw}$ ) whose sum overall VMs placed between DCs  $i$  and  $j$  cannot exceed  $B_{ij}$ . Furthermore, there is a required latency ( $l_{vw}$ ), such that VMs  $v$  and  $w$  cannot be placed in data centers  $i$  and  $j$  if the required latency is greater than the respective data center latency. Finally, there is a set  $U$  of users who access the system. Each user  $u$  is located at a data center  $d(u)$  and has a required latency  $t_{uv}$  for each VM  $v$ .

Figure 1 shows a representation of the input data components: the data centers (Figure 1a), and the virtual machines (Figure 1b). The first component is composed by three data centers (rounded rectangles). Each data center has a number of users and a capacity (represented as a number of spots where VMs can be placed). The connection between each pair of DCs represents the bandwidth capacity, latency, and cost. The second component is composed by eight virtual machines, where each link represents the required bandwidth and the required latency.

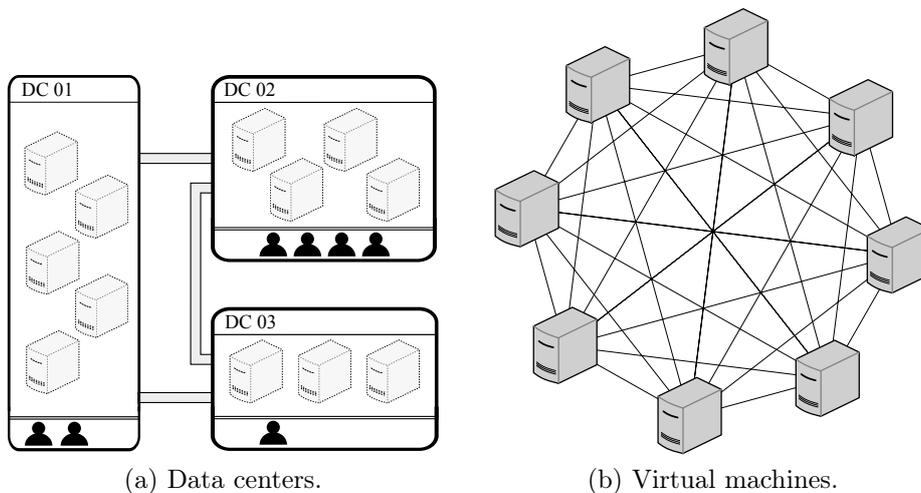


Figure 1: Input data representation.

In the next subsections, we present three mathematical formulations for the VMPlacement problem. We first present a quadratic formulation, followed by two mixed integer linear formulations. Since VMPlacement is a generalization of the NP-hard Generalized Quadratic Assignment Problem, we extend the linear mathematical models proposed for the GQAP in Lee and Ma (2004) to VMPlacement. The models in Lee and Ma (2004) were also extended from the mixed integer linear programming formulation from Kaufman and Broeckx (1978) and Frieze and Yadegar (1983) to the Quadratic Assignment Problem, all of them based on the formulation for the QAP described in Koopmans and Beckmann (1957).

### 2.1. Quadratic mathematical model

A natural formulation for VMPlacement is based in a quadratic formulation as a generalization of GQAP. In what follows we summarize the parameters and present a quadratic mathematical model for VMPlacement (QMVMP) introduced in [Stefanello et al. \(2015a\)](#).

#### Parameters:

- $N$  : set of data centers;
- $K$  : set of virtual machines;
- $U$  : set of users;
- $a_i$  : capacity in number of VMs that DC  $i$  can host;
- $B_{ij}$  : bandwidth between DCs  $i$  and  $j$ ;
- $L_{ij}$  : latency between DCs  $i$  and  $j$ ;
- $C_{ij}$  : cost of transferring a data unit between DCs  $i$  and  $j$ ;
- $b_{vw}$  : required bandwidth between VMs  $v$  and  $w$ ;
- $l_{vw}$  : required latency between VMs  $v$  and  $w$ ;
- $d(u)$  : DC which hosts user  $u$ ;
- $t_{vu}$  : required latency between user  $u$  and VM  $v$ ;
- $c_{iv}$  : cost of placing a VM  $v$  in a DC  $i$ .

Equations (1)-(7) present the quadratic mathematical model for the VMPlacement (QMVMP), where the binary decision variable  $x_{iv} = 1$  if VM  $v$  is located into DC  $i$ , and  $x_{iv} = 0$  otherwise.

$$\min \sum_{i \in N} \sum_{v \in K} c_{iv} x_{iv} + \sum_{i \in N} \sum_{j \in N} \sum_{v \in K} \sum_{w \in K} x_{iv} x_{jw} C_{ij} b_{vw} \quad (1)$$

subject to:

$$\sum_{v \in K} x_{iv} \leq a_i \quad \forall i \in N, \quad (2)$$

$$\sum_{i \in N} x_{iv} = 1 \quad \forall v \in K, \quad (3)$$

$$\sum_{v \in K} \sum_{w \in K} x_{iv} x_{jw} b_{vw} \leq B_{ij} \quad \forall i, j \in N, \quad (4)$$

$$\sum_{i \in N} \sum_{j \in N} x_{iv} x_{jw} L_{ij} \leq l_{vw} \quad \forall v, w \in K, \quad (5)$$

$$\sum_{i \in N} x_{iv} L_{i,d(u)} \leq t_{vu} \quad \forall u \in U, \forall v \in K, \quad (6)$$

$$x_{iv} \in \{0, 1\} \quad \forall i \in N, \forall v \in K. \quad (7)$$

Objective function (1) minimizes the cost of placing each pair of virtual machines  $v$  and  $w$  to DCs  $i$  and  $j$ . Constraints (2) require that the number of VMs in each DC must not exceed the DC capacity. Constraints (3) require that each VM must be assigned to exactly

one DC. Constraints (4) require that the given bandwidth between each pair  $i$  and  $j$  of DCs should not be surpassed by the total sum of bandwidth required among the virtual machines placed in these DCs. Constraints (5) assure that the latency required between each pair of VMs should be respected, i.e, if VMs  $v$  and  $w$  are placed respectively to DCs  $i$  and  $j$ , then the latency between DCs  $i$  and  $j$  should not exceed the required latency between VMs  $v$  and  $w$ . Constraints (6) require that the latency between a VM  $v$  and the DC where the user  $u$  is located be respected, i.e, a VM  $v$  can be only placed on a DC  $i$  if the latency between  $i$  and  $d(u)$  is less than or equal to a given latency between the VM  $v$  and the user  $u$ . Finally, constraints (7) define the variables domain.

The performance of mixed integer linear programming solvers has improved considerably over the last few years. CPLEX<sup>2</sup> is a general-purpose black-box solver based on simplex and a branch-and-bound algorithm with the state-of-the-art exact algorithms for integer programming and has been successfully applied in many combinatorial optimization problems. To analyze the CPLEX performance and provide baseline results for comparison of heuristic methods, the following subsections present two linear mathematical models for VMPlacement problem.

## 2.2. Linear mathematical model I: LMVMP

Based on model L3 from Lee and Ma (2004) for GQAP, and from Frieze and Yadegar (1983) for QAP, we present a mixed-integer linear model for the VMPlacement. Let  $y_{ivjw} = x_{iv}x_{jw}$ ,  $\forall i, j = \{1, \dots, N\}$  and  $v, w = \{1, \dots, K\}$ , the mixed-integer linear mathematical model for VMPlacement named as LMVMP can be formulated as the following:

---

<sup>2</sup>[www.cplex.com](http://www.cplex.com)

$$\min \sum_{i \in N} \sum_{v \in K} c_{iv} x_{iv} + \sum_{i \in N} \sum_{j \in N} \sum_{v \in K} \sum_{w \in K} y_{ivjw} C_{ij} b_{vw} \quad (8)$$

subject to:

$$\sum_{v \in K} x_{iv} \leq a_i \quad \forall i \in N, \quad (9)$$

$$\sum_{i \in N} x_{iv} = 1 \quad \forall v \in K, \quad (10)$$

$$\sum_{i \in N} y_{ivjw} = x_{jw} \quad \forall v, w \in K, \forall j \in N, \quad (11)$$

$$y_{ivjw} = y_{jwiv} \quad \forall v, w \in K, \forall i, j \in N, \quad (12)$$

$$\sum_{v \in K} \sum_{w \in K} y_{ivjw} b_{vw} \leq B_{ij} \quad \forall i, j \in N, \quad (13)$$

$$\sum_{i \in N} \sum_{j \in N} y_{ivjw} L_{ij} \leq l_{vw} \quad \forall v, w \in K, \quad (14)$$

$$\sum_{i \in N} x_{iv} L_{i,d(u)} \leq t_{vu} \quad \forall u \in U, \forall v \in K, \quad (15)$$

$$x_{iv} \in \{0, 1\} \quad \forall i \in N, \forall v \in K, \quad (16)$$

$$0 \leq y_{ivjw} \leq 1 \quad \forall i, j \in N, \forall v, w \in K. \quad (17)$$

The LMVMP is obtained by replacing the product  $x_{iv}x_{jw}$  by  $y_{ivjw}$  from QMVMP. In addition four extra sets of constraints are considered. Constraints (11) and (12) define the relation between variables  $x$  and  $y$ . Constraints (12) impose the symmetry relation to variables  $y$ . Finally, constraints (17) define the domain of variables  $y$ .

We note that the model QMVMP has quadratic constraints, while LMVMP not. The objective function also changes from a quadratic function in QMVMP to linear in LMVMP. However, the mixed-integer linear problem LMVMP has a considerable higher number of variables, having variables  $y_{ivjw}$  in addition to the previous variables  $x_{iv}$ . Thus, the number of variables changes from  $O(NK)$  in QMVMP to  $O(N^2K^2)$  in LMVMP. We note that if the optimal solution of LMVMP is  $(x_{iv}^*, y_{ivjw}^*)$ , then  $(x_{iv}^*)$  is the optimal solution for QMVMP. The proof that both models are equivalent can be obtained by extending the proof for QAP provided in Lee and Ma (2004).

### 2.3. Linear mathematical model II LMVMP-II

Next we present a second linearization for VMPlacement problem (LMVMP-II). This linear model is derived from Kaufman and Broeckx (1978) for QAP, which is probably the linearization for QAP with a lower number of variables and constraints. In Lee and Ma (2004) the authors extend the formulation for GQAP.

In this model, each binary decision variable  $x_{iv}$  is set to one when VM  $v$  is located into DC  $i$ , and zero otherwise. The auxiliary variables  $y_{iv}$  aggregate the cost for each placed VM  $v$  in a DC  $i$ , and  $n_{ijb}$  aggregate the bandwidth from VM  $v$  between data centers  $i$  and  $j$ .

$$\min \sum_{i \in N} \sum_{v \in K} c_{iv} x_{iv} + \sum_{i \in N} \sum_{v \in K} y_{iv} \quad (18)$$

subject to:

$$\sum_{v \in K} x_{iv} \leq a_i \quad \forall i \in N, \quad (19)$$

$$\sum_{i \in N} x_{iv} = 1 \quad \forall v \in K, \quad (20)$$

$$\sum_{j \in N} \sum_{w \in K} C_{ij} b_{vw} x_{jw} - y_{iv} \leq m_{iv} (1 - x_{iv}) \quad \forall v \in K, \forall i \in N, \quad (21)$$

$$\sum_{w \in K} b_{vw} x_{jw} - n_{ijv} \leq M(1 - x_{iv}) \quad \forall v \in K, \forall i, j \in N, \quad (22)$$

$$\sum_{v \in K} n_{ijv} \leq B_{ij} \quad \forall i, j \in N, \quad (23)$$

$$L_{ij} x_{iv} \leq l_{vw} + L_{ij} (2 - x_{iv} - x_{jw}) \quad \forall v, w \in K, \forall i, j \in N, \quad (24)$$

$$\sum_{i \in N} x_{iv} L_{i,d(u)} \leq t_{vu} \quad \forall u \in U, \forall v \in K, \quad (25)$$

$$x_{iv} \in \{0, 1\} \quad \forall i \in N, \forall v \in K, \quad (26)$$

$$y_{iv} \geq 0 \quad \forall i \in N, \forall v \in K, \quad (27)$$

$$n_{ijv} \geq 0 \quad \forall v, w \in K, \forall i, j \in N. \quad (28)$$

where

$$m_{iv} \geq \sum_{j \in N} \sum_{w \in K} C_{ij} b_{vw}, \quad \forall i \in N, \forall v \in K.$$

Constraints (21) impose the cost between the data centers  $i$  and  $j$  to the variables  $y_{ij}$ . Constraints (22) and (23) impose the bandwidth constraints while constraints (24) impose the latency constraints. The constraints (24) can be replaced by

$$x_{iv} + x_{jw} \leq 1 \quad \forall i, j \in N, \forall v, w \in K \text{ if } L_{ij} > l_{vw}. \quad (29)$$

We observe that CPLEX converts (24) into (29) in the pre-processing phase. Finally, constraints (26), (27), and (28) define the domain of variables.

This model is not used in practice since it uses big-M constraints (21), and the root-node bound is always zero. However, a stronger formulation can be obtained by adding the following cuts

$$y_{iv} \geq Y_{iv} x_{iv} \quad \forall i \in N, \forall v \in K, \quad (30)$$

where  $Y_{iv}$  is defined as the optimal value of the following generalized assignment model:

$$Y_{iv} = \min \sum_{j \in N} \sum_{w \in K} C_{ij} b_{vw} x_{jw} \quad (31)$$

subject to:

$$\sum_{w \in K} x_{jw} \leq a_j \quad \forall j \in N, \quad (32)$$

$$\sum_{j \in N} x_{jw} = 1 \quad \forall w \in K, \quad (33)$$

$$\sum_{j \in N} \sum_{w \in K} b_{vw} x_{jw} \leq B_{ij} \quad \forall w \in K, \quad (34)$$

$$(24), (25) \quad (35)$$

$$x_{iv} = 1 \quad (36)$$

$$x_{jw} \in \{0, 1\} \quad \forall j \in N, \forall w \in K. \quad (37)$$

These additional cuts were applied for the 3-dimensional assignment problem in [Mittelman and Salvagnin \(2015\)](#). Note that these cuts can also be applied for GQAP suppressing the constraints (22)–(25) and (28) from LMVMPII, and (34)–(35) from the assignment model, which are the specific constraints for VMPlacement.

### 3. Heuristic Approaches

In this section we propose two metaheuristics approaches to solve the VMPlacement problem. We first describe the local search procedures and two path-relinking strategies used in the proposed metaheuristics. Finally, we describe the Greedy Randomized Adaptive Search Procedure (GRASP), followed by the Biased Random-Key Genetic Algorithm (BRKGA).

Place a virtual machine into a data center can violate some of the constraints. To deal with this situation, we use a penalization strategy to minimize the number of violated constraints. Thus, the cost of placing a VM  $v$  in DC  $i$  is calculated by the regular placement cost added by a sufficiently large number  $M$  for each violated constraint. This penalization strategy is applied whenever a solution is evaluated. In our experiments we use  $M = 10^{10}$ . Also, the notation for a solution  $S$  represents a vector of size  $|K|$  with  $S = \{n_1, \dots, n_{|K|}\}$ , with  $n_v$  representing the label (or index) of the data center where the virtual machine  $v$  is placed.

#### 3.1. Local search procedures

Local search is a general approach for finding and improving solutions to hard combinatorial optimization problems. The most basic strategy of local search algorithms is to start from an initial solution and iteratively replace the current solution by a better neighbor solution, until no improvement can be reached. A neighbor solution can be obtained by applying moves defined by a neighborhood structure. In this work we define four neighborhood structures to obtain neighbor solutions, namely *shift*, *swap*, *chain2L*, and *chain3L*.

**Shift.** A shift operation moves a virtual machine from the current data center to a different data center. In a *shift search*, we select a virtual machine  $v$  in a circular order of their indexes (starting from index zero), and calculate the cost to move it from the current data center to each other data center. A virtual machine  $v$  is moved to the data center that produces the greatest improvement in the objective function. This procedure is repeated until reaching a local minimum.

**Swap.** A swap operation interchanges the positions of two virtual machines. In a *swap search*, we evaluate the cost of all swap moves between two virtual machines  $v$  and  $w$  in a circular order of their indexes (starting from index zero). When an improvement is reached, the virtual machine positions are interchanged, and the search continues by selecting a next pair of virtual machines. The procedure ends when no swap move can improve the solution. The evaluation of symmetrical movements is avoided.

**Chain2L.** A chain operation is a composition of two shift moves applied sequentially. In a *chain2L search*, the first shift move selects a virtual machine  $v$  from the data center  $i$  to move to a data center  $j$ . In the second shift move, a virtual machine  $w$  from  $j$  is moved to a data center  $k$ . We evaluate all compositions of two shift moves. When an improvement is reached, the moves are applied to the solution and the process is restarted with the new solution. Note that this neighborhood includes the *swap search* when  $k$  is equal to  $i$ .

