Metaheuristics for the Minimum Gap Graph Partitioning Problem

Maurizio Bruglieri

Dipartimento di Design, Politecnico di Milano

Roberto Cordone*

Dipartimento di Informatica, Università degli Studi di Milano

Abstract

The *Minimum Gap Graph Partitioning Problem* (*MGGPP*) consists in partitioning a vertex-weighted undirected graph into a given number of connected subgraphs with the minimum difference between the largest and the smallest weight in each subgraph. We propose a multi-level Tabu Search algorithm and an Adaptive Large Neighborhood Search algorithm to solve the *MGGPP* in reasonable time on instances with several hundreds of vertices. The quality of the heuristic solutions is assessed comparing them with the solutions of a polynomially solvable combinatorial relaxation.

Keywords: Graph partitioning, Tabu Search, Adaptive Large Neighborhood Search

1. Introduction

Let $G = (V, E)$ be an undirected graph, with $|V| = n$ vertices and $|E| = m$ edges, let $p \in \mathbb{N}$ with $1 < p < n$ and let w be a *weight vector* which associates a weight $w_v \in \mathbb{N}$ with each vertex $v \in V$. For each vertex subset $U \subseteq V$, we denote by $m_U = \min_{u \in U} w_u$ and $M_U = \max_{u \in U} w_u$ the minimum and maximum value of the weights in U, respectively. We denote by *gap* the difference $\gamma_U =$

*Corresponding author

Email addresses: maurizio.bruglieri@polimi.it (Maurizio Bruglieri), roberto.cordone@unimi.it (Roberto Cordone)

$M_U - m_U$, that corresponds to the maximum difference between the weights of the vertices in U.

The *Minimum Gap Graph Partitioning Problem* (*MGGPP*) consists in finding a partition of G into p vertex-disjoint connected subgraphs $G_r = (V_r, E_r)$, with $r = 1, \ldots, p$. Here, we consider the *nondegenerate* problem (*MGGPPnd*), in which all subgraphs must have at least two vertices, and we focus on the *min-sum* version that minimizes the sum of all gaps, i.e., $\sum_{r=1}^{p} \gamma_{V_r}$.

The *MGGPP* is motivated by applications in different fields. First, the standard approach to manage large Water Distribution Networks is to sectorize them into subnetworks called District Metered Areas (*DMA*s) [22]. This allows to localize leakages more accurately, by monitoring the input and the output discharges for each district. Moreover, it also achieves a better management of pressure through valves and turbines that produce energy, reducing the amount of existing losses and limiting the occurrence of new damages. The optimal design of *DMA*s takes into account, among other objectives, the minimization of the difference between the required heads within the *DMA*s. The purpose is to establish a unique target pressure value in each *DMA*, and consequently, to achieve an efficient pressure regulation also in networks with strong variations in ground elevation [12]. If the network is represented as a graph where the edges correspond to pipes, the vertices to their intersections and the weight of a vertex to the ground elevation of an intersection, the *MGGPP* models the search for an optimal sectorization.

A second possible application concerns the segmentation of images, that is the clustering of pixels into objects [28]. While the classical approach is based on similarity weights associated with each pair of adjacent pixels, in other more recent applications pixels are also characterized by scalar values, such as their color intensity or texture. The goal in these problems is to collect within each segmented object only pixels with a limited variability in their range of values. Correspondingly, the aim of the problem is to partition a graph that represents an image, modelling the pixels as vertices and the pairs of adjacent pixels with edges between them, into connected components that

represent meaningful objects, with a limited range of weights, that is values. This fits well into the *MGGPP* framework.

Finally, in the agriculture field, the levelling of farmlands is an important foundation of modern ground irrigation systems, as it improves the uniformity of irrigation and soil salt distribution, thus controlling weeds, saving water and increasing yield [32]. In general, it might be impractical or too expensive to flatten a sloping land as a whole. In this case, it is necessary to divide the land into parcels, and build a flat terrace on each parcel by suitable earthworks. Choosing parcels with a small ground elevation difference helps to reduce the amount of ground to be moved, and therefore the cost associated with the following earthworks. Height is not the only significant property in land partitioning: in [1], for example, the problem of partitioning an agricultural field into the minimum number of rectangular zones with an upper bound on the variance of a suitable soil property is proposed, in a multiobjective and stochastic framework. The *MGGPP* allows to model the division of a land into parcels with a limited difference in height, or any other land attribute: the vertices of the graph correspond to sampled locations in the land, the weights to their heights, while the edges link adjacent locations.

In this paper we propose a multi-level Tabu Search (mTS) and an Adaptive Large Neighborhood Search ($ALNS$) metaheuristic to solve the *MGGPP* in reasonable time also on instances of large size. The former approach is characterized by the clustering of vertices into macro-vertices that are handled as single objects during the search: many small macro-vertices on the lower levels, few large macro-vertices on the higher levels. This has been made for two reasons. First, moving a single vertex from a component to another one is often unfeasible, because the resulting components of the partition can be disconnected. Second, moving a single vertex may yield a solution with the same total gap, because the gap of a component changes only if the vertex moved has the minimum or the maximum weight. On the other hand, the $ALNS$ approach is based on the alternated application of removal and insertion heuristics to destroy and rebuild part of the solution so as to avoid the limitations of reassigning single vertices.

Finally, in order to evaluate the quality of the solutions found, we compute a combinatorial bound by relaxing the connectivity constraint on the subgraphs and applying the polynomial procedure presented in [6].

The paper is organized as follows: Section 2 surveys the related literature; Section 3 gives the problem statement and a Mixed Integer Linear Programming (*MILP*) formulation of the *MGGPP*; Section 4 describes the Tabu Search metaheuristic and its multi-level extension; Section 5 describes the *ALNS* algorithm; Section 6 shows the numerical results obtained on instances of different size and topology. Finally, we draw some conclusions in Section 7.

2. Literature

The *MGGPP* has been introduced in [4], where it is proved to be $\mathcal{NP}$-hard and a couple of its special cases are investigated. Other special cases have been characterized in [6], while a basic Tabu Search heuristic procedure (without the multi-level extension) is tested on small random instances in [5].

The *MGGPP* clearly falls within the large field of graph partitioning problems [2, 3], but it is distinguished by a rather uncommon objective function. Most graph partitioning problems, in fact, optimize the cost of the edges linking different subgraphs of the partition (see for example the classical *Minimum k-Cut Problem* [11, 17]), whereas the *MGGPP* does not consider any edge cost. The graph partitioning problems whose definition includes vertex weights usually take into account the *total* weight of the vertices in each subgraph, and require it to be as uniform as possible across different subgraphs. This is obtained either by minimizing the difference between the total weights of different subgraphs [7] or by imposing bounds on each of them [15, 16]. For example, in [19], the authors consider the problem of partitioning a vertex-weighted path into subpaths, such that the total weight of every subpath lies between two given values.

Some similarity with the *MGGPP* can be found in the sectorization problem more recently proposed in [30], where a graph must be divided into districts and

the total weight of each district should differ as little as possible from a given reference value. In [18], a graph must be partitioned into components each including a given special vertex (*center*), so that the sum of the assignment costs of the other vertices to the centers is minimized. This work also considers uniform partitioning problems, whose goal is to balance as much as possible the total cost or the total weight of the components. Finally, the *Minimum-Diameter Partitioning Problem* (*MDPP*) introduced by [14] partitions a set of objects into a given number of clusters, such that the largest dissimilarity between objects in the same cluster is minimized. If the objects have weights and the dissimilarity between two objects is defined as the difference of their weights, one obtains the *MGGPP* on a complete graph. The two problems, therefore, share a common case. This case can be solved in polynomial time, as discussed in [13], together with other clustering problems dealing with different objective functions.

3. Statement of the problem and formulation

In this section we introduce a *MILP* formulation of the *MGGPP* where a feasible solution is represented as a spanning tree on an auxiliary directed graph $\tilde{G} = (V \cup \{r\}, A)$, being r a super-root node and A a set of arcs including both arcs (i, j) and (j, i) for each edge $[i, j] \in E$, and an arc (i, r) for each vertex $i \in V$. Each subtree appended to r identifies a connected subgraph of the solution (i.e., a component of the partition of G), and the predecessor of r is the minimum weight node in each subtree. Therefore, the resulting model is a multi-commodity flow formulation based on three families of variables: binary flow variables y_{ijk}, for all $(i, j) \in A$ and $k \in V$, equal to 1 if the flow of commodity k passes along arc (i, j), 0 otherwise; binary arc variables x_{ij}, for all $(i, j) \in A$), equal to 1 if arc (i, j) belongs to the spanning tree, 0 otherwise; real variables γ_i, for all $i \in V$, equal to the gap of the subgraph linked to r through node i, 0

if i is not linked to r.

$$\min \sum_{i \in V} \gamma_i \tag{1}$$

$$\sum_{(i,j) \in \delta^+(i)} y_{ijk} = \sum_{(i,j) \in \delta^-(i)} y_{jik} \quad \forall i \in V, k \in V : i \neq k \tag{2}$$

$$\sum_{(i,j) \in \delta^+(i)} y_{ijk} = 1 \quad \forall i \in V, k \in V : i = k \tag{3}$$

$$\sum_{i \in V} y_{irk} = 1 \quad \forall k \in V \tag{4}$$

$$\sum_{(i,j) \in \delta^+(i)} x_{ij} = 1 \quad \forall i \in V \tag{5}$$

$$x_{ij} + x_{ji} \leq 1 \quad \forall (i,j) \in A \tag{6}$$

$$y_{ijk} \leq x_{ij} \quad (i,j) \in A, k \in V \tag{7}$$

$$y_{irk} = 0 \quad \forall i \in V, k \in V : w_k < w_i \tag{8}$$

$$w_k y_{irk} - w_i x_{ir} \leq \gamma_i \quad \forall i \in V, k \in V \tag{9}$$

$$\sum_{(j,i) \in \delta^-(i)} x_{ji} \geq x_{ir} \quad \forall i \in V \tag{10}$$

$$\sum_{i \in V} x_{ir} = p \tag{11}$$

$$x_{ij} \in \{0,1\} \quad \forall (i,j) \in A \tag{12}$$

$$y_{ijk} \in \{0,1\} \quad \forall (i,j) \in A, k \in V \tag{13}$$

where $\delta^+(i)$ and $\delta^-(i)$ denote the forward and the backward star of vertex i, respectively.

The objective function (1), to be minimized, is the sum of the gaps of each subgraph in the solution. Constraints (2) guarantee the flow conservation for each commodity k in each node $i \neq k$; (3) ensure that the total flow of each commodity k emitted by node k is 1; (4) impose that for each commodity k the super-root r can be reached from one and only one node $i \in V$. Constraints (5) guarantee that for each node $i \in V$ one and only one arc of its forward star is selected. Constraints (6) prevent the two-cycles. Consistency constraints (7) impose that no commodity k can flow through arc (i,j) if this is not in the

solution. Constraints (8) guarantee that the last node that sends commodity k directly to the super-root r has weight lower than w_k, and thus is the minimum weight node of the subgraph. Constraints (9) enforce variables γ_i to model the gap of the subgraph linked to r through node i and (10) ensure that no subgraph in the solution is a singleton. Finally, constraint (11) guarantees that the solution is made up of p subgraphs and constraints (12) and (13) specify the nature of the variables.

4. Tabu Search

This procedure consists of two phases: the first builds an initial solution through a *construction procedure*, while the second improves it by *Tabu Search* [9]. The construction procedure is inspired by Prim's algorithm for the minimum spanning tree problem [24] and it returns a spanning forest of p trees identifying a feasible solution for the *MGGPP*. The procedure is detailed in Algorithm 1. First, it builds a forest, extracting for p times an edge with the minimum weight difference, avoiding the edges adjacent to the previously selected ones. Then, adjacent edges are iteratively selected and appended to one of the p trees, so as to keep the minimum total gap and not to merge different trees, until the forest spans the whole graph.

Tabu Search is a well known metaheuristic approach introduced in [9] to enhance the performance of local search. Our implementation for the *MGGPP* exploits a very simple neighborhood, based on moving a single vertex from a subgraph of the current solution to another one. This neighborhood consists of at most $n(p-1)$ solutions. The feasibility requires three conditions: 1) the moved vertex must not be an *articulation vertex* for the original subgraph[1]; 2) the destination subgraph must be adjacent to the vertex; 3) the original subgraph must include at least three vertices. The feasible moves can be identified in $O(m)$ time for the whole neighborhood, thanks to the identification of the articulation

[1] An articulation vertex is a vertex whose removal disconnects the graph it belongs to.

Algorithm 1 Constructive procedure

Order the edges $[i, j] \in E$ by non decreasing values of the gaps $|w_i - w_j|$;
Let $X := \emptyset$;
while $|X| \leq p$ **do**
    Let $[i, j] := extract\_first(E \setminus X)$;
    if $[i, j]$ is not adjacent to any vertex in X **then**
        Let $X := X \cup \{[i, j]\}$;
    end if
end while
while (there are vertices uncovered by X) **do**
    for each edge e in the cut generated by the vertices covered by X **do**
        Compute the total gap obtained if e were added to X;
    end for
    Select the edge e^* that minimizes the total gap;
    Let $X := X \cup \{e^*\}$;
end while

vertices and of the subgraphs adjacent to each vertex.

The *landscape* [29] of the *MGGPP* is characterized by large *plateaus*, that is, subsets of feasible solutions with the same objective function value. This is challenging for any neighbourhood-based procedure, because it may hide the most promising search directions. Therefore, we consider four criteria that are optimized in a lexicographic way in order to discriminate the feasible solutions in the plateau and to drive the search in a less miopic direction. The first criterium is the original objective function f. The second and the third one allow a partial look-ahead, considering not only the impact of a single move, but also its combination with the next ones. Indeed, although a move may have no immediate effect on the objective function, the subsequent transfer of a second vertex from the same subgraph could allow a strong improvement. Consequently, the second criterium to evaluate the transfer of a vertex v from its current subgraph G_s to a different one is defined as the better of the following two values: (i) the gap of G_s after removing both v and the vertex with the maximum weight; (ii) the gap of G_s after removing both v and the vertex with the minimum weight. If v coincides with either the maximum weight vertex or the minimum one, the

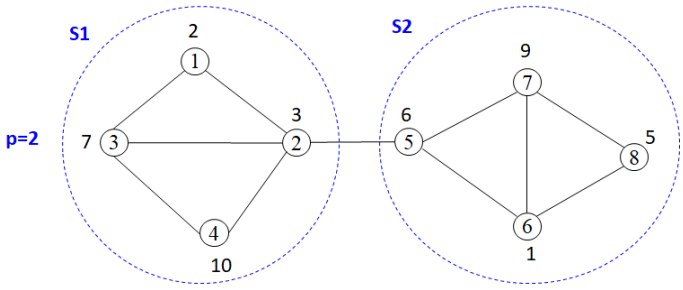

Figure 1: Example of the application of the lexicographic optimization to escape from plateaus

previous measures consider the second maximum or minimum weight vertex. The third criterium does the same, but moves two vertices in addition to v, corresponding to (i) the two maximum weights, (ii) the two minimum ones, or (iii) the maximum and the minimum one. Finally, the fourth criterium is the index of the most recent iteration in which the vertex v has been transferred in the past. This aims to favour the oldest moves upon the most recent ones. See an example in Figure 1: the two only feasible moves transfer, respectively, vertex 2 from subgraph S1 to S2 and vertex 5 from subgraph S2 to S1. Both moves yield solutions with a total gap equal to 16. Applying the second criterium, the gap of S1 after moving $v = 2$ and the maximum weight vertex (4) is 5, whereas its gap after moving $v = 2$ and the minimum weight vertex (1) is 3: the better of the two values is 3. On the other hand, the gap of S2 after moving $v = 5$ and the maximum weight vertex (7) is 4; its gap after moving $v = 5$ and the minimum weight vertex (6) is again 4. The lexicographic function therefore selects the first move. In this way, at the next iteration the best move consists in moving also vertex 1 to S2, thus leading to the optimal solution.

The Tabu Search mechanism adopted uses as an attribute the inclusion of a vertex v in a subgraph G_s for $s = 1, \ldots, p$. A suitable matrix L stores the most recent iteration l_{vs} in which vertex v has left G_s. A move that transfers vertex v into G_s is tabu until iteration $l_{vs} + l_{in}$, where l_{in} is a parametric value known as *tabu tenure*. At the beginning, all elements of matrix L are set to a very large

negative value ($l_{vs} = -\infty$), so that all moves are nontabu. In order to intensify or diversify the search, the tenure parameter changes in an adaptive way based on the value assumed by the objective in the previous iteration: l_{in} decreases by 1 if the objective has improved, increases by 1 otherwise, remaining inside a prescribed range $\{l_{\min}, ..., l_{\max}\}$. The rationale is to intensify the search in more promising regions of the solution space and to diversify it in less promising ones.

4.1. Multilevel Tabu Search

In spite of the lexicographic approach and the Tabu Search mechanism, plateaus remain a challenging feature of the problem. A well-known approach to deal with problems in which a graph must be partitioned into connected components is to apply an iterative coarsening procedure, where adjacent vertices are merged to create macro-vertices that can be moved as a whole [26]. The procedure can be applied at several levels, creating a sequence of graphs with different numbers $\tilde{n}$ of macro-vertices, starting from $\tilde{n} = n/2$ macro-vertices made of two vertices each, down to $p+1$ macro-vertices. At each level, the *MG-GPP* can be solved, dividing the aggregated graph into p subgraphs at minimum total gap: the resulting solution can be transformed back into a solution of the original problem, by simply replacing each macro-vertex with the corresponding subset of vertices. In fact, the components obtained are connected subgraphs of macro-vertices and each macro-vertex induces a connected subgraph on the original graph G. This solution is refined by a further application of the Tabu Search procedure.

