

---

# Tabu Search and GRASP for the Capacitated Clustering Problem

ANNA MARTÍNEZ-GAVARA

Departamento de Estadística e Investigación Operativa, Universidad de Valencia, Spain  
Gavara@uv.es

VICENTE CAMPOS

Departamento de Estadística e Investigación Operativa, Universidad de Valencia, Spain  
Vicente.Campos@uv.es

MICAEL GALLEGO

Departamento de Ciencias de la Computación, Universidad Rey Juan Carlos, Spain.  
Micael.Gallego@urjc.es

MANUEL LAGUNA

Leeds School of Business, University of Colorado at Boulder, USA  
laguna@colorado.edu

RAFAEL MARTÍ

Departamento de Estadística e Investigación Operativa, Universidad de Valencia, Spain  
Rafael.Marti@uv.es

---

## ABSTRACT

The Capacitated Clustering Problem (CCP) consists of forming a specified number of clusters or groups from a set of elements in such a way that the sum of the weights of the elements in each cluster is within some capacity limits, and the sum of the benefits between the pairs of elements in the same cluster is maximized. This problem— which has been recently tackled with a GRASP/VNS approach— arises in the context of facility planners at mail processing and distribution. We propose a tabu search and several GRASP variants to find high quality solutions to this NP-hard problem. These variants are based on several neighborhoods, including a new one, in which we implement a one-for-two swapping strategy. We also hybridize both methodologies to achieve improved outcomes.

The Maximally Diverse Grouping Problem (MDGP) is a special case of the CCP in which all the elements have a weight of one unit. This problem has been recently studied in the academic context when forming student groups, and we adapt the best method reported in the literature, a tabu search with strategic oscillation (TS\_SO), to the CCP. Our empirical study with 50 large instances shows the superiority of the new GRASP with tabu search for the CCP with respect to both the previous GRASP/VNS approach and the adapted TS\_SO.

---

**Keywords:** Capacitated Clustering, Diversity problems, tabu search, metaheuristics.

**Draft version:** July 29, 2013.

## 1. Introduction

Given a graph  $G = (V, E)$  where  $V$  is a set of  $n$  nodes and  $E$  is a set of edges, let  $w_i \geq 0$  be the weight of node  $i \in V$  and let  $c_{ij}$  be the benefit of edge  $(i, j) \in E$ . The Capacitated Clustering Problem (CCP) consists of partition  $V$  into  $p$  clusters in such a way that the sum of the weights of the elements in each cluster is within some integer capacity limits,  $L$  and  $U$ , and the sum of the benefits between the pairs of elements in the same cluster is maximized.

The CCP can be formulated as a quadratic integer program with binary variables  $x_{ik}$  that take the value of 1 if element  $i$  is in cluster  $k$  and 0 otherwise.

$$\begin{aligned}
 \text{(CCP) Maximize} \quad & \sum_{k=1}^p \sum_{i=1}^{n-1} \sum_{j>i}^n c_{ij} x_{ik} x_{jk} \\
 \text{subject to} \quad & \sum_{k=1}^p x_{ik} = 1 \quad i = 1, 2, \dots, n \\
 & L \leq \sum_{i=1}^n w_i x_{ik} \leq U \quad k = 1, 2, \dots, p \\
 & x_{ik} \in \{0, 1\} \quad i = 1, \dots, n \quad k = 1, \dots, p
 \end{aligned}$$

The objective function adds the total benefit of all pairs of elements that belong to the same cluster. The first set of constraints forces the assignment of each element to a cluster. The second set of constraints forces the sum of the weights of the pairs of elements in the same cluster to be between  $L$  and  $U$ .

The literature on the CCP and related problems is vast and a recent paper (Deng and Bard, 2011) summarizes previous heuristics and formulations for this problem. We therefore do not duplicate the work here and refer the reader to Deng and Bard's excellent survey. These authors proposed a GRASP (Greedy Randomized Adaptive Search Procedure) with VNS (Variable Neighborhood Search) that, according to their computational study, outperforms previous approaches. An interesting application of the problem, as mentioned by Deng and Bard, arises in the context of facility planners at mail processing and distribution centers within the US Postal Service. In particular, the design of the zones to help rationalize the bulk movement of mail can be framed as a CCP.

An interesting characteristic of the GRASP in Deng and Bard (2011) is that it constructs the customary *restricted candidate list* (RCL) using both nodes and edges, which to the best of our knowledge is an unconventional strategy (Festa and Resende, 2009). This design triggered our interest on this problem, to compare this unconventional GRASP implementation with a standard one in which the RCL is formed with either nodes or edges but not both simultaneously. Deng and Bard (2011) computational study does not consider general CCP instances and limits itself to the case in which the node weights,  $w_i$ , are equal to 1. This means that although they proposed a procedure for the CPP, their tests were performed on the special case of the MDGP. We extend their testing to include general CCP instances to compare several GRASP designs.

When reviewing the recent clustering literature, we found an interesting connection not previously reported. In particular, the maximally diverse grouping problem (MDGP) consists of grouping a set of elements into  $p$  mutually disjoint groups in such a way that the diversity among the elements in each group is maximized. The diversity among the elements in a group is calculated as the sum of the individual distance between each pair of elements. The objective of the problem is to maximize the overall diversity, i.e., the sum of the diversity of all groups, when the size of each group is within a specified range. Clearly, the MDGP is a special case of the CCP for which  $w_i = 1$  for all node  $i$ , and the distance between each pair of nodes  $(i, j)$  is the benefit  $c_{ij}$ . Therefore, from the CCP formulation above, the MDGP can be formulated as:

$$\begin{aligned}
 \text{(MDGP) Maximize} \quad & \sum_{k=1}^p \sum_{i=1}^{n-1} \sum_{j>i}^n c_{ij} x_{ik} x_{jk} \\
 \text{subject to} \quad & \sum_{k=1}^p x_{ik} = 1 \quad i = 1, 2, \dots, n \\
 & L \leq \sum_{i=1}^n x_{ik} \leq U \quad k = 1, 2, \dots, p \\
 & x_{ik} \in \{0, 1\} \quad i = 1, \dots, n \quad k = 1, \dots, p
 \end{aligned}$$

The MDGP is called the  $k$ -partition problem in Feo et al. (1992) and the equitable partition problem in O'Brien and Mingers (1995). It belongs to the family of diversity problems (Gallego et al. 2009, Martí and Sandoya 2013). One of the most popular MDGP applications appears in the academic context of forming student groups (Weitz and Jelassi, 1992). Gallego et al. (2013) proposed a tabu search with strategic oscillation, TS\_SO, for the MDGP that outperforms the previous approaches for this problem. Recognizing the connection between MDGP and CCP, as part of our work, we also extend the TS\_SO to tackle the general CCP.

The main contributions of our current development are:

- A new (simplified) GRASP for the CCP and a thorough comparison with Deng's and Bard's GRASP
- An adaptation of TS\_SO to handle instances of the CCP
- A new TS for the CCP to be hybridized with GRASP
- Testing of all the methods on a new set of CCP instances

The next section summarizes the Deng and Bard GRASP for the CCP and the Gallego et al. TS\_SO for the MDGP. Section 3 describes the proposed new procedures. This is followed by the description of our extensive experimentation that includes 50 new CCP instances with  $n = 82, 240$ , and 480. Statistical analysis shows the merit of our approach when compared to existing methods (including the adaptation of TS\_SO to the CCP).

## 2. Existing CCP and Adapted MDGP Approaches

Deng and Bard (2011) proposed a GRASP for the CCP, which as customary in GRASP implementations, it alternates between two phases: construction and improvement. In the construction phase, the  $p$  clusters are first seeded with the heaviest weight edges algorithm (HWE), and then completed with a greedy randomized procedure. Specifically, the HWE identifies the  $p$  nodes with the largest weights and assigns them, separately, to the  $p$  clusters. The heaviest edges incident to these nodes are then identified, and their endpoints are assigned to the corresponding clusters. An alternative constructive method, labeled CMC, proposed by Deng and Bard was shown to be inferior to HWE and therefore we do not consider it here.

HWE produces clusters containing two nodes. At this point, a candidate list  $CL$  of elements is built to continue the construction process according to the GRASP methodology. In particular,  $CL$  is formed with the nodes and edges (pairs of nodes) that can be inserted into a solution cluster without exceeding the upper capacity limit  $U$ . For each node  $i \in CL$ , let  $I(i, k)$  be the increase in the objective function if node  $i$  is added to cluster  $k$ . In other words, the sum of the benefits  $c_{ij}$  for all the nodes  $j$  already in cluster  $k$ . Similarly, let  $I(e, k)$  be the increase in the objective function when edge  $e$  is added to cluster  $k$  (i.e., when its two endpoint nodes are in the cluster). The restricted candidate list  $RCL$  is formed with the best elements in  $CL$  according to this evaluation. This set of top candidates is computed as a fraction  $\alpha$  of the number of elements in  $CL$ , where  $\alpha \in ]0,1]$ . Deng and Bard implemented a so-called reactive GRASP in which effective values for this search parameter are self-adjusted as the iterations progress.

An important characteristic of this method is that the  $CL$  consists of both, nodes and edges. That is, the next element to be included in the solution may be either a single node or a pair of nodes (if an edge is selected). This is somewhat unconventional in GRASP, if one considers the implementations in the annotated bibliography by Festa and Resende (2009). We will compare this design with a more conventional design in which the  $CL$  consists of either nodes or edges, but not both.

In the second phase of their GRASP, Deng and Bard proposed three different neighborhoods to improve a constructed solution  $x$ :  $N_1(x)$ ,  $N_2(x)$ , and  $N_3(x)$ . Let  $V_k$  be the set of nodes in cluster  $k$  of this solution, and let  $W_k$  be the sum of the weights of the nodes in  $V_k$  (i.e.,  $W_k = \sum_{i \in V_k} w_i$ ); then,  $W_k$  must be within the capacity limits:  $L \leq W_k \leq U$  for  $k = 1, 2, \dots, p$ .

$N_1(x)$  is the result of extended insertion moves, in which a node  $i$  is moved from a cluster  $k$  to a cluster  $s$ . The node  $i$  is only removed from  $V_k$  if this cluster remains feasible after the move (i.e., if  $L \leq W_k - w_i$ ), but if there is not enough capacity in cluster  $s$  for node  $i$  (i.e., if  $W_s + w_i > U$ ), instead of discarding the move, a node  $j$  in  $V_s$  is moved from cluster  $s$  to a different one, say cluster  $t$ , with enough capacity for node  $j$ . The move is performed if the final solution is feasible, in other words, if  $L \leq W_k - w_i$ ,  $L \leq W_s + w_i - w_j \leq U$ , and  $W_t + w_j \leq U$ .

$N_2(x)$  consists of edge insertions. Given an edge  $(i, j) \in E$ , if both nodes are in the same cluster  $k$  (i.e.,  $i, j \in V_k$ ), an edge insertion considers moving the nodes to another cluster  $s$ , as long as the resulting solution remains feasible. In mathematical terms, we move  $(i, j)$  from cluster  $k$  to cluster  $s$  if

$$L \leq W_k - w_i - w_j \text{ and } W_s + w_i + w_j \leq U$$

This neighborhood also contains moves of edges with endpoints in two different clusters. In particular, given an edge  $(i, j) \in E$  with  $i \in V_k$  and  $j \in V_s$ , it examines moving  $i$  to cluster  $s$ , moving  $j$  to cluster  $k$ , or moving both nodes to another cluster  $t$ . In any case, only capacity-feasible moves are considered.

$N_3(x)$  implements a classical swap move, that is, one in which a node  $i$  is moved from a cluster  $k$  to a cluster  $s$ , and simultaneously a node  $j$  is moved from the cluster  $s$  to the cluster  $k$ . As in the other neighborhoods, the move is performed if the resulting solution is feasible. In mathematical terms:

$$L \leq W_k - w_i + w_j \leq U \text{ and } L \leq W_s + w_i - w_j \leq U.$$

Deng and Bard proposed two improvement methods based on these three neighborhoods. The first one, called CNS, examines the neighborhoods in a sequential/cyclical fashion, terminating when no improvement is possible. The second one implements the so-called variable neighborhood descent, VND, in which the method switches from the current neighborhood  $N_u$  to the next one,  $N_{u+1}$ , when  $N_u$  cannot produce an improvement of the current solution. Moreover, if a better solution is obtained with  $N_u$  with  $u > 1$ , the method switches back to  $N_1$ . The VND terminates when  $N_3$  is applied and no further improvement is possible. Finally, the authors implemented a Randomized VND, called RNVD, in which the neighborhood to be searched in the next iteration is probabilistically selected, where the probability of selection is linked to the merit of each neighborhood as determined by the quality of the solutions found during the search.

In the computational testing, Deng and Bard (2011) compared their designs and concluded that the combination of HWE with RVND resulted in the best overall performance. We use this variant for the purpose of comparison and call it Prev\_GRASP. As a side note, Deng and Bard tested a Path Relinking post-processing strategy and concluded that its inclusion did not result in a significant improvement and therefore we have left it out of our experiments.

Gallego et al. (2013) proposed a tabu search with strategic oscillation, TS\_SO, for the maximally diverse grouping problem (MDGP). TS\_SO is the state-of-the-art solution method for MDGPs. Empirical evidence shows that TS\_SO outperforms the improvement method LCW (Weitz and Lakshminarayanan, 1998) and the hybrid genetic algorithm LSGA (Fan et al., 2011). As described in the previous section, the MDGP is a special case of the CCP. We briefly describe TS\_SO, which is based on three steps: 1) construction of the initial solution, 2) neighborhood search and 3) strategic oscillation, and our adaptation of TS\_SO to the CCP.

The construction step (GC, for greedy construction) for the MDGP starts by randomly selecting  $p$  nodes and assigning each of these elements to a separate group (cluster in our case). Therefore, at the end of the first step, each group has one node assigned to it. Then, the procedure performs  $n - p$  iterations to assign the remaining unassigned elements to groups. In order to generate a feasible solution, the iterations are divided into two phases. In the first phase, the elements are assigned to groups that currently contain fewer elements than the desired minimum number of elements. In the second phase, the remaining elements are assigned to groups with a number of elements that is smaller than the desired maximum number of elements.

In our adaptation of GC to the CCP, the first phase adds a randomly selected node  $i$  to the cluster  $k$  that maximizes  $I(i, k)$  and for which  $W_k < L$ . When all the clusters have a sum of weights of their nodes larger than or equal to  $L$ , the second phase proceeds in a similar way, by adding randomly selected elements to the cluster  $k$  that maximizes  $I(i, k)$ , without exceeding the upper capacity limit  $U$ .

Once a solution  $x$  has been constructed, the neighborhood search (phase 2 of the TS\_SO method) is applied. The search neighborhood in TS\_SO consists of all node insertions and swaps. Note that insertions and swaps are considered simultaneously instead of sequentially, as done by Deng and Bard (2011). Moreover, the TS\_SO implements simple insertions, involving only two groups, instead of the extended version (involving three groups) implemented in  $N_1(x)$ . We denote the simple insertion neighborhood as  $N_0(x)$ . In mathematical terms, at each step the method selects the best feasible move in  $N_0(x) \cup N_3(x)$ . TS\_SO includes a short-term tabu memory (Glover and Laguna, 1997) that enables the method to search beyond the first local optimal point. Specifically, when the local search reaches a point where no improving moves are available, the best non-improving move is selected and executed. At this point, elements that are moved from their current group to another are not allowed to move again for  $tabuTenure = n/10$  iterations. The process terminates when  $maxIter = n/2$  consecutive iterations have been performed without improving the best solution found during the search. TS\_SO can be straightforwardly adapted to the CCP by simply checking the feasibility of moves in terms of the sum of the weights of the nodes in each cluster instead of the cluster cardinality.

So far, we have assumed that feasibility is maintained through the search. However, the third phase of TS\_SO —the *strategic oscillation* (SO) phase— explores solutions for which the group cardinality restrictions may be violated. In particular, the method applies the neighborhood search described above but the number of the elements in a group is allowed to be outside the specified limits by a certain amount  $so$  that ranges between 0 and  $so_{max}$ . This means that when  $so > 0$  the search is allowed to visit infeasible solutions. To create the oscillation pattern, the value of  $so$  is reset to one after every successful application of the improvement method, otherwise  $so$  is increased by one unit until it reaches  $so_{max}$ .

In our adaptation of the SO phase to the CCP, we consider that the sum of the weights of the elements in a cluster has to be within the following limits:

$$L - so \leq W_k \leq U + so$$

Note that this method does not guarantee that the final solution is feasible. When this happens, we apply a repair mechanism that consists of removing elements from clusters  $k$  for which  $W_k > U$  and adding elements to clusters  $s$  for  $W_s < L$ . The elements are selected at random and the process continues until the weight constraints of the clusters are satisfied. We refer to the entire adaptation of TS\_SO to the CPP as CCP\_TS\_SO.

### 3. A Simplified GRASP for the CCP

In our view, the GRASP implementation of Deng and Bard (2011) is overly complicated and a simplified version should be able to obtain solutions of similar quality. The best way to describe their

method is that it is the result of hybridizing Reactive GRASP with VND with three different neighborhoods that are combined in a probabilistic fashion, where the neighborhoods consist of a set of compound moves that may change the node assignment of up to three clusters. In addition, the candidate list in the construction process includes both nodes and edges. In contrast, we propose a simplified GRASP implementation in which only nodes are candidates in the construction process and two simple neighborhoods are combined into a deterministic VND design.

Our GRASP starts by seeding the  $p$  clusters  $V_1, V_2, \dots, V_p$  with  $p$  randomly selected nodes. Then, we explore the clusters in lexicographical order assigning elements until all of them satisfy the lower bound constraint (i.e., until the sum of the weights of the nodes already assigned to the cluster is larger than or equal to  $L$ ). To do so, the candidate set  $CL$  is formed with all the unassigned nodes and the value  $I(i, k)$  is calculated for all pairs  $(i, k)$  of nodes and clusters.  $RCL_k$ —that is, the restricted candidate list of nodes for cluster  $k$ —is formed with all nodes  $i$  for which  $I(i, k)$  is within a percentage  $\alpha \in ]0,1]$  of the maximum value  $I_{max}$  in  $CL$ :

$$RCL_k = \{i \in CL : I(i, k) \geq \alpha I_{max}\} \text{ where } I_{max} = \max_{i \in CL} I(i, k)$$

Our constructive method randomly selects an element in  $RCL_k$ , and performs the corresponding assignment. The method proceeds in this way, starting with cluster  $k = 1$ , and assigning elements to it until the lower bound constraint is satisfied ( $L \leq W_1$  where  $W_1 = \sum_{j \in V_1} w_j$ ). Then, the process moves to cluster 2 and proceeds in this way until the sum of the weights of the elements assigned to each cluster is larger than or equal to  $L$ .

In the following steps, the candidate set  $CL$  is formed with the pairs  $(i, k)$  with unassigned nodes  $i$  and those clusters  $k$  such that the sum of the weights of the elements already assigned to the cluster plus the weight of  $i$  is lower than or equal to  $U$ .

$$CL = \{(i, k) : 1 \leq k \leq p, W_k + w_i \leq U\} \text{ where } W_k = \sum_{j \in V_k} w_j$$

The method proceeds to evaluate  $I(i, k)$  for all  $(i, k)$  in  $CL$  build  $RCL$  with the  $(i, k)$  pairs with an evaluation within a percentage  $\alpha \in ]0,1]$  of the maximum value in  $CL$ , and select one pair at random. It stops when all the nodes have been assigned to clusters.

Once a solution  $x$  is obtained, we apply our improvement method, which consists of a deterministic VND (Variable Neighborhood Descent) based on two neighborhoods,  $N_0(x)$  and  $N_3(x)$ . The method determines first a best neighbor  $x'$  of  $x$  in  $N_0(x)$ . If  $x'$  is better than  $x$ , then  $x$  is replaced with  $x'$  and the method searches now for the best neighbor in  $N_0(x')$ , thus performing a local search in  $N_0$  while it improves the current solution. When the current solution  $x$  cannot be improved in  $N_0$ , then the method resorts to  $N_3$  and determines the best neighbor  $x'$  of  $x$  in  $N_3(x)$ . If  $x'$  is better than  $x$ , then the method comes back to search in  $N_0(x')$ ; otherwise the VND finishes. In short, the algorithm performs a local search for the best solution in  $N_0$  and only resorts to searching  $N_3$  when the process is trapped in a local optimum found in  $N_0$ . The improvement method considers only feasible moves.

## 4. Tabu Search for the CCP

As an alternative to the methods described in the previous sections, we engaged in the development of a tabu search specifically designed for the CCP. In other words, we studied the characteristics of the CCP and instead of adapting an existing procedure, we designed an original one. The main characteristic of this procedure consists of a new neighborhood structure based on 2-1 exchanges that we refer to as  $N_4(x)$ .

$N_4(x)$  may be considered a variant of  $N_2(x)$  in the sense that it explores exchanges of two nodes, say  $i$  and  $j$ , belonging to the same cluster  $k$  for a node  $l$  that belongs to a cluster  $s$ . In this neighborhood, unlike in  $N_2(x)$ , nodes  $i$  and  $j$  may or may not be connected. That is,  $N_4(x)$  does not require that  $(i, j) \in E$ . The neighborhood includes moves that are contained neither in  $N_0(x)$  nor in  $N_3(x)$ . For instance, consider the following situation.

Cluster A contains nodes 1 and 2 with weights  $w_1 = w_2 = 3$  and has remaining capacity of 2 units. Cluster B contains node 3 with  $w_3 = 6$  and remaining capacity of 1. Under  $N_0(x)$ , the insertions of node 1 to cluster B, node 2 to cluster B and node 3 to cluster A are considered. However, none of these moves are feasible. Likewise, the swaps of nodes 1 with 3 and 2 with 3 under  $N_3(x)$  result in infeasible solutions. However, the 2-1 exchange that moves nodes 1 and 2 from cluster A to B and node 3 from cluster B to A is feasible.

Our short-term memory tabu search based on  $N_4(x)$  operates as follows. All 2-1 exchanges are evaluated and the best move (according to the objective function value) is selected. The three nodes participating in the exchange are made tabu-active. In subsequent iterations, a move is classified tabu if it contains one or more tabu-active nodes. Nodes remain tabu-active for *tenure* iterations. The tabu status of a move is waved if the exchange leads to a solution that improves upon the incumbent.

## 5. Computational Experiments

This section describes the computational experiments that we performed to test the effectiveness and efficiency of the procedures discussed above. The GRASP by Deng and Bard (2011), Prev\_GRASP, and our adaptation to the CCP of the TS\_SO originally proposed for the MDGP by Gallego et al. (2013), CCP\_TS\_SO, were implemented in Java SE 6. Our GRASP (Section 3), simply called GRASP, and the TS (Section 4) for the CCP were implemented in C. All experiments were conducted on an Intel Core 2 Quad CPU Q 8300 with 6 GB of RAM.

We employed 50 instances in our experimentation. This benchmark set of instances, referred to as CCPLIB, is available at <http://www.opticom.es/ccp>. The set is divided into two subsets:

1. *RanReal* — This set, originally proposed by Gallego et al. (2013) for the MDGP, consists of 40  $n \times n$  matrices in which the benefit values  $c_{ij}$  are real numbers generated using a Uniform distribution  $U(0,1000)$ . We have adapted these instances to the CCP by generating the node weights with a Uniform distribution  $U(0,10)$ . There are 20 instances with  $n = 240$ ,  $p = 12$ ,  $L = 75$ , and  $U = 125$ . The remaining 20 instances have  $n = 480$ ,  $p = 20$ ,  $L = 100$ , and  $U = 150$ .



2. *DB* — This set is based on the problem with  $n = 82$  and  $p = 8$  introduced by Deng and Bard (2011) in the context of mail delivery. We keep the benefit values as they appear in that original instances. However, instead of considering all the node weights equal to 1, as in Deng’s and Bard’s article, we randomly generate the node weights with a Uniform distribution  $U(0,10)$  and thus changing the instances from MDGP to CCP. The set consists of 10 instances with  $n = 82$ ,  $p = 8$ ,  $L = 25$ , and  $U = 75$ .

We have selected 15 representative instances with different characteristics to perform a preliminary experimentation to identify effective values for critical search parameters and to compare different designs. Specifically, we select 6 *RanReal* instances with  $n = 240$ , 6 *RanReal* instances with  $n = 480$ , and 3 *DB* instances.

In our first preliminary experiment we test four different variants of the simplified GRASP method for the CCP described in Section 3. In particular, we consider four different values of the parameter  $\alpha$ , namely, 0.2, 0.4, 0.6, and 0.8. We generated 100 different solutions to each instance with each GRASP variant. Table 1 shows, for each variant, the average objective function value, the average percent deviation from the best solutions obtained within this experiment, and the number of instances in which the method is able to match the best solutions obtained within this experiment (out of 15 instances). These measures are “local” in the sense that the best solutions are those found within the experiment. They are used because they allow us to discriminate among the procedures being tested and identify the better alternatives. However, as a point of reference, Table 1 also includes the average deviation achieved by each variant against the best-known solutions. This is a “global” measure in the sense that the best-known solutions are those found across all experiments in this paper, which as far as we know represents the best-known published solutions.

Additionally, Table 1 includes the Score statistic described in Ribeiro et al. (2002). For each problem instance in an experiment, the Score of a method is the number of methods that found a solution that is strictly better than the one that the method being scored found. For a set of instances, the Score is the sum of the individual scores. Therefore, while the overall best score in any experiment is zero, the worst score depends on the number of alternative procedures and the size of the test set. Since there are 15 instances in the test set for this experiment and 4 different procedures (one for each value of alpha) then the worst possible score is 45 (i.e., 15 instances times 3 competing procedures).

Metric	$\alpha$			
	0.2	0.4	0.6	0.8
Average objective function value	233397.5	233203.9	233392.9	233226.3
Number of best solutions	3	3	5	4
Local average deviation	0.36%	0.44%	0.36%	0.43%
Global average deviation	14.95%	14.91%	14.95%	14.96%
Score (worst = 45)	24	23	19	24

Table 1. GRASP with different values of  $\alpha$ .

Table 1 shows the slightly better performance (in terms of both number of best solutions and deviation) achieved by GRASP with the parameter  $\alpha$  set to 0.6. This table also shows that, in

general, the differences observed among these statistics related to the four variants tested are relatively small. This is particularly true in the deviations with respect to the overall best known solutions, which confirms that GRASP constructions on their own are not able to produce high quality solutions. The local measures, including the Score all favor the parameter setting with  $\alpha = 0.6$ .

An important question in the context of multi-start procedures in general, and GRASP in particular, relates to the origin of the high-quality solutions obtained after applying the improvement method. Specifically, we are interested in measuring the correlation between the objective function value of the constructed solutions and the value of the improved solutions (i.e., the local optima found after applying the improvement method). We would like determine whether high-quality local optima are related to high-quality constructions. If so, this could lead us to implementing filter mechanisms to save the computational time invested on attempting to improve somewhat poor constructions. To this end, we generated 100 solutions for each problem in our set of 15 representative instances (testing set) and computed the correlation coefficient between both the initial solution and the local optimum. The resulting correlation coefficient of 0.01 indicates that good local optima can come from initial solutions of any quality. Figure 1 shows this lack of correlation for 100 solutions of a representative instance in our test set. We observe that the points are scattered in the plane without any discernible pattern.




Figure 1. Constructed and Improved values of 100 solutions

In our second preliminary experiment we compare GRASP (with  $\alpha$  set to 0.6) with the Prev\_GRASP on the test set of 15 instances. We run both methods for 60 seconds on each instance and report the results in Table 2.

Metric	Prev_GRASP	GRASP(0.6)
Average objective function value	194109.8	233841.7
Number of best solutions	6	9
Local average deviation	10.99%	2.84%
Global average deviation	20.68%	14.38%
Score (worst = 15)	9	6

Table 2. Comparison of GRASP methods.

Table 2 shows that the proposed simplified GRASP seems to perform slightly better than the one in the literature. GRASP(0.6) is able to match 9 out of 15 best solutions and exhibits a 2.84% deviation from the best solutions found within this experiment. The global average deviation (i.e., the one against the overall best-known solutions) of 14.38% compares well with 20.68% achieved by Prev\_GRASP. Note that in a comparison of two methods, the Score provides the same information obtained by the “number of best solutions found.” Although not shown in Table 2, it is interesting to point out that Prev\_GRASP performs better than GRASP(0.6) on the medium size RanReal instances ( $n = 240$ ), while GRASP(0.6) outperforms Prev\_GRASP on the large RanReal instances ( $n = 480$ ). We apply a statistical, non-parametric, test to confirm that both methods perform similarly. We employ the *Wilcoxon test* to make a pairwise comparison of Prev\_GRASP and GRASP(0.6). The results of this test (with a  $p$ -value of 0.233) determined that the solutions obtained by the two methods could come from the same population.

To finish the analysis of GRASP methods, we compare three different designs for the construction phase: 1) the one used in Prev\_GRASP in which the candidate list  $CL$  consists of both, nodes and edges (N&E\_GRASP), 2) one in which the  $CL$  consists of nodes (N\_GRASP), and 3) one in which the  $CL$  consists of edges (E\_GRASP). To this end, we generated 100 solutions for each problem in our set of 15 representative instances with each of these three constructive methods and report, in Table 3, the results. They clearly shows the superiority of the conventional GRASP construction, N\_GRASP, in which the next element to be included in the solution is a single node, instead of an edge, or either a node or an edge.

Metric	N_GRASP	E_GRASP	N&E_GRASP
Average objective function value	195495.8	189815.8	181338.4
Number of best solutions	15	0	0
Local average deviation	0.0%	5.4%	8.7%
Global average deviation	28.1%	32.3%	34.6%
Score (worst = 30)	0	16	29

Table 3. Comparison of GRASP constructive variants.

In our fourth preliminary experiment, we test the tabu search method, TS, described in Section 4, with several *tenure* values. We set this parameter to 5, 10, 15, and  $\sqrt{n}$ . Table 4 shows the results (i.e., average values of the five statistics) on the 15 test instances. These results show that the best outcomes are obtained with *tenure* = 5 since this variant is able to achieve 8 best solutions and exhibits a local average percent deviation of 0.29%.

Metric	<i>tenure</i>			
	5	10	15	$\sqrt{n}$
Average objective function value	282305.3	282739.3	281972.7	281532.8
Number of best solutions	8	4	2	1
Local average deviation	0.29%	0.88%	2.05%	1.03%
Global average deviation	2.40%	2.47%	3.61%	2.88%
Score (worst = 45)	14	23	26	21

Table 4. TS with different values of *tenure*.

In the final experiment of this block, we add to the comparison set our adaptation to the CCP of the tabu search developed by Gallego et al. (2013) for the MDGP. As described in Section 2, this method has the ability to search in the space of infeasible solutions by way of a strategic oscillation component. To test the merit of this feature, we have considered two variants of this adaptation: AdTS, without the SO, and AdTS\_SO, with the oscillation. We also compare both methods with our TS with  $tenure = 5$ . As in the previous preliminary experiments, we run each method for 60 seconds on each instance. The results associated with this experiment are in Table 5.

<b>Metric</b>	<b>AdTS</b>	<b>AdTS_SO</b>	<b>TS</b>
Average objective function value	256406.2	269549.5	281203.4
Number of best solutions	1	2	12
Local average deviation	6.81%	2.89%	1.06%
Global average deviation	8.02%	4.16%	2.37%
Score (worst = 30)	26	14	6

Table 5. Tabu search methods

The results in Table 5 indicate that, as expected, our tabu search method specifically designed for the CCP obtains better solutions than the two versions adapted from the MDGP to the CCP. This table also shows that the strategic oscillation component (AdTS\_SO) is effective when compared to the version without it. TS obtains 12 best solutions out of the 15 instances, while AdTS and AdTS-SO obtain 1 and 2, respectively. Additionally, both the local and global average deviations as well as the score favor TS.

We applied the Friedman test for paired samples to the data used to generate Table 5. The resulting  $p$ -value of 0.001 obtained in this experiment indicates that there are statistically significant differences among the three methods tested (we are using the typical significance level of  $\alpha = 0.05$  as the threshold between rejecting or not the null hypothesis). A typical post-test analysis consists of ranking the methods under consideration by their average rank values. The result is that the TS method is the best with an average rank of 2.60, followed by AdTS\_SO (2.10), while AdTS ranks third (1.30). Finally, we compare the best two methods in this experiment, TS and AdTS\_SO, with the pairwise Wilcoxon test. The resulting  $p$ -value of 0.002 obtained confirms that the two methods are significantly different.

In our final experiment we compare the best methods on the entire set of 50 instances. Specifically, we compare:

- Prev\_GRASP. The GRASP method by Deng and Bard (2011) described in Section 2
- AdTS\_SO. Our adaptation of the tabu search by Gallego et al. (2013) described in Section 2.
- GRASP. Our GRASP method described in Section 3.
- TS. Our tabu search method described in Section 4.
- GRASP+TS. A combination of the two methods that we have developed.

The GRASP+TS hybrid consists of running GRASP(0.6) for half of the total running time in the experiment, and then, for the rest of the running time, applying TS starting from the best solution found by GRASP.

Tables 6, 7, and 8, show respectively the results of the 20 *RanReal* instances with  $n = 240$ , 20 *RanReal* instances with  $n = 480$ , and the 10 DB instances. We run each method for 60 seconds on each instance.

Metric	Prev_GRASP	AdTS_SO	TS	GRASP	GRASP+TS
Global average deviation	9.43%	2.79%	0.15%	14.55%	0.61%
Number of best solutions	0	0	16	0	4
Score (worst = 200)	79	40	4	100	16

Table 6. Comparison of best methods on the 20 *RanReal* instances with  $n = 240$

Metric	Prev_GRASP	AdTS_SO	TS	GRASP	GRASP+TS
Global average deviation	41.78%	7.54%	1.99%	19.85%	2.07%
Number of best solutions	0	0	10	0	10
Score (worst = 200)	100	37	10	80	13

Table 7. Comparison of best methods on the 20 *RanReal* instances with  $n = 480$

Metric	Prev_GRASP	AdTS_SO	TS	GRASP	GRASP+TS
Global average deviation	0.33%	0.02%	4.52%	0.08%	0.13%
Number of best solutions	0	6	0	2	2
Score (worst = 200)	41	11	50	20	21

Table 8. Comparison of best methods on the 10 DB instances

The results in Tables 6, 7, and 8 indicate that the performance of the competing methods varies with the data sets. For instance, Prev\_GRASP has difficulties finding solutions of reasonable quality for *RanReal* problems with  $n = 480$ , producing a large average deviation of 41.78%. However, the same method performs much better on the DB set, where it obtains an average deviation of 0.33%. On the other hand, TS shows better performance than the other procedures in the *RanReal* sets but comes out last in the DB set.

We apply statistical tests to the data obtained from this last experiment in order to detect differences among the Prev\_GRASP, AdTS\_SO and GRASP+TS. From our three designs, GRASP, TS and GRASP+TS, we chose GRASP+TS because of its better average overall performance. We first applied the Friedman test to test the hypothesis that there are differences among the procedures. The hypothesis that they are the same is rejected with a  $p$ -value lower than 0.001. The resulting ranking is GRASP+TS (2.74), AdTS\_SO (2.19), and Prev\_GRASP (1.07). We then applied the pairwise Wilcoxon test to GRASP+TS and AdTS\_SO, with the outcome indicating a better performance of GRASP+TS. This conclusion is reached with a  $p$ -value lower than 0.001.

In an attempt to understand how GRASP and TS interact within GRASP+TS, we recorded the objective function value of the best solution found during the search when applying GRASP+TS to

the first *RanReal* instance with  $n = 240$ . The profile of this run is depicted in Figure 2. The plot shows that TS is able to provide a boost after the GRASP search stalls.




Figure 2. Search profile of GRASP+TS on the RanReal240\_01 instance

## 6. Conclusions

The CCP is a difficult combinatorial optimization problem that is closely related to the MDGP, although this connection had not been previously discussed. Of particular interest in our work has been testing the effects of a variety of search strategies —such as those that allow the search to move outside the feasible region— and neighborhood structures within the GRASP framework. We also explored a tabu search design with a 2-1 exchange neighborhood that proved effective. Our work is the first one in which the experiments are performed on general instances of the CCP, that is, those for which the node weights are different from one. These instances and associated tests establish a benchmark that currently is not in the literature.

## Acknowledgments

This research has been partially supported by the *Ministerio de Educación Cultura y Deporte* of Spain (Grant Ref. PRX12/00016), the *Ministerio de Economía y Competitividad* of Spain (Grant Ref. TIN2012-35632) and the Generalitat Valenciana (Prometeo 2013/049).

## References

- Deng, Y. and J.F. Bard (2011) "A reactive GRASP with path relinking for the capacitated clustering," *Journal of Heuristics*, vol. 17, pp. 119-152.
- Fan, Z. P., Y. Chen, J. Ma and S. Zeng (2011) "A hybrid genetic algorithmic approach to the maximally diverse grouping problem", *Journal of the Operational Research Society*, vol. 62, pp. 92-99.
- Feo, T., O. Goldschmidt and M. Khellaf (1992) "One-half approximation algorithms for the k-partition problem," *Operations Research*, vol. 40, pp. S170-S173.
- Festa P. and M. G. C. Resende (2009) "An annotated bibliography of GRASP, Part I: Algorithms," *International Transactions in Operational Research*, vol. 16, pp. 1-24.
- Gallego, M., A. Duarte, M. Laguna, and R. Martí (2009) Hybrid heuristics for the maximum diversity problem. *Computational Optimization and Applications*, 44(3):411, 2009.
- Gallego, M., M. Laguna, R., Martí, and A. Duarte (2013) "Tabu search with strategic oscillation for the maximally diverse grouping problem," *Journal of the Operational Research Society*. To appear.
- Glover, F. and M. Laguna (1997) *Tabu Search*, Kluwer Academic Publisher: Boston.
- O'Brien, F. A. and J. Mingers (1995) "The equitable partitioning problem: a heuristic algorithm applied to the allocation of university student accommodation," Warwick Business School, *Research Paper* no. 187.
- Martí, R. and F. Sandoya (2013) "The equitable dispersion problem," *Computers and Operations Research*, Available online 21 April 2012, forthcoming.
- Ribeiro, C.C., E. Uchoa, and R.F. Werneck. (2002) "A hybrid GRASP with perturbations for the Steiner problem in graphs". *INFORMS Journal on Computing*, 14:228–246.
- Weitz, R. R. and M. T. Jelassi (1992) "Assigning students to groups: a multi-criteria decision support system approach," *Decision Sciences*, vol. 23, no. 3, pp. 746-757.
- Weitz, R. R. and S. Lakshminarayanan (1998) "An empirical comparison of heuristic methods for creating maximally diverse groups," *Journal of the Operational Research Society*, vol. 49, no. 6, pp. 635-646.