**Chain3L.** The last neighborhood proposed extends the concept of the chain moves to a composition of three shift moves applied sequentially. A *chain3L search* involves four virtual machines and up to four data centers. This search also comprises three shift moves between two data centers, or a circle move among three data centers.

Given a solution  $S$ , the worst case time complexity to evaluate the cost to insert or remove a virtual machine  $v$  in  $S$  is  $O(|K|)$ . This complexity is reached because is necessary to evaluate whether the latency requirement is satisfied between  $v$  and  $w \in S$ . Also, it is necessary evaluate the bandwidth cost between  $v$  and  $w \in S$ . Capacity constraints, uniqueness in the placement, and user latency requirements can be evaluated in constant time. Therefore, for the *shift search* a number of  $|K|$  removals and  $|K| * |N|$  insertions are evaluated, resulting in a complexity of  $O(|K|^2|N|)$  (for the case of no improvement during the search). For the *swap search* the time complexity is  $O(|K|^3)$ , while for *chain2L search* is  $O(|K|^3|N|)$ . Finally, the *chain3L search* has complexity of  $O(|K|^4|N|)$ .

Naturally, in all searches described above, the processing time can be reduced avoiding to evaluate moves that lead to a worse solution. Also, heuristic strategies do not need to explore the whole neighborhood, and only evaluate a subset of the neighbor solutions. This is done in the *chain3L search* prohibiting the insertion of a virtual machine in some data centers, thus reducing the search space.

Two heuristic strategies to reduce the amount of explored neighbor solutions in the *chain3L search* are proposed. The first reduce heuristic strategy is based in the cost  $Y_{iv}$  described in the Model (31)-(37) from Subsection 2.3. Given a parameter  $\alpha_{3L} \in [0, 1]$ , we allow to insert a virtual machine  $v$  in a data center  $i$  only if  $i$  belongs to the  $\lfloor N * \alpha_{3L} \rfloor$  data centers with the lowest cost  $Y_{iv}$ . A low value of  $\alpha_{3L}$  indicates a restricted search only in data centers with low values of  $Y_{iv}$ . A high value of  $\alpha_{3L}$  indicates a more broad search. An analyzis for the instances described in Section 4.1 with their respective best known solutions

corroborate for this heuristic strategy. In these solutions, more than 50% of the virtual machines are placed in the 30% data centers with lowest cost  $Y_{iv}$ , and approximately 75% are placed in the 50% data centers with the lowest cost.

The second reduce heuristic strategy is based on the communication cost assigned to the data centers. Given a solution  $\mathcal{S}$ , let  $C_i$  be the sum of the communication cost from  $i$  to all other data centers. We only evaluate compositions of movements that start from a virtual machine that belongs to the  $\beta_{3L}$  data center with the highest cost  $C_i$ . In summary, the idea is to limit the search to movements that reduce the communication cost of the most required data centers. Combining both reduction heuristics, the search space is considerably reduced, and the heuristics are still able to visit solutions in a promising search space. By default, we use  $\alpha_{3L} = 0.3$  and  $\beta_{3L} = 3$ .

The neighborhoods are integrated based on the idea of Variable Neighborhood Descent (Hansen et al., 2010). Every neighborhood is applied sequentially starting from the lowest to highest complexity neighborhood (*shift*, *swap*, *chain2L*, and *chain3L*). However, the search is restarted with the first neighborhood if an improvement is reached. We name each local search strategy using two characters. The first character is a number that indicates the highest complex neighborhood applied. The character V indicates that we use the VND strategy to integrate the neighborhoods. For example, the name 3V indicates that we consider local search procedures *shift search*, *swap search*, and *chain2L search*.

Note that using a penalization strategy described in the previous subsection, the local search is also applied to infeasible solutions. In this case, the local search also works as a repair procedure.

### 3.2. Path-relinking

Path-relinking (PR) is an approach to integrate intensification and diversification in the search. It consists in exploring trajectories that connect high-quality solutions. 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.

The path-relinking was first suggested for tabu search in Glover (1989) and then formalized in Glover and Laguna (1993). Since then, this strategy was applied for a large number of combinatorial optimization problems and related studies as Glover (1997), Glover et al. (2000), Resende and Ribeiro (2005), Resende et al. (2010), Festa and Resende (2013), and Glover (2014), just to name a few.

The Algorithm 1 shows the pseudo-code for the path-relinking operator between solutions  $S$  and  $T$ , where  $S$  is the initial solution and  $T$  the guiding solution. In line 2 the incumbent solution is initialized. The loop between the lines 3 and 10 is repeated while the distance between both solutions is greater than an input parameter  $\delta_{lim}$ . The distance  $\Delta\{S, T\}$  is the minimum number of moves needed to transform  $S$  into  $T$  or vice-versa. We use  $\delta_{lim} = \Delta\{S, T\}/2$  defined at the beginning of the procedure, since we observed that the probability to find an improved solution is greater in the first steps of the path. In line 4 a movement is applied to the solution  $S$  in direction to the solution  $T$ . Basically, we analyze

all changes in  $S$  to  $T$  that reduce the distance  $\Delta\{S, T\}$  by one, and apply in  $S$  the change with the least cost.

Several studies have experimentally found that it is convenient to add a local search exploration from some of the generated solutions in the path (Martí et al., 2006). Since two consecutive solutions obtained by a relinking step are often very similar, it is generally not efficient to apply the local search at every step of the relinking process. Thus, in line 7 the local search is applied at each  $n$  iterations, where  $n = (|K| * 0.5)/(4 + |K|/50)$ . The parameter  $n$  is rounded up to the first even number. This ensures that the local search is applied to  $S$  and  $T$  each time, since, in line 10, the solutions are interchanged to use *back-and-forward* strategy (Festa and Resende, 2013). Finally, in line 8 the incumbent solution is updated and returned in line 11.

---

**Algorithm 1:** Pseudo-code for a greedy path-relinking operator.

---

```

1 Procedure PathRelinkingOperator( $S, T$ )
2    $S^* \leftarrow S$ ;
3   while  $\Delta\{S, T\} \geq \delta_{lim}$  do
4      $S \leftarrow \text{MakeMovement}(S, T)$ ;
5      $S' \leftarrow S$ ;
6     if  $LSCcondition$  then
7        $S' \leftarrow \text{LocalSearch}(S')$ ;
8     if  $f(S') < f(S^*)$  then
9        $S^* \leftarrow S'$ ;
10     $\text{Swap}(S, T)$ ;
11  return  $S^*$ ;

```

---

Algorithm 2 and 3 describe two frameworks of how to administrate the elite set  $\mathcal{E}$ , and how solutions are selected for the path-relinking operator. In the first approach, the path-relinking operator is applied between a provided solution  $S$  and a selected solution  $T$  from the elite set. In the second approach, the path-relinking operator is applied multiple times, between a provided solution  $S$  and every solution from the elite set.

To simplify the notation, we name by path-relinking the general process to administrate the elite and make the path-relinking operator between two solutions. Regarding to the specific administration procedures, we name the first approach by PRS and the second approach by PRM. The first two letters indicate that the algorithm is related to the path-relinking and the last letter indicates the framework for administrate the path-relinking operator. In the first, S indicates the application of a single path-relinking operator and in the second, M indicates the application of a multiple path-relinking operator.

In the beginning of the Algorithm 2, the elite set  $\mathcal{E}$  is empty. The maximum size  $\delta_{max}$  of the elite set is an input parameter. A solution  $S$  is added to  $\mathcal{E}$  if it improves the best solution or is sufficiently different from each solution in  $\mathcal{E}$  (lines 3, 6, and 11). We define that two solutions are sufficiently different if  $\Delta\{S, T\} > \delta_{dis}$ . We denote by  $S \not\approx \mathcal{E}$  if  $\Delta\{S, T'\} > \delta_{dis}, \forall T' \in \mathcal{E}$ . In our case, we use  $\delta_{dis} = \lfloor |K| * 0.1 \rfloor$ . If the elite set has a

minimum number of solutions ( $\delta_{min} = 2$  in line 2), and  $S$  is sufficiently different, a solution  $T$  is selected for the path-relinking operator between  $S$  and  $T$ .

In line 4, solutions are ordered by a linear rank  $r(T_i)$ ,  $\forall i = 1 \dots |\mathcal{E}|$ , where the best solution has rank  $r(T_1) = |\mathcal{E}|$  and the worst solution has rank  $r(T_{|\mathcal{E}|}) = 1$ . A solution  $T \in \mathcal{E}$  is selected using the roulette wheel criterion, i.e,  $T$  is selected from the elite set with probability of  $r(T)/\sum_{i \in r(T_i)}$ . With this criterion, better solutions have more chance to be chosen. Since  $T$  is selected, the path-relinking operator is applied between  $S$  and  $T$ .

Lines 6 to 9 update the elite set after the path-relinking operator is applied. Line 6 evaluates whether  $S$  is sufficiently different or improves the best solution. If this is the case,  $S$  is added to  $\mathcal{E}$ . Line 8 is responsible for maintaining the cardinality of the elite set.  $T'$  is the most similar solution to  $S$  with worst solution than  $S$ , i.e,  $T' \leftarrow \{T'' \in \mathcal{E} \mid f(T'') > f(S) \text{ and } \Delta\{T'', S\} \leq \Delta\{T'', S'\}, \forall S' \in \mathcal{E}\}$ . Finally, in line 13 the solution  $S$  is returned.

---

**Algorithm 2:** Pseudo-code for the path-relinking administration framework PRS.

---

```

1 Procedure PathReLinking( $S$ )
2   if  $|\mathcal{E}| > \delta_{min}$  then
3     if  $S \not\approx \mathcal{E}$  then
4       Select  $T \in \mathcal{E}$  ;
5        $S \leftarrow \text{PathReLinkingOperator}(S, T)$ ;
6       if  $S \not\approx \mathcal{E}$  then
7          $\mathcal{E} \leftarrow \mathcal{E} \cup \{S\}$ ;
8       if  $|\mathcal{E}| > \delta_{max}$  then
9         Remove  $T' \in \mathcal{E}$ ;
10  else
11    if  $S \not\approx \mathcal{E}$  then
12       $\mathcal{E} \leftarrow \mathcal{E} \cup \{S\}$ ;
13  return  $S$ ;

```

---

Another approach is outlined in Algorithm 3. The main difference from the previous approach is that when a solution  $S$  is added to  $\mathcal{E}$ , the path-relinking operator is applied between all pair of solutions from  $\mathcal{E}$  to which the operator was not previously applied. This is done in the loop in lines 5 and 13. The flag *hasPairsST* can be easily updated using memory structures, and it indicates whether there are  $S'$  and  $T'$  in the elite set that the path-relinking operator was not applied. In line 6 two solutions are selected. Pairs with better objective functions are selected first. The insertion operation (line 11) and the removal operation (line 13) follow the same rules than in the previous algorithm. In line 14 a solution is removed from the elite set in case it is full. We also use roulette wheel criterion by rank, where the best solution has rank equal to zero, and the worst solution has rank  $|\mathcal{E}| - 1$ . This ensures that the best solution are kept in the pool, and the worst solution has a high probability to be removed. Finally, in line 18 the solution  $S$  is returned.

---

**Algorithm 3:** Pseudo-code for the path-relinking administration framework PRM.

---

```

1 Procedure PathRelinking( $S$ )
2   if  $|\mathcal{E}| > \delta_{min}$  then
3     if  $S \not\approx \mathcal{E}$  then
4        $\mathcal{E} \leftarrow \mathcal{E} \cup \{S\}$ ;
5       while hasPairsST do
6         Select  $S'$  and  $T'$ ;
7          $S' \leftarrow \text{PathRelinkingOperator}(S', T')$ ;
8         if  $f(S') < f(S)$  then
9            $S \leftarrow S'$ ;
10        if  $S' \not\approx \mathcal{E}$  then
11           $\mathcal{E} \leftarrow \mathcal{E} \cup \{S'\}$ ;
12        if  $|\mathcal{E}| > \delta_{max}$  then
13          Remove  $T' \in \mathcal{E}$ ;
14        Remove  $T' \in \mathcal{E}$ ;
15   else
16     if  $S \not\approx \mathcal{E}$  then
17        $\mathcal{E} \leftarrow \mathcal{E} \cup \{S\}$ ;
18   return  $S$ ;

```

---

Note that Algorithm 2 and 3 are designed in independent components, requiring only few configuration parameters (as local search strategy and elite size), and a procedure that provides different solutions. For this reason, these path-relinking strategy can be easily included in many different heuristic and metaheuristic frameworks. In the next subsections, we describe the GRASP and BRKGA metaheuristics which use the path-relinking strategies.

### 3.3. Greedy Randomized Adaptive Search Procedure - GRASP

GRASP (Greedy Randomized Adaptive Search Procedure) is a multi-start metaheuristic for combinatorial optimization problems (Feo and Resende, 1989; 1995; Festa and Resende, 2009a;b; Resende and Ribeiro, 2010; Martí et al., 2013; Resende and Ribeiro, 2014). The algorithm is basically composed of two phases: construction and local search. The construction phase builds a feasible solution following a greedy randomized criterion, whose neighborhood is investigated until a local minimum is found during the local search phase. These phases are repeated over the iterations and the best local optimum found is returned as the heuristic solution. Its basic implementation is memoryless, because it does not make use of information collected in previous iterations. One way to add memory to GRASP is its hybridisation with path-relinking. A large review of this technique is presented in Festa and Resende (2013).

A high-level description of GRASP is presented in Algorithm 4. GRASP iterations are carried out in lines 3 to 8. In line 4, the procedure attempts to build a greedy randomized

solution (see algorithm 5). A solution  $S$  is used as the starting solution for the local search in line 5. In line 6 the path-relinking is applied if a hybrid approach is used. If the local optimum  $S$  is better than the incumbent solution, then, in line 8 the incumbent solution is updated. In line 9, the best solution found overall GRASP iterations is returned as the GRASP solution.

---

**Algorithm 4:** Pseudo-code for GRASP.

---

```

1 Procedure GRASP()
2    $S^* \leftarrow \emptyset$ ;
3   while stopping criterion is not satisfied do
4      $S \leftarrow \text{GreedyRandomized}()$ ;
5      $S \leftarrow \text{LocalSearch}(S)$ ;
6      $S \leftarrow \text{PathRelinking}(S)$ ;
7     if  $f(S) < f(S^*)$  then
8        $S^* \leftarrow S$ ;
9   return  $S^*$ ;

```

---

The Algorithm 5 shows a general pseudo-code for the greedy randomized constructive heuristic used in the GRASP procedure. The construction builds a solution  $S$ , one element at a time. In line 2, solution  $S$  is initialized empty. In line 3, the set of candidate list is initialized, as well as, the cost of placing each candidate element in a data center is calculated in line 4. The solution is built in the loop in lines 5 to 10. In line 6, a restricted candidate list ( $RCL$ ) is build, and in line 7 a candidate is selected. In line 8 this element is added to the partial solution. In line 9, the candidate is removed from the set of candidates  $\mathcal{C}$ . In line 10 the placement costs for each candidate is re-evaluated. Finally, in line 11 the solution  $S$  is returned.

---

**Algorithm 5:** Pseudo-code for a greedy randomized constructive solution procedure.

---

```

1 Procedure GreedyRandomized()
2    $S \leftarrow \emptyset$ ;
3   Initialize set of candidates  $\mathcal{C}$ ;
4   Evaluate the incremental cost of candidates;
5   while  $\mathcal{C} \neq \emptyset$  do
6     Build the  $RCL$ ;
7     Select  $c \in RCL$ ;
8     Add  $c$  to solution:  $S \leftarrow S \cup \{c\}$ ;
9     Update candidates:  $\mathcal{C} \leftarrow \mathcal{C} \setminus \{c\}$ ;
10    Re-evaluate the incremental costs;
11  return  $S$ ;

```

---

Based on this framework, four constructive heuristics were developed to generate initial solutions in line 4 of GRASP procedure.

**Random.** In the first constructive heuristic, on each placement step a non-placed virtual machine is randomly selected and placed in a random data center. Note that this constructive approach is not a greedy heuristic. We decided to evaluate this approach since it is one of the simplest constructive methods, and we use it as a baseline for comparing more sophisticated constructive methods.

**Greedy.** In the second constructive heuristic, on each placement step a non-placed virtual machine is randomly selected and placed in the data center that produces the least increase in the objective function. In this case, the *RCL* is composed only by one candidate that is the best data center for a specific randomly selected virtual machine.

**DC-Greedy.** In the third constructive heuristic, on each placement step a non-placed virtual machine is randomly selected and placed in one of the  $n$  data centers from the *RCL*. The *RCL* contains the candidates that produce lowest incremental placement cost. By default, we use  $n = \max\{3, |N| * 0.2\}$ . The value of  $n$  can variate on each iteration since we consider some additional conditions. First, candidates with the same cost to the first  $n$  candidates are also added to *RCL*. Second, if at least one candidate can be added to the solution keeping it feasible, then candidates that generate infeasibilities are ignored.

**VM-Greedy.** In the last constructive heuristic, on each placement step, all possible placements of virtual machines to data centers are evaluated, and one of the  $n$  best candidates is randomly selected. In this constructive heuristic, we use  $n = \max\{5, |K| * 0.1\}$  for the *RCL* size, but the effective value of  $n$  can variate since we also include both additional conditions described in the previous constructive heuristic.

In these constructive heuristic, any random selection is based on uniform distribution. Note that these constructive heuristic can generate infeasible solutions, and the repair process is made by the local search procedure.

### 3.4. Biased random-key genetic algorithm - BRKGA

A biased random-key genetic algorithm (BRKGA) is a populational metaheuristic that recently has been successfully applied to several hard combinatorial optimization problems (Gonçalves and Resende, 2011). A BRKGA is a class of genetic algorithms which solutions are encoded as vectors of random keys, i.e. randomly generated real numbers from uniform distribution in the interval  $[0, 1)$ . A vector of random keys is translated into a solution of the optimization problem by *decoders*. A decoder is a deterministic algorithm that takes as input a vector of random keys and returns a solution of the optimization problem as well as its cost (or *fitness*).

Algorithm 6 presents a general scheme of a BRKGA. The algorithm starts with a set of  $p$  random vectors of size  $n$  (*population*). Parameter  $n$  depends on the encoding while parameter  $p$  is user-defined. Starting from the initial population, the algorithm generates a series of populations. Each iteration of the algorithm is called a *generation*. The algorithm evolves the population over the generations by combining pairs of solutions from one generation to produce offspring solutions to the following generation.

At each generation  $g$ , the decoder is applied to all newly created random keys, and the population is partitioned into a smaller set of  $p_e$  elite solutions, i.e., the best fittest  $p_e$  solutions in the population, and another larger set of  $p - p_e > p_e$  non-elite solutions

---

**Algorithm 6:** Pseudo-code of BRKGA.

---

```
1 Procedure BRKGA()  
2   Generate  $p$  vectors of random keys;  
3   Decode each vector of random keys;  
4   while stopping criterion is not satisfied do  
5     Sort solutions by their fitness;  
6     Classify solutions as elite and non-elite;  
7     Copy elite solutions to next population;  
8     Generate mutants in the next population;  
9     Combine an elite and a non-elite and add offspring to next population;  
10    Decode each vector of random keys;  
11  return best solution  $S$ ;
```

---

(line 6). Population  $g + 1$  is generated as follows. All  $p_e$  elite solutions of population  $g$  are copied without changing to population  $g + 1$  (line 7). This elitist strategy maintains the best solutions on hand. To ensure that mutation is present in the evolution,  $p_m$  *mutants* are added to population  $g + 1$ . A mutant is simply a vector of random keys, generated in the same fashion than an initial solution (line 8).

With  $p_e + p_m$  solutions accounted for population  $g + 1$ , other  $p - p_e - p_m$  additional solutions must be generated to complete the  $p$  solutions that make up population  $g + 1$ . This is done through *mating* or *crossover* (line 9). A parent- $A$  is selected randomly from the elite solutions, and the parent- $B$  is selected randomly among the set of non-elite solutions. A child  $C$  is produced by combining the parents using parameterized uniform crossover. Let  $\rho_A > 1/2$  be the probability that the offspring solution inherits the key of parent- $A$  and  $\rho_B = 1 - \rho_A$  be the probability that it inherits the key of parent- $B$ , i.e.  $c_i = a_i$  with probability  $\rho_A$  or  $c_i = b_i$  with probability  $\rho_B = 1 - \rho_A$ , where  $a_i$  and  $b_i$  are, respectively, the  $i$ -th key of parent- $A$  and parent- $B$ , for  $i = 1, \dots, n$ .

Random key genetic algorithms were first introduced in [Bean \(1994\)](#). BRKGA differs from the original approach in the way parents are chosen from the population and how the child inherits the key ([Gonçalves and Resende, 2011](#)). In that work, the authors describes the BRKGA as a *problem-independent* and *problem-dependent* components. The independent components are applied without any acknowledgement of the problem. The problem-dependent components are the only connection with the problem. Thus, to describe a BRKGA, one needs only to show how solutions are encoded and decoded, what choice of parameters  $p$ ,  $p_e$ ,  $p_m$ , and  $\rho_A$  was made, and how the algorithm stops. We next describe the encoding and decoding procedures for the proposed algorithm, and attribute values for parameters and stopping criterion in Section 4.

### 3.4.1. Decoders

Solutions of the optimization problem are encoded as a vector  $\mathcal{X}$  with  $n = |K|$  random keys. To translate this vector into the solution of the VMPlacement problem, we propose two decoders, as described next.

**Greedy Ordered Decoder - D1:** In this decoder, the keys provide the order of placement. Following this order, a greedy strategy is used, placing each VM to the DC which produces the least increase in the objective function.

The decoder starts with a list whose each element is composed of the random key and the index of the virtual machine. The list is sorted by the keys. Now, the sorted list of index of virtual machines is used as an order in which virtual machines should be placed. Following this order, the next step is to place each virtual machine  $v$  to a DC. This is done by placing virtual machine  $v$  in DC  $i$  which produces the least increase in the objective function. Note that the cost to insert the VM  $v$  in each DC considers the previous virtual machines placed. When all VMs are placed, the decoder returns the fitness value for the respective vector  $\mathcal{X}$  of random keys.

**Location Decoder - D2:** In this decoder, each key is decoded as the data center in which the virtual machine should be placed. Let  $k_i$  be the key corresponding to the VM of index  $i$  in  $\mathcal{X}$ , then this decoder simply places the VM of index  $i$  to DC  $\lfloor k_i * |N| \rfloor$ . Figure 2 shows an example of decoding for the decoder D2 for 3 data centers, 8 virtual machines and a vector  $\mathcal{X}$  of random keys.

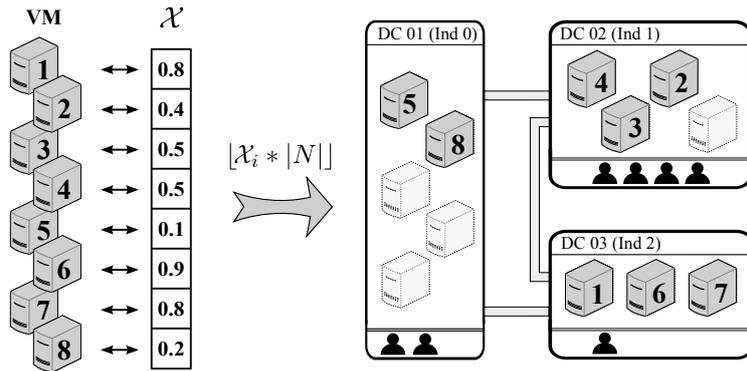


Figure 2: Example of decoding for decoder D2.

Since both decoders are deterministic, we always obtain the same solution  $S'$  from a vector of random keys  $\mathcal{X}$ . In the case where a deterministic local search strategy is applied to  $S'$ , a solution  $S''$  can be obtained using the decoders and the local search, and thus,  $S''$  can be associated to  $\mathcal{X}$ . However, depending on the decoder, it is possible to use a process called *recode* in  $\mathcal{X}$ , to obtain  $S''$  only by using the decoder (without local search). In decoder D1, the recode can be computationally expensive and then we decided to maintain the local search as part of the decoder. In decoder D2, the recode can be applied by calculating the value of the key that will correspond to a data center in  $S''$ . Let  $i'$  the index of DC  $i$ ,  $lb = i' / |N|$  and  $lu = (i + 1)' / |N|$ , a key that corresponds to DC can be given by  $lb + (lu - lb) / 2$  or any number in the interval  $[lb, ub)$ . Experiments including the recode process are presented in the Subsection 4.4. To simplify the notation, we denote by D3 the decoder D2 with recode process.

### 3.4.2. Hybrid BRKGA and path-relinking

Path-relinking is an approach to integrate intensification and diversification in the search and can be incorporated on many metaheuristic frameworks. Path-relinking is frequently used with GRASP algorithms (Laguna and Martí, 1999; Oliveira et al., 2004; Mateus et al., 2010; Festa and Resende, 2013), but references can be found for other metaheuristics as Scatter Search (Glover et al., 2003), Tabu Search (Armentano et al., 2011), VNS (Pérez et al., 2007). Regarding GAs, we find several papers where the path-relinking technique is applied, as for example Basseur et al. (2005), Zhang and Lai (2006), and Vallada and Ruiz (2010), just to name a few.

In this work, we propose a new approach that hybridise the path-relinking with BRKGA algorithm. In our approach, we include the procedure 2 or 3 at  $n$  generations of BRKGA. By default we use  $n = 1$ , i.e., we apply the path-relinking each time a new generation is produced. Since both path-relinking procedures require a solution  $S$ , we randomly select  $S$  by uniform distribution among all elite solutions of the BRKGA population. In the case that the path-relinking returns a solution  $S'$  better than  $S$ , then the corresponding vector  $\mathcal{X}$  of  $S$  is recoded to be adjusted to  $S'$ .

In Section 4, experiments with both path-relinking strategies in comparison with the standard approach are presented in detail. Since the path-relinking uses the recode process, experiments are performed only with decoder D3.

## 4. Computational results

The experiments were conducted on a computer with an AMD FX-8150 Eight-Core 3.6 GHz CPUs, with at least 32 GB of RAM running GNU/Linux, except for experiments with CPLEX which we use a computer with quad-core Intel Xeon E5530 2.4 GHz CPUs, with at least 48 GB of RAM running GNU/Linux. Algorithms are implemented in C++, with optimization flag `-O3`. For BRKGA we use the API described in Toso and Resende (2014). The commercial solver IBM ILOG CPLEX Optimizer version 12.6.0.0 (C++ API) was used to evaluate the mathematical models. All experiments used a single thread but multiple experiments were run in parallel.

Experiments were conducted with two main objectives. The first was to investigate the CPLEX performance considering the different mathematical models described in Section 2. The second was to evaluate the performance of the heuristic approaches described in Section 3 and some hybridisation variants, including metaheuristic approaches, local search strategies, and path-relinking procedures.

Initially we describe the method to generate the dataset used in the experiments. Following we report results for CPLEX, GRASP, BRKGA, and then we report the best results of each metaheuristic method. Finally, we use our best approaches on a set of instances of GQAP, comparing with the state-of-the-art results for the problem described in Mateus et al. (2010).

#### 4.1. Data set

In this subsection we present an instance generator we proposed to generate the data set used in the computational experiments reported by this work. For each instance the generator receives as input four parameters:  $|N|$ ,  $|K|$ ,  $|U|$ , and  $P$  (the latter represents the percentage of the overall data center occupation).

To ensure the generator creates instances that admit feasible solutions, we generate the data for each instance based on  $n$  sets of pre-placed virtual machines to data centers (by default  $n = 3$ ). Given a capacity of each data center, each set of pre-placed  $s \in \mathcal{S}$  is generated by randomly placing each virtual machine to a data center with probability proportional to the data center capacity, ensuring the capacity is not violated. Biased on these pre-placements, we generate the remaining data respecting the constraints of the problem, ensuring at least  $n$  feasible solutions for each instance.

Considering a random number generator using uniform distribution, and  $M'$  be a sufficiently large number, we generate the parameter data for each instance using the following steps.

**Data center capacity:** The total number of available virtual machines is given by  $n' = \max\left(|N|, \left\lceil \frac{|K|}{|P|} \right\rceil\right)$ . Thus, to define the values of  $a_i$  for each DC  $i$ , we start with all  $a_i = 0$  and select  $n'$  times a random data center  $i$ , and increase  $a_i$  by one. We also ensure that each data center has a capacity  $a_i$  greater or equal to one. At this step, the  $n$  pre-placements described before are generated.

**Required virtual machine bandwidth:** For each pair of virtual machine  $v \in K$  and  $w \in K$  the bandwidth  $b_{vw}$  is a random number in the interval  $[0 : 9]$ . This matrix is symmetric, i.e,  $b_{vw} = b_{wv}$ , and  $b_{vv} = 0$ .

**Data center bandwidth:** Having defined the bandwidth between each pair of virtual machines in the previous step, we generate the values of data center bandwidth based in the pre-placements  $\mathcal{S}$ . Let  $b_{ij}^s$  be the sum of bandwidth between all virtual machines pre-placed to  $i$  and  $j$  in  $s \in \mathcal{S}$ . For each pair of data centers  $ij$ , we associate a bandwidth  $B_{ij} = \max\{b_{ij}^s\}$ ,  $\forall s \in \mathcal{S}$ . This matrix also is symmetric, i.e  $B_{ij} = B_{ji}$ , with  $B_{ii} = M'$ .

**Data center latency:** For each pair of data center  $ij$ , the latency  $L_{ij}$  is a random number selected in the interval  $[5 : 20]$ . This matrix is symmetric, i.e,  $L_{ij} = L_{ji}$ , with  $L_{ii} = 0$ .

**Required virtual machine latency:** Let  $l_{vm}^s$  be the latency  $L_{ij}$  between the data centers  $ij$  where  $v$  is placed in  $i$  and  $w$  is placed in  $j$  in the pre-placement  $s \in \mathcal{S}$ . We randomly select  $n = |K| * 2$  distinct pairs  $vw$  to associate a required latency  $l_{vw} = \max\{l_{vm}^s\}$ ,  $\forall s \in \mathcal{S}$ . The remaining latency  $l_{vw}$  is defined as  $M'$ , indicating that no latency is required. We also ensure that  $l_{vw} = l_{wv}$ , and  $l_{vv} = 0$ .

**Users in data centers:** Users are allocated at random to data centers chosen with probability proportional to their capacity. More than one user can be located at the same data center.

**Required user latency:** For each user, we randomly select a virtual machine  $v$  to define a required latency. Let  $i(s)$  be the data center where  $v$  is placed in  $s \in \mathcal{S}$ , thus the

required latency between  $u$  and  $v$  is given by  $t_{vu} = \max\{L_{d(u),i(s)}\}, \forall s \in \mathcal{S}$ . The remaining user required latency  $t$  is set to  $M'$ .

**Transferring data center cost:** For each pair of data center  $ij$ , the cost  $C_{ij}$  is a random number in the interval  $[10.00 : 100.00]$ . This matrix is symmetric, i.e  $C_{ij} = C_{ji}$ , and  $C_{ii} = 0$ .

The parameters  $|N|$  and  $|K|$  can be used to define the instance size, while parameters  $|U|$  and  $P$  can be used to adjust how the problem should be restricted. Finally, we generate a set of instances by combining values from  $N = \{5, 10, 25\}$ ,  $K = \{15, 20, 25, 50, 100, 150, 200\}$ ,  $U = \{K_i * 0.5, K_i, K_i * 1.5\}$ , and  $P = \{70, 90\}$ . Tables 1 and 3 contain all instances we generated, and the values of  $|N|$ ,  $|K|$ ,  $|U|$  and  $P$  are encoded in the name of instance in this respective order. All instances and their best known solutions are available at [www.inf.ufsm.br/~stefanello/instances/](http://www.inf.ufsm.br/~stefanello/instances/).

Part of the objective function from Equation (1) is in fact not used for these instances since we are not considering the fix cost  $c_{iv} \in \mathbb{R}$  of place a VM to a DC. Thus, without loss of generality, for these instances of the VMPlacement problem we may assume  $c_{iv} = 0, \forall i \in N$  and  $v \in K$ . However, these values can be different for GQAP instances, or a more general case of VMPlacement instances that consider installation cost.

## 4.2. CPLEX results

In this section we investigate the CPLEX performance for both linear mathematical models described in Section 2. We carried out an empirical study to analyze which size of instances the CPLEX can handle and prove optimality, and which model provides better integer solutions and lower bounds when CPLEX ends by a time limit. The objective is also to obtain baseline results for comparison of heuristic methods. We first analyze the results for the small size instances, followed by the set of the medium size instances and large size instances. We run CPLEX for QMVMP model. Since the performance of linear models was considered better than for the quadratic model, we report only results for the mixed-integer linear model.

### 4.2.1. Results for small size instances

In the first experiment we evaluated the performance of CPLEX with the mathematical models described in Section 2. We used the standard CPLEX solvers for models LMVMP and LMVMPII. The time limit was set to three hours per run (10,800 seconds) and the number of threads was set to one. The remaining parameters were maintained on the default values.

In Fischetti and Monaci (2014) the authors exploited erraticism in search and how to take advantage of this behavior. Also, the authors show that CPLEX can have a wide contrast in its behavior due randomized initial conditions. Thus, to better evaluate the CPLEX performance we perform ten runs for each instance, each one with a different random seed defined by the CPLEX parameter `RandomSeed`.

Table 1 shows CPLEX results. The first column shows the name of the instances. The second column (BKS) shows the objective function of the best known solution value for each instance (optimal solutions are showed in boldface). The next two columns show the average

running times of CPLEX to solve each instance for each mathematical model. Numbers in parenthesis denote the number of runs that the solver did not prove optimality within the time limit. The next two columns show the average running times that CPLEX spent to find the BKS value. The average is over ten runs or the number of times that CPLEX found BKS (indicated by the number in parenthesis). A signal ‘–’ indicates that no run found a solution as good as BKS within the time limit. Finally, the last two columns show the best running times over all runs that CPLEX spent to find the BKS.

Table 1: CPLEX detailed results for small instances.

Instance	BKS	AVG Time(s)		AVG BKS Time(s)		MIN BKS Time(s)	
		LMVMP	LMVMPII	LMVMP	LMVMPII	LMVMP	LMVMPII
05_015_007_70	<b>25,844.02</b>	2.9	6.2	2.3	2.2	1.1	0.4
05_015_007_90	<b>23,557.30</b>	26.0	61.9	17.4	15.4	10.2	1.4
05_015_015_70	<b>10,904.78</b>	0.4	1.0	0.3	0.4	0.2	0.3
05_015_015_90	<b>24,354.96</b>	4.2	6.9	3.1	1.8	2.3	0.8
05_015_022_70	<b>14,163.60</b>	0.2	0.7	0.1	0.1	0.1	0.0
05_015_022_90	<b>32,318.02</b>	22.4	21.3	21.2	5.2	12.7	0.9
05_020_010_70	<b>38,572.62</b>	292.8	701.7	232.4	138.1	6.4	1.2
05_020_010_90	<b>64,710.80</b>	68.3	107.3	58.5	24.7	31.2	1.9
05_020_020_70	<b>55,288.76</b>	134.5	437.6	100.9	212.4	10.5	37.8
05_020_020_90	<b>57,574.90</b>	1.7	2.1	1.3	0.4	0.8	0.4
05_020_030_70	<b>28,433.34</b>	11.1	17.4	7.1	3.7	2.8	0.4
05_020_030_90	<b>66,088.70</b>	1.9	2.6	1.8	0.7	1.2	0.3
05_025_012_70	<b>43,300.76</b>	1271.5	4,025.5	1,063.8	862.6	195.0	2.2
05_025_012_90	100,865.02	10,800 (10)	10,800 (10)	9,902.6 (4)	6,165.6 (8)	9,746.0	28.7
05_025_025_70	<b>42,890.40</b>	58.4	126.7	52.5	39.0	26.9	1.3
05_025_025_90	103,791.96	10,800 (10)	10,800 (10)	9,764.3 (1)	–	9,764.3	–
05_025_037_70	<b>97,335.12</b>	161.4	592.8	117.0	161.2	19.6	2.8
05_025_037_90	<b>73,363.56</b>	10,262 ( 7)	10,800 (10)	8,993.2	3,782.8 (9)	5,916.3	33.8

We can draw three main observations from this experiment. First, for both models, the solver tends to increase significantly the time spent to prove optimality for instances with 5 or more data centers, and 25 or more virtual machines. This indicates a baseline for the size of instances that this version of CPLEX can handle and prove optimality using these models in this computer. The second observation is that CPLEX performance for both models was relatively similar. CPLEX was able to prove or not the optimality in the same instances, except for instance 05\_025\_037\_90 which CPLEX proved optimality in three runs for LMVMP, while for LMVMPII CPLEX has an average gap of 8%. However, considering the time that each model spent to find a solution with objective function equal to BKS, in most cases CPLEX with LMVMPII spent less time than considering the model LMVMP. Finally, the last observation is about the variance of time to find the BKS. With the randomized initial conditions, CPLEX presented a large variance in its behavior. For example, for the instance 05\_025\_012\_70 and LMVMPII, in the best case CPLEX found BKS in 2.2 seconds, while in the worst case it took 3,458.01 seconds to find the same solution. For this instance, the time to prove optimality also have large variance, alternating between a minimum of 3,685.67 and a maximum of 4,828.05.

We also evaluate the model LMVMPII without including the additional set of cuts (30). In comparison with the model with the cuts, we observed a higher decrease in the performance of CPLEX. For six instances CPLEX was not able to prove optimality in any of the

ten runs. For these cases, the gap of CPLEX is still high. For example, for the instance 05\_025\_012\_90, the average gap was 53%, confirming that the relaxation quality of this model is low. Furthermore, the time to find BKS increased around 8 times, and the number of nodes explored within the time limit or to prove optimality increased around 7 times in comparison with the model that includes the set of cuts (30).

#### 4.2.2. Results for medium and large size instances

In this subsection we present results for CPLEX when applied to a set of median and large instances. In this experiment, we evaluate CPLEX for each model and each instance with one run for a time limit of one day (86,400 seconds).

Table 2 shows for each instance the BKS obtained over all experiments, including the reported in this work for heuristic methods and additional non-reported experiments used to evaluate the previous version of the algorithms. The column FO shows, for each model, the objective function value of the best integer solution found by CPLEX. Column CPLEX GAP shows the gap returned by CPLEX. The last column shows the gap from BKS to the best integer solution returned by CPLEX.

The first observation from this experiment is that for instances with 10 data centers, CPLEX was able to start solve the models, but still with a high gap even after 24 hours of computation. However, we observed that the gap returned by CPLEX as well as the gap to BKS are lower for LMVMPII in comparison with the values of LMVMP. The average gap returned by CPLEX with LMVMPII was around half of the average gap for LMVMP. The average gap to BKS was 2.19 for LMVMPII, while for LMVMP the gap was 4.96.

The second observation is that for instances with 25 data centers and LMVMP model, CPLEX spent the whole time in the presolve phase without solving the root relaxation node. Instead for LMVMPII, CPLEX was able to start a node exploration and found feasible solutions. For instances with 25 data centers, CPLEX explored an average of approximately 34800 nodes for instances with 100 VMs, 5700 nodes for instances with 150 VMs, and 1000 nodes for instances with 200 VMs. Even CPLEX maintaining a high gap, the gap from BKS was relatively small, considering the size of the instance and the difficulty to find feasible solutions (Stefanello et al., 2015a).

#### 4.3. GRASP results

This subsection presents results and analyzes feasibility for different constructive heuristic combined with different local search procedures embedded in the GRASP framework. Also, we report experiments with the path-relinking procedure.

The main goals of the following experiments are show how difficult is to find feasible solutions using a constructive heuristic, and explore the effect of embedding a local search procedure to obtain a feasible solution. Finally, the last main objective is to evaluate the impact of each component added to the framework, i.e., constructive heuristic, local search method, and path-relinking strategy.

In the first part of our experiment, we evaluate the performance of each constructive heuristic described in Subsection 3.3, embedded in the GRASP algorithm. We performed one run for each instance, combination of constructive heuristic, and local search method.

Table 2: CPLEX detailed results for median and large instances.

Instance	FO			CPLEX GAP		BKS GAP	
	BKS	LMVMP	LMVMPII	LMVMP	LMVMPII	LMVMP	LMVMPII
10_025_012_70	114,582.50	116,264.44	115,734.58	42.37	24.29	1.47	1.01
10_025_012_90	84,461.30	88,087.12	85,814.72	53.92	28.31	4.29	1.60
10_025_025_70	90,997.90	93,729.48	92,407.94	29.59	19.30	3.00	1.55
10_025_025_90	124,763.66	125,365.26	125,335.92	31.18	16.56	0.48	0.46
10_025_037_70	100,801.80	104,350.38	100,801.80	24.16	15.04	3.52	0.00
10_025_037_90	106,617.94	107,558.00	107,471.04	24.06	20.06	0.88	0.80
10_050_025_70	414,535.12	442,548.40	435,929.26	83.22	39.98	6.76	5.16
10_050_025_90	458,879.74	480,146.00	477,439.06	75.60	30.79	4.63	4.04
10_050_050_70	360,102.12	374,071.02	366,972.40	68.85	43.23	3.88	1.91
10_050_050_90	400,233.16	420,173.88	416,615.52	72.88	36.32	4.98	4.09
10_050_075_70	349,135.78	362,853.92	353,979.20	54.03	35.20	3.93	1.39
10_050_075_90	498,190.58	513,161.02	510,062.50	54.45	25.23	3.01	2.38
10_100_050_70	1,647,975.00	1,884,262.62	1,732,985.94	91.55	42.59	14.34	5.16
10_100_050_90	1,792,257.68	1,916,126.76	1,826,484.56	89.56	34.79	6.91	1.91
10_100_100_70	1,463,498.00	1,546,897.78	1,500,763.50	86.22	47.12	5.70	2.55
10_100_100_90	2,126,993.26	2,256,408.46	2,169,460.06	82.73	31.03	6.08	2.00
10_100_150_70	1,563,152.40	1,702,573.74	1,586,082.86	78.44	45.02	8.92	1.47
10_100_150_90	1,847,076.66	1,968,341.96	1,883,289.84	70.74	33.59	6.57	1.96
<b>Average</b>	-	-	-	61.86	31.58	4.96	2.19
25_100_050_70	1,887,688.24	-	1,976,649.48	-	44.15	-	4.71
25_100_050_90	2,116,849.00	-	2,184,190.20	-	34.89	-	3.18
25_100_100_70	1,953,155.20	-	2,056,428.04	-	44.55	-	5.29
25_100_100_90	2,021,228.76	-	2,091,723.56	-	35.72	-	3.49
25_100_150_70	1,967,364.52	-	2,061,181.36	-	43.91	-	4.77
25_100_150_90	2,160,014.54	-	2,235,936.02	-	35.35	-	3.51
25_150_075_70	4,603,163.50	-	4,768,514.02	-	42.99	-	3.59
25_150_075_90	4,618,491.80	-	4,786,179.50	-	36.54	-	3.63
25_150_150_70	3,882,650.94	-	4,085,463.10	-	46.34	-	5.22
25_150_150_90	4,706,129.66	-	4,899,667.66	-	36.82	-	4.11
25_150_225_70	4,340,090.18	-	4,530,866.46	-	46.00	-	4.40
25_150_225_90	4,523,393.44	-	4,708,994.42	-	35.51	-	4.10
25_200_100_70	6,937,008.98	-	7,539,073.36	-	50.79	-	8.68
25_200_100_90	9,034,147.98	-	9,391,760.06	-	34.93	-	3.96
25_200_200_70	7,146,330.32	-	7,579,453.48	-	49.13	-	6.06
25_200_200_90	8,578,620.94	-	8,909,349.84	-	36.53	-	3.86
25_200_300_70	7,638,447.16	-	8,033,640.92	-	42.32	-	5.17
25_200_300_90	8,195,152.30	-	8,579,652.20	-	38.94	-	4.69
<b>Average</b>	-	-	-	-	40.86	-	4.58

The stopping criterion of each run was a time limit of  $|K| * |N| * \theta$  seconds, where  $\theta = 0.8$ . We use this stopping criterion in all experiments in order to provide a fair comparison for all strategies evaluated.

Table 3 shows the average of the percentage of feasible solutions found on each run of the algorithm for all instances and combinations of local search (rows on first column) and constructive heuristic (columns 2 to 5).

Table 3: Percentage of feasible solutions found by GRASP for different constructive heuristic and local search strategy.

<b>LSType</b>	<i>Random</i>	<i>Greedy</i>	<i>DC-Greedy</i>	<i>VM-Greedy</i>	<b>Average</b>
NoLS	0.00	1.02	1.96	0.57	0.89
1V	23.70	22.12	27.88	21.96	23.91
2V	77.26	78.02	79.72	78.41	78.35
3V	82.93	82.66	84.06	83.20	83.21
4V	84.08	83.61	84.98	84.19	84.22
Average	53.59	53.49	<b>55.72</b>	53.67	54.12

The first observation from Table 3 is that the probability of finding a feasible solution using only a constructive heuristic without local search is low (row NoLS). In average, the percentage of feasible solutions was less than 1%, with an average of less than 2% for the constructive heuristic *DC-Greedy*. However, the main observation is that the constructive heuristic has a lower impact on the percentage of feasible solution in comparison with the impact of any local search method. When a local search procedure is considered (also used as repair procedure), the percentage of feasible solutions increase considerable. For example, using 2V the percentage of feasible solutions is higher than 77% or higher than 82% for 3V. This occurs even with the random constructive heuristic, showing that the local search has a performance that is not high-dependent of the initial solution. We also analyze the quality of the feasible solutions found on each strategy. The conclusions are similar to the percentage of feasible solutions, where the solution quality are similar for all constructive heuristic and tends to be better when we consider a more complex neighborhood search. Thus, in the remaining experiments we adopted the *DC-Greedy* constructive heuristic as default since it presents the highest percentage of feasible solutions.

Table 4 shows the percentage of feasible solutions for each instance for the constructive heuristic *DC-Greedy*.

Even with a significant increase in the percentage of feasible solutions found when the local search is applied, some instances still have a low percentage of feasible solutions. This shows that the set of instances used to evaluate the algorithms is diverse and not trivial to solve. We also observed that, as expected, most instances with a higher percentage of data center occupation has a low percentage of feasible solutions in comparison with the respective instance with a low percentage of occupation.

In the next set of experiments, the main objective is to evaluate the path-relinking component in the GRASP procedure. For each experiment, we performed five independent

Table 4: Percentage of feasible solutions found by GRASP with the *DC-Greedy* constructive heuristic.

Instance	NoLS	1V	2V	3V	4V	Instance	NoLS	1V	2V	3V	4V
10_025_012_70	10.90	60.50	90.88	95.75	96.24	25_100_050_70	8.10	77.11	95.89	98.03	98.46
10_025_012_90	1.27	24.57	92.25	95.68	96.18	25_100_050_90	0.00	13.16	84.37	87.90	89.16
10_025_025_70	25.75	92.23	99.44	99.78	99.86	25_100_100_70	0.56	40.15	88.12	91.93	91.82
10_025_025_90	0.08	6.09	52.23	62.96	64.72	25_100_100_90	0.00	3.08	66.90	74.69	75.48
10_025_037_70	2.91	59.80	81.60	86.55	88.67	25_100_150_70	1.23	46.98	93.80	96.42	97.31
10_025_037_90	0.06	4.93	44.24	57.58	63.35	25_100_150_90	0.00	2.78	66.30	73.97	75.14
10_050_025_70	4.38	42.79	94.45	95.80	96.17	25_150_075_70	5.84	60.39	95.35	96.03	97.12
10_050_025_90	0.13	21.10	88.80	92.63	93.11	25_150_075_90	0.01	6.60	84.39	87.73	88.13
10_050_050_70	0.34	23.77	71.09	76.97	78.78	25_150_150_70	0.77	31.90	79.68	84.00	85.29
10_050_050_90	0.05	13.24	56.53	65.21	65.81	25_150_150_90	0.00	5.23	64.75	70.49	69.27
10_050_075_70	1.50	49.55	72.63	77.26	78.96	25_150_225_70	0.53	31.81	80.36	83.15	83.33
10_050_075_90	0.07	10.02	68.59	72.29	74.02	25_150_225_90	0.00	1.59	47.46	57.90	56.01
10_100_050_70	3.34	65.63	98.49	98.96	98.62	25_200_100_70	0.21	31.49	85.27	86.11	90.86
10_100_050_90	0.01	3.07	80.60	86.54	86.99	25_200_100_90	0.00	21.90	95.36	97.76	97.36
10_100_100_70	1.66	39.19	87.45	91.28	93.38	25_200_200_70	0.07	22.82	75.00	80.57	81.38
10_100_100_90	0.00	7.45	84.02	86.81	86.66	25_200_200_90	0.00	21.36	93.51	95.41	95.36
10_100_150_70	0.82	41.94	83.35	86.01	86.89	25_200_300_70	0.00	13.36	66.92	72.53	73.55
10_100_150_90	0.00	1.46	73.33	74.18	75.63	25_200_300_90	0.00	4.49	86.54	89.12	90.08
Average	2.96	31.52	78.89	83.46	84.67	Average	0.96	24.23	80.55	84.65	85.28

runs totalling an amount of 180 samples for each experiment. We also use  $\theta = 0.8$  as stopping criterion. The local search strategy used in the path-relinking is always the same used after the construction phase in GRASP procedure. Experiments are chosen to analyze the performance for different combination of algorithms and to show the impact of each component.

To compare the results with respect to the best solution found on each run, we first scale the objective function value to the range  $[0, 1]$  since each instance can have very different values. The scale is a simple transformation where for each instance, the largest cost over all analyzed experiments is scaled to 1 and the lowest is scaled to 0.

Figure 3 shows the distribution of the scaled cost for each algorithm. The box plots show the location of the minimum value, lower quartile, median, upper quartile, and maximum value of each experiment. The dots are the outliers. Experiments are represented on the horizontal axis and are labelled as composition of main algorithm name, path-relinking type, local search strategy, and size of elite set parameter. Experiments ended with “\*” use a full exploration on the neighborhood *chain3L*, i.e,  $\alpha_{3L} = 1$  and  $\beta_{3L} = |N|$ . Also, note that the first three experiments are a simple GRASP procedure and do not include the path-relinking component.

The first observation is about the relation between the intensification of the neighborhood exploration and restarting the search. On the one hand, from Table 3 and 4 we observe that the local search that includes more complex neighborhoods overcomes in most cases the searches with less completely neighborhoods in the percentage of feasible solutions. Since the strategies are executed for the same time, is natural that a less complex search has a higher number of iterations and a higher numbers of restarts. This is an indication that

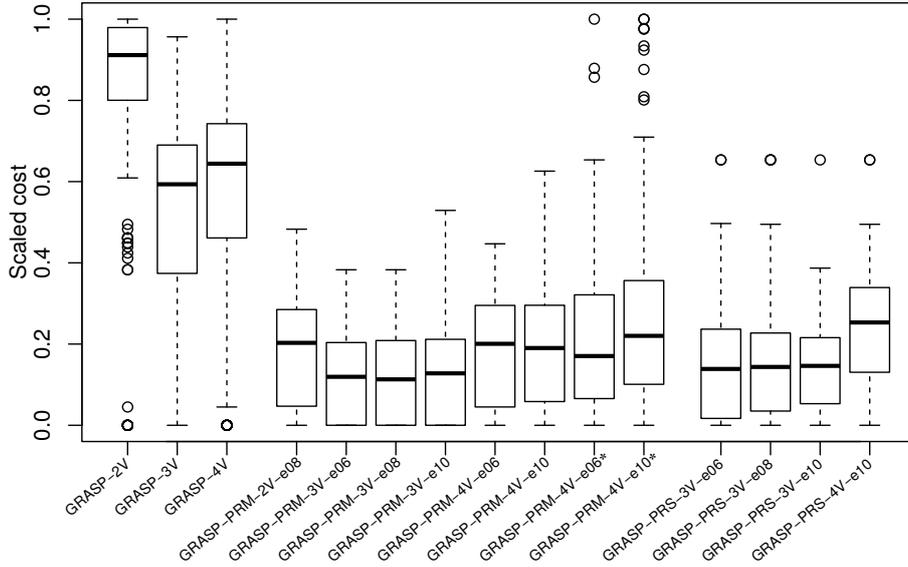


Figure 3: Dispersion of scaled cost for each algorithm.

for this problem is preferable to invest in the exploration of the neighborhood of a high-quality solution instead of making a full restart to a new point of the search space. On the other hand, even with a higher ability of the local search 4V in producing feasible solutions and exploring a large neighborhood, when we consider the best solution found in the experiments, the local search 3V produces better results than the search 4V. This indicates that for the proposed neighborhood strategies, there is a trade-off between intensification of the neighborhood exploration and time spent on each local search strategy. This can be used to guide the search and select the best strategies for a general proposed method.

The second observation is that the path-relinking component improves significantly the results in comparison with the case without this strategy (case GRASP-2V, GRASP-3V, and GRASP-4V). We also observe that both path-relinking strategies have similar performance. Furthermore, the results have a low influence by the values of elite size in the tested interval.

To confirm the results presented in Figure 3, we use the R package to test the normality of these distributions using the Shapiro-Wilk test and apply the Mann-Whitney-Wilcoxon U test. For all tests, we assume a confidence interval of 99%. Shapiro-Wilk tests indicate that no cost distribution fits a normal distribution since the p-values for all tests are less than 0.01. Therefore, we applied the U test which assumes as null hypothesis that the location statistics are equal in both distributions. We also use a p-value correction procedure based on false discovery rate (FDR) to minimize the number of false positives.

Table 4.3 shows U test results for each pair of algorithms. The structure of this table is as follows: Each row and column is indexed by one algorithm. Each element in the diagonal (bold) is the median of the scaled cost of the corresponding algorithm. The upper-right diagonal elements are the differences in location statistics for each pair of algorithms. A negative difference indicates that the “row algorithm” has its location statistics lower

(better) than the “column algorithm”, and the positive difference is the opposite. The bottom-left diagonal elements are the p-values of each test. Math signals indicate when  $p < 0.01$  for a U test between “row algorithm” and “column algorithm” for the respective signal alternative hypothesis. The case “less” indicates that the “row algorithm” overcomes the “column algorithm”, or the opposite in the case “greater”.

Table 5: Values of medians, p-values, and difference in median location for cost distributions using a confidence interval of 99% for GRASP algorithm.

	GRASP			GRASP-PRM								GRASP-PRS			
	2V	3V	4V	2V-e08	3V-e06	3V-e08	3V-e10	4V-e06	4V-e10	5V-e06	5V-e10	3V-e06	3V-e08	3V-e10	4V-e10
2V	<b>0.911</b>	0.311	0.265	0.683	0.755	0.756	0.751	0.683	0.679	0.676	0.642	0.730	0.728	0.734	0.647
3V	<	<b>0.593</b>	-0.048	0.379	0.452	0.456	0.448	0.387	0.379	0.360	0.332	0.424	0.422	0.427	0.342
4V	<	>	<b>0.644</b>	0.426	0.496	0.498	0.489	0.434	0.421	0.416	0.389	0.468	0.465	0.476	0.393
PRM-2V-e08	<	<	<	<b>0.203</b>	0.075	0.081	0.069	-0.001	0.003	-0.014	-0.051	0.056	0.053	0.061	-0.035
PRM-3V-e06	<	<	<	<	<b>0.119</b>	0.002	-0.006	-0.077	-0.075	-0.085	-0.121	-0.022	-0.029	-0.017	-0.118
PRM-3V-e08	<	<	<	<	0.839	<b>0.113</b>	-0.010	-0.077	-0.082	-0.086	-0.125	-0.027	-0.028	-0.022	-0.119
PRM-3V-e10	<	<	<	<	0.568	0.353	<b>0.128</b>	-0.072	-0.075	-0.082	-0.116	-0.020	-0.017	-0.012	-0.110
PRM-4V-e06	<	<	<	0.930	>	>	>	<b>0.201</b>	-0.007	-0.014	-0.049	0.051	0.054	0.058	-0.039
PRM-4V-e10	<	<	<	0.846	>	>	>	0.582	<b>0.190</b>	-0.015	-0.045	0.056	0.055	0.058	-0.037
PRM-5V-e06	<	<	<	0.409	>	>	>	0.376	0.294	<b>0.171</b>	-0.041	0.063	0.058	0.067	-0.031
PRM-5V-e10	<	<	<	>	>	>	>	>	>	>	<b>0.220</b>	0.100	0.093	0.096	0.006
PRS-3V-e06	<	<	<	<	0.075	0.040	0.089	<	<	<	<	<b>0.139</b>	-0.004	0.003	-0.097
PRS-3V-e08	<	<	<	<	0.016	0.010	0.091	<	<	<	<	0.737	<b>0.144</b>	0.009	-0.097
PRS-3V-e10	<	<	<	<	0.065	0.066	0.243	<	<	<	<	0.774	0.384	<b>0.146</b>	-0.102
PRS-4V-e10	<	<	<	>	>	>	>	>	>	0.056	0.774	>	>	>	<b>0.253</b>

Table 4.3 confirms that the difference between the results of the standard GRASP procedure and the procedure that includes the path-relinking component are statistically significant for a confidence interval of 99%. Tests between cases with local search 3V and different elite sizes indicate that the difference is not significant, even for both path-relinking strategies. However analyzing the values of median and median location, we conclude that there is a low advantage for the PRM strategy with elite size equal to 8.

A report for the solution quality obtained with this strategy is showed in the Subsection 4.5. Next we report results for the BRKGA.

#### 4.4. BRKGA results

This subsection presents results for the biased random-key genetic algorithm. The objective is to analyze different procedures, decoders, the recode process, and the path-relinking procedure introduced in the algorithm. The experiments follow the rules and the objectives explained in the previous subsection.

BRKGA parameters was set to the same parameters as described in Stefanello et al. (2015a) which are similar to the values suggested by Gonçalves and Resende (2011), i.e., the elite size  $p_e = 0.24p$ , the set of mutants  $p_m = 0.12p$ , and the probability of inheriting  $\rho_A = 0.6$ . The restart parameter was disable and the population size  $p = 75$ . The stopping criterion for all runs was a time set to  $|K| * |N| * \theta$  seconds (where  $\theta = 0.8$ ).

We run experiments that include different local searches, decoders and path-relinking procedures. In each experiment, we performed five independent runs with different random seeds for each instance, given a total of 180 runs. To compare the results we scale the values of the objective function for the best solution found on each run to the range  $[0, 1]$ , as described in the previous subsection.

Figure 4 shows the box plot for the distribution of the scaled cost for each algorithm. Algorithms are represented on the horizontal axis and are labelled as a composition of main algorithm name, path-relinking strategy, decoder method, local search strategy, and size of elite set parameter.

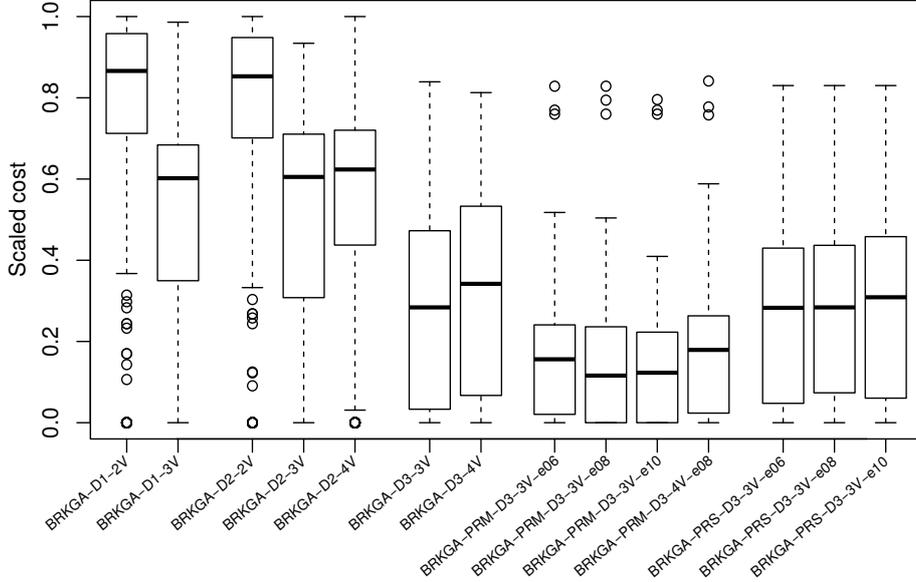


Figure 4: Dispersion of scaled cost for each algorithm.

The first observation is that the recode process helps to improve the results, since all experiments with decoder D3 have medians lower than the respective correspondent algorithm with decoders D1 and D2. Note for example that without the recode, one modification in a random key can influence in one or more placements (since the local search can apply many changes). Using the recode, when one key is changed, the modification is in only one placement. This well-defined behavior helps to keep a more accurate information of each individual. The second observation is that for the local search strategies, the same conclusions from the GRASP can be taken for BRKGA. Experiments with local search 3V provides better results than the other local search strategies. Finally, strategies that include the path-relinking procedure are more effective to produce better results than the standard counterpart. However, differently as observed in GRASP, the difference between PRM and PRS is significant. We attribute the worst performance of PRS due to the small number of applications of the path-relinking since it is limited to the number of generations. Also, after some iterations of BRKGA, the solutions in the elite set of BRKGA tend to be similar, and fail in the similarity check with the solution in the elite set of path-relinking, and thus, the path-relinking operator is not performed.

Table 6 shows U test results for each pair of algorithms. We present the values of medians, p-values, and the difference in median location for the scaled cost distributions using a confidence interval of 99% for all experiments. The organization of this table follows the description of Table 4.3.

Table 6: Values of medians, p-values, and difference in median location for cost distributions using a confidence interval of 99% for BRKGA algorithm.

	BRKGA							BRKGA-PRM				BRKGA-PRS		
	D1-2V	D1-3V	D2-2V	D2-3V	D2-4V	D3-3V	D3-4V	3V-e06	3V-e08	3V-e10	4V-e08	3V-e06	3V-e08	3V-e10
D1-2V	<b>0.866</b>	0.251	0.000	0.242	0.211	0.512	0.449	0.678	0.697	0.701	0.645	0.533	0.520	0.515
D1-3V	<	<b>0.602</b>	-0.247	0.000	-0.028	0.250	0.185	0.413	0.437	0.440	0.397	0.275	0.263	0.261
D2-2V	0.987	>	<b>0.853</b>	0.243	0.200	0.505	0.440	0.669	0.691	0.692	0.640	0.524	0.512	0.506
D2-3V	<	0.979	<	<b>0.605</b>	-0.036	0.250	0.188	0.424	0.443	0.449	0.397	0.272	0.265	0.260
D2-4V	<	0.016	<	>	<b>0.623</b>	0.275	0.216	0.449	0.470	0.477	0.424	0.300	0.291	0.287
D3-3V	<	<	<	<	<	<b>0.284</b>	-0.062	0.179	0.208	0.202	0.142	0.031	0.014	0.018
D3-4V	<	<	<	<	<	>	<b>0.342</b>	0.233	0.266	0.263	0.206	0.091	0.073	0.075
PRM-3V-e06	<	<	<	<	<	<	<	<b>0.156</b>	0.020	0.023	-0.026	-0.151	-0.163	-0.169
PRM-3V-e08	<	<	<	<	<	<	<	0.074	<b>0.116</b>	0.001	-0.042	-0.176	-0.186	-0.190
PRM-3V-e10	<	<	<	<	<	<	<	0.040	0.891	<b>0.123</b>	-0.046	-0.179	-0.194	-0.193
PRM-4V-e08	<	<	<	<	<	<	<	0.047	>	>	<b>0.179</b>	-0.118	-0.129	-0.133
PRS-3V-e06	<	<	<	<	<	0.011	<	>	>	>	>	<b>0.283</b>	-0.015	-0.015
PRS-3V-e08	<	<	<	<	<	0.171	<	>	>	>	>	0.102	<b>0.284</b>	-0.003
PRS-3V-e10	<	<	<	<	<	0.071	<	>	>	>	>	0.118	0.810	<b>0.309</b>

This table confirms the considerations described above showing that the best approaches uses the path-relinking component (PRM). Also, as mentioned for GRASP in the previous subsection, the experiments for different elite sizes indicate that the differences are not statistical significant for a confidence interval of 99%. However, analyzing the values of the median, and median location we conclude that using the size of elite set equal to 8, the procedure tends to produce better results.

In [Stefanello et al. \(2015a\)](#) this problem was first introduced, and the first version of BRKGA is presented. From the original approach, three main improvements are made. First, we added a new neighborhood search. In the original approach, we described a local search that includes only shift and swap moves (namely LSW that correspond to 2V in this work). We observe from the results in [Figure 4](#) and in the [Table 6](#) that the inclusion of this new neighborhood strategy (*chain search*) helps to improve significantly the results (comparison between BRKGA-D2-2V and BRKGA-D2-3V). Second, the inclusion of the recode procedure in the decoder D2 (defined as decoder D3), also improves significantly the results (comparison between BRKGA-D2-2V and BRKGA-D3-2V, as well as, BRKGA-D2-3V and BRKGA-D3-3V). Finally, we included a path-relinking component that in the best of our knowledge, is the first experiment that considers this hybrid approach between BRKGA and path-relinking.

A final observation about running times of BRKGA is that results are reported using single-thread in order to provide a fair comparison with the GRASP and CPLEX results. However, the BRKGA API provides an efficient multi-thread decoding [Toso and Resende \(2014\)](#), that could be used to reduce substantially the running times when multiple processors are available.

#### 4.5. Additional comparison

In this subsection we report an overview over the best strategies for BRKGA and GRASP, providing addition comparisons and informations for both algorithms. [Figure 5](#) shows an overview for the experiments for both strategies reported in the previous subsections.

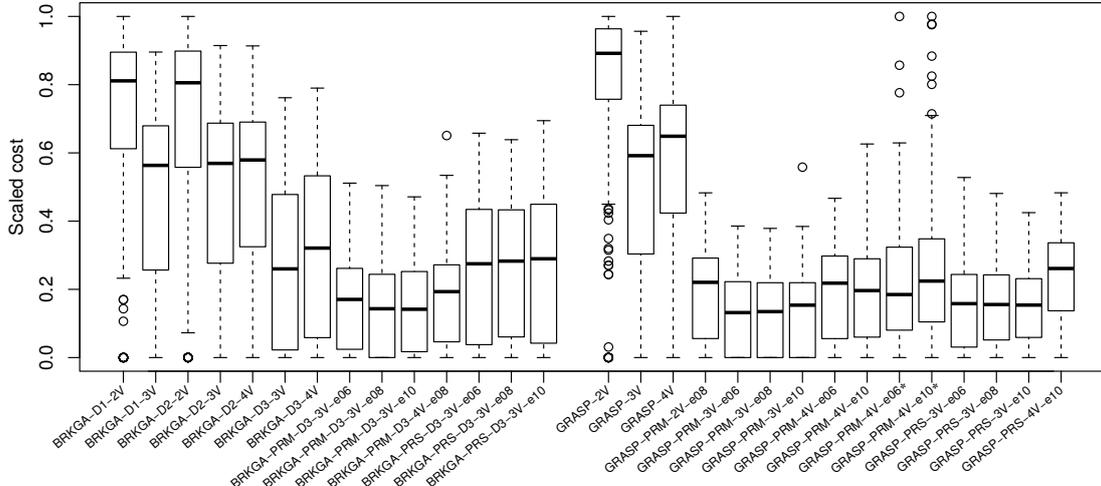


Figure 5: Dispersion of scaled cost for each algorithm.

As depicted in Figure 5 and reported in subsections 4.3 and 4.4, the best approach for both algorithms is consider the local search 3V and to include the path-relinking component with elite size equal to 8. Thus, Table 7 presents supplementary information for the experiments with both strategies. The second column shows the best-known solution value (BKS) for each instance. The next columns show the minimum, average and maximum percentage gap for both algorithms. The percentage gap is calculated by the formula  $(FO - BKS) * 100 / BKS$ , were FO is the respective value of the objective function for the considered algorithm. Finally, columns Time shows the minimum and average time in seconds in which each algorithm obtained the last improvement in the objective function.

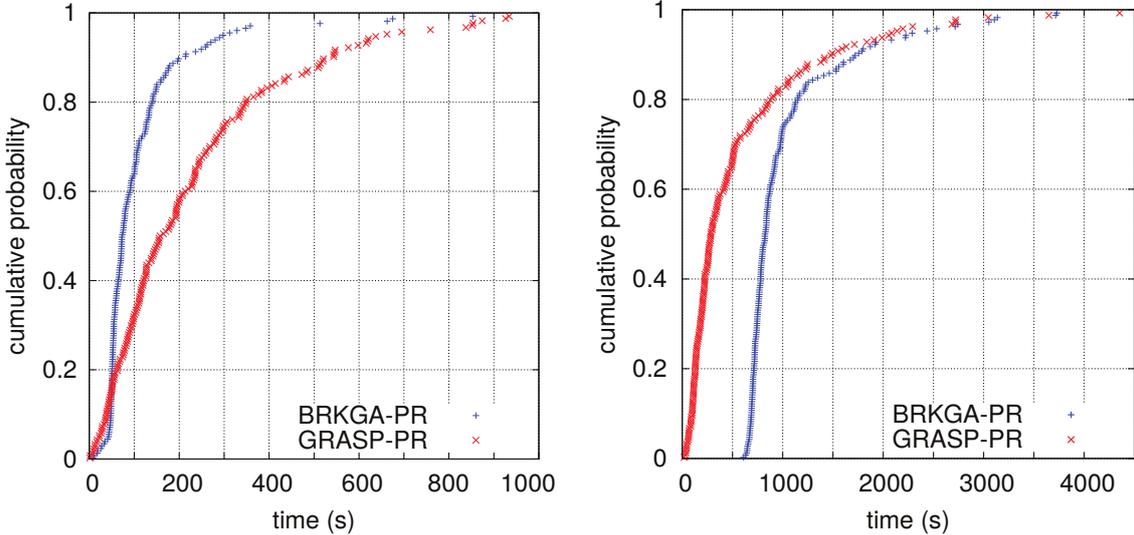
From this table, we observe that the percentage gap is relatively small for both algorithms, showing that the approaches are efficient to find good quality solutions. For instances with 10 data centers, in most cases both algorithms found the BKS. Also, considering that each run was stopped by time limit ( $|N| * |K| * 0.8$ ), from columns Time we observe that a significant parcel of the execution was spent without improvements. This indicates that the algorithm has a fast convergence, and will probably provide a good solution if stopped before the used time limit. Also, we observe that BRKGA-PR is slightly faster than GRASP-PR for the instance set with 10 data centers. However, the opposite happens for the instances with 25 data centers.

Finally, the last experiment uses the Time-To-Target (TTT) plots to display the running time distribution for the algorithm to find a solution at least as good as a given target value. TTT plots were proposed by Aiex et al. (2007) and have been advocated by Hoos and Stützle (1998b;a) as a way to characterize the running times of stochastic algorithms for combinatorial optimization problems. The experiment consists in performing 200 runs of BRKGA-PR and GRASP-PR algorithms for two instances until a target is reached. The target was set to the worst solution found for both strategies on the previous experiment, i.e, 1,468,973 for 010\_100\_100\_70 and 9,077,672 for 025\_200\_100\_90. The instances chosen

Table 7: Comparison of percentage gap and last improve time for BRKGA-PR and GRASP-PR.

Instance	BKS	BRKGA-PR					GRASP-PR				
		GAP (%)			Time (s)		GAP (%)			Time (s)	
		min	Avg	Max	Min	Avg	min	Avg	Max	Min	Avg
10_025_012_70	114,582.50	0.00	0.00	0.00	1.6	22.7	0.00	0.00	0.00	1.3	38.7
10_025_012_90	84,461.30	0.00	0.32	0.54	1.8	35.9	0.00	0.00	0.00	16.1	49.2
10_025_025_70	90,997.90	0.00	0.00	0.00	0.5	0.8	0.00	0.00	0.00	48.5	111.3
10_025_025_90	124,763.66	0.00	0.15	0.18	1.9	45.1	0.00	0.04	0.18	19.1	59.8
10_025_037_70	100,801.80	0.00	0.00	0.00	0.7	0.8	0.00	0.00	0.00	0.4	1.0
10_025_037_90	106,617.94	0.00	0.00	0.00	0.7	13.3	0.00	0.00	0.00	1.1	2.3
10_050_025_70	414,535.12	0.11	0.26	0.47	12.5	86.5	0.00	0.08	0.13	47.9	190.9
10_050_025_90	458,879.74	0.00	0.02	0.07	38.3	76.0	0.00	0.09	0.16	87.5	193.6
10_050_050_70	360,102.12	0.00	0.10	0.49	11.4	36.0	0.00	0.00	0.00	4.5	92.4
10_050_050_90	400,233.16	0.00	0.37	0.64	21.2	32.9	0.00	0.28	0.62	32.9	185.6
10_050_075_70	349,135.78	0.00	0.00	0.00	4.7	29.8	0.00	0.00	0.00	7.7	62.1
10_050_075_90	498,190.58	0.39	0.58	0.87	47.9	153.7	0.00	0.26	0.47	18.1	147.8
10_100_050_70	1,647,975.00	0.00	1.98	3.32	82.8	395.4	1.12	2.26	3.37	457.5	652.2
10_100_050_90	1,792,257.68	0.01	0.18	0.31	98.1	380.7	0.14	0.31	0.55	292.5	493.0
10_100_100_70	1,463,498.00	0.00	0.08	0.37	50.9	271.7	0.00	0.03	0.14	518.3	624.6
10_100_100_90	2,126,993.26	0.00	0.23	0.48	433.7	656.5	0.12	0.28	0.53	60.6	483.5
10_100_150_70	1,563,152.40	0.00	0.13	0.40	190.1	437.5	0.06	0.30	0.59	198.7	537.4
10_100_150_90	1,847,076.66	0.01	0.12	0.32	137.2	429.0	0.08	0.26	0.39	252.6	575.8
Average		0.03	0.25	0.47	63.1	172.5	0.08	0.23	0.40	114.7	250.0
25_100_050_70	1,887,688.24	0.18	0.43	0.78	160.3	1,128.9	0.12	0.30	0.47	953.2	1,464.6
25_100_050_90	2,116,849.00	0.03	0.52	0.81	904.3	1,615.4	0.26	0.37	0.46	741.2	1,326.4
25_100_100_70	1,953,155.20	0.36	0.62	0.83	284.5	770.2	0.00	0.35	0.75	253.9	1,122.9
25_100_100_90	2,021,228.76	0.35	0.72	1.24	976.0	1,516.7	0.44	0.55	0.65	847.5	1,207.5
25_100_150_70	1,967,364.52	0.30	0.57	0.82	276.5	1,141.6	0.00	0.42	0.70	160.5	1,377.1
25_100_150_90	2,160,014.54	0.56	0.80	1.08	776.3	1,509.1	0.44	0.62	0.77	528.9	1,156.7
25_150_075_70	4,603,163.50	0.20	0.37	0.49	531.1	1,877.3	0.24	0.33	0.41	992.6	1,567.2
25_150_075_90	4,618,491.80	0.21	0.33	0.49	1,223.7	1,807.9	0.21	0.31	0.40	1,986.0	2,548.8
25_150_150_70	3,882,650.94	0.19	0.36	0.54	788.5	2,183.6	0.19	0.34	0.49	1,560.7	2,156.1
25_150_150_90	4,706,129.66	0.41	0.53	0.68	597.5	1,984.0	0.32	0.46	0.58	889.2	2,067.9
25_150_225_70	4,340,090.18	0.13	0.34	0.53	1,657.6	2,347.1	0.47	0.59	0.80	1,648.5	2,321.8
25_150_225_90	4,523,393.44	0.44	0.54	0.69	271.5	1,684.5	0.17	0.37	0.53	543.2	1,607.5
25_200_100_70	6,937,008.98	0.22	0.28	0.35	3,123.8	3,938.7	0.07	0.26	0.44	1,855.8	2,841.7
25_200_100_90	9,034,147.98	0.15	0.29	0.48	2,017.7	3,352.2	0.15	0.29	0.44	1,371.5	2,688.1
25_200_200_70	7,146,330.32	0.13	0.24	0.44	1,271.1	2,948.6	0.10	0.35	0.52	1,718.1	3,201.1
25_200_200_90	8,578,620.94	0.00	0.15	0.26	1,170.0	2,669.4	0.18	0.22	0.30	326.3	1,681.0
25_200_300_70	7,638,447.16	0.00	0.25	0.35	1,763.6	3,216.9	0.19	0.28	0.39	819.1	2,745.9
25_200_300_90	8,195,152.30	0.05	0.18	0.27	1,523.7	2,689.1	0.13	0.27	0.37	913.1	2,796.6
Average		0.22	0.42	0.62	1,073.2	2,132.3	0.20	0.37	0.53	1,006.1	1,993.3

was one at each group of data centers. We select the instance with the lowest difference of average percentage gap between both approaches (excluding equal values).



(a) TTT plot for 10\_100\_100\_70.

(b) TTT plot for 25\_200\_100\_90.

Figure 6: Cumulative probability distribution.

Figures 6a and 6b illustrate the cumulative probability plot obtained by using BRKGA-PR and GRASP-PR for two instances. For the instance 10\_100\_100\_70, we observe that BRKGA-PR has a higher probability to find a solution at least as good as the target in less time than GRASP-PR. For instance 25\_150\_225\_70 the probability is the opposite.

In summary, both approaches have good performance according to specific scenarios. The main performance of BRKGA-PR is in small and median instances while GRASP-PR has a better performance in large instances. The path-relinking component provides a higher improvement for both strategies. We attribute this performance to a rugged landscape for this problem and the fact that is hard to find feasible solutions. Since in general the path-relinking is performed between two feasible solutions, this component explores a search space where solutions are feasible or has a lower number of infeasibilities. Thus, when the local search is applied to a solution in the path, the search tends to be faster and leads to a new minimum local, that can be better than the initial and the guided solution.

Finally, we evaluate our algorithm considering a set of instances which sizes were not address before in the literature for this problem. We test in instances with up to 25 data centers and 200 virtual machines. Even that in real world application this size can be considered small, previous works on related problems in the literature have addressed instances up to 50 facilities (virtual machines). Thus, we consider a significant step since we develop solutions for instances five times larger than the previously considered size. Furthermore, we provide a new database for a future comparison of new algorithms.

In the next subsection we evaluate our proposed method in an adapted set of instances from a similar problem and compare our approach with a method described in the literature.

#### 4.6. Results for the Generalized Quadratic Assignment Problem

The VMPlacement is a generalization of GQAP and, our proposed method can be easily adapted for GQAP. Since we evaluate the infeasibilities in the objective function, we only need to change the cost evaluation function. However, for convenience, we adapted the instances of GQAP for the VMPlacement problem, just given a sufficiently large value for the bandwidth capacity between each pair of data centers and required latency between each pair of virtual machines. Finally, we also define an empty set of users.

The main objective of this experiment is to provide a brief comparison with a method described in the literature to show that our approach has a competitive performance. We first perform an experiment with CPLEX to provide a baseline comparison with exact methods. After, we compare our approaches to a GRASP with path-relinking proposed by [Mateus et al. \(2010\)](#).

We use a set of instances proposed by [Cordeau et al. \(2006\)](#). These instances have 20 to 50 facilities (virtual machines) and 6 to 20 locations (data centers). Instances are listed in Table 8 and are labelled respectively by the number of facilities (virtual machines), locations (data centers), and a parameter that controls the tightness of the capacity constraints. The higher the value of the parameter, the higher is the tightness of the capacity constraints. Since this set comprises the largest instance set available for GQAP, and we observe a good performance of our approach, we evaluate our strategies only in this set of instances. These experiments are carried out on a computer with an Intel(R) Core(TM) i7 CPU 930 2.80 GHz with 12GB of main memory.

[Cordeau et al. \(2006\)](#) evaluate the CPLEX (version 8.1) exploring the emphasizes parameter on the branch-and-bound tree exploration for a time limit of two hours. CPLEX was able to prove optimality only for the instance 30-08-55. [Pessoa et al. \(2010\)](#) proposed exact algorithms that combine branch-and-bound with a new Lagrangean decomposition and the Reformulation-Linearization Technique and prove optimality for another 13 instances of the previously described instance set.

In the first experiment we evaluate the performance of CPLEX with LMVMP and LMVMPII described in Section 2. We run CPLEX for each instance with a time limit of one day (86,400 seconds). We also set the tree memory parameter (`TreLim`) to 5000 to stop the execution when this memory limit is exceeded. The remaining parameters are kept to their default values.

Table 8 shows the results for each instance listed in the first column. The second column shows the best-known solutions for each instance (optimal values are given in boldface). The next two sets of columns show the results for each linear mathematical model. Column Nodes shows the number of nodes solved by CPLEX. Column Integer Sol shows the objective function of best solution found during the execution. Column gap shows the percentage gap between the lower bound and the best integer solution found by CPLEX. Column Time shows the running time in seconds to CPLEX prove optimality or reach the stopping criterion. Instances that CPLEX found the optimal solution or proved the optimality are highlighted in boldface.

Observing the results in [Pessoa et al. \(2010\)](#) and comparing with the results in Table 8, we note that the branch-and-bound approach in [Pessoa et al. \(2010\)](#) tends to be more

Table 8: CPLEX results for GQAP instances.

Instance	BKS	LMVMP				LMVMPII			
		Nodes	Integer Sol	gap (%)	Time (s)	Nodes	Integer Sol	gap (%)	Time (s)
20-15-35	<b>1,471,896</b>	2,990	<b>1,471,896</b>	<b>0.00</b>	3,788.3	30,709,073	<b>1,471,896</b>	7.95	41,981.1
20-15-55	<b>1,723,638</b>	10,839	<b>1,723,638</b>	<b>0.00</b>	9,149.2	30,269,870	<b>1,723,638</b>	6.58	39,265.6
20-15-75	<b>1,953,188</b>	5,725	<b>1,953,188</b>	<b>0.00</b>	4,546.3	1,867,883	<b>1,953,188</b>	<b>0.00</b>	3,979.6
30-06-95	<b>5,160,920</b>	24,889	<b>5,160,920</b>	<b>0.00</b>	9,471.3	16,820,694	<b>5,160,920</b>	<b>0.00</b>	35,086.4
30-07-75	<b>4,383,923</b>	47,884	<b>4,383,923</b>	<b>0.00</b>	22,429.5	26,644,434	<b>4,383,923</b>	4.50	63,524.3
30-08-55	<b>3,501,695</b>	793	<b>3,501,695</b>	<b>0.00</b>	370.1	14,090,023	<b>3,501,695</b>	<b>0.00</b>	36,009.5
30-10-65	<b>3,620,959</b>	41,020	<b>3,620,959</b>	<b>0.00</b>	65,805.4	19,972,814	<b>3,620,959</b>	13.32	58,100.3
30-20-35	<b>3,379,359</b>	1,872	3,605,129	23.96	86,400.0	14,857,869	<b>3,379,359</b>	35.36	50,023.8
30-20-55	<b>3,593,105</b>	1,680	3,865,716	31.23	86,400.0	14,117,019	<b>3,593,105</b>	30.27	54,253.0
30-20-75	<b>4,050,938</b>	1,946	4,245,753	26.43	86,400.0	13,553,603	<b>4,050,938</b>	21.13	59,363.0
30-20-95	<b>5,710,645</b>	3,575	5,840,934	12.96	86,400.0	6,869,516	<b>5,710,645</b>	2.71	86,400.2
35-15-35	<b>4,456,670</b>	3,290	<b>4,456,670</b>	13.23	86,400.0	15,797,188	4,457,348	34.76	40,190.5
35-15-55	<b>4,639,128</b>	3,078	4,723,959	19.34	86,400.0	13,903,550	<b>4,639,128</b>	27.52	40,295.3
35-15-75	6,301,723	1,139	6,395,402	27.83	86,400.0	13,348,410	6,327,723	28.02	59,428.9
35-15-95	<b>6,670,264</b>	1,546	7,370,866	32.09	86,400.0	9,664,857	6,689,421	20.38	86,400.0
40-07-75	7,405,793	34,769	<b>7,405,793</b>	3.60	86,400.0	20,466,362	<b>7,405,793</b>	12.94	35,228.8
40-09-95	7,667,719	8,471	7,941,330	18.26	86,400.0	16,718,721	7,694,904	18.61	38,337.7
40-10-65	7,265,559	6,790	7,305,381	10.80	86,400.0	16,475,064	<b>7,265,559</b>	22.79	49,526.1
50-10-65	10,513,029	5,467	<b>10,513,029</b>	4.15	86,400.0	14,617,831	<b>10,513,029</b>	19.39	52,238.3
50-10-75	11,217,503	2,209	11,415,840	19.89	86,400.0	14,018,304	11,251,072	24.42	54,332.5
50-10-95	12,845,598	1,960	13,242,115	18.57	86,400.0	13,205,512	<b>12,845,598</b>	19.25	57,175.5

efficient than CPLEX. In general, the running time tends to be smaller than the required by CPLEX. However, both methods are not strictly comparable since the experiments use different stop criterion, and are reported over different computers. In the end, the main observation from these results is that even with a high increase in the computational power and the development of many algorithms and improvements over each CPLEX version, this set of instances is still hard to be solved by exact methods.

In a second experiment, we compare our strategies similarly as described for GRASP-PR proposed by [Mateus et al. \(2010\)](#), the state-of-the-art algorithm for GQAP. For each instance we performed 200 independent runs. Each run stopped when a solution value as good as the best-known solution was found (column BKS in Table 8).

Two main differences are observed between our approaches and GRASP-PR proposed by [Mateus et al. \(2010\)](#). First, we performed a wide exploration with the local search strategy. [Mateus et al. \(2010\)](#) uses an approximate local search using shift and swap moves. In our approach, we explore all neighbors using three neighborhood structures. Second, we use a penalization strategy to deal with infeasible solutions. This allows passing by through infeasible solutions to reach new feasible solutions, which can be difficult to be reached only exploring a feasible search space.

Table 9 shows a comparison of GRASP-PR from [Mateus et al. \(2010\)](#) and our best two approaches described in the previous subsection. Columns Min, Avg, Max, Sd give the minimum, average, and maximum times, as well as the standard deviation of these runs to find a solution with values equal to BKS. Finally, column 0.95 shows the time in which 95% of the runs find the BKS.

The results described in [Mateus et al. \(2010\)](#) are reported over a computer Dell PE1950 with dual quad core 2.66 GHz Intel Xeon processors. Unfortunately, there is no specific information about the processor model, and then we suppose the model Intel Xeon X5355 2.66 GHz, based on server model, the number of cores and clock frequency. We run our algorithms for this set of experiments on a computer with an Intel(R) Core(TM) i7 CPU 930 2.80 GHz. The machine that we use is approximately 1.2 times faster than the one used in [Mateus et al. \(2010\)](#) (based on Single Thread Rating data from [www.cpubenchmark.net](http://www.cpubenchmark.net)). The execution times for the two implementations are not strictly comparable since the languages and compilers used are different. However, at a higher level, this comparison provides an idea of the behavior of the algorithms.

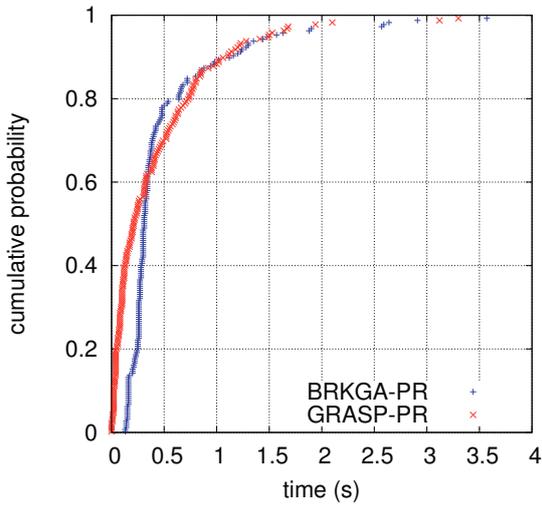
Table 9: Comparison algorithms for GQAP.

Instance	Mateus et al. (2010)					BRKGA-PR					GRASP-PR				
	Min	Avg	Max	Sd	0.95	Min	Avg	Max	Sd	0.95	Min	Avg	Max	Sd	0.95
20-15-35	0.16	7.05	38.87	6.47	21.04	0.06	0.13	0.28	0.04	0.17	0.01	0.08	0.57	0.09	0.22
20-15-55	0.24	2.87	14.42	2.18	7.69	0.06	0.12	0.17	0.03	0.16	0.01	0.13	0.75	0.14	0.36
20-15-75	0.26	2.01	12.82	1.72	5.25	0.08	0.39	3.44	0.50	1.32	0.01	0.41	3.30	0.51	1.41
30-06-95	0.55	2.59	23.81	2.22	6.44	0.16	0.41	2.23	0.25	0.88	0.01	0.48	1.78	0.36	1.16
30-07-75	0.50	7.80	38.47	5.47	18.18	0.15	0.66	5.22	0.63	1.86	0.01	0.48	4.10	0.49	1.28
30-08-55	0.18	1.61	4.89	0.95	3.60	0.13	0.27	0.40	0.06	0.35	0.01	0.03	0.24	0.03	0.10
30-10-65	2.75	121.94	1,032.80	146.06	514.82	0.18	1.58	12.27	1.83	5.35	0.01	1.59	8.83	1.53	4.41
30-20-35	1.08	79.03	4,441.40	312.62	166.21	0.24	0.73	1.73	0.29	1.27	0.01	1.81	8.60	1.60	5.14
30-20-55	1.28	25.16	150.11	21.19	66.82	0.25	1.35	9.84	1.37	4.05	0.01	1.73	7.43	1.37	4.58
30-20-75	2.11	41.43	759.81	68.39	148.43	0.29	0.71	1.47	0.23	1.15	0.01	1.01	4.84	0.90	2.67
30-20-95	833.99	543,019.01	2,533,608.00	747,962.39	2,186,440.80	8.82	514.61	2,780.71	491.24	1,556.91	1.40	1,880.75	7,363.47	1,643.47	5,563.42
35-15-35	8.41	306.11	1,717.94	242.49	775.25	0.24	5.62	27.82	5.64	17.71	0.09	12.26	53.11	11.26	35.16
35-15-55	4.33	21.13	75.69	11.95	42.47	0.32	0.97	4.10	0.54	1.71	0.01	0.38	1.56	0.31	0.97
35-15-75	5.18	68.23	621.83	74.17	183.19	0.35	1.09	3.30	0.48	1.99	0.07	1.66	6.85	1.39	4.63
35-15-95	6.61	1,454.00	19,171.48	3,057.43	6,949.08	3.75	152.66	693.83	130.20	406.25	0.44	126.32	822.98	132.03	386.36
40-07-75	4.53	59.37	377.06	51.21	159.00	0.26	0.64	2.46	0.27	1.02	0.01	0.57	3.29	0.58	1.64
40-09-95	6.18	417.00	5,017.56	610.28	1,490.31	1.65	58.30	398.04	58.47	148.19	0.27	16.87	99.82	16.26	48.26
40-10-65	0.84	17.87	115.06	15.88	52.73	0.37	0.79	1.12	0.21	1.03	0.01	0.14	0.71	0.13	0.38
50-10-65	2.52	24.56	84.64	16.34	64.04	0.67	1.42	2.06	0.34	1.88	0.01	0.12	0.58	0.09	0.28
50-10-75	22.79	1,352.41	24,507.34	3,085.42	4,404.50	1.70	299.78	1,595.95	311.53	965.27	0.19	38.58	195.59	40.51	120.84
50-10-95	9.97	89.36	1,059.59	91.95	200.20	1.98	22.36	145.65	25.65	73.16	0.14	9.38	39.38	7.72	25.55
Average	43.55	26,053.36	123,470.17	35,989.85	104,843.81	1.03	50.69	271.05	49.04	151.98	0.13	99.75	410.85	88.61	295.66
Median	2.52	41.43	377.06	51.21	148.43	0.26	0.97	3.44	0.50	1.86	0.01	1.01	4.84	0.90	2.67

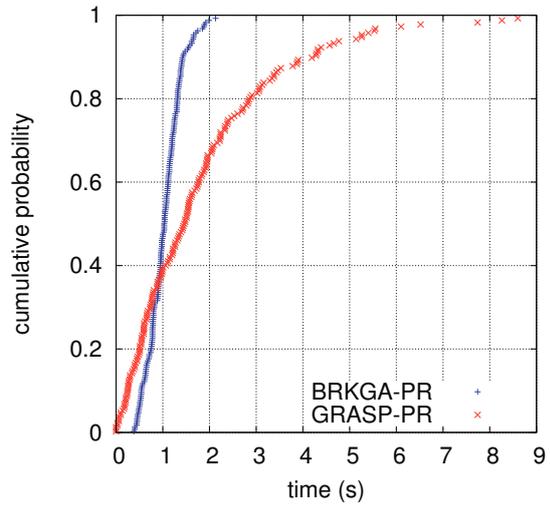
Even considering that the running time is over different hardware, the data from Table 9 shows that our version of GRASP-PR, as well as the BRKGA-PR approach have a significant lower running time in comparison with the results reported in [Mateus et al. \(2010\)](#). The average time reduction for each instance for the minimum time to solve is near to two orders of magnitude. This evidence that our proposed strategies has competitive performance since they are able to find the BKS in all runs in running times significantly shorter than the reported in [Mateus et al. \(2010\)](#). We also observe that our strategies are robust in the sense the algorithms were developed for a specific problem and they also perform well in a general application without any modification in the algorithm. Finally, statistical tests similar to the one described in previous subsections comparing our two approaches shows that GRASP-PR is slightly faster than BRKGA-PR, especially for the instances with 35 or more facilities. For instances with up to 30 facilities, the behaviour is the opposite.

Figure 4.6 shows the Time-To-Target plots that shows the running time distribution for BRKGA-PR and GRASP-PR to find a target solution defined as BKS. We choose the

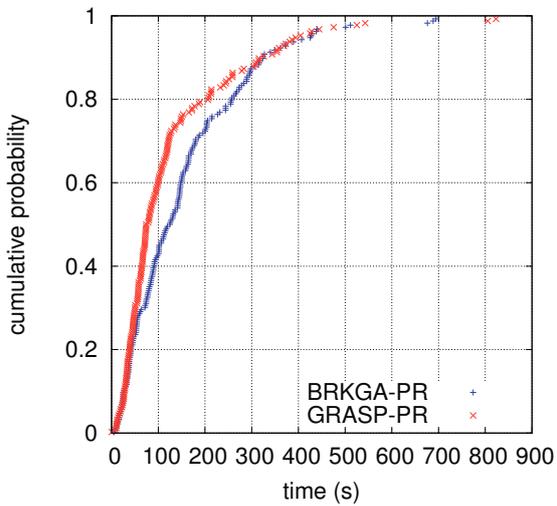
same instances used to show the time-to-target for GRASP-PR in [Mateus et al. \(2010\)](#). We observe that the behaviour, in general, is similar for BRKGA-PR and GRASP-PR, but the performance of one algorithm can be better than the other depending on the instance.



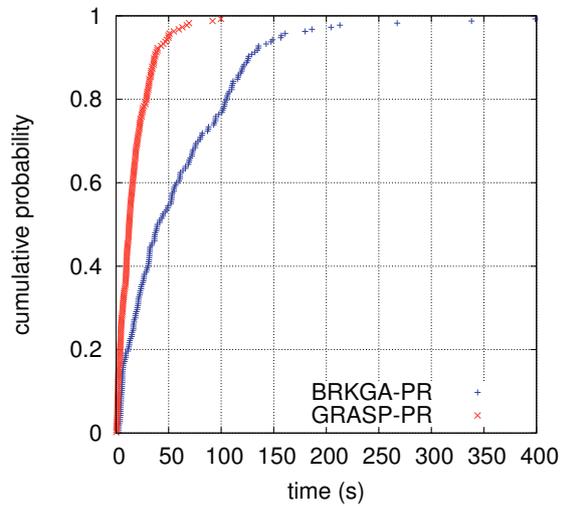
(a) TTT plot for 20-15-75.



(b) TTT plot for 30-20-35.



(c) TTT plot for 35-15-95.



(d) TTT plot for 40-09-95.

Figure 7: Cumulative probability distribution for BRKGA-PR and GRASP-PR running times for instances 20-15-75, 30-20-35, 35-15-95, and 40-09-95.

## 5. Concluding remarks

In this paper, we presented the problem of minimizing the cost of placement virtual machines across geo-separated data centers. A quadratic and two linear mathematical formulations were presented. Moreover, we presented several heuristics for solving this problem. In the experiments, we evaluate the performance of CPLEX using the proposed mathematical formulations, and the proposed heuristic methods.

The results of CPLEX show that by adding the set of cuts the solver improves significantly the quality of the lower bounds for the LMVMPII model. We also observed that for this model, CPLEX can handle larger instances than considering LMVMP model, obtaining better lower bounds and feasible solutions in less computational time. However, as an exact method, CPLEX is limited to solve small instances, showing that heuristic methods are required on larger size instances.

Two metaheuristics approaches are used to solve the problem, the GRASP, and BRKGA. In both methods, we use a path-relinking procedure as an intensification approach. Also, we use a local search method that includes many neighborhood strategies. Both strategies had similar performance, with the best performance achieved using local search 3V and path-relinking PRM. A slight advantage for BRKGA was observed in small instances while GRASP has a slight advantage in larger instances. The same performance was observed when our strategies were applied to GQAP instances. In this case, our methods outperforms the previous state-of-the-art results.

Finally, considering the high cost involved in this kind of problem, and the difficulty to obtain feasible solutions when considering several constraints and limited resources, the proposed algorithms are good alternative to reduce costs, while maintaining the reliability and the demand requirements of data centers.

## 6. Acknowledgments

This work has been partially supported by CAPES, CNPq project 462425/2014-2, PRH PB-217 Petrobras S.A. from Brazil. Also thanks to K. Joshi at AT&T Labs-Research for many helpful discussions.

## References

- Renata M. Aiex, Mauricio G. C. Resende, and Celso Carneiro Ribeiro. TTT plots: A perl program to create time-to-target plots. *Optimization Letters*, 1(4):355–366, 2007. ISSN 18624472. doi: 10.1007/s11590-006-0031-4. URL <http://dx.doi.org/10.1007/s11590-006-0031-4>.
- Vinícius Amaral Armentano, A.L. Shiguemoto, and A. Løkketangen. Tabu search with path relinking for an integrated production–distribution problem. *Computers & Operations Research*, 38(8):1199–1209, aug 2011. ISSN 03050548. doi: 10.1016/j.cor.2010.10.026. URL <http://www.sciencedirect.com/science/article/pii/S0305054810002558>.
- Hitesh Ballani, Paolo Costa, Thomas Karagiannis, and Ant Rowstron. Towards predictable datacenter networks. *ACM SIGCOMM Computer Communication Review*, 41(4):242, oct 2011. ISSN 01464833. doi: 10.1145/2043164.2018465. URL <http://doi.acm.org/10.1145/2018436.2018465>.

- Matthieu Basseur, Franck Seynhaeve, and El-Ghazali Talbi. Path Relinking in Pareto Multi-objective Genetic Algorithms. In Carlos A. Coello Coello, Arturo Hernández Aguirre, and Eckart Zitzler, editors, *Evolutionary Multi-Criterion Optimization SE - 9*, volume 3410, pages 120–134. Springer Berlin Heidelberg, 2005. ISBN 978-3-540-24983-2. doi: 10.1007/978-3-540-31880-4\_9. URL [http://dx.doi.org/10.1007/978-3-540-31880-4\\_9](http://dx.doi.org/10.1007/978-3-540-31880-4_9).
- J. C. Bean. Genetic Algorithms and Random Keys for Sequencing and Optimization. *INFORMS Journal on Computing*, 6(2):154–160, 1994. ISSN 1091-9856. doi: 10.1287/ijoc.6.2.154. URL <http://dx.doi.org/10.1287/ijoc.6.2.154>.
- Jean François Cordeau, Manlio Gaudioso, Gilbert Laporte, and Luigi Moccia. A memetic heuristic for the generalized quadratic assignment problem. *INFORMS Journal on Computing*, 18(4):433–443, 2006. ISSN 10919856. doi: 10.1287/ijoc.1040.0128. URL <http://dx.doi.org/10.1287/ijoc.1040.0128>.
- Thomas A. Feo and Mauricio G. C. Resende. A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters*, 8(2):67–71, 1989. ISSN 0167-6377. doi: 10.1016/0167-6377(89)90002-3. URL <http://www.sciencedirect.com/science/article/pii/0167637789900023>.
- Thomas A. Feo and Mauricio G. C. Resende. Greedy Randomized Adaptive Search Procedures. *Journal of Global Optimization*, 6(2):109–133, 1995. ISSN 0925-5001. doi: 10.1007/BF01096763. URL <http://dx.doi.org/10.1007/BF01096763>.
- Paola Festa and Mauricio G. C. Resende. An annotated bibliography of GRASP - Part I: Algorithms. *International Transactions in Operational Research*, 16(1):1–24, jan 2009a. ISSN 09696016. doi: 10.1111/j.1475-3995.2009.00663.x. URL <http://doi.wiley.com/10.1111/j.1475-3995.2009.00663.x>.
- Paola Festa and Mauricio G. C. Resende. An annotated bibliography of GRASP-Part II: Applications. *International Transactions in Operational Research*, 16(2):131–172, mar 2009b. ISSN 09696016. doi: 10.1111/j.1475-3995.2009.00664.x. URL <http://doi.wiley.com/10.1111/j.1475-3995.2009.00664.x>.
- Paola Festa and Mauricio G. C. Resende. Hybridizations of GRASP with Path-Relinking. In El-Ghazali Talbi, editor, *Hybrid Metaheuristics*, volume 434 of *Studies in Computational Intelligence*, pages 135–155. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-30670-9. doi: 10.1007/978-3-642-30671-6\_5. URL [http://dx.doi.org/10.1007/978-3-642-30671-6\\_5](http://dx.doi.org/10.1007/978-3-642-30671-6_5).
- Andreas Fischer, Juan Felipe Botero, Michael Till Beck, Hermann de Meer, and Xavier Hesselbach. Virtual Network Embedding: A Survey. *IEEE Communications Surveys & Tutorials*, 15(4):1888–1906, jan 2013. ISSN 1553-877X. doi: 10.1109/SURV.2013.013013.00155. URL <http://dx.doi.org/10.1109/SURV.2013.013013.00155>.
- Matteo Fischetti and Michele Monaci. Exploiting Erraticism in Search. *Operations Research*, 62(1):114–122, 2014. ISSN 0030-364X. doi: 10.1287/opre.2013.1231. URL <http://dx.doi.org/10.1287/opre.2013.1231>.
- A.M. Frieze and J. Yadegar. On the quadratic assignment problem. *Discrete Applied Mathematics*, 5(1): 89–98, jan 1983. ISSN 0166218X. doi: 10.1016/0166-218X(83)90018-5. URL [http://dx.doi.org/10.1016/0166-218X\(83\)90018-5](http://dx.doi.org/10.1016/0166-218X(83)90018-5).
- Fred Glover. Tabu Search—Part I. *ORSA Journal on Computing*, 1(3):190–206, 1989. doi: 10.1287/ijoc.1.3.190. URL <http://dx.doi.org/10.1287/ijoc.1.3.190>.
- Fred Glover. Tabu Search and Adaptive Memory Programming - Advances, Applications and Challenges. In Richard S. Barr, Richard V. Helgason, and Jeffery L. Kennington, editors, *Interfaces in Computer Science and Operations Research*, volume 7 of *Operations Research/Computer Science Interfaces Series*, pages 1–75. Springer US, 1997. ISBN 978-1-4613-6837-3. doi: 10.1007/978-1-4615-4102-8\_1. URL [http://dx.doi.org/10.1007/978-1-4615-4102-8\\_1](http://dx.doi.org/10.1007/978-1-4615-4102-8_1).
- Fred Glover. Exterior Path Relinking for Zero-One Optimization. *International Journal of Applied Metaheuristic Computing*, 5(3):1–8, jan 2014. ISSN 1947-8283. doi: 10.4018/ijamc.2014070101. URL <http://dx.doi.org/10.4018/ijamc.2014070101>.
- Fred Glover and Manuel Laguna. Tabu Search. In Colin R Reeves, editor, *Modern Heuristic Techniques for Combinatorial Problems*, pages 70–150. Blackwell Scientific Publishing, Oxford, England, 1993. ISBN 0-470-22079-1.

- Fred Glover, Manuel Laguna, and Rafael Martí. Fundamentals of scatter search and path relinking. *Control and Cybernetics*, Vol. 29, n(3):653–684, 2000.
- Fred Glover, Manuel Laguna, and Rafael Martí. Scatter Search and Path Relinking: Advances and Applications. In Fred Glover and Gary A. Kochenberger, editors, *Handbook of Metaheuristics*, volume 57 of *International Series in Operations Research & Management Science*, pages 1–35. Springer US, 2003. ISBN 978-1-4020-7263-5. doi: 10.1007/0-306-48056-5\_1. URL [http://dx.doi.org/10.1007/0-306-48056-5\\_1](http://dx.doi.org/10.1007/0-306-48056-5_1).
- José Fernando Gonçalves and Mauricio G. C. Resende. Biased random-key genetic algorithms for combinatorial optimization. *Journal of Heuristics*, 17(5):487–525, aug 2011. ISSN 13811231. doi: 10.1007/s10732-010-9143-1. URL <http://link.springer.com/10.1007/s10732-010-9143-1>.
- Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David a. Maltz, Parveen Patel, and Sudipta Sengupta. VL2: A Scalable and Flexible Data Center Network. *ACM SIGCOMM Computer Communication Review*, 39(4):51, 2009. ISSN 01464833. doi: 10.1145/1594977.1592576. URL <http://doi.acm.org/10.1145/1594977.1592576>.
- Chuanxiong Guo, Guohan Lu, Helen J. Wang, Shuang Yang, Chao Kong, Peng Sun, Wenfei Wu, and Yongguang Zhang. SecondNet: A Data Center Network Virtualization Architecture with Bandwidth Guarantees. In *International Conference on - Co-NEXT, 10*, page 12, New York, 2010. ACM. ISBN 9781450304481. doi: 10.1145/1921168.1921188. URL <http://dx.doi.org/10.1145/1921168.1921188>.
- Pierre Hansen, Nenad Mladenović, and José A. Moreno Pérez. Variable neighbourhood search: Methods and applications. *Annals of Operations Research*, 175(1):367–407, 2010. ISSN 02545330. doi: 10.1007/s10479-009-0657-6. URL <http://dx.doi.org/10.1007/s10479-009-0657-6>.
- Holger H. Hoos and Thomas Stützle. Evaluating Las Vegas Algorithms: Pitfalls and Remedies. In *Conference on Uncertainty in Artificial Intelligence, 14*, pages 238–245, San Francisco, 1998a. Morgan Kaufmann Pub. ISBN 1-55860-555-X. URL <http://dl.acm.org/citation.cfm?id=2074094.2074122>.
- Holger H. Hoos and Thomas Stützle. On the empirical evaluation of Las Vegas algorithms. Technical report, Computer Science Department, University of British Columbia, 1998b. URL <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.33.132>.
- L. Kaufman and F. Broeckx. An algorithm for the quadratic assignment problem using Bender’s decomposition. *European Journal of Operational Research*, 2(3):207–211, may 1978. ISSN 03772217. doi: 10.1016/0377-2217(78)90095-4. URL [http://dx.doi.org/10.1016/0377-2217\(78\)90095-4](http://dx.doi.org/10.1016/0377-2217(78)90095-4).
- Tjalling C. Koopmans and Martin J. Beckmann. Assignment problems and the location of economic activities. *Econometrica*, 25(1):53–76, 1957.
- Manuel Laguna and Rafael Martí. GRASP and Path Relinking for 2-Layer Straight Line Crossing Minimization. *INFORMS Journal on Computing*, 11(1):44–52, 1999. doi: 10.1287/ijoc.11.1.44. URL <http://dx.doi.org/10.1287/ijoc.11.1.44>.
- C G Lee and Z Ma. The generalized quadratic assignment problem. Technical report, Department of Mechanical and Industrial Engineering at the University of Toronto, Toronto, Ontario, M5S 3G8, Canada, 2004. URL <http://www.mie.utoronto.ca/labs/ilr/papers/GQAP.pdf>.
- Rafael Martí, Manuel Laguna, and Fred Glover. Principles of scatter search. *European Journal of Operational Research*, 169(2):359–372, mar 2006. ISSN 03772217. doi: 10.1016/j.ejor.2004.08.004. URL <http://dx.doi.org/10.1016/j.ejor.2004.08.004>.
- Rafael Martí, Mauricio G. C. Resende, and Celso Carneiro Ribeiro. Multi-start methods for combinatorial optimization. *European Journal of Operational Research*, 226(1):1–8, 2013. ISSN 03772217. doi: 10.1016/j.ejor.2012.10.012. URL <http://dx.doi.org/10.1016/j.ejor.2012.10.012>.
- Geraldo R. Mateus, Mauricio G. C. Resende, and Ricardo M. A. Silva. GRASP with path-relinking for the generalized quadratic assignment problem. *Journal of Heuristics*, 17(5):527–565, sep 2010. ISSN 1381-1231. doi: 10.1007/s10732-010-9144-0. URL <http://dx.doi.org/10.1007/s10732-010-9144-0>.
- Hans D. Mittelmann and Domenico Salvagnin. On solving a hard quadratic 3-dimensional assignment problem. *Mathematical Programming Computation*, pages 1–16, 2015. ISSN 1867-2949. doi: 10.1007/s12532-015-0077-3. URL <http://dx.doi.org/10.1007/s12532-015-0077-3>.
- Carlos A. S. Oliveira, Panos M. Pardalos, and Mauricio G. C. Resende. GRASP with Path-Relinking for the Quadratic Assignment Problem. In Celso C. Ribeiro and Simone L. Martins, editors, *Experimental*

- and *Efficient Algorithms*, volume 3059 of *Lecture Notes in Computer Science*, pages 356–368. Springer Berlin Heidelberg, 2004. ISBN 978-3-540-22067-1. doi: 10.1007/978-3-540-24838-5\_27. URL [http://dx.doi.org/10.1007/978-3-540-24838-5\\_27](http://dx.doi.org/10.1007/978-3-540-24838-5_27).
- Melquíades Pérez Pérez, Francisco Almeida Rodríguez, and J Marcos Moreno-Vega. A hybrid VNS–path relinking for the p-hub median problem. *IMA Journal of Management Mathematics*, 18(2):157–171, apr 2007. doi: 10.1093/imaman/dpm013. URL <http://dx.doi.org/10.1093/imaman/dpm013>.
- Artur Alves Pessoa, Peter M. Hahn, Monique Guignard, and Yi-Rong Zhu. Algorithms for the generalized quadratic assignment problem combining Lagrangean decomposition and the Reformulation-Linearization Technique. *European Journal of Operational Research*, 206(1):54–63, oct 2010. ISSN 03772217. doi: 10.1016/j.ejor.2010.02.006. URL <http://dx.doi.org/10.1016/j.ejor.2010.02.006>.
- Mauricio G. C. Resende and Celso Carneiro Ribeiro. GRASP with Path-Relinking: Recent Advances and Applications. In *Metaheuristics: Progress as Real Problem Solvers*, volume 1, pages 29–63. Springer-Verlag, New York, 2005. doi: 10.1007/0-387-25383-1\_2. URL [http://dx.doi.org/10.1007/0-387-25383-1\\_2](http://dx.doi.org/10.1007/0-387-25383-1_2).
- Mauricio G. C. Resende and Celso Carneiro Ribeiro. Greedy Randomized Adaptive Search Procedures: Advances, Hybridizations, and Applications. In Michel Gendreau and Jean-Yves Potvin, editors, *Handbook of Metaheuristics*, volume 146 of *International Series in Operations Research & Management Science*, pages 283–319. Springer US, 2010. ISBN 978-1-4419-1663-1. doi: 10.1007/978-1-4419-1665-5\_10. URL [http://dx.doi.org/10.1007/978-1-4419-1665-5\\_10](http://dx.doi.org/10.1007/978-1-4419-1665-5_10).
- Mauricio G. C. Resende and Celso Carneiro Ribeiro. GRASP: Greedy Randomized Adaptive Search Procedures. In Edmund K Burke and Graham Kendall, editors, *Search Methodologies*, pages 287–312. Springer US, 2014. ISBN 978-1-4614-6939-1. doi: 10.1007/978-1-4614-6940-7\_11. URL [http://dx.doi.org/10.1007/978-1-4614-6940-7\\_11](http://dx.doi.org/10.1007/978-1-4614-6940-7_11).
- Mauricio G. C. Resende, Celso Carneiro Ribeiro, Fred Glover, and Rafael Martí. Scatter search and path-relinking: Fundamentals, advances, and applications. In *Handbook of Metaheuristics*, pages 87–107. 2010. ISBN 978-1-4419-1665-5. doi: 10.1007/978-1-4419-1665-5\_4. URL [http://dx.doi.org/10.1007/978-1-4419-1665-5\\_4](http://dx.doi.org/10.1007/978-1-4419-1665-5_4).
- Fernando Stefanello, Vaneet Aggarwal, Luciana S. Buriol, José Fernando Gonçalves, and Mauricio G. C. Resende. A Biased Random-key Genetic Algorithm for Placement of Virtual Machines across Geo-Separated Data Centers. In *Conference on Genetic and evolutionary computation, 15*, pages 1–8, Madrid, 2015a. ACM. ISBN 9781450326629. doi: dx.doi.org/10.1145/2739480.2754768. URL <http://dx.doi.org/10.1145/2739480.2754768>.
- Fernando Stefanello, Luciana S. Buriol, Vaneet Aggarwal, and Mauricio G. C. Resende. A New Linear Model for Placement of Virtual Machines across Geo-Separated Data Centers. In *Simpósio Brasileiro de Pesquisa Operacional, 47*, pages 1–11, Porto de Galinhas, PE, 2015b. Sociedade Brasileira de Pesquisa Operacional. doi: 10.13140/RG.2.1.2236.8088. URL <http://dx.doi.org/10.13140/RG.2.1.2236.8088>.
- Rodrigo Franco Toso and Mauricio G. C. Resende. A C++ application programming interface for biased random-key genetic algorithms. *Optimization Methods and Software*, 30(1):1–15, 2014. ISSN 10556788. doi: 10.1080/10556788.2014.890197. URL <http://dx.doi.org/10.1080/10556788.2014.890197>.
- Eva Vallada and Rubén Ruiz. Genetic algorithms with path relinking for the minimum tardiness permutation flowshop problem. *Omega*, 38(1-2):57–67, feb 2010. ISSN 03050483. doi: 10.1016/j.omega.2009.04.002. URL <http://dx.doi.org/10.1016/j.omega.2009.04.002>.
- Di Xie and Y Charlie Hu. The Only Constant is Change: Incorporating Time-Varying Network Reservations in Data Centers. In *Sigcomm, 12*, SIGCOMM ’12, pages 199–210, New York, 2012. ACM. ISBN 9781450314190. doi: 10.1145/2342356.2342397. URL <http://doi.acm.org/10.1145/2342356.2342397>.
- G.Q. Zhang and K.K. Lai. Combining path relinking and genetic algorithms for the multiple-level warehouse layout problem. *European Journal of Operational Research*, 169(2):413–425, mar 2006. ISSN 03772217. doi: 10.1016/j.ejor.2004.08.007. URL <http://dx.doi.org/10.1016/j.ejor.2004.08.007>.