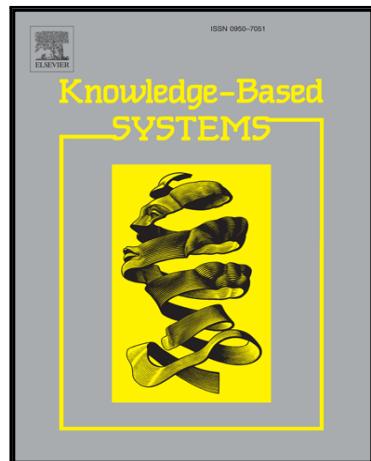


Accepted Manuscript

Multilevel approach for combinatorial optimization in bipartite network

Alan Valejo, Maria Cristina Ferreira de Oliveira, Geraldo P.R. Filho,
Alneu de Andrade Lopes

PII: S0950-7051(18)30153-9
DOI: [10.1016/j.knosys.2018.03.021](https://doi.org/10.1016/j.knosys.2018.03.021)
Reference: KNOSYS 4270



To appear in: *Knowledge-Based Systems*

Received date: 12 July 2017
Revised date: 12 March 2018
Accepted date: 13 March 2018

Please cite this article as: Alan Valejo, Maria Cristina Ferreira de Oliveira, Geraldo P.R. Filho, Alneu de Andrade Lopes, Multilevel approach for combinatorial optimization in bipartite network, *Knowledge-Based Systems* (2018), doi: [10.1016/j.knosys.2018.03.021](https://doi.org/10.1016/j.knosys.2018.03.021)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- A novel multilevel optimization method applicable to problems modeled as bipartite networks. To the extent of our knowledge, the proposal is the first for bipartite network.
- The method has the capability of handling layers independently while executing the multilevel process.
- The implementation of the multilevel framework incorporates two novel efficient matching algorithms, as well as novel contracting and uncoarsening algorithms.
- Applications of the general-purpose method to solve two problems: community detection, in which the method is employed to scale a known algorithm, and dimensionality reduction, in which it is employed to define a novel algorithm.
- A comprehensive experimental evaluation of the proposed solution on real and synthetic bipartite networks that demonstrates it scales the original algorithm and preserves solution quality.
- A test case on dimensionality reduction in text classification, with promising results in terms of runtime and accuracy, is presented.
- A discussion on the underlying features of the framework and its applicability to solving various practical network problems are presented.



Multilevel approach for combinatorial optimization in bipartite network

Alan Valejo^{a,*}, Maria Cristina Ferreira de Oliveira^a, Geraldo P. R. Filho^a, Alneu de Andrade Lopes^a

^a*Institute of Mathematical and Computer Sciences (ICMC), University of São Paulo (USP)
P.O. Box 668, P.C. 14560-970, São Carlos, SP, Brazil*

Abstract

Multilevel approaches aim at reducing the cost of a target algorithm over a given network by applying it to a coarsened (or reduced) version of the original network. They have been successfully employed in a variety of problems, most notably community detection. However, current solutions are not directly applicable to bipartite networks and the literature lacks studies that illustrate their application for solving multilevel optimization problems in such networks. This article addresses this gap and introduces a multilevel optimization approach for bipartite networks and the implementation of a general multilevel framework including novel algorithms for coarsening and uncorsening, applicable to a variety of problems. We analyze how the proposed multilevel strategy affects the topological features of bipartite networks and show that a controlled coarsening strategy can preserve properties such as degree and clustering coefficient centralities. The applicability of the general framework is illustrated in two optimization problems, one for solving the Barber's modularity for community detection and the second for dimensionality reduction in text classification. We show that the solutions thus obtained are statistically equivalent, regarding accuracy, to those of conventional approaches, whilst requiring considerably lower execution times.

© 2017 Published by Elsevier Ltd.

Keywords: Complex Networks, Bipartite Networks, Combinatorial Optimization, Meta-heuristic, Multilevel Optimization, Large-scale Networks

1. Introduction

Bipartite networks comprise a particular class of network models in which the set of vertices is split into two disjoint subsets, with edges connecting vertices placed in different sets. Also known as two-layer networks, they provide a powerful representation of relationships in many real-world systems, including document-word [45], protein-ligand [19], actor-movie [64], georeferenced user-location [63] and paper co-authorship or citation networks [38]. Bipartite network models have been widely employed in hard combinatorial optimization problems that require finding a minimum (or maximum) cost, wherein the number of possible states is finite and usually exponential. Many such problems, e.g., biclique, matching, vertex cover, community structure, traveling salesman and network coloring [2] have proven to be \mathcal{NP} -complete or \mathcal{NP} -hard.

*Corresponding author

Email addresses: alanvalejo@gmail.com (Alan Valejo), cristina@icmc.usp.br (Maria Cristina Ferreira de Oliveira), geraldoprfilho@gmail.com (Geraldo P. R. Filho), alneu@icmc.usp.br (Alneu de Andrade Lopes)

10 Multilevel techniques are being investigated as a global strategy to handle decision-making and optimization problems in a variety of application domains. We refer the reader interested in applications to management problems to
 12 recent literature on the topic [67, 18, 34]. In this paper we investigate the multilevel strategy to handle computationally expensive optimization problems in bipartite networks. In this context, the approach consists of iteratively
 14 coarsening an original network into a hierarchy of smaller sized approximations. A starting solution is obtained in the coarsest network and successively projected back and refined over the inverse sequence of coarsened networks, until
 16 the original one. Previous studies demonstrated the strategy enables running computationally expensive algorithms on large networks with no significant loss in solution quality [10, 22, 50, 41, 53, 54, 32]. Sciences [58] and Walshaw
 18 [59] argued over the relevance and feasibility of the multilevel strategy for solving combinatorial optimization problems. Empirical evidence has been shown by Walshaw [59] that the coarsening process filters the solution space by
 20 gradually removing irrelevant high-cost solutions and drastically reducing the search space, and hence, optimization convergence times.

22 Multilevel algorithms have been applied to many classic network problems, including network coloring [55], traveling salesman [57], network drawing [55], network partitioning [8] and computation of centrality measures [11].
 24 However, current multilevel approaches are not directly applicable to bipartite networks and, to the best of our knowledge, the multilevel strategy has not been considered in this context.

26 We address this gap and introduce a novel multilevel optimization approach applicable to bipartite networks. Furthermore, we describe an implementation of this approach as a general-purpose multilevel framework that incorporates
 28 two novel efficient matching algorithms, as well as novel contracting and uncoarsening algorithms.

30 In order to illustrate its potential, we employed the framework to handle two distinct problems defined in bipartite
 32 networks, namely community detection and dimensionality reduction. In the community detection problem, tests on a large set of synthetic networks demonstrated that, combined with a proper local search strategy, it yields good
 34 speedups and preserves solution quality. When employed to perform dimensionality reduction in text classification it yielded encouraging results in terms of both runtime and accuracy as compared with a standard dimensionality reduction technique.

36 The remainder of the paper is organized as follows: Section 2 reviews some basic concepts on networks and
 38 provides an overview of the standard multilevel approach. Section 3 discusses previous work and application of multilevel strategies on combinatorial optimization problems. Section 4 introduces a multilevel formulation for bipartite
 40 networks and its implementation. Section 5 reports our empirical results, which include (i) an analysis of how the multilevel representation impacts the topological features of a real bipartite network; (ii) an empirical study of instantiating
 42 the framework to solve the community detection problem on a large synthetic test suite; and (iii) an empirical study of its application to dimensionality reduction in text classification. Section 6 briefly discusses how the general
 44 framework can be tuned for application in other types of problems. Finally, Section 7 summarizes our findings and
 46 discusses future work.

44 2. Background

2.1. Basic concepts

46 Let $G = (V, E, \sigma, \omega)$ be an undirected weighted network, where $V = \{1, \dots, n\}$ denotes the set of vertices and
 48 $E \subseteq V \times V$ denotes the set of edges, such that $(v, u) = \{(u, v) = (v, u) \mid u, v \in V\}$. Let $n = |V|$ be the total number of
 50 vertices and $m = |E|$ be the total number of edges, where operator “ $|.$ ” stands for the cardinality of a set. The weight
 52 of an edge (u, v) is represented by $\omega(u, v)$ with $\omega : V \times V \rightarrow \mathbb{R}^*$ and the weight of a vertex v is represented by $\sigma(v)$ with
 $\sigma : V \rightarrow \mathbb{R}^*$.

54 A network $G = (V, E, \sigma, \omega)$ is bipartite (two-layer network) if V is partitioned into two sets V_1 and V_2 , such that
 56 $V_1 \cap V_2 = \emptyset$ and $E \subseteq V_1 \times V_2$. Hereafter, each vertex subset is called a layer. A bipartite network thus has two layers
 58 so that vertices in the same layer are not connected.

54 The degree of a vertex $v \in V$, denoted κ_v , is given by the total weight of its adjacent edges, i.e. $\kappa_v = