Discrete Optimization

Advanced Greedy Randomized Adaptive Search Procedure for the Obnoxious p-Median problem☆

J. Manuel Colmenar[a], Peter Greistorfer[b], Rafael Martí[c], Abraham Duarte[a,*]

[a] *Departamento de Ciencias de la Computación, Universidad Rey Juan Carlos, Spain*
[b] *Institut für Produktion und Logistik, Karl-Franzens-Universität Graz, Austria*
[c] *Departamento de Estadística e Investigación Operativa, Universidad de Valencia, Spain*

A B S T R A C T

The Obnoxious p-Median problem consists in selecting a subset of p facilities from a given set of possible locations, in such a way that the sum of the distances between each customer and its nearest facility is maximized. The problem is $\mathcal{NP}$-hard and can be formulated as an integer linear program. It was introduced in the 1990s, and a branch and cut method coupled with a tabu search has been recently proposed. In this paper, we propose a heuristic method – based on the Greedy Randomized Adaptive Search Procedure, GRASP, methodology – for finding approximate solutions to this optimization problem. In particular, we consider an advanced GRASP design in which a filtering mechanism avoids applying the local search method to low quality constructed solutions. Empirical results indicate that the proposed implementation compares favorably to previous methods. This fact is confirmed with non-parametric statistical tests.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Facility location has been of practical and theoretical interest for more than the half of a century (Cohen, 1973; Klose & Drexl, 2005). First linear programming (LP) formulations origin in the late fifties and were soon followed by solution techniques, which later on became well-known as Branch and Bound (B&B) or Mixed Integer Programing (MIP). The classical so-called warehouse location problem (WLP), also well-known as (simple) facility or plant location problem, is an uncapacitated optimization problem, which can be modeled as an LP (Balinski, 1965) or network problem (Drezner & Hamacher, 2004; Melkote & Daskin, 2001). It is a site-selecting location-allocation model with a min–sum objective, where from a number of potential facility or warehouse sites the set of costumers has to be serviced, while minimizing the total fixed site-costs (location) plus the total variable customer assignment costs (allocation). The WLP is $\mathcal{NP}$-hard (Garey & Johnson, 1979; Papadimitriou & Yannakakis, 1991) and, as such, has been in the focus of researchers interested in developing specialized LP-approaches (Körkel, 1999) or approximative constructive

and local search or improvement algorithms, including standard add- or drop-heuristics, and Lagrangian approaches (Beasley, 1993; Kuehn & Hamburger, 1963), which then were followed by more advanced metaheuristics like genetic algorithms or tabu search and its derivatives (Greistorfer & Rego, 2006; Kratica, Tošić, Filipović, & Ljubic, 2001; Michel & Hentenryck, 2004). Comparing the WLP with the p-median problem, pMP (Hakimi, 1964; 1965), one identifies two differences: (1) there are no fixed site-costs involved and (2) the number of finally opened sites, p, is no longer a decision variable, but becomes included in the model. The pMP can be solved in polynomial time for fixed values of p, but is strongly $\mathcal{NP}$-hard for variable values of p (Current, Daskin, and Shilling, 2004, chap. 3; Garey & Johnson, 1979; Megiddo & Supowit, 1984). Consequently, any pMP is a special case of the general class of WLPs. Such as for the WLP, a variety of solution procedures, exact approaches, primal-dual approaches and metaheuristics, have been introduced for the pMP (Mladenović, Brimberg, Hansen, & Moreno-Pérez, 2007; Reese, 2006; Tansel, Francis, & Lowe, 1983). A most recent paper, Batta, Lejeuneb, and Prasad (2014), is recommended for a classical as well as modern view (dispersion, population, and equity criteria) on locational developments, offering an additional special focus on pMP.

This work relates to location type problems like the WLP and pMP, additionally taking account of the so-called obnoxious or semi-obnoxious effects. Such effects often occur when interesting services of some provider are based on unwanted, but inevitably necessary locations, which do not add additional value to the

Table 1
Matrix of distances. Last row shows the sum of all the client distances for each facility.

	j_1	j_2	j_3	j_4	j_5	j_6
i_1	3	2	10	13	4	8
i_2	14	14	11	3	15	12
i_3	5	4	6	12	14	6
i_4	3	1	1	3	10	9
i_5	13	9	2	10	9	4
i_6	9	5	11	14	4	11
i_7	4	3	4	13	4	2
i_8	12	2	9	3	15	13
i_9	9	11	1	2	7	4
$\sum$	72	51	55	73	82	69

Table 2
Example of solutions with $p = 3$.

(a) $S = \{j_2, j_5, j_6\}$			(b) $S' = \{j_2, j_3, j_4\}$		
	j_2	j_5	j_6		j_2	j_3	j_4
i_1	**2**	4	8	i_1	**2**	10	13
i_2	14	15	**12**	i_2	14	11	**3**
i_3	**4**	14	6	i_3	**4**	6	12
i_4	**1**	10	9	i_4	**1**	1	3
i_5	9	9	**4**	i_5	9	**2**	10
i_6	5	**4**	11	i_6	**5**	11	14
i_7	3	4	**2**	i_7	**3**	4	13
i_8	**2**	15	13	i_8	**2**	9	3
i_9	11	7	**4**	i_9	11	**1**	2
		$f(S) = 35$				$f(S) = 23$	

product. Quite contrary, pure obnoxiousness clearly devalues the production process. Obnoxious problems were firstly coined in Church and Garfinkel (1978) who located a facility in a network using an exact method and also introduced the term of the so-called bottleneck points for a network.

Erkut and Neuman (1989) used the terms *disservice* and *service* of the people in the vicinity of an (semi-)obnoxious facility that occurs during the processing of a product. In general, such services include a type of hazardous material, waste disposal, water treatment, nuclear power or chemical plants as well as big public facilities like airports. Regularly, these problems arise in the context of urban settings when the network may consist of noisy or polluting roads, transportation corridors, or rail lines (Segal, 2003).

The present work deals with the Obnoxious p-Median (OpM) problem. It can be formally defined as follows. Let I be a set of clients, J a set of facilities, and d_{ij} the distance between the client $i \in I$ and the facility $j \in J$. The OpM problem consists in finding a set S with p facilities (with $S \subseteq J$ and $p < |J|$), such that the sum of the minimum distance between each client and the set of facilities is maximized. In mathematical terms:

$$\max \sum_{i \in I} \min\{d_{ij} : j \in S\}$$

subject to
$$S \subseteq J , \ |S| = p$$

Facilities in S are called open facilities, while facilities in $J \backslash S$ are known as closed or unopened facilities.

Table 1 shows an example of pair-client distances, where there are 9 clients $I = \{i_1, i_2, i_3, i_4, i_5, i_6, i_7, i_8, i_9\}$ and 6 potential facilities $J = \{j_1, j_2, j_3, j_4, j_5, j_6\}$. Suppose that $p = 3$. Then, a solution of the OpM consists of selecting 3 facilities out of 6 in J. Table 2a shows a solution with $S = \{j_2, j_5, j_6\}$. The minimum distance between each client and the corresponding facility is highlighted with bold font. For example, the distances from client i_1 to each facility are: $d(i_1, j_2) = 2$, $d(i_1, j_5) = 4$, and $d(i_1, j_8) = 8$. Then, j_2 is the closest

facility to client i_1. The value of the objective function of this solution, denoted as $f(S)$, is the sum of the minimum distances from each client to the set of facilities: $f(S) = 2 + 12 + 4 + 1 + 4 + 4 + 2 + 2 + 4 = 35$.

Table 2 b shows a different solution $S' = \{j_2, j_3, j_4\}$, where the value of the objective function is $f(S') = 23$. Considering that the OpM is a maximization problem, solution S is better than solution S'. In other words, the minimum distances among the clients and facilities in S has a sum that is larger than the one in S'.

In this paper we explore the adaptation of the Greedy Randomized Adaptive Search Procedure methodology, GRASP, introduced in Feo and Resende (1989), to solve the OpM problem. Each GRASP iteration consists of constructing a trial solution and then applying an improvement procedure to find a local optimum. We explore different designs for both phases, construction and improvement. In particular, in Section 3 we propose two different constructive methods and in Section 4, we describe two local search algorithms. Additionally, we propose an efficient strategy to update the objective function value that considerably reduces the running time of GRASP. In Section 5, we present a filtering strategy intended to discard low-quality solutions and to selectively apply the local search only to promising solutions. This mechanism reduces the computing time without deteriorating the quality of the final solution. Section 6 reports on an extensive computational experience to validate the proposed algorithm by comparing its performance with those in the current state of the art. Finally, Section 7 summarizes the main conclusions of our research.

2. Literature review

There exists an extensive literature treating obnoxious and semi-obnoxious situations, modeled on the plane or in a graph, using max–min, max–max and often bi-criteria objectives (Batta et al., 2014; Cappanera, 1999; Conceição Fonseca & Captivo, 2007; Erkut & Neuman, 1989; Farahani, SteadieSeifi, & Asgari, 2010; Segal, 2003). This goes along with the vast research on (obnoxious) p-median, p-center and similar location problems. Some of these sources are highlighted subsequently.

In Drezner and Wesolowsky (1980) a planar, Euclidean model of a 1-facility-problem is presented in which the shortest weighted distance to a point is maximized. Simultaneously, a side constraint must hold, i.e., the facility must be within a pre-specified distance from each point. The problem is solved using a graphical circles-based approach. A continuation of this work can be found in Drezner and Wesolowsky (1983), where a constrained obnoxious problem is considered for which a weighted rectangular-metric objective has to be maximized. The authors propose two algorithms, a so-called boundary and segment search and an LP-based approach. A semi-obnoxious scenario is solved in Melachrinoudis (1985). It uses a max–min model and seeks for a point on a convex two-dimensional bounded region, which maximizes the minimum weighted distance from that point to a given set of existing points in the region to be considered. As before, a circles-based, straight-forward geometrical approach as well as an exact algorithm, using Kuhn–Tucker boundary points, is introduced. Complexity issues regarding the placement of several facilities in an obnoxious setting, a p max–min problem, were considered by Tamir (1991). It is shown that even the finding of an approximate solution is $\mathcal{NP}$-hard. Again, Drezner and Wesolowsky (1996) deal with an Euclidean network and an obnoxious urban situation in which a location has to be determined inside the convex hull of a number of nodes in order to maximize the minimum weighted distance between this point and the nodes and arcs of the network. Among others, the authors offer an ϵ-approximation scheme, which was later improved in Segal (2003). Welch and Salhi (1997) present three heuristics to solve the max–min formulation for siting p

facilities on a network considering pollution dispersion equation and additionally facility interaction. Starting with $p = 1$, a graphical method approximates the pollution dispersion by the use of polygons. Afterwards, the general case for p and the combination of both the p max–min and p max–sum objectives, using a lexicographic approach, are investigated. Moreover, a simulated annealing algorithm is used to evaluate the base-algorithms.

Ohsawa and Tamura (2003) study the placement of a semi-obnoxious facility in a continuous plane underlying the bi-objective of maximizing the distance to the nearest inhabitant and minimizing the sum of distances to all users. In doing so, an elliptic max–min criterion and the rectangular min–sum criterion are used to model the push- and pull-aspect, respectively. For determining an efficient set and according trade-off curves, polynomial-time algorithms are presented. The obnoxious topic is considerably extended in Cappanera, Gallo, and Maffioli (2004), who additionally consider a routing component, in which sites and transportation links may be affected by a hazardous single commodity. The resulting flow model of this obnoxious facility location problem is then solved by a Lagrangian heuristic and possibly improved by a subsequent B&B-algorithm. A multi-objective model for the location of landfills is discussed in Rakas, Teodorović, and Kim (2004). They use a combined weighted objective and solve the overall problem based on a composition of several MIP-solutions. Highlighted are the use of fuzzy arithmetic rules with triangular fuzzy numbers to cover the aspect of uncertainty and the coverage of a real-case-study. Tamir (2006) focuses on the problem of finding the location of two new and obnoxious facilities in the plane. The objective is to maximize the minimum of all weighted distances between the existing customer facilities and the two new facilities, and the weighted distance between the pair of new facilities. For the Euclidean and rectilinear versions of this 2-obnoxious facility location problem efficient logarithmic running time algorithms are presented, and the rectilinear case is extended to also handle general (non-convex) planar and compact polygonal sets. Yapicioglu, Smith, and Dozier (2007) use a single- and a bi-objective particle swarm optimizer, well-known from evolutionary computation for their semi-obnoxious facility location problem. The idea is to simultaneously minimize transportation costs and undesirable effects, represented by a piecewise distance function depicting the degree of obnoxiousness. A computational analysis approximates the linear complexity in effort with increasing problem size. Berman and Wang (2007; 2008) consider the problem of locating single and multiple semi-obnoxious facilities while expropriating the nearest demand nodes under the restriction of a given total budget hold by the developer. The objective is to maximize the minimum weighted distance to the non-expropriated demand nodes. The algorithms introduced are based on dominating sets, Lagrangian relaxation and B&B. In Berman, Drezner, Wang, and Wesolowsky (2008) the base topic is extended to the possible expropriation of population centers in the vicinity of selected routes of a transportation model, thus introducing a routing effect, e.g., when avoiding the transportation of hazardous material in populated resident zones. Solution approaches for the single- and multiple-flow-case are based on a greedy heuristic and on column generation or branch and price, respectively.

A multi-objective obnoxious facility location model is treated in Bhattacharya (2011), in which an obnoxious facility interacts with a given set of existing costomers on the plane and has to be located according to a max–min criterion, while at the same time maximizing the number of existing facility points covered. The algorithm proposed is a rectangle decomposition, a step-wise analytical approach that embeds linear as well as non-linear programming sub-procedures. Ortigosa, Hendrix, and Redondo (2011) describe heuristic methods, based on Pareto-approaches, for

generating semi-obnoxious locations. In doing so, they use a bi-objective convex/non-convex approach and several randomized sampling methods to optimize the desirable and the non-desirable aspect of the model, i.e. the min–sum pull-objective and the obnoxious push-objective. Coutinho-Rodrigues, Tralhão, and Alçada-Almeida (2012) deal with an empiric case of an urban waste collection problem in which a number of waste containers has to be located in a way that combines push and pull aspects of being serviced by a location too near (visual aesthetics, smell, nuisance, attractiveness) and by a location too far away (walking effort to deposit waste), respectively. Again, a bi-objective model approach is used, this time to minimize both, the container investment cost and a so-called weighted average customer dissatisfaction. The solution approach is a MIP-model, letting the decision-maker the final choice from non-dominated locations. Lozano, Duarte, Gortázar, and Martí (2012) study the so-called separation or dual bandwidth problem. This $\mathcal{NP}$-hard problem uses a max–min objective function to optimize transmitters in radio frequency assignment by assigning different frequencies in such a way that physically neighboring transmitters have as different frequencies as possible. To solve this type of layout problem, different meta-heuristic approaches are studied, e.g., ejection chains, tabu search and variable neighborhood search. The interesting point in Heydari and Melachrinoudis (2012), who describe a semi-obnoxious facility with elliptic max–min and network min–sum objectives, is the use of two different metric measures. To find the efficient set of this bi-objective problem three phases are implemented: a network redefinition (of the transportation network due to bottleneck points and Voronoi-clustering), elimination (of inefficient edges) and construction (of the efficient set and its non-dominated trade-off curve). In Plastria, Gordillo, and Carrizosa (2013) a continuous, i.e., Euclidean, semi-obnoxious facility location problem is in the focus of interest. The analytical method offered, based on machine learning, gives the necessary conditions for optimality, and, using the latter, develops a polynomial enumeration on the set of dominating solutions.

In this paper, we target the Obnoxious p-Median (OpM) problem, which consists of selecting a subset of p facilities from a given set of possible locations, in such a way that the sum over all customers of the distances between each customer and its nearest facility is maximized. OpMP was introduced in the 1990s (Cappanera, 1999; Eiselt & Laporte, 1995; Welch & Salhi, 1997) and, as of today, has only gained relatively few attention. It is $\mathcal{NP}$-hard (Tamir, 1991) and can be formulated as a binary LP. Note that Burkard, Fathali, and Taghizadeh Kakhki (2007) proved that the special case of the OpM problem on a tree can be solved in linear time. A branch and cut method coupled with a tabu search has been recently suggested by Belotti, Labbé, Maffioli, and Ndiaye (2007). We include both methods, the branch and cut and the tabu search in our computational experimentation reported in Section 6.

3. Constructive methods

The GRASP construction phase is an iterative, greedy, randomized and adaptive procedure that constructs solutions from scratch. It is iterative since the solution is built by considering one element at a time. It is greedy because the addition of each element to the solution under construction is guided by a greedy function. It is randomized because it performs a random selection from a list of good candidates. Finally, the method is adaptive in the sense of updating relevant information from a construction step to the next one.

A complete solution of the OpM is a set $S \subset J$ with exactly p elements ($|S| = p$). Let $S \subset J$ be a *partial solution* to the OpM problem (i.e., $0 \leq |S| < p$). For each client $i \in I$, we define the

Table 3
Greedy evaluation with respect to the partial solution $S = \{j_5\}$.

	$g(S, j_1)$	$g(S, j_2)$	$g(S, j_3)$	$g(S, j_4)$	$g(S, j_6)$
i_1	3−4 = −1	2−4= −2	10−4= 6	13−4= 9	8−4 = 4
i_2	14−15= −1	14−15= −1	11−15=−4	3−15= −12	12−15= −3
i_3	5−14=−9	4−14=−10	6−14=−8	12−14=−2	6−14=−8
i_4	3−10= −7	1−10= −9	1−10= −9	3−10= −7	9−10= −1
i_5	13−9=4	9−9=0	2−9=−7	10−9=1	4−9=−5
i_6	9−4=5	5−4=1	11−4=7	14−4=10	11−4=7
i_7	4−4 =0	3−4 = −1	4−4 =0	13−4= 9	2 −4 = −2
i_8	12−15=−3	2−15=−13	9−15=−6	3−15=−12	13−15=−2
i_9	9−7=2	11−7=4	1−7=−6	2−7=−5	4−7=−3
	−**21**	−36	−40	−38	−24

Table 4
Greedy evaluation with respect to the partial solution $S = \{j_1, j_5\}$.

	$g(S, j_2)$	$g(S, j_3)$	$g(S, j_4)$	$g(S, j_6)$
i_1	2−3= −1	10−3= 7	13−3= 10	8−3 = 5
i_2	14−14= 0	11−14=−3	3−14= −11	12−14= −2
i_3	4−5=−1	6−5=1	12−5=7	6−5=1
i_4	1−3= −2	1−3= −2	3−3= 0	9−3= 6
i_5	9−9=0	2−9=−7	10−9=1	4−9=−5
i_6	5−4=1	11−4=7	14−4=10	11−4=7
i_7	3−4 = −1	4−4 =0	13−4= 9	2 −4 = −2
i_8	2−12=−10	9−12=−3	3−12=−9	13−12=1
i_9	11−7=4	1−7=−6	2−7=−5	4−7=−3
	−15	−21	−25	−**12**

minimum distance between i and the facilities in S, denoted as δ_i, as follows:

$$\forall i \in I \rightarrow \delta_i = \min_{j \in S} d_{ij} \qquad (1)$$

The value of this partial solution, $f(S)$, can be directly computed as:

$$f(S) = \sum_{i \in I} \delta_i \qquad (2)$$

We propose a greedy function that calculates the change in the objective function value when a new facility is added to a partial solution. In particular, if the facility j is included in the partial solution S, producing a new solution $S' = S \cup \{j\}$, the corresponding δ'-value for each client is updated as follows:

$$\forall i \in I \rightarrow \delta'_i = \min\{\delta_i, d_{ij}\}$$

Instead of directly computing the change of the objective function value, we determine the contribution of the new facility to the objective function. This is obtained with the greedy function $g(S, j)$:

$$g(S, j) = \sum_{i \in I} \min\{0, (d_{ij} - \delta_i)\} \qquad (3)$$

where $(d_{ij} - \delta_i) < 0$ indicates that the facility j is closer to i than the remaining facilities in S. Then, the inclusion of j reduces δ_i. On the other hand, if $(d_{ij} - \delta_i) \geq 0$, the inclusion of the facility j does not affect client i. We use $g(S, j)$ to determine the best potential facility, $j^\star$, to be included in the partial solution. In mathematical terms:

$$j^\star = \arg \max_{j \in J \setminus S}\{g(S, j)\}$$

Let us illustrate the computation of the proposed greedy function with the example shown in Section 1. Considering that the OpM problem consists of finding a subset of facilities distant from the clients, we first identify the facility with the largest sum of distances. In this example, the sum of distances from all clients to each facility is 72, 51, 55, 73, 82 and 69. Then, the most distant facility with respect to all clients is j_5 with a sum of distances equal to 82. Therefore, the greedy strategy selects this facility, obtaining the partial solution $S = \{j_5\}$ and the minimum distance from each client to j_5 is [4, 15, 14, 10, 9, 4, 4, 15, 7], respectively.

Table 3 calculates the objective function change when we try to include a new facility, showing for each client i, $d_{ij} - \delta_i$. For example, if we selected j_1 to be included in the partial solution, the situation of clients i_1, i_2, i_3, i_4, and i_8 would be deteriorated since j_1 is closer to these clients than j_5. On the other hand, i_5, i_6, i_7, and i_9 are not affected for the inclusion of this new facility. The facility which presents the best (higher) evaluation of the greedy function is, in fact, j_1, with a value equal to −21 according to Eq. (3). We then update the partial solution to $S = \{j_1, j_5\}$ and the array of distances to [3, 14, 5, 3, 9, 4, 4, 12, 7], respectively.

Similarly, Table 4 shows the values of the greedy function when we try to add a new facility to the current solution. The best option now is to include j_6, which obtains a value of −12. We then include j_6 in the partial solution, obtaining $S = \{j_1, j_5, j_6\}$ and the distances to elements: [3, 12, 5, 3, 4, 4, 2, 12, 4]. Considering that $p = 3$, S is a feasible solution with an objective function value of 49 (i.e., the sum of minimum distances).

Now we describe our two constructive algorithms $C1$ and $C2$ for the OpM problem. $C1$ implements a typical GRASP construction where each candidate element is initially evaluated by a greedy function to construct a Restricted Candidate List (RCL), and one element is selected at random from the RCL. Algorithm 1 shows the pseudo-code for $C1$. It initially creates a list of candidates (CL) which contains the elements that can be added to the partial solution under construction (Step 2). In order to favor the diversity of the constructed solutions, the method randomly selects the first facility from CL (Step 3) and includes it in the partial solution (Step 4). The method thus iterates until it obtains a solution with p facilities (Steps 6–13). In each iteration, $C1$ calculates the minimum (g_{min}) and maximum (g_{max}) values of the greedy function $g(S, j)$ (Steps 7 and 8). After that, $C1$ constructs a restricted candidate list (RCL) with all the elements (facilities) whose greedy value exceeds a percentage α of the best greedy value (Step 9). Finally, in Step 10, the method selects at random one facility from the RCL and adds it to the solution, updating CL (Steps 11–12).

We now describe $C2$, based on the construction strategy introduced in Resende and Werneck (2004), in which randomization takes place before the greedy selection in each construction step. This strategy has been successfully applied in Campos, Martí, Sánchez-Oro, and Duarte (2014), Duarte, Martí, Resende, and Silva (2011) and Duarte, Sánchez-Oro, Resende, Glover, and Martí (2015). In $C2$, we first randomly choose candidates and then evaluate them

---

Algorithm 1: Greedy constructive algorithm ($C1$).

```
1:  S ← ∅
2:  CL ← J
3:  j₀ ← SelectRandom(CL)
4:  S ← S ∪ {j₀}
5:  CL ← CL \ {j₀}
6:  while |S| < p do
7:      g_min ← min g(S, j)
              j∈CL
8:      g_max ← max g(S, j)
              j∈CL
9:      RCL ← {j ∈ CL | g(S, j) ≥ g_min + α · (g_max − g_min)}
10:     j ← SelectRandom(RCL)
11:     S ← S ∪ {j}
12:     CL ← CL \ {j}
13: end while
14: return  S
```

Algorithm 2: Greedy constructive algorithm (*C2*).

```
1:  S ← ∅
2:  CL ← J
3:  j₀ ← SelectRandom(CL)
4:  S ← S ∪ {j₀}
5:  CL ← CL \ {j₀}
6:  while |S| < p do
7:      size ← max{⌊α · |CL|⌋, 1}
8:      RCL ← SelectRandom(CL, size)
9:      j* ← arg max g(S, j)
             j∈RCL
10:     S ← S ∪ {j*}
11:     CL ← CL \ {j*}
12: end while
13: return  S
```

Algorithm 3: Random selection local search (*LS1*).

```
1:  improve ← true
2:  while improve do
3:      improve ← false
4:      Jc ← RandomlySort(S)
5:      J'c ← RandomlySort(J \ S)
6:      while (|Jc| ≠ ∅) & (improve ≠ true) do
7:          j ← next(Jc)
8:          while (|J'c| ≠ ∅) & (improve ≠ true) do
9:              j' ← next(J'c)
10:             S' ← exchange(S, j, j')
11:             if f(S') > f(S) then
12:                 S ← S'
13:                 improve ← true
14:             else
15:                 Jc ← Jc \ {j}
16:                 J'c ← J'c \ {j'}
17:             end if
18:         end while
19:     end while
20: end while
21: return  S
```

---

to make the greedy choice. *C2* first constructs a restricted candidate list RCL with a fraction (with $0 \leq \alpha \leq 1$) of the elements in the CL selected at random. Then, it evaluates all of them, computing the greedy function for all elements in the RCL, and selects the best one. Algorithm 2 shows the pseudo-code for *C2*. As it can be seen, the initial steps are the very same than those in Algorithm 1. The main difference between both methods resides on the instructions in the main loop (Steps 6–12 in Algorithm 2). In particular, it randomly selects a number of facilities from the CL. This number is determined by α, the percentage of the current size of the CL. Notice that we ensure that the value of *size* ranges from 1 to |CL| (see Step 7). The restricted candidate list is built with *size* elements of the CL, selected at random (Step 8). Then, the method selects the element j^*, which presents the best value according to the greedy function g (Step 9). Similarly to *C1*, the constructive procedure adds the selected element to S (Step 10) and removes it from CL (Step 11).

The α parameter controls the greediness/randomness of the GRASP constructive procedures. Specifically, if $\alpha = 0$, the corresponding method would be totally random. On the other hand, if $\alpha = 1$, the constructive procedure would be a pure greedy method. In Section 6 we investigate the influence of this parameter on the performance of *C1* and *C2*, selecting the value that obtains the best result.

4. Local search

The second stage of a GRASP algorithm consists in improving the constructed solution by applying a local search method to obtain a local optimum. Given a solution S, a neighbor solution S' can be obtained by applying a move, which can be viewed as a perturbation in the solution S. The set of all neighbor solutions of S, called neighborhood, is denoted as $N(S)$. In particular, we define a move operator that interchanges a facility $j \in S$ with a facility $j' \in J \setminus S$, producing a new feasible solution $S' = S \setminus \{j\} \cup \{j'\}$. For the sake of simplicity, we denote this move as $S' \leftarrow exchange(S, j, j')$. Therefore, the neighborhood of a solution can be defined as follows:

$$N(S) = \{S' \leftarrow exchange(S, j, j') : j \in S, j' \in J \setminus S\}$$

The neighborhood can be explored with two different strategies: best improvement, in which all the solutions are examined (i.e., the associated neighborhood is completely explored), to identify the best; and first improvement, that tries to avoid the time complexity of exploring the whole neighborhood by performing the first improving move encountered during the exploration of the corresponding neighborhood. Therefore, the order in which the neighbors are inspected has a significant influence on the search.

Notice that the order of exploration in the best improvement strategy is irrelevant since the corresponding neighborhood is fully explored.

We propose two different local search procedures based on the first improvement strategy. The difference between them resides in the neighborhood scanning strategy. In particular, *LS1* explores the neighborhood at random. Algorithm 3 shows the pseudo-code of this procedure. It receives as input argument a feasible solution S, which is improved (Steps 2–20). The local search uses working copies of the set of facilities J, which are given as J_c (opened sites, Step 4) and J'_c (unopened sites, Step 5), which are both randomly sorted. The neighborhood is, therefore, explored at random by selecting an opened facility (Step 7) and an unopened facility (Step 9). Specifically, in Step 10 a neighbor solution is visited. Then, *LS1* tests whether S' improves upon S (Step 11) or not (Step 14). If so, the exploration of the current neighborhood is abandoned, updating the incumbent solution (Steps 12 and 13); otherwise, the copied sets are updated (Steps 15 and 16). The local search ends when no further improvement is found.

The second local search, *LS2*, sorts the facilities in the solution in ascending order of its contribution to the solution. Here, the contribution of each facility is the sum of the δ-values for each client in the solution (see Section 3). For the sake of simplicity we do not include the pseudo-code of this method since is similar to the one presented in Algorithm 3. The only difference is that facilities j and j' are selected according to the contribution to a solution (instead of selecting them at random). We will study the performance of both methods in the computational experience.

One of the key elements in designing an effective local search method is the definition of the move and the associated move value (change in the objective function value). A move produces a change in the objective function value, called *move_value*. More precisely, let S be a solution whose associated cost is $f(S)$, let $j_1 \in S$ be an opened facility, and let $j_2 \in J \setminus S$ be an unopened facility. Then, $exchange(S, j_1, j_2)$ produces a new solution $S' = S \setminus \{j_1\} \cup \{j_2\}$ with cost $f(S')$. The change in the objective function value (due to the move) can be computed as:

$$move\_value = f(S') - f(S)$$

We say that $exchange(S, j_1, j_2)$ is an improving move if $move\_value > 0$, since the OpM is a maximization problem.

Table 5
List of distances for each client considering the example in Table 1 and solution $S = \{j_2, j_5, j_6\}$.

(a) Facilities in S		(b) Facilities in $J\backslash S$	
$\Delta_S(i_1)$	$\{(2, j_2), (4, j_5), (8, j_6)\}$	$\Delta_{J\backslash S}(i_1)$	$\{(3, j_1), (10, j_3), (14, j_4)\}$
$\Delta_S(i_2)$	$\{(12, j_6), (14, j_2), (15, j_5)\}$	$\Delta_{J\backslash S}(i_2)$	$\{(3, j_4), (11, j_3), (14, j_1)\}$
$\Delta_S(i_3)$	$\{(4, j_2), (6, j_6), (14, j_5)\}$	$\Delta_{J\backslash S}(i_3)$	$\{(5, j_1), (6, j_3), (12, j_4)\}$
$\Delta_S(i_4)$	$\{(1, j_2), (9, j_6), (10, j_5)\}$	$\Delta_{J\backslash S}(i_4)$	$\{(1, j_3), (3, j_1), (3, j_4)\}$
$\Delta_S(i_5)$	$\{(4, j_6), (9, j_2), (9, j_5)\}$	$\Delta_{J\backslash S}(i_5)$	$\{(2, j_3), (10, j_4), (13, j_1)\}$
$\Delta_S(i_6)$	$\{(4, j_5), (5, j_2), (11, j_6)\}$	$\Delta_{J\backslash S}(i_6)$	$\{(9, j_1), (11, j_3), (14j_4)\}$
$\Delta_S(i_7)$	$\{(2, j_6), (3, j_2), (4, j_5)\}$	$\Delta_{J\backslash S}(i_7)$	$\{(4, j_1), (4, j_3), (13, j_4)\}$
$\Delta_S(i_8)$	$\{(2, j_2), (13, j_6), (15, j_5)\}$	$\Delta_{J\backslash S}(i_8)$	$\{(3, j_4), (9, j_3), (12, j_1)\}$
$\Delta_S(i_9)$	$\{(4, j_6), (7, j_5), (11, j_2)\}$	$\Delta_{J\backslash S}(i_9)$	$\{(1, j_3), (2, j_4), (9, j_1)\}$

The local search performs a sequence of moves to reach the final solution (local optimum). To perform a single move, it evaluates many candidates in the neighborhood. Therefore, the move evaluation is a critical part when designing a GRASP, since it requires a considerable computational effort. In a straightforward implementation, once the improvement method executes a move, it has to be evaluated, computing its corresponding *move_value*. Normally, this computation requires to scan all the clients and all the facilities (i.e., the computational complexity of such a straightforward move evaluation is $\Theta(|I||J|)$). Therefore, the larger the number of clients or facilities, the longer the execution time. However, we can considerably improve and reduce this complexity by maintaining a list for each client, which stores the currently best link to the closest facility. We conveniently separate these distances in two independent data structures. The first one contains the distances from all clients to the opened facilities. In particular, given a solution S, we maintain for each client i the list $\Delta_S(i)$ with those distances from i to each facility $j \in S$. This list is sorted is ascending order. For example, given the matrix of distances shown in Table 1 and the solution $S = \{j_2, j_5, j_6\}$, the corresponding data structure is depicted in Table 5a, where each row contains the information of each client. For instance, the list of distances for client i_3 is $\Delta_S(i_3) = \{(4, j_2), (6, j_6), (14, j_5)\}$, which means that the closest facility is j_2 (where the distance from i_3 to j_2 is 4), the second closest facility is j_6 (distance 6), and the farthest facility is j_5 (distance 14).

We construct a similar structure with the distances from clients to unopened facilities (i.e., those in $J\backslash S$). Specifically, we store for each client i the ordered list $\Delta_{J\backslash S}(i)$ with the distances from i to each facility $j \in J\backslash S$. Table 5b shows the list of distances for each client considering again the example described above (i.e. $J \backslash S = \{j_1, j_3, j_4\}$).

The evaluation of a move is considerably improved by using these two data structures. In particular, if we execute the move $S' \leftarrow exchange(S, j, j')$, the value $f(S')$ can be obtained in $\mathcal{O}(|I|)$, since it is only required to access the first element of each list. In addition, the aforementioned data structures can be updated in $\mathcal{O}(|I| \log(|J|))$, since this operation requires to insert the elements j and j' in $2*|I|$ ordered lists (that can be performed in $\mathcal{O}(\log(|J|))$). Therefore, *move_value* can be computed in $\mathcal{O}(|I| \log(|J|)) + \mathcal{O}(|I|) = \mathcal{O}(|I| \log(|J|))$, which is considerably better than the straightforward implementation. In Section 6 we will study the influence of this strategy that we denote as incremental evaluation.

5. GRASP with filtering

GRASP repeatedly applies a construction followed by a local search. However, as mentioned in Section 4, the local search is time consuming, so we implement a filtering mechanism to discard low-quality constructed solutions and skip the local search. Filtering low-quality solutions was proposed in the early paper by Feo, Resende, and Smith (1994). We implement here the specific design proposed in Laguna and Martí (1999) and applied in Martí, Pantrigo, Duarte, Campos, and Glover (2011).

The percentage of improvement achieved by the application of the local search can be estimated as follows:

$$P(i) = \frac{f(S'_i) - f(S_i)}{f(S'_i)}$$

where i indicates the ith iteration, S_i is the solution generated by the constructive algorithm at iteration i, and S'_i is the solution obtained after applying the local search method to S_i. Considering that the OpM problem is a maximization problem the value of this percentage ranges from $P(i) = 0$ (i.e., the improvement method is not able to improve the constructed solution) to $P(i) < 1$.

The filtering strategy requires a "warming up" phase to compute the average improvement obtained by the local search, as well as the standard deviation of the improvement. This phase corresponds to the first k iterations of the GRASP algorithm, where k is an input parameter of the algorithm. After that phase, the mean μ_P and standard deviation σ_P of the percentage of improvement P can be estimated as:

$$\hat{\mu}_P = \frac{\sum_{i=1}^{k} P(i)}{k}$$

$$\hat{\sigma}_P = \sqrt{\frac{\sum_{i=1}^{k} (P(i) - \hat{\mu}_P)^2}{k-1}}$$

These two statistics can be used to determine whether it is "likely" that the constructed solution at iteration i (with $i > k$) can be improved enough to outperform the current best solution, $S^\star$. To do this, we calculate the minimum percentage improvement that a constructed solution S_i needs to be improved in order to be better than $S^\star$. This value, denoted as $imp(i)$, is obtained with the following expression:

$$imp(i) = \frac{f(S^\star) - f(S_i)}{f(S_i)}$$

If this value is close to $\hat{\mu}_P$, we can assume that the improvement method could improve the current best solution. On the other hand, if $imp(i)$ is considerably different than the estimation of the mean, it is unlikely that the local search method is able to obtain a solution better than the current one. In particular we only apply the local search method if and only if $imp(i)$ is smaller than the estimated mean plus q times the estimation of the standard deviation (to consider a conservative rule). In mathematical terms:

$$imp(i) < \hat{\mu}_P + q \cdot \hat{\sigma}_P$$

where q is a search parameter representing a threshold on the number of standard deviations away from the estimated mean percentage improvement. It is well-known that in a GRASP implementation the local search method consumes most of the running time. Therefore, we can considerably reduce it if we only improve promising constructed solutions. In Section 6, we will study the effect of different q values on solution quality and running time.

The pseudo-code of the GRASP approach with the filtering strategy is shown in Algorithm 4 . It performs N independent iterations as it is customary in GRASP. The algorithm starts each iteration by constructing a solution (Step 3) with any of the procedures described in Section 3 (C1 or C2). If the number of iterations i is lower than k, the constructed solution is improved (Step 5) with any of the local search procedures introduced in Section 4 (LS1 or LS2). Additionally, the percentage of improvement with respect to the best solution found is also updated (Steps 6 and 7). After k iterations (out of N), the method decides whether to improve a new

Algorithm 4: GRASP with filtering strategy.

1: $S^\star \leftarrow \emptyset$
2: **for** $i = 1$ **to** N **do**
3:    $S_i \leftarrow \texttt{Constructive}()$
4:    **if** $(i < k)$ **then**
5:       $S_i' \leftarrow \texttt{LocalSearch}(S_i)$
6:       $P(i) = \dfrac{f(S_i') - f(S_i)}{f(S_i)}$
7:       $S^\star \leftarrow \texttt{UpdateBest}(S_i')$
8:    **else**
9:       $\hat{\mu}_P = \dfrac{\sum\limits_{x=1}^{i} P(x)}{i}$
10:      $\hat{\sigma}_P = \sqrt{\dfrac{\sum\limits_{x=1}^{i} (P(x) - \hat{\mu}_P)^2}{i-1}}$
11:      $imp(i) = \dfrac{f(S^\star) - f(S_i)}{f(S_i)}$
12:      **if** $(imp(i) < (\hat{\mu}_P + q \cdot \hat{\sigma}_P))$ **then**
13:         $S_i' \leftarrow \texttt{LocalSearch}(S_i)$
14:         $P(i) = \dfrac{f(S_i') - f(S_i)}{f(S_i)}$
15:         $S^\star \leftarrow \texttt{UpdateBest}(S_i')$
16:      **end if**
17:    **end if**
18: **end for**
19: **return** $S^\star$

constructed solution or not. In particular, the average mean, standard deviation and the percentage of improvement are computed in Steps 9–11. Notice that these computations are performed in an incremental way. Additionally, these three values are updated during the entire execution of the procedure. This strategy has the effect of being more restrictive in the application of the local search procedure as soon as the number of iterations increases.

6. Experimental results

In this section we first describe the preliminary experiments that guided us in the selection of the values of the different key search parameters and strategies of both constructive and local search algorithms. Then, we show the results of our GRASP method compared with the state-of-the-art methods, i.e., branch & cut and tabu search procedures described in Belotti et al. (2007). All the experiments, including the algorithms reported in Belotti et al. (2007), were executed on the same computer, an Intel i5 660 processor running at 3.3 gigahertz with 8 gigabytes of RAM using GNU/Linux. In addition, all the results related to execution times will be displayed in seconds.

We have used a set of 72 instances in our experimentation. In particular, they were generated by considering 24 instances (from pmed17 to pmed40) of the well-known p-median problem,[1] where the number of nodes ranges from 400 to 900. In order to transform a p-median instance into an obnoxious p-median one, Belotti et al. (2007) described the following procedure. Given the original instance with n nodes, this method selects $n/2$ nodes at random to be the set of clients. The remaining $n/2$ become the set of facilities. Additionally, for each original instance, Belotti et al. (2007) derived three new instances for the OpM by considering three values of p: $\lfloor n/2 \rfloor$, $\lfloor n/4 \rfloor$ and $\lfloor n/8 \rfloor$. Table 6 reports the main characteristics of the new set of 72 instances, where n indicates the number of nodes, $|I|/|J|$ represents the number of clients/facilities, and p the number of required facilities.

Table 6
Instances generated from the OR-Library (Beasley, 1990).

Instance	n	$|I|/|J|$	p	Instance	Nodes	$|I|/|J|$	p
pmed17-p100	400	200	100	pmed29-p150	600	300	150
pmed17-p25	**400**	**200**	**25**	pmed29-p37	600	300	37
pmed17-p50	400	200	50	pmed29-p75	600	300	75
pmed18-p100	400	200	100	pmed30-p150	600	300	150
pmed18-p25	400	200	25	pmed30-p37	600	300	37
pmed18-p50	400	200	50	pmed30-p75	600	300	75
pmed19-p100	400	200	100	pmed31-p175	700	350	175
pmed19-p25	400	200	25	pmed31-p43	700	350	43
pmed19-p50	400	200	50	pmed31-p87	700	350	87
pmed20-p100	400	200	100	pmed32-p175	700	350	175
pmed20-p25	400	200	25	pmed32-p43	700	350	43
pmed20-p50	**400**	**200**	**50**	pmed32-p87	700	350	87
pmed21-p125	500	250	125	pmed33-p175	700	350	175
pmed21-p31	500	250	31	pmed33-p43	700	350	43
pmed21-p62	500	250	62	**pmed33-p87**	**700**	**350**	**87**
pmed22-p125	500	250	125	pmed34-p175	700	350	175
pmed22-p31	500	250	31	pmed34-p43	700	350	43
pmed22-p62	**500**	**250**	**62**	pmed34-p87	700	350	87
pmed23-p125	500	250	125	pmed35-p100	800	400	100
pmed23-p31	500	250	31	pmed35-p200	800	400	200
pmed23-p62	500	250	62	pmed35-p50	800	400	50
pmed24-p125	500	250	125	**pmed36-p100**	**800**	**400**	**100**
pmed24-p31	500	250	31	pmed36-p200	800	400	200
pmed24-p62	500	250	62	pmed36-p50	800	400	50
pmed25-p125	500	250	125	pmed37-p100	800	400	100
pmed25-p31	500	250	31	pmed37-p200	800	400	200
pmed25-p62	500	250	62	pmed37-p50	800	400	50
pmed26-p150	600	300	150	pmed38-p112	900	450	112
pmed26-p37	600	300	37	pmed38-p225	900	450	225
pmed26-p75	600	300	75	pmed38-p56	900	450	56
pmed27-p150	600	300	150	**pmed39-p112**	**900**	**450**	**112**
pmed27-p37	600	300	37	pmed39-p225	900	450	225
pmed27-p75	600	300	75	pmed39-p56	900	450	56
pmed28-p150	600	300	150	pmed40-p112	900	450	112
pmed28-p37	600	300	37	**pmed40-p225**	**900**	**450**	**225**
pmed28-p75	**600**	**300**	**75**	pmed40-p56	900	450	56

We have designed two constructive algorithms based on the same greedy function (see Section 3). These procedures, namely $C1$ and $C2$ are parameterized by α, which controls the trade-off between randomness and greediness. In the first experiment, we evaluate the influence of this parameter over the performance of $C1$ and $C2$ by considering four different values: 0.25, 0.5, 0.75, and random, where random indicates that the method randomly selects an α value in the range $[0, 1]$ for each construction. In order to avoid the over-fitting of our methods, we consider a representative subset of 10 percent of them (i.e., 8 instances) from the whole set of instances. In this way, we select, at least, one instance for each considered size, and an intermediate value of p on each case. We do not consider additional features of the instances since the coordinates of the nodes (facilities and clients) in the corresponding instances were generated by a random uniform distribution (Beasley, 1990). Notice that the remaining 90 percent instances (64 out of 72) are reserved for the final comparison with the state-of-the-art procedures. Table 6 shows with bold font the 8 instances that form the representative subset.

Table 7 shows the corresponding results when generating 100 independent constructions averaged over the subset of 8 instances. We report the average best cost (Avg. cost); deviation with respect to the best result in this experiment (Dev. (percent)); average computing time (Avg. time); the diversity of the solutions (Div.), computed as the sum of the different facilities found in the constructions generated on each run, divided by the number of constructions; and the number of best results (#Best), computed as the number of times the algorithm is able to match the best solution in this experiment. Attending to these results, we can assert that $C2$ outperforms $C1$ in terms of effectiveness, because it

Table 7
Comparison of different variants of the constructive algorithms. The best results are highlighted with bold font.

Constr. (α)	Avg. cost	Dev. (percent)	Avg. time	Div.	#Best
C1(0.25)	3851.63	7.9	7.56	**6164.82**	0
C1(0.5)	4063.75	8.17	13.33	6072.53	0
C1(0.75)	4745.13	7.77	10.91	5362.49	0
C1(random)	5837.38	30.38	7.02	5912.39	0
C2(0.25)	5892.38	2.98	6.46	1495.35	2
C2(0.5)	5896.75	2.65	8.86	1422.15	1
C2(0.75)	**5904.5**	**2.53**	11.53	1355.42	**3**
C2(random)	5902.38	3.89	**4.69**	1740.97	**3**

Table 8
Local search methods coupled with two constructive procedures. The best results are highlighted with bold font.

GRASP	Avg. cost	Dev. (percent)	Avg. time	#Best
C1(0.25) + LS1	5870.82	1.29	3581.15	**4**
C1(0.25) + LS2	5822.89	1.93	730.11	1
C2(0.75) + LS1	**5892.81**	0.91	2207.63	**4**
C2(0.75) + LS2	5888.66	**0.84**	**461.63**	1

obtains the best results regarding the average cost, deviation and number of best solutions found, emerging *C*2 with $\alpha = 0.75$ as the best variant. On the other hand, *C*1 presents better results in terms of diversity, being *C*1 with $\alpha = 0.25$ the best variant. Finally, *C*2 with random α is the fastest algorithm, although the average cost and deviation are worse than *C*2(0.75).

At this point, we thought that it could be interesting to select both, the most effective constructive algorithm, *C*2(0.75), and the algorithm that produces the most diverse solutions, *C*1(0.25). The first one feeds the local search with high-quality solutions, although they share most of their facilities. On the contrary, the second one provides diverse solutions to the local search to drive the search to different areas in the search space, but the quality of this solutions is considerably worse. It is well-known that the design of efficient metaheuristics mainly relies on a balance between search intensification and diversification. Therefore, we cannot anticipate which would result in better outcomes when coupled with the local search method.

In the next experiment we analyzed this combination by considering the two local search procedures defined in Section 4, namely *LS*1 and *LS*2 with the best two constructive procedures aforementioned (i.e., *C*1(0.25) and *C*2(0.75)). Table 8 shows the results of the four derived GRASP variants in terms of cost, deviation, time, and number of best solutions. We do not include in this experiment the diversity computation since it is not useful for further improvements.

As shown in Table 8, the GRASP variants that use *C*2(0.75) as constructive procedure obtain better results than the ones that consider *C*1(0.25). We can then conclude that in this problem it is more relevant to produce high quality solutions, although it requires to sacrifice the diversity (at least in the set of instances considered in this experiment). Additionally, *LS*1 obtains slightly better results both in average cost and number of best results, spending longer computing times than *LS*2. This fact can be partially explained because *LS*2 sorts the elements in the solution according to the δ-values (see Section 4), while *LS*1 scans the elements in a random way. Given that the local search is based on the first improvement strategy, *LS*2 reaches improved solutions faster than *LS*1.

In the next experiment, we study the computing time of the best GRASP method when (1) using a direct cost computation (DCC) and (2) using the incremental cost computation (ICC) presented in Section 4. Fig. 1 depicts a bar diagram where the *X*-axis represents the name of the 8 instances considered in the preliminary experiments and the *Y*-axis gives the computing time required to obtain a local optimum for both methods, respectively, (DCC and ICC) in the corresponding instance. The figure clearly shows that the saving in computing time is significant for ICC. Specifically, for these 8 instances we have calculated the geometric mean of the times required to obtain a local optimum in both methods. For DCC, the geometric mean value is 775.36 seconds, while the value for ICC is 139.72 seconds to obtain the same optimum value, which corresponds to a reduction of a 81 percent.

In the next preliminary experiment we test the influence of the filter strategy described in Section 5. This mechanism skips the improvement method when the value of the constructed solution does not reach a minimum quality. The parameter *q* controls the number of standard deviations away from the estimated mean percentage improvement. In order to determine the effectiveness of this strategy and to select the best value of *q*, we configure the GRASP algorithm with $k = 20$ (warming up phase) and
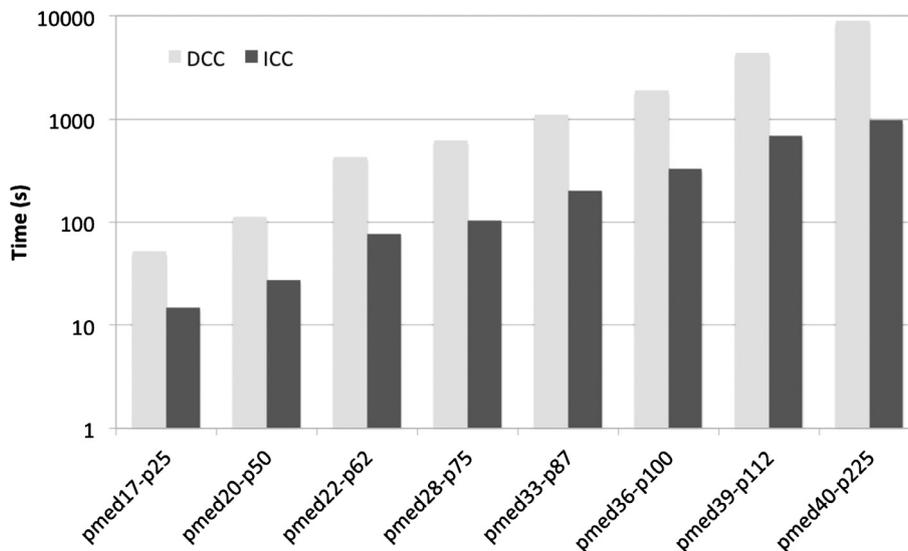

Fig. 1. Computing time comparison for direct (DCC) and incremental (ICC) cost computation, in logarithmic scale.

Table 9

Performance of the GRASP when filtering solutions with different values of q. Two GRASP versions without the filtering strategies are included as a baseline. The best results are highlighted with bold font.

q (#Constr.)	Avg. cost	Dev. (percent)	Avg. time (percent vs max)	#Best	Avg. #skip
0.5 (100)	5889.51	0.074	1.51 (0.76)	38	**53.17**
1.0 (100)	5889.51	0.077	1.41 (0.71)	38	41.04
1.5 (100)	5889.22	0.085	**1.38 (0.69)**	39	30.56
2.0 (100)	5890.75	**0.051**	1.70 (0.85)	47	21.01
2.5 (100)	5890.21	0.063	1.65 (0.83)	41	14.06
GRASP(20)	5884.53	0.176	1.72 (0.86)	27	–
GRASP(100)	**5890.81**	0.053	198.99 (100)	45	–

$N = 100$, applying the filter strategy from iteration 21 to 100 (see Algorithm 4).

Table 9 shows the results of this experiment, where five values of q, from 0.5 to 2.5, are studied. In order to analyze the improvement with respect to traditional GRASP implementations, we include two versions without the filter strategy (last two rows of the table). The number between parentheses indicates the number of iterations of each variant.

We additionally include the percentage time reduction with respect to the slowest algorithm (percent vs max), and the average number of skipped improvements (Avg. #skip). As expected, the lower the value of q, the larger the number of skipped improvements, and the larger the reduction in the computing time. It is important to remark that the quality of the obtained results (in terms of the average percentage deviation and the number of best solutions found) does not present significant changes across different values of q, thus indicating the robustness of the filter (i.e. solutions skipped for improvement hardly modify the final result).

The reduction in the execution time is also remarkable. In particular, the slowest algorithm, GRASP(100), employs 198.99 seconds on average, while the filter variants need less than 2 seconds to find similar results. The filter strategy described in Section 5 reduces the computing time in more than 99 percent. We select the variant with $q = 2$ since it presents the best results in terms of quality and with a really competitive computing time (1.70 seconds). Notice that a meaningful consequence of the filter approach is the fact that it allows us to run the GRASP algorithm for a larger number of iterations.

As described in Section 2, the best identified procedures in the related literature are the branch & cut (B&C) and the tabu search method (XTS) described in Belotti et al. (2007).[2] As it was aforementioned, the B&C is a finite algorithm that obtains the optimal solution of a given problem. However, when solving a large instance, it usually requires very long computational times, which can even make it unpractical. Considering that we are proposing heuristic procedures, we set the maximum computing time of this exact method to 3600 seconds. If after that time the B&C has not found the optimal solution, we interrupt its execution, returning the best solution found. The XTS is configured with the best parameters described in Belotti et al. (2007). We compare these methods with our best GRASP variant (constructive $C2$ with $\alpha = 0.75$, local search $LS1$, incremental computation of the cost, and filter strategy with $q = 2$) executed for two time horizons (1000 and 5000 iterations).

Table 10 presents the performance of each algorithm over the whole group of 72 instances. We report the same statistics previously used. These results show the merit of both GRASP approaches. In particular, the fastest version (1000 iterations) clearly outperforms B&C and XTS in all the considered metrics. In

Table 10

Comparison of the best GRASP method over the whole set of instances.

Algorithm	Avg. cost	Dev. (percent)	Avg. time	#Best
GRASP(1000)	5893.15	0.76	**237.34**	34
GRASP(5000)	**5894.51**	**0.73**	749.09	**47**
B&C	5764.35	3.27	3600.00	11
XTS	5847.08	1.57	511.35	25

particular, it obtains the best results in 34 instances (out of 72), while the competitors obtain 11 and 25, respectively. In the same line, GRASP(1000) presents a remarkable average percentage deviation of 0.76 percent, which compares favorably to the 3.27 percent and 1.57 percent, achieved by B&C and XTS. Notice that OpM instances used in this experimentation present a large value of the objective function (see the column Avg. cost in Table 10). Therefore, improving the value of the cost in hundreds barely affects to the value of the average deviation (this problem is even aggravated in large instances). Then, the reduction of 0.81 percent in this set of instances can be considered as a success. GRASP(1000) obtains these results in considerably shorter computing time. It is important to remark that all algorithms were executed in the same computer.

In order to show how the performance of the GRASP procedure is affected by the number of iterations, we include a version executed for 5000 iterations. It improves upon the results of the fastest version, obtaining the best solution in 47 instances (out of 72) with an average percentage deviation of 0.73 percent. In this case, GRASP(5000) is slower than XTS but faster than B&C. Notice that the inclusion of the B&C solution in this table has the objective of providing a baseline in our heuristic comparison, and it is not meant to compare the B&C itself with other heuristics.

We finally compare our best method, GRASP(1000), with the best previous heuristic, XTS, with two well-known non-parametric tests for pairwise comparisons: the Wilcoxon test and the Sign test. The Wilcoxon test answers the question: Do the two samples (solutions obtained with both methods in our case) represent two different populations? The resulting p-value of 0.017 clearly indicates that the values compared come from different methods (using a typical significance level $\alpha = 0.05$ as the threshold for reject or not the null hypothesis). On the other hand, the Sign test computes the number of instances on which an algorithm supersedes another one. The resulting p-value of 0.013 indicates that there are significant differences between both algorithms, confirming the superiority of our GRASP method.

7. Conclusions

In this work we have studied the application of a GRASP approach to solve the Obnoxious p-Median problem. In particular, we have developed two constructive algorithms based on a greedy function that calculates the contribution of a facility to the current solution. Our experimentation showed that one of them emphasizes search intensification, while the other one mainly diversifies the search into different regions.

We also proposed two different local search procedures to be executed after the constructive phase. Knowing that the local search changes only one facility in a given solution, an incremental move evaluation of the solutions was proposed, which geometrically averages an 81 percent reduction in the execution time in relation to a full evaluation of the solutions. In line with our objective of designing an efficient method, we implemented a filtering mechanism that is able to skip the application of the improvement phase in not promising constructed solutions. In this way,

---

[2] The authors kindly provided us with the source code of their algorithms.

this technique saves an average of 91.4 percent of the computation time for the GRASP with 100 constructions, obtaining similar results in terms of cost than the approach with no filter.

Finally, before engaging in competitive testing, we performed a series of scientific preliminary tests to determine the contribution of the various elements that we have designed. We believe that the reader can find them very useful since valuable lessons can be learned from them, and applied to other problems.

The experimental results show that our GRASP algorithm is able to outperform the current state-of-the-art methods in both short and long time horizons.

As future work, we will incorporate additional cost functions to enhance the quality of the solutions. The idea is to deal with a multi-objective optimization scenario for the OpM problem.

Appendix A. Best known solutions

In order to contribute to a further research in this problem, we detail in this appendix the results of our experiments breaking down them to the instance level. Table 11 shows the best cost found for each one of the 72 instances we have dealt with. In addition, it is also indicated the algorithm/s that obtained that solution, as well as the time spent to reach the solution. In the cases where two or more algorithms found the same best solution, the reported time corresponds to the faster one, which will be shown in bold font.

Table 11
Best cost obtained by each one of the instances, algorithms that obtained the value, and time spent by the faster algorithm (in bold).

Instance	Best cost	Method	Best time (seconds)
pmed17-p100	4054	GRASP(1000); **GRASP(5000)**	94.37
pmed17-p25	7317	**GRASP(1000)**; GRASP(5000)	37.29
pmed17-p50	5563	B&C; **XTS**	92.67
pmed18-p100	4215	**GRASP(1000)**; GRASP(5000)	48.47
pmed18-p25	7662	B&C; **XTS**	42.64
pmed18-p50	5852	B&C; **XTS**	81.74)
pmed19-p100	4033	**GRASP(1000)**; GRASP(5000)	72.85
pmed19-p25	7020	**GRASP(1000)**; GRASP(5000)	38.77
pmed19-p50	5408	B&C; **XTS**	105.76
pmed20-p100	4063	GRASP(5000)	309.3
pmed20-p25	7648	**GRASP(1000)**; GRASP(5000)	39.2
pmed20-p50	5872	**GRASP(1000)**; GRASP(5000)	73.15
pmed21-p125	4153	**GRASP(1000)**; GRASP(5000)	317.81
pmed21-p31	7331	B&C; **XTS**	148.57
pmed21-p62	5865	B&C; **XTS**	149.89
pmed22-p125	4347	GRASP(5000)	118.51
pmed22-p31	7900	**GRASP(1000)**; GRASP(5000)	66.34
pmed22-p62	6201	**XTS**	220.97
pmed23-p125	4107	GRASP(5000)	421.86
pmed23-p31	7841	**GRASP(1000)**; GRASP(5000)	94.4
pmed23-p62	5785	**GRASP(1000)**; GRASP(5000)	157.41
pmed24-p125	4079	GRASP(5000)	160.48
pmed24-p31	7425	**GRASP(1000)**; GRASP(5000)	92.95
pmed24-p62	5705	B&C; **XTS**	142.98
pmed25-p125	4220	XTS	152.73
pmed25-p31	7552	**GRASP(1000)**; GRASP(5000)	72.62
pmed25-p62	5767	**GRASP(1000)**; GRASP(5000)	190.83
pmed26-p150	4341	GRASP(5000)	626.95
pmed26-p37	8112	**GRASP(1000)**; GRASP(5000)	118.36
pmed26-p75	5899	XTS	260.9
pmed27-p150	4140	XTS	230.84
pmed27-p37	7556	**GRASP(1000)**; GRASP(5000)	161.23
pmed27-p75	5836	B&C; **XTS**	287.88
pmed28-p150	4096	**GRASP(1000)**; GRASP(5000)	329.76
pmed28-p37	7388	XTS	275.08
pmed28-p75	5675	GRASP(5000)	382.31
pmed29-p150	4132	GRASP(5000)	86.85
pmed29-p37	7520	XTS	226.71
pmed29-p75	5880	**GRASP(1000)**; GRASP(5000)	241.91

(continued on next page)

Table 11 (*continued*)

Instance	Best cost	Method	Best time (seconds)
pmed30-p150	4383	GRASP(5000)	773.79
pmed30-p37	8048	B&C; **XTS**	337.7
pmed30-p75	6189	**GRASP(1000)**; GRASP(5000)	268.42
pmed31-p175	4136	GRASP(5000)	294.53
pmed31-p43	7424	**GRASP(1000)**; GRASP(5000)	214.56
pmed31-p87	5905	**GRASP(1000)**; GRASP(5000)	337.72
pmed32-p175	4233	GRASP(5000)	170.96
pmed32-p43	7899	B&C; **XTS**	693.09
pmed32-p87	5925	GRASP(5000)	549.05
pmed33-p175	4110	XTS	497.59
pmed33-p43	7598	**GRASP(1000)**; GRASP(5000)	233.38
pmed33-p87	5815	XTS	711.83
pmed34-p175	4283	**GRASP(1000)**; GRASP(5000)	69.48
pmed34-p43	7725	**GRASP(1000)**; GRASP(5000)	221.65
pmed34-p87	5849	**GRASP(1000)**; GRASP(5000)	310.09
pmed35-p100	5845	**GRASP(1000)**; GRASP(5000)	715.93
pmed35-p200	4085	XTS	795.62
pmed35-p50	7570	XTS	1220.7
pmed36-p100	6461	**GRASP(1000)**; GRASP(5000)	332.06
pmed36-p200	4312	**GRASP(1000)**; GRASP(5000)	565.87
pmed36-p50	8179	**GRASP(1000)**; GRASP(5000)	308.91
pmed37-p100	6203	**GRASP(1000)**; GRASP(5000)	123.87
pmed37-p200	4583	XTS	957.92
pmed37-p50	8375	B&C; **XTS**	1270.58
pmed38-p112	5912	**GRASP(1000)**; GRASP(5000)	816.44
pmed38-p225	4439	XTS	1181.82
pmed38-p56	7535	XTS	1334.35
pmed39-p112	6163	XTS	2157.19
pmed39-p225	4355	GRASP(5000)	844.55
pmed39-p56	7712	**GRASP(1000)**; GRASP(5000)	587.88
pmed40-p112	6271	**GRASP(1000)**; GRASP(5000)	485.01
pmed40-p225	4565	GRASP(5000)	578.68
pmed40-p56	8211	**GRASP(1000)**; GRASP(5000)	338.48

References

Balinski, M. (1965). Integer programming: Methods, uses, computation.. *Management Science, 12*, 253–313.

Batta, R., Lejeuneb, M., & Prasad, S. (2014). Public facility location using dispersion, population, and equity criteria. *European Journal of Operational Research, 234*(3), 819–829.

Beasley, J. (1993). Lagrangian heuristics for location problems. *European Journal of Operational Research, 65*, 383–399.

Beasley, J. E. (1990). OR-Library: Distributing test problems by electronic mail. *Journal of the Operational Research Society, 41*(11), 1069–1072.

Belotti, P., Labbé, M., Maffioli, F., & Ndiaye, M. (2007). A branch-and-cut method for the obnoxious *p*-median problem. *4OR, 5*(4), 299–314.

Berman, O., Drezner, Z., Wang, Q., & Wesolowsky, G. (2008). The route expropriation problem. *IIE Transactions, 40*, 468–477.

Berman, O., & Wang, Q. (2007). Locating a semi-obnoxious facility with expropriation. *Journal of Operational Research Society, 58*, 378–390.

Berman, O., & Wang, Q. (2008). Locating a semi-obnoxious facility with expropriation. *Computers & Operations Research, 35*(2), 392–403.

Bhattacharya, U. (2011). A multi-objective obnoxious facility location model on a plane. *American Journal of Operations Research, 1*, 29–45.

Burkard, R., Fathali, J., & Taghizadeh Kakhki, H. (2007). The *p*-maxian problem on a tree. *Operations Research Letters, 35*(3), 331–335.

Campos, V., Martí, R., Sánchez-Oro, J., & Duarte, A. (2014). GRASP with path relinking for the orienteering problem. *Journal of the Operational Research Society, 65*, 1800–1813.

Cappanera, P. (1999). *A Survey on obnoxious facility location problems. Technical report*. Dipartimento di Informatica, Università di Pisa.

Cappanera, P., Gallo, G., & Maffioli, F. (2004). Discrete facility location and routing of obnoxious activities. *Discrete Applied Mathematics, 133*, 3–28.

Church, R., & Garfinkel, R. (1978). Locating an obnoxious facility on a network. *Transportation Science, 12*(2), 107–118.

Cohen, J. (1973). *A survey on the warehouse location problem. Operations Research Center working paper, OR 022-73*. Massachusetts Institute of Technology (MIT).

Conceição Fonseca, M., & Captivo, M. (2007). *Analysis of some models for semiobnoxious facility location. CIO working paper 16*. Universidade de Lisboa, Faculdade de Ciências, Centro de Investigação Operacional.

Coutinho-Rodrigues, J., Tralhão, L., & Alçada-Almeida, L. (2012). A bi-objective modeling approach applied to an urban semi-desirable facilitylocation problem. *European Journal of Operational Research, 223*(1), 203–213.

Current, J., Daskin, M., & Shilling, D. (2004). Discrete network location models. In Z. Drezner, & H. W. Hamacher (Eds.), *Facility location: Applications and theory*. In *Springer series in operations research* (pp. pp.83–120). Springer Science & Business Media.

Drezner, Z., & Hamacher, H. (2004). *Facility location: Applications and theory*. Springer.

Drezner, Z., & Wesolowsky, G. (1980). A maximin location problem with maximum distance constraints. *AIIE Transactions, 12*(3), 249–252.

Drezner, Z., & Wesolowsky, G. (1983). The location of an obnoxious facility with rectangular distances. *Journal of Regional Science, 23*(2), 241–248.

Drezner, Z., & Wesolowsky, G. (1996). Obnoxious facility location in the interior of a planar network. *Journal of Regional Science, 35*(4), 675–688.

Duarte, A., Martí, R., Resende, M., & Silva, R. (2011). GRASP with path relinking heuristics for the antibandwidth problem. *Networks, 58*, 171–189.

Duarte, A., Sánchez-Oro, J., Resende, M., Glover, F., & Martí, R. (2015). Grasp with exterior path relinking for differential dispersion minimization. *Information Sciences, 296*, 46–60. doi:10.1016/j.ins.2014.10.010.

Eiselt, H., & Laporte, G. (1995). Objectives in location problems. . In Z. Drezner (Ed.), *Facility location. a survey of applications and methods.* (pp. 151–180). New York: Springer.

Erkut, E., & Neuman, S. (1989). Analytical models for locating undesirable facilities. *European Journal of Operational Research, 40*(3), 275–291.

Farahani, R., SteadieSeifi, M., & Asgari, N. (2010). Multiple criteria facility location problems: A survey. *Applied Mathematical Modelling, 34*, 1689–1709.

Feo, T., & Resende, M. (1989). A probabilistic heuristic for a computationally difficult set covering problem. *Operations Research Letters, 8*, 67–71.

Feo, T., Resende, M., & Smith, S. (1994). A greedy randomized adaptive search procedure for maximum independent set. *Operations Research, 42*, 860–878.

Garey, M., & Johnson, D. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. New York: W.H. Freeman and Co., New York.

Greistorfer, P., & Rego, C. (2006). A simple filter-and-fan approach to the facility location problem. *Computers & Operations Research, 33*(9), 2590–2601.

Hakimi, S. (1964). Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research, 12*(3), 450–459.

Hakimi, S. (1965). Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Operations Research, 13*(3), 462–475.

Heydari, R., & Melachrinoudis, E. (2012). Location of a semi-obnoxious facility with elliptic maximin and network minisum objectives. *European Journal of Operational Research, 223*, 452–460.

Klose, A., & Drexl, A. (2005). Facility location models for distribution system design. *European Journal of Operational Research, 162*(1), 4–29.

Körkel, M. (1999). *Effiziente Verfahren zur Lösung unkapazitierter Standort-Probleme*. Berlin: Verlag für Wissenschaft und Forschung.

Kratica, J., Tošić, D., Filipović, V., & Ljubic, I. (2001). Solving the simple plant location problems by genetic algorithm. *RAIRO Operations Research, 35*, 127–142.

Kuehn, A., & Hamburger, M. (1963). A heuristic program for locating warehouses. *Management Science, 9*, 643–666.

Laguna, M., & Martí, R. (1999). GRASP and path relinking for 2-layer straight line crossing minimization. *INFORMS Journal on Computing, 11*(1), 44–52.

Lozano, M., Duarte, A., Gortázar, F., & Martí, R. (2012). Variable neighborhood search with ejection chains for the antibandwidth problem. *Journal of Heuristics, 18*, 919–938.

Martí, R., Pantrigo, J., Duarte, A., Campos, V., & Glover, F. (2011). Scatter search and path relinking: A tutorial on the linear arrangement problem. *International Journal of Swarm Intelligence Research, 2*(2), 1–21.

Megiddo, N., & Supowit, K. (1984). On the complexity of some common geometric location problem. *SIAM Journal on Computing, 13*(1), 182–196.

Melachrinoudis, E. (1985). Determining an optimum location for an undesirable facility in a workroom environment. *Applied Mathematical Modelling, 9*, 365–369.

Melkote, S., & Daskin, M. S. (2001). An integrated model of facility location and transportation network design. *Transportation Research Part A, 35*, 515–538.

Michel, L., & Hentenryck, P. V. (2004). A simple tabu search for warehouse location. *European Journal of Operational Research, 157*(3), 576–591.

Mladenović, N., Brimberg, J., Hansen, P., & Moreno-Pérez, J. (2007). The *p*-median problem: A survey of metaheuristic approaches. *European Journal of Operational Research, 179*(3), 927–939.

Ohsawa, Y., & Tamura, K. (2003). Efficient location for a semi-obnoxious facility. *Annals of OR, 123*, 173–188.

Ortigosa, P., Hendrix, E., & Redondo, J. (2011). *On methods for generating semi-obnoxious locations heuristically. WASS working paper No. 2*. Wageningen School of Social Sciences.

Papadimitriou, C., & Yannakakis, M. (1991). Optimization, approximation and complexity classes. *Journal of Computer and System Sciences, 43*, 425–440.

Plastria, F., Gordillo, J., & Carrizosa, E. (2013). Locating a semi-obnoxious covering facility with repelling polygonal regions. *Discrete Applied Mathematics, 161*(16–17), 2604–2623.

Rakas, J., Teodorović, D., & Kim, T. (2004). Multi-objective modeling for determining location of undesirable facilities. *Transportation Research Part D, 9*, 125–138.

Reese, J. (2006). Solution methods for the p-median problem: An annotated bibliography. *Networks, 48*(3), 125–142.

Resende, M., & Werneck, R. (2004). A hybrid heuristic for the *p*-median problem. *Journal of Heuristics, 10*(1), 59–88.

Segal, M. (2003). Placing an abnoxious facility in geometric networks. *Nordic Journal of Computing, 10*(3), 224–237.

Tamir, A. (1991). Obnoxious facility location on graphs. *SIAM Journal on Discrete Mathematics, 2*(4), 550–567.

Tamir, A. (2006). Locating two obnoxious facilities using the weighted maximin criterion. *Operations Research Letters, 34*, 97–105.

Tansel, B., Francis, R., & Lowe, T. (1983). Location on networks: A survey. Part I: The *p*-center and *p*-median problems. *Management Science, 29*(4), 482–497.

Welch, S., & Salhi, S. (1997). The obnoxious P facility network location problem with facility interaction. *European Journal of Operational Research, 102*, 302–319.

Yapicioglu, H., Smith, A., & Dozier, G. (2007). Solving the semi-desirable facility location problem using bi-objective particle swarm. *European Journal of Operational Research, 177*, 733–749.