In more detail, for each possible value of $\tilde{n}$, the *mTS* algorithm goes through three phases. The *building phase* applies the Prim-like constructive and the Tabu Search procedure to the original graph, with the number of subgraphs set to $\tilde{n}$, instead of p. The solution of this phase is converted into an aggregate graph, contracting each subgraph into a macro-vertex. In fact, the components of a good solution to the *MGGPP* tend to have small gaps, and this makes them promising candidates for the role of macro-vertices. The *improving phase* applies the constructive and the Tabu Search procedure to the aggregate graph,

and returns a partition into p components. The Tabu Search applied in this phase differs from the one described in the previous section for two aspects. First, it allows singleton components, because it works on macro-vertices that correspond by construction to subsets of at least two vertices. Second, each macro-vertex is associated with a range of weights $\left[w_v^{\min}, w_v^{\max}\right]$, instead of a single weight w_v. The whole procedure, however, can be trivially extended to handle both these aspects. The solution of the improving phase corresponds to a solution on the original graph. The *refinement phase* applies the Tabu Search procedure to this solution on the original graph, in a final attempt to further improve it.

5. Adaptive Large Neighborhood Search

The *ALNS* is a metaheuristic introduced by [25]. It has been successfully applied to several problems, as surveyed in [23], mainly in the field of vehicle routing. The applications to other contexts are less numerous: see [8] for the scheduling of technicians and tasks in a telecommunications company, [21] for the resource-constrained project scheduling problem and [20] for the lot-sizing problem. To the best of our knowledge, this is the first application of *ALNS* to a graph partitioning problem.

ALNS is an extension of Large Neighborhood Search (*LNS*). The basic idea of *LNS*, introduced by [27], is to search in large neighborhoods, which may contain more and potentially better solutions compared to small neighborhoods. The neighborhood of a solution is defined implicitly by a destroy method, i.e., a *removal heuristic*, and a repair method, i.e., an *insertion heuristic*. The former disrupts part of the current solution while the latter rebuilds the destroyed solution. In *ALNS*, several removal heuristics and insertion heuristic are employed. In each iteration, the removal/insertion heuristics to be applied to the current solution are selected according to their *scores*, which are adaptively adjusted based on the performance of these heuristics in the previous iterations: the heuristics that have performed better in the previous iterations have larger

scores and are more likely to be chosen in the current one.

Algorithm 2 reports the pseudocode of our implementation of *ALNS*. First, we initialize probability P_{RH} for each removal heuristic $RH \in \mathcal{RH}$ and probability P_{IH} for each insertion heuristic $IH \in \mathcal{IH}$, using a uniform probability distribution (lines 1–2). The initial solution X_0 is computed with the Prim-like constructive heuristic (line 3). The algorithm runs as long as the maximum number of iterations and the maximum execution time have not been reached (line 7). We randomly select a removal heuristic and an insertion heuristic according to the current values of the probabilities (lines 8 and 9) and we apply them to the current solution X_{cur}, obtaining the new solution X_{new} (line 10). To accept a solution, we use the same criterion defined in the Simulated Annealing algorithm. If the total gap of the new solution, $\gamma(X_{new})$, is smaller than that of the current one, $\gamma(X_{cur})$, then the new solution is accepted (lines 11–12). In the opposite case (lines 13–19), the new solution is accepted with a probability equal to $e^{(-(\gamma(X_{new})-\gamma(X_{cur}))/\gamma(X_0)\tau)}$, where τ denotes the current value of a *temperature* parameter that starts at a suitable value τ_0 (line 5) and decreases at each iteration through a *cooling rate* $h \in [0, 1]$ (line 23).

5.1. Removal heuristics

In our implementation of the *ALNS* for the *MGGPP*, we consider three families of removal heuristics:

1. *Random removal*: heuristic RH1 randomly eliminates a fraction q of the n vertices from the current solution. In order to ensure that the insertion procedure will be able to rebuild a feasible solution, we compute the connected components for each cluster from which vertices have been removed and keep only the component of maximum cardinality (and minimum gap, in case of ties). If all components of a cluster are singletons, we remove all the vertices. In general, therefore, the number of removed vertices is $\geq q\,n$.

Algorithm 2 Pseudocode for the *ALNS* procedure

Set $P_{RH} := 1/|\mathcal{RH}|$ for each removal heuristic $RH \in \mathcal{RH}$;
Set $P_{IH} := 1/|\mathcal{IH}|$ for each insertion heuristic $IH \in \mathcal{IH}$;
Generate an initial solution X^0 by applying the Prim-like constructive heuristic;
Set $X_{best} := X_{cur} := X^0$;
Set $\tau := \tau^0$;
Set $N := 0$;
while the termination condition is not satisfied **do**
    Select a removal heuristic $RH \in \mathcal{RH}$ with probability P_{RH};
    Select an insertion heuristic $IH \in \mathcal{IH}$ with probability P_{IH};
    Let X_{new} be the solution obtained applying RH to X_{cur} and IH to the result;
    if $\gamma(X_{new}) < \gamma(X_{cur})$ **then**
        Set $X_{cur} := X_{new}$;
    else
        Set $\nu := e^{(-(\gamma(X_{new})-\gamma(X_{cur}))/\gamma(X^0)\tau)}$;
        Generate a random number $\epsilon \in [0, 1]$;
        if $\epsilon < \nu$ **then**
            Set $X_{cur} := X_{new}$;
        end if
    end if
    if $\gamma(X_{cur}) < \gamma(X_{best})$ **then**
        Set $X_{best} := X_{cur}$;
    end if
    Set $\tau := h\tau$;
    Update P_{RH} and P_{IH} using $AWAP$;
    Let $N := N + 1$;
end while

2. *Worst cluster removal*: heuristic RH2 removes all the vertices of the cluster with the maximum gap. We also consider heuristic RH3, which randomly selects either the cluster with the second maximum gap or that with the third maximum gap, and removes all its vertices.

3. *Shaw removal*: this heuristic, introduced by [27], removes vertices that are somewhat similar, on the basis of a specific relatedness indicator, since we expect it to be easier to create better solutions by reassigning similar vertices. Given two vertices i and j, we define their *relatedness* $R(i,j)$ as a convex combination of their weight difference, $|w_i - w_j|$, and of the number of edges, l_{ij}, of the minimum cardinality path between them:

$$R(i,j) = \omega \frac{|w_i - w_j|}{\Gamma} + (1-\omega)\frac{l_{ij}}{L}$$

where $\omega \in [0,1]$ tunes the relative weight of the two components, $\Gamma = \max_{i,j \in V} |w_i - w_j|$, and $L = \max_{i,j \in V} l_{ij}$. We obtain three heuristics of this kind that we call RH4, RH5 and RH6 considering the following values of weight, respectively: $\omega = 0$, $\omega = 0.5$, and $\omega = 1$.

Algorithm 3 Pseudocode for the Shaw removal heuristic

```
Let v̄ be a vertex randomly selected from V;
Set S := {v̄};
while |S| < q n do
    Let v be a vertex randomly selected from S;
    Let α be an array containing all vertices of V \ S and n_α = |V \ S|;
    Sort α in such a way that i < j ⇒ R(v, α[i]) ≤ R(v, α[j]);
    Choose a random number y from the interval [0,1);
    Set S := S ∪ {α[⌊y^ρ n_α⌋]};
end while
Remove the vertices of S from X;
```

We report the detailed pseudocode of the Shaw removal heuristic, based on that of [25], in Algorithm 3. Starting from a random vertex $\overline{v}$, the procedure removes a total number of $q\,n$ vertices from the current solution X. They are selected choosing one of the previously removed vertices at random, v, sorting

the vertices that still belong to the solution by increasing relatedness to v and selecting one of them with a biased random extraction. The real number $\rho \geq 1$ tunes the randomness in the selection. In fact, $\rho = 1$ corresponds to a completely random removal algorithm, whereas very high values of ρ let y^ρ tend to 0, so that line 8 selects the first vertex of array α, i.e., the most related to v.

5.2. Insertion heuristics

We consider four insertion heuristics, IH1-IH4, to reinsert the removed vertices back into the current solution X. The heuristic IH1 applies the Prim-like procedure: for each removed vertex that can be feasibly added to one of the current subgraphs of X (that is, the adjacent ones), it computes the resulting total gap; then, it adds the best one until all the removed vertices are inserted. Heuristic IH2 is a randomized version of the same procedure: instead of choosing the best adjacent vertex, it randomly selects one of the k best vertices, where k is set to one third of the total number of adjacent vertices. Heuristic IH3 does the same, with a candidate list including two thirds of the adjacent removed vertices. Finally, IH4 adds an adjacent vertex at random.

5.3. Adaptive Weight Adjustment Procedure

The Adaptive Weight Adjustment Procedure ($AWAP$), used at line 24 of Algorithm 2, allows to adjust automatically the probabilities of using the removal and insertion heuristics on the basis of their success in the previous iterations. Its pseudocode is shown in Algorithm 4. A *score* s_H is associated with each heuristic $H \in RH \cup IH$, in order to measure how well the heuristic has performed recently. A counter θ_H tracks the number of applications of each heuristic. The whole search is divided into a number of *segments*, that is sequences of NS iterations of the $ALNS$. At the beginning of each segment, when the remainder of the division of iteration index N by NS is equal to zero (line 1), all scores and counters are initialized to zero. At each iteration, the counters of the removal and of the insertion heuristic applied are increased by 1 and the corresponding scores are increased by one of three possible values,

depending on the quality of the solution achieved. If a new overall best solution is found (line 4), the two heuristics are rewarded with an increase of σ_1. If the solution found has not been visited before and it is accepted by the *ALNS* (line 7), then the heuristics are rewarded with an increase of σ_2 (line 7) or σ_3 (line 10), according to the fact that the new solution is better or worse than the previous one. To check whether a solution has never been visited before, we can memorize all the new solutions accepted in a hash table. Notice that, as we shall see in Section 6, the experimental results suggest to accept only improving solutions, so that the case of line 10 will never occur and the hash table is not actually required.

At the end of each segment, when $mod(N, NS) = NS - 1$, the probability of each heuristic is updated with a convex combination of the previous probability and the average score s_H/θ_H achieved applying the heuristic (see line 18). The *reaction factor*, r, controls how quickly the weight-adjustment algorithm reacts to changes in the effectiveness of the heuristics: if $r = 0$, the scores are not used and thus the probabilities do not change; if $r = 1$, only the scores obtained in the last segment determine the new probabilities. We introduce a slight modification in the original formula of [25] in order to address the relative order of magnitude of probabilities and average scores: first, we normalize the average scores s_H/θ_H before combining them; second, we renormalize the probabilities after the update. In both cases, we proceed separately for the removal and for the insertion heuristics. This allows a much easier tuning of the scores.

6. Experimental campaign

Our experimental campaign was performed coding all algorithms in C, compiling them with `gcc 7.3.1` and running them on an Intel Xeon E5-2620 2.1GHz server with 16GB of RAM.

We built instances with five different sizes (n ranging from 100 to 500 by steps of 100 vertices) and three different structures with respect to the density and topology: random graphs with $m = 2n(n-1)/3$ edges and $m = n(n-1)/3$ edges,

Algorithm 4 Pseudocode for the *AWAP*

if $mod(N, NS) = 0$ **then**
    Set $s_H := 0$, $\theta_H := 0$ for each $H \in \mathcal{RH} \cup \mathcal{IH}$;
end if
if applying RH and IH yields a new best known solution **then**
    Set $s_{RH} := s_{RH} + \sigma_1$;
    Set $s_{IH} := s_{IH} + \sigma_1$;
else if applying RH and IH yields a solution that has never been accepted and is better than the current one **then**
    Set $s_{RH} := s_{RH} + \sigma_2$;
    Set $s_{IH} := s_{IH} + \sigma_2$;
else if applying RH and IH yields a solution that has never been accepted, is not better than the current one, but has been accepted **then**
    Set $s_{RH} := s_{RH} + \sigma_3$;
    Set $s_{IH} := s_{IH} + \sigma_3$;
end if
Set $\theta_{RH} := \theta_{RH} + 1$;
Set $\theta_{IH} := \theta_{IH} + 1$;
if $mod(N, NS) = NS - 1$ **then** ▷ Update the probabilities at the end of each segment
    for each $H \in \mathcal{RH} \cup \mathcal{IH}$ **do**
        Set $P_H := (1 - r)\, P_H + r\, Norm(\frac{s_H}{\theta_H})$;
    end for
    for each $H \in \mathcal{RH} \cup \mathcal{IH}$ **do**
        Set $P_H := Norm(P_H)$;
    end for
end if

and planar graphs obtained with a greedy triangulation of n points uniformly distributed at random in a square. In the following, we will denote these three classes as *dense*, *sparse* and *planar* instances. For each of them, three subclasses were obtained setting the number of subgraphs p equal to $\ln(n)$, $\sqrt{n}$ or $n/\ln(n)$ (rounded to the closest integer), in order to obtain solutions with few large subgraphs, balanced subgraphs, or many small subgraphs, respectively. The vertex weights are integer numbers uniformly distributed in $\{1, \ldots, n\}$. For each combination we generated 5 instances, thus obtaining a benchmark set of $5 \cdot 3 \cdot 3 \cdot 5 = 225$ instances of the *MGGPP*. This set is available at `https://homes.di.unimi.it/cordone/research/research.html`.

6.1. Indicators of the solution quality

The quality of a heuristic solution is usually expressed by the percent gap between its value z and the optimum z^*. However, the *MGGPP* is strongly $\mathcal{NP}$-hard [4], and Formulation (1-13) allows to compute the optimum with a general-purpose solver only for very small instances ($n \leq 20$) [5]. We therefore have to replace z^* with estimates. A lower bound $z_{LB} \leq z^*$ is computed neglecting the connectivity constraint and then solving the problem on the resulting complete graph with the polynomial procedure described in [6]. An upper bound $z_{UB} \geq z^*$ is provided by the best solution found during the whole experimental campaign (i.e., including also the results found in the parameter tuning phases). Correspondingly, we measure the quality of solutions with two indicators: i) the gap $\delta_{UB}(z) = (z - z_{UB})/z_{UB}$ with respect to the upper bound z_{UB}; ii) the gap $\delta_{LB}(z) = (z - z_{LB})/z_{LB}$ with respect to the lower bound z_{LB}. The former underestimates the actual gap between z and z^*, while the latter overestimates it. Our experiments show that both estimates are very tight on the dense instances with few subgraphs[2]. Moving to sparser instances and larger values of

[2]Sometimes, the two bounds even coincide; notice that, however, in most of these cases the bounding procedure provides an unfeasible partition, though with the same total gap as the optimal solution.

	p		
δ	$\ln n$	$\sqrt{n}$	$n/\ln n$
Dense	0.12%	0.59%	5.57%
Sparse	0.29%	1.90%	28.85%
Planar	11.39%	40.03%	296.83%

Table 1: Relative difference $\delta = (z_{UB} - z_{LB})/z_{LB}$ between the best known value z_{UB} and the combinatorial lower bound z_{LB} for the nine classes of instances: the difference strongly increases as the graph becomes sparser and the number of subgraphs increases

p, the lower bound strongly degrades because the solution of the combinatorial relaxation mostly consists of disconnected subgraphs, while the upper bound becomes less robust as different parameter tunings often yield different results. Table 1 shows how the density of the graph and the cardinality of the partition influence the distance between the two bounds, measured by the relative difference $\delta = (z_{UB} - z_{LB})/z_{LB}$.

6.2. Tuning of the Tabu Search algorithm

Preliminary investigations showed that the effective tuning of the Tabu Search parameters strongly depends on the features of the single instance: the number of vertices n, the number of subgraphs p and (more weakly) the density of the graph. On the one hand, large instances with many subgraphs require large tenures to avoid cyclic behaviors; on the other hand, small instances with few subgraphs require small tenures to avoid forbidding all moves.

Tabu tenure. First, we considered different ranges of values for the tabu tenure of the basic Tabu Search procedure. Inspired by a remark by Glover and Laguna [10], we decided to adopt a tenure proportional to the square root of the number of attributes, that is $\sqrt{np}$. We therefore let the tabu tenure vary in the range $[l_{\min}, l_{\max}]$, setting $l_{\min} = \alpha_{\min}\sqrt{np}$ and $l_{\max} = \alpha_{\max}\sqrt{np}$, where the coefficients $\alpha_{\min}$ and $\alpha_{\max}$ assume all pairs of values in $\{0.5, 0.75, 1, 1.25, 1.5\}$ such that $\alpha_{\min} \leq \alpha_{\max}$. This allows to consider both fixed tenures (when

$\alpha_{\min} = \alpha_{\max}$) and variable tenures, with a more or less wide range of variation. The Tabu Search procedure was then run on the benchmark instances for $T = 10\,000$ iterations with no aggregation of vertices into macro-vertices.

Table 2 reports the results: rows and columns correspond to the possible values of $\alpha_{\min}$ and $\alpha_{\max}$, respectively. Each cell reports the average values of δ_{UB} and δ_{LB} obtained by the corresponding configuration. Both estimates of the gap exhibit a smooth profile, with a minimum value (marked in bold) surrounded by gradually worsening results as $\alpha_{\min}$ and $\alpha_{\max}$ move farther away. Unfortunately, the optimal tunings for δ_{UB} and δ_{LB} are different. Since this is mainly due to the planar instances, where the quality of the lower bound is particularly bad, we adopted the first configuration ($\alpha_{\min} = 0.75$ and $\alpha_{\max} = 1$) and compared its results with those of the other ones with Wilcoxon's signed rank test [31]. The cells shaded in grey indicate the configurations for which the difference with respect to the chosen one is not statistically significant (i. e., the P-value of the test is $\geq 5\%$). Overall, the quality of these results looks disappointing, motivating first a deeper analysis of the algorithm, then the search for alternative, more effective mechanisms.

Plateau moves. Our second batch of experiments focused on the use of the lexicographic objective function to discriminate between alternative moves. We compared the basic version of the Tabu Search algorithm, considering only the actual objective function, an intermediate version, which includes also the second and third look-ahead criteria, and the full-fledged version, which also breaks ties choosing the less recent move. All three versions were run for 10 000 iterations to determine whether the additional information exploited to choose among moves is actually useful. Table 3 reports the average values of the gap estimates δ_{UB} and δ_{LB}: the full-fledged version proves better than the intermediate one, which dominates the basic version. Wilcoxon's test suggests that, in fact, the basic version is significantly dominated by the intermediate one ($P = 3.69 \cdot 10^{-13}$) and this is dominated by the full-fledged one ($P = 8.95 \cdot 10^{-3}$). In the following we adopt the full-fledged version.

$\alpha_{\min}\backslash\alpha_{\max}$	0.50	0.75	1.00	1.25	1.50
0.50	34.05%	33.35%	33.58%	33.99%	34.17%
	137.48%	134.24%	134.93%	136.24%	136.61%
0.75	-	31.14%	**30.59%**	31.32%	31.51%
	-	125.45%	123.71%	125.76%	126.02%
1.00	-	-	30.73%	30.84%	31.83%
	-	-	**122.75%**	123.69%	126.78%
1.25	-	-	-	31.37%	32.55%
	-	-	-	124.45%	128.30%
1.50	-	-	-	-	33.33%
	-	-	-	-	131.87%

Table 2: Average gap with respect to upper and lower bounds of the results obtained after 10 000 iterations of the basic Tabu Search with a tenure varying in $\{\alpha_{\min}\sqrt{np}, \ldots, \alpha_{\max}\sqrt{np}\}$

	Basic	Intermediate	Full-fledged
δ_{UB}	41.08%	30.99%	30.59%
δ_{LB}	140.35%	123.83%	123.71%

Table 3: Average gap with respect to the upper and lower bounds of the results obtained after 10 000 iterations of the basic (actual objective), the intermediate (three-level objective) and the full-fledged Tabu Search (three-level objective and recency)

Computational time for different classes of instances. Figure 2 reports the average computational time required to perform 10 000 Tabu Search iterations on each group of 5 benchmark instances with the same size n, topology (planar, sparse and dense) and value of p ($\ln n$, $\sqrt{n}$ and $n/\ln n$). A least-square interpolating profile is drawn for each class of instances to estimate the dependence of the computational time on the size of the instance. Quite clearly, the density of the graph strongly influences the computational time (denser instances take more time), while the overall number of clusters has a weaker, and less consistent, effect (more clusters, but in some cases also less clusters, imply more time).

The interpolation suggests that for planar instances the computational time is slightly less than linear in n, especially for large numbers of clusters. In these instances, in fact, each vertex tends to have a constant number of adjacent vertices. Since some vertices cannot be moved, because they would disconnect a subgraph or they are surrounded by vertices of the same subgraph, the overall number of neighbor solutions increases slightly less than linearly with n. For the sparse and dense instances, on the contrary, each vertex has on average a linear number of adjacent vertices, part of them belonging to different subgraphs, and the computational time increases quadratically with the number of vertices.

Restart. Even with the best tuning of the tabu tenure, the basic Tabu Search procedure is clearly unable to obtain good solutions in short time. We have then evaluated the possible improvements obtained by increasing the computational time and restarting the search when it appears to be unpromising. Since the most challenging instances are the planar ones and the complexity of the neighbourhood exploration is nearly linear on these instances, we have set an overall running time of $n/100$ minutes (i.e., from one to five minutes according to the size), and we have imposed a restart from a random solution after different numbers of nonimproving iterations, specifically $T_{ni} \in \{10^2, 10^3, 10^4, 10^5, 10^6\}$, where the first value corresponds to a very frequent restart, with a search phase possibly shorter than the tabu tenure, and the last value to at most 10 restarts

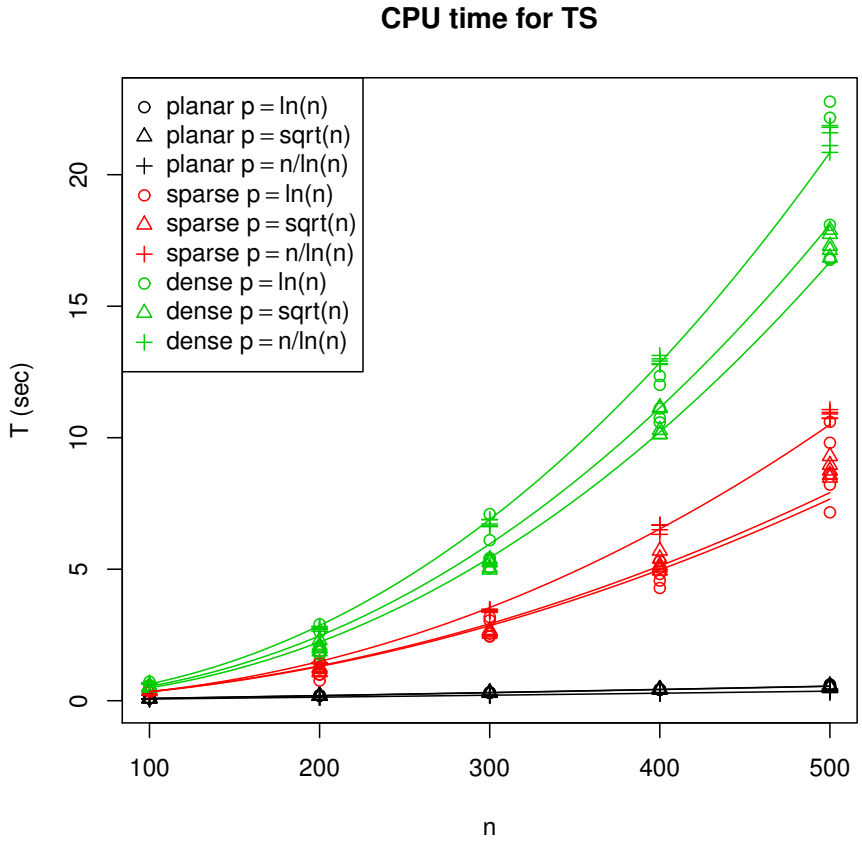

Figure 2: Average computational time (in seconds) required to perform 10 000 iterations of the Tabu Search algorithm on instances of the nine classes

T_{ni}	100	1 000	10 000	100 000	1000 000
δ_{UB}	19.63%	18.08%	**16.80%**	19.65%	22.28%
δ_{LB}	104.06%	98.53%	**93.28%**	96.69%	99.83%

Table 4: Average gap with respect to the best known upper and lower bounds of the results obtained by the Tabu Search, with a restart after different numbers T_{ni} of nonimproving iterations

during the whole computation. Table 4 reports the results, as usual providing the average gaps with respect to the upper and lower bound. As in the previous tables, the best result is bolded: it corresponds to a restart of the search after 10 000 nonimproving iterations. The other tunings are significantly worse according to Wilcoxon's test ($P = 0.035$ for $T_{ni} = 1,000$ and $P < 10^{-5}$ for the other values). Though better, the results are still not satisfactory.

Multilevel Tabu Search. The bad performance of the basic Tabu Search procedure mainly concerns the planar instances: the overall 16.80% average gap actually combines a 5.67% average gap on the sparse and dense instances and a 39.20% gap on the planar ones. This is the effect of two main limitations of the local search: first, the objective function, even if enhanced by the lexicographic objective, is unable to clearly discriminate between promising and unpromising moves; second, in planar graphs it is difficult to move vertices from cluster to cluster while keeping them connected. Random restart can move the search to different regions of the solution space, but does not allow to explore it effectively. The search can be strongly improved by moving suitably identified macro-vertices, as discussed in Section 4.1. We have therefore divided the the overall computational time of $n/100$ minutes into $n/2 - p$ equal intervals of length $\tau = n/(100(n/2-p))$, one for each possible value of the number of macro-vertices $\tilde{p}$ (from $p + 1$ to $n/2$). Then, we have divided each interval into three sequential phases. The building phase applies for $\alpha_b\tau$ minutes the constructive and Tabu Search procedures to the original graph in order to build $\tilde{p}$ clusters, that are converted into macro-vertices of an aggregated graph. Then, the ex-

change phase applies for $\alpha_e\tau$ minutes the greedy and Tabu Search procedures to partition the aggregated graph into p clusters. Finally, a refinement phase applies for $(1-\alpha_b-\alpha_e)\tau$ minutes the Tabu Search procedure to the original graph, so as to refine the p clusters obtained from the exchange phase. We have generated 15 configurations extracting the values of the two coefficients α_b and α_e from $\{0, 0.25, 0.5, 0.75, 1\}$ so that $\alpha_b + \alpha_e \leq 1$. When $\alpha_b = 0$ (respectively, $\alpha_e = 0$), the building (respectively, the exchange) phase applies only the greedy procedure; when $\alpha_b+\alpha_e = 1$, the refinement phase is skipped. The Tabu Search procedure is always applied with the best tabu tenure tuning identified above.

Table 5 reports the results of the *mTS*: rows and columns correspond to the possible values of α_b and α_e, respectively. Each cell reports the average values of δ_{UB} and δ_{LB} obtained by the corresponding configuration. The best configuration with respect to both indicators is $\alpha_b = 0.25$ and $\alpha_e = 0.50$, but the symmetric configuration with $\alpha_b = 0.50$ and $\alpha_e = 0.25$ is nearly as effective and the difference is not statistically significant (both cells, in fact, are shaded in grey in the table). Since a third configuration ($\alpha_b = 0.25$ and $\alpha_e = 0.25$) is only slightly worse, with a P-value of 2.98%, it is reasonable to conclude that all three phases of the algorithm (building good macro-vertices, collecting them into good subgraphs and refining the solution thus obtained) have an important impact on the final result, and that any "balanced" distribution of the computational time among them yields an acceptable behaviour.

The quality of the results is clearly much better than that of the basic Tabu Search, even with a restart mechanism (notice that the number of levels, $n/2-p$, is of the same order of magnitude of the best number of random restarts) We conclude that the aggregation into macro-vertices actually plays a role in building an effective search.

6.3. Tuning of the ALNS *algorithm*

The *ALNS* algorithm has a large number of parameters, as it combines several different insertion and removal procedures. In our case, the insertion procedures have no parameters: each of them is completely characterized by the

$\alpha_b \backslash \alpha_e$	0.00	0.25	0.50	0.75	1.00
0.00	9.44%	6.12%	6.05%	6.09%	11.19%
	72.63%	62.14%	61.87%	62.02%	78.44%
0.25	5.26%	3.39%	**3.37%**	4.16%	-
	57.06%	53.49%	**53.44%**	55.92%	-
0.50	5.11%	3.37%	4.13%	-	-
	56.43%	53.48%	55.91%	-	-
0.75	5.27%	4.21%	-	-	-
	57.00%	56.12%	-	-	-
1.00	9.17%	-	-	-	-
	66.24%	-	-	-	-

Table 5: Average gap with respect to upper and lower bounds of the results obtained in $n/100$ minutes by the multi-level Tabu Search with different distributions of the running time among the building, exchange and refinement phase

given cardinality of the list of candidate vertices. Our first round of experiments, therefore, was focused on the removal procedures.

Tuning of the removal parameters. The random removal heuristic RH1 is characterized by the fraction q of vertices removed. In order to tune it, we have provisionally disabled the other removal heuristics and assigned a fixed uniform probability to each of the insertion heuristics. For the sake of simplicity, we have also set to 0 the parameter τ_0 that controls solution acceptance, so that all nonimproving solutions are rejected. Consequently, the cooling parameter h becomes inactive and the results only depend on q. We have considered all values of q in $\{0.05, 0.10, 0.15, 0.20\}$. Notice that q strongly affects the time required by each iteration of the algorithm, because removing more vertices slows down both the removal and the insertion heuristic. In order to make a fair comparison, with stable results, but in a reasonable time, we set a time limit of $n/1\,000$ minutes (from six to thirty seconds), that is comparable to the time used to tune the *TS* algorithm. Table 6 reports the results of this experiment. It has

q	0.05	0.10	0.15	0.20
δ_{UB}	9.34%	**7.27%**	8.49%	8.28%
δ_{LB}	60.99%	**57.21%**	62.07%	62.78%

Table 6: Average gap with respect to the best known upper and lower bounds of the results obtained in $n/1\,000$ minutes by the *ALNS* removing different fractions of the vertices, uniformly chosen at random

the following structure: the first row provides the values tested for parameter q, the second one the percent gap with respect to the best known result and the third one the gap with respect to the lower bound. The best average gaps (reported in bold) are obtained by setting $q = 0.10$, that is the tuning adopted in all following experiments, but Wilcoxon's test suggests that the difference with respect to $q = 0.15$ is not significant (both results are shaded in grey).

The worst clusters removal heuristics RH2 and RH3 have no parameters to tune. The Shaw removal heuristics RH4, RH5, RH6 have two parameters: the fraction q of vertices chosen to be removed and the real parameter ρ that tunes the randomness of their choice. We have considered all pairs of values (q, ρ) with $q \in \{0.05, 0.10, 0.15, 0.20\}$ and $\rho \in \{5, 10, 20, 50, 100\}$. The experimental framework adopted is the same: $n/1\,000$ minutes of computation, $\tau_0 = 0$ and equal probabilities for the four insertion heuristics; also the three Shaw removal heuristics have been assigned the same probability. Table 7 reports the results: the rows correspond to the values of q and the columns to the values of ρ. The best results with respect to the upper bound are obtained setting $q = 0.15$ and $\rho = 50$, and these are the values that will be employed in all following experiments. However, there is a rather large region in the parameter space that does not differ in a statistically significant way from the chosen tuning.

Tuning of parameter τ_0. After determining a good value for the specific parameters of each removal heuristic, we have tuned the parameter τ_0 that controls the probability to accept a worsening solution. This parameter rules the balance between the intensification and diversification aspects of the search. We remind

q \ ρ	5	10	20	50	100
0.05	4.39%	4.75%	3.38%	4.11%	4.37%
	49.91%	53.99%	49.50%	52.35%	51.10%
0.10	4.44%	3.61%	2.72%	2.50%	3.22%
	53.47%	52.77%	49.54%	**48.65%**	51.37%
0.15	4.54%	2.91%	2.50%	**2.42%**	2.48%
	54.60%	51.10%	50.19%	49.41%	49.53%
0.20	4.69%	3.64%	3.04%	3.46%	2.89%
	57.32%	54.77%	52.65%	54.08%	51.56%

Table 7: Average gap with respect to upper and lower bounds of the results obtained in $n/1\,000$ minutes by the *ALNS* with the Shaw removal heuristics with different fractions of removed vertices (q), chosen more or less randomly (increasing ρ reduces randomness)

that the *ALNS* algorithm always accepts an improving solution, whereas it accepts a solution with a relative increase of the total gap equal to δ with a probability equal to $e^{-\delta/\tau_0}$. An intuitive way to describe this mechanism is to say that solutions with a relative gap increase equal to $\tau_0 \ln 2$ have the same probability to be accepted or rejected (50%), solutions with a smaller increase are accepted more often and solutions with a larger increase are more likely to be rejected. By $\tau_0 = 0$ we mean that all nonimproving solutions are rejected. In our experiments, we have considered all values of τ_0 in $\{0.00, 0.05/\ln 2, \ldots, 0.20/\ln 2\}$, so that the relative gap increase accepted with 50% probability ranges from 0 (only improving solutions) to 0.20 (solutions with a gap 1.2 times the current one). The algorithm has been run for $n/1\,000$ minutes, as in the previous experiment, assigning the same probability to each of the four insertion heuristics and of the six removal heuristics. Table 8 reports, as usual, the average values of δ_{UB} and δ_{LB} obtained by each configuration. The best one is $\tau_0 = 0$, that is, accepting only improving solutions; the quality of the best found result becomes gradually worse as the probability of accepting worse solutions increases.

τ_0	0.00	$\frac{0.05}{\ln 2}$	$\frac{0.10}{\ln 2}$	$\frac{0.15}{\ln 2}$	$\frac{0.20}{\ln 2}$
δ_{UB}	**1.50%**	3.36%	10.69%	18.78%	23.10%
δ_{LB}	**46.75%**	49.06%	79.51%	111.97%	128.22%

Table 8: Average gap with respect to the best known upper and lower bounds of the results obtained by the *ALNS*, with different values of the probability to accept worsening solutions

This is a not very common tuning for *ALNS*, but it has actually already been applied to other problems with objective functions characterized by *plateaus*, such as the Resource-constrained Project Scheduling Problem [21]. According to Wilcoxon's test, the difference with respect to all other tunings is statistically significant. In the following, therefore, we have set $\tau_0 = 0$, which also makes it no longer necessary to tune the cooling rate h.

Computational time for different classes of instances. Figure 3 reports the average computational time required to perform 1 000 *ALNS* iterations on each group of 5 benchmark instances with the same size n, topology (planar, sparse and dense) and value of p ($\ln n$, $\sqrt{n}$ and $n/\ln n$). A least-square interpolating profile is drawn for each class of instances to estimate the dependence of the computational time on the size of the instance. As for the Tabu Search algorithm, denser instances take a longer time, while the effect of the overall number of clusters is in this case unambiguous: more clusters clearly imply more time.

The interpolation suggests that the dependence of the computational time on the number of vertices is between linear and quadratic. This is consistent with the structure of the insertion mechanisms, that are Prim-like $O\left(n^2 + m\right)$ procedures, and of the removal mechanisms, that are $O\left(n\right)$ (worst and random-worst), $O\left(m\right)$ (random) or $O\left(n^2\right)$ (Shaw removal).

Parameter configuration for the ALNS. Finally, we introduce the adaptive mechanism *AWAP* described in Section 5.3. We increase the overall computational time to $n/100$ minutes, that is 10 times larger than in the previous experiment and equal to the time used in the final test on the Multilevel Tabu Search algo-

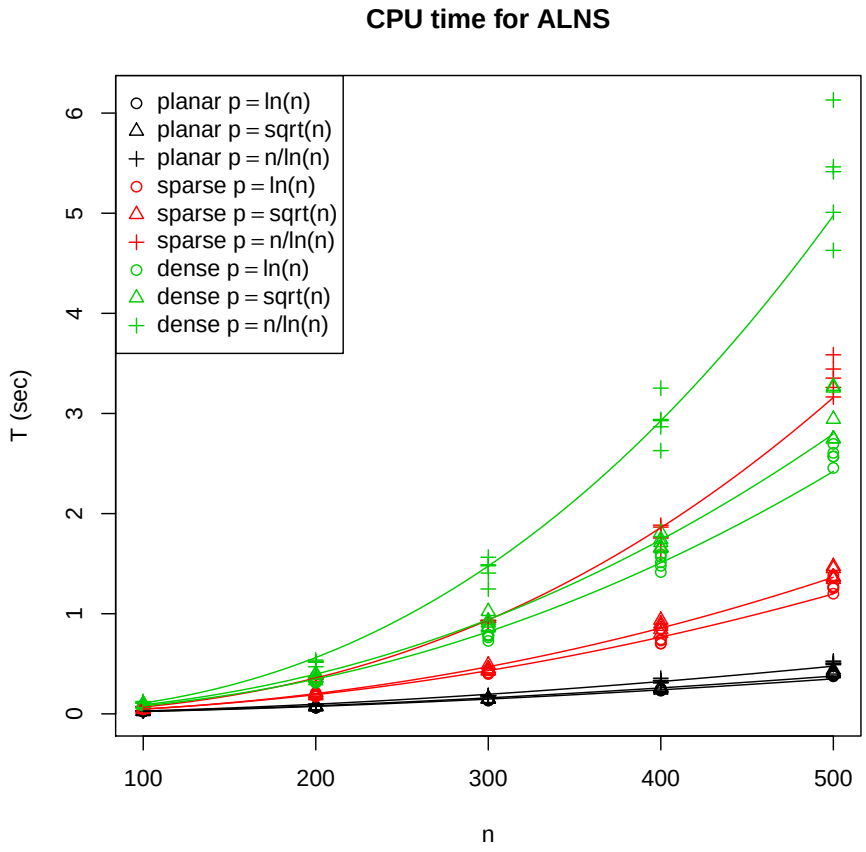

Figure 3: Average computational time (in seconds) required to perform 1 000 iterations of the *ALNS* algorithm (with equal probability for each combination of insertion and removal procedures) on instances of the nine classes

rithm. Since this time limit allows to perform millions of iterations on all the instances of the benchmark, we set the number of iterations for each segment to $NS = 1\,000$. This guarantees both the generation of a large number of segments and several applications of the $4 \cdot 6 = 24$ possible combinations of insertion and removal heuristics, in each segment. Setting $\tau_0 = 0$ avoids tuning the scores σ_2 and σ_3, because the corresponding cases never occur: only solutions that improve the best known result are accepted. As for σ_1, this score has been set to 1, so that the ratio s_H/θ_H used to update the probability of each heuristic H is equal to the frequency with which that heuristic has returned an improved solution in the last segment. The only remaining parameter is the reaction factor, that is set to $r = 0.1$ following [25].

Comparison between the mTS *and the* ALNS. Detailed results for the *mTS* and the *ALNS* on the single instances can be found at `https://homes.di.unimi.it/cordone/research/research.html`. Table 9 reports the average values of δ_{LB} and δ_{UB} on the whole benchmark. There are three blocks of rows, for planar, sparse and dense graphs, respectively. Each block consists of three rows, corresponding to different values of the number of subgraphs p ($\ln(n)$, $\sqrt{n}$ and $n/\ln(n)$). We emphasize in boldface the better between the values obtained by the two algorithms. In most cases, the *ALNS* outperforms the *mTS*: this occurs when the number of subgraphs is larger ($p = \sqrt{n}$ or $p = n/\ln(n)$). When this number is smaller, *mTS* is slightly better than *ALNS*. A possible explanation of this different performance could be that the removal and insertion heuristics of *ALNS* become less effective due to the low number of subgraphs, since most removed vertices are simply reinserted in their original position. The last row of the table provides the overall average values, according to which the performance of the *ALNS* is superior. This is confirmed by the number of instances in which it finds the overall best known result: 157 versus 94, out of 225. Since both algorithms find the best known result in 74 cases, there are 83 best known results found by the *ALNS* approach that are not detected by the *mTS*, while the opposite occurs only for 20 instances.

		mTS		*ALNS*	
	p	δ_{UB}	δ_{LB}	δ_{UB}	δ_{LB}
Planar	$\ln(n)$	**0.00%**	**11.39%**	**0.00%**	**11.39%**
	$\sqrt{n}$	1.49%	42.11%	**0.00%**	**40.03%**
	$n/\ln(n)$	19.75%	380.01%	**4.11%**	**315.11%**
Sparse	$\ln(n)$	**0.04%**	**0.33%**	0.52%	0.82%
	$\sqrt{n}$	1.29%	3.22%	**0.05%**	**1.95%**
	$n/\ln(n)$	3.63%	33.39%	**1.51%**	**30.73%**
Dense	$\ln(n)$	**0.02%**	**0.14%**	0.23%	0.35%
	$\sqrt{n}$	0.93%	1.46%	**0.04%**	**0.57%**
	$n/\ln(n)$	3.16%	8.85%	**0.34%**	**5.92%**
Average		3.37%	53.44%	**0.75%**	**45.21%**

Table 9: Average gap with respect to the best known upper and lower bounds of the results obtained by the *mTS* and the *ALNS* on different classes of instances

Figures 4 and 5 show the behavior of δ_{UB} varying the size of the instances, for *mTS* and *ALNS*, respectively. The general behavior of the *mTS* consists in an increase in δ_{UB} as the size of the instance becomes larger. The *ALNS* has a less clear behavior, since for some classes δ_{UB} decreases as the instance size grows when the number of subgraphs p is small, whereas it tends to increase in the other cases. In fact, these are exactly the only instances where *mTS* performs slightly better than the *ALNS*.

7. Conclusions

In this paper, we investigate the *Minimum Gap Graph Partitioning Problem* (*MGGPP*), that is a partitioning problem on vertex-weighted graphs, aiming to produce connected subgraphs with a narrow internal distribution of weights. We tackle the problem with two competing approaches. Due to the misleading nature of the objective function and the strong limitation imposed by the connectivity constraint, a basic Tabu Search algorithm exchanging vertices between

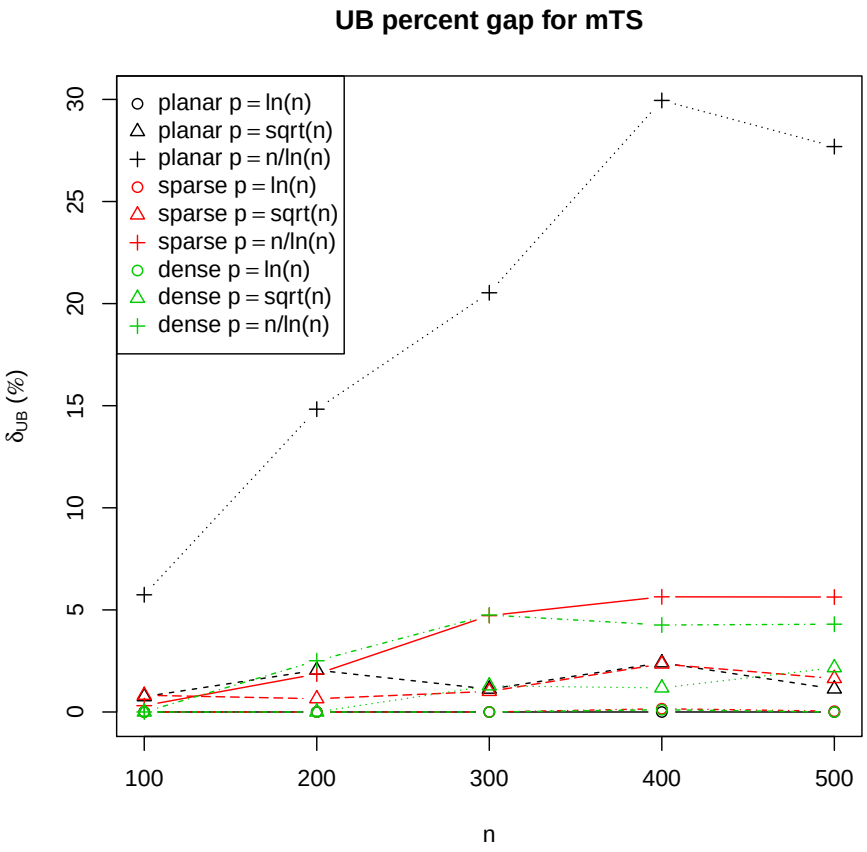

Figure 4: Average behavior of δ_{UB} over different sizes of the instances for mTS

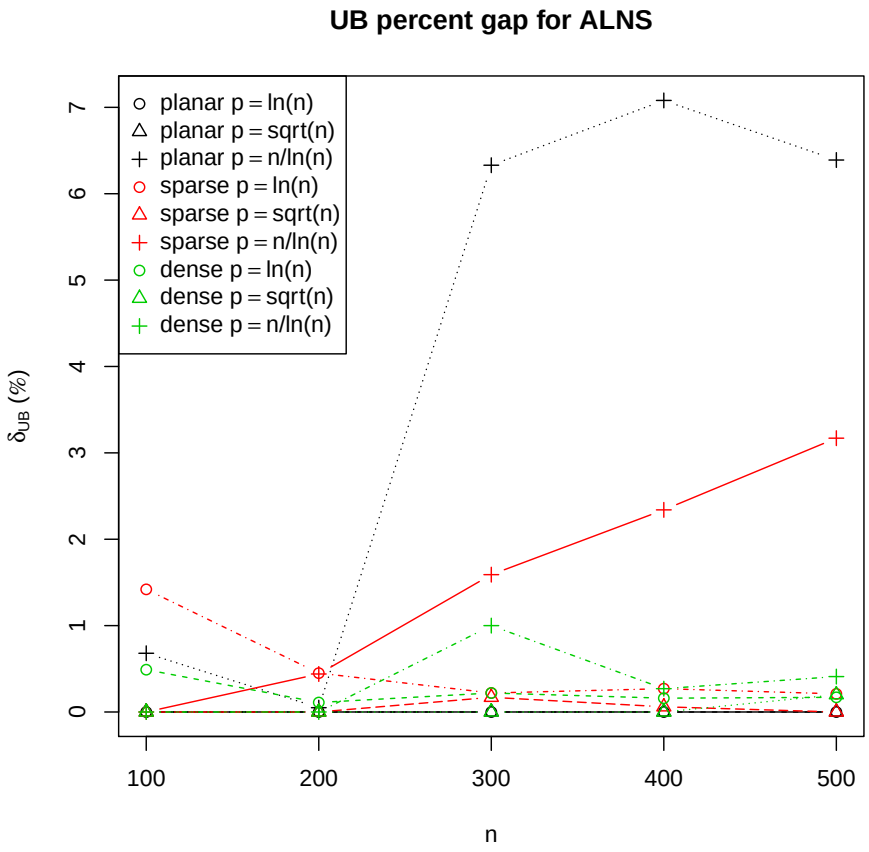

Figure 5: Average behavior of δ_{UB} over different sizes of the instances for $ALNS$

subgraphs proves little effective. We therefore enhance it with a multi-level approach that builds groups of vertices and moves them jointly as macro-vertices. On the other hand, we develop an Adaptive Large Neighbourhood Search *ALNS* algorithm that alternatively removes vertices with different strategies and reinserts them with greedy randomized procedures, and progressively adapts the frequency of each strategy based on the quality of the results obtained. The adaptive tuning mechanism and the experimental advantage to accept only improving solutions allow to use a particularly small number of parameters. Computational experiments have been performed on a large benchmark of instances of different size (from 100 to 500 vertices), topology and density (planar, sparse and dense graphs), and cardinality of the partition ($p = \ln(n)$, $p = \sqrt{n}$ and $p = n/\ln(n)$). In general, the *ALNS* algorithm outperforms the *mTS* algorithm, since the former finds the overall best known result in 157 cases versus 94 of the latter. The solution quality of *ALNS* is very good on dense graphs and on sparse ones for small and medium values of p (i.e., $\ln(n)$ and $\sqrt{n}$), since in these cases the average gap with the lower bound is smaller than 1%, whereas it becomes almost 6% on dense graphs for $p = n/\ln(n)$ and it is very high on planar graphs. This is probably due to the fact that the lower bound is very weak on planar graphs since it gets tighter as the graph becomes denser, providing exactly the optimum for complete graphs. For this reason, our future work will concern the development of methods to find tighter bounds also for sparser graphs and high values of p. For instance this could be obtained considering the relaxation of different formulations of the problem or stronger combinatorial bounding techniques.

References

[1] V.M. Albornoz, M.I. Véliz, R.P. Ortega, and V. Ortíz-Araya. Integrated versus hierarchical approach for zone delineation and crop planning under uncertainty. *Annals of Operational Research*, 286:617–634, 2020.

[2] D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner, editors. *Graph*

Partitioning and Graph Clustering, volume 588 of *Contemporary Mathematics*. American Mathematical Society, Providence, Rhode Island, 2013.

[3] C.-E. Bichot and P. Siarry, editors. *Graph Partitioning*. Wiley-ISTE, 2013.

[4] M. Bruglieri and R. Cordone. Partitioning a graph into minimum gap components. *Electronic Notes in Discrete Mathematics*, 55:33–36, 2016.

[5] M. Bruglieri, R. Cordone, and V. Caurio. A metaheuristic for the minimum gap graph partitioning problem. In *Proceedings of the 15th Cologne-Twente Workshop on Graphs and Combinatorial Optimization*, pages 23–26, June 6–8th 2017.

[6] M. Bruglieri, R. Cordone, I. Lari, F. Ricca, and A. Scozzari. A metaheuristic for the minimum gap graph partitioning problem. In *Proceedings of the 16th Cologne-Twente Workshop on Graphs and Combinatorial Optimization*, pages 107–110, June 18–20th 2018.

[7] J. Chlebikova. Approximability of the maximally balanced connected partition problem in graphs. *Information Processing Letters*, 60:225–230, 1996.

[8] J. F. Cordeau, G. Laporte, F. Pasin, and Røpke S. Scheduling technicians and tasks in a telecommunications company. *Journal of Scheduling*, 13:393–409, 2010.

[9] F. Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13:533–549, 1986.

[10] F. W. Glover and M. Laguna. *Tabu Search*. Kluwer, 1997.

[11] O. Goldschmidt and D. S. Hochbaum. Polynomial algorithms for the k-cut problem. In *Proceedings of the 29th annual IEEE Symposium on Foundations of Computer Sciences (*FOCS*)*, pages 444–451, 1988.

[12] R. Gomes, J. Sousa, J. Muranho, and A. Sá Marques. Different design criteria for district metered areas in water distribution networks. *Procedia Engineering*, 119:1221–1230, 2015.

[13] D. Hochbaum. Algorithms and complexity of range clustering. *Networks*, 73(2):170–186, March 2019.

[14] L. J. Hubert. Some applications of graph theory to clustering. *Psychometrika*, 39(3):283–309, 1974.

[15] T. Ito, T. Nishizeki, M. Schröder, T. Uno, and X. Zhou. Partitioning a weighted tree into subtrees with weights in a given range. *Algorithmica*, 62:823–841, 2012.

[16] T. Ito, X. Zhou, and T. Nishizeki. Partitioning a weighted graph to connected subgraphs of almost uniform size. *IEICE Transactions on Information and Systems*, E90-D(2):449–456, February 2007.

[17] R. Krauthgamer, J. Naor, and R. Schwartz. Partitioning graphs into balanced components. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '09, pages 942–949, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics.

[18] I. Lari, F. Ricca, J. Puerto, and A. Scozzari. Partitioning a graph into connected components with fixed centers and optimizing cost-based objective functions or equipartition criteria. *Networks*, 67(1):69–81, 2016.

[19] M. Lucertini, Y. Perl, and B. Simeone. Most uniform path partitioning and its use in image processing. *Discrete Applied Mathematics*, 42:227–256, 1993.

[20] L. F. Muller, S. Spoorendonk, and D. Pisinger. A hybrid adaptive large neighborhood search heuristic for lot-sizing with setup times. *European Journal of Operational Research*, 218(3):614 – 623, 2012.

[21] L.F. Muller. An adaptive large neighborhood search algorithm for the resource-constrained project scheduling problem. In *Proceedings of the VIII Metaheuristics International Conference (MIC), Hamburg, Germany*, 2009.

[22] F. De Paola, N. Fontana, E. Galdiero, M. Giugni, G. Sorgenti degli Uberti, and M. Vitaletti. Optimal design of district metered areas in water distribution networks. *Procedia Engineering*, 70:449–457, 2014.

[23] D. Pisinger and S. Røpke. Large neighborhood search. In M. Gendreau and J.-Y. Potvin, editors, *Handbook of Metaheuristics*. Springer, 2010.

[24] R. C. Prim. Shortest connection networks and some generalizations. *Bell System Technical Journas*, 36:1389–1401, 1957.

[25] S. Røpke and D. Pisinger. An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation science*, 40(4):455–472, 2006.

[26] P. Sanders and C. Schulz. Engineering multilevel graph partitioning algorithms. In C. Demetrescu and Magnús M. Halldòrsson, editors, *Proceedings of the 19th European Symposium on Algorithms*, volume 6942 of *Lecture Notes in Computer Science*, pages 469–480, Saarbrücken, Germany, September 5–9th 2011. Springer.

[27] P. Shaw. Using constraint programming and local search methods to solve vehicle routing problems. In *International Conference on Principles and Practice of Constraint Programming*, pages 417–431. Springer, 1998.

[28] P. Sivakumar and S. Meenakshi. A review on image segmentation techniques. *International Journal of Advanced Research in Computer Engineering & Technology*, 5(3):641–647, 2016.

[29] P. F. Stadler. Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20:1–45, 1996.

[30] X. Tang, A. Soukhal, and V. T'kindt. Preprocessing for a map sectorization problem by means of mathematical programming. *Annals of Operations Research*, 222:551–569, 2014.

[31] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics*, 1:80–83, 1945.

[32] L. Xiao, L. Hongpeng, N. Dongling, W. Yan, and L. Gang. Optimization of GNSS-controlled land leveling system and related experiments. *Transactions of the Chinese Society of Agricultural Engineering*, 31(3):48–55, 2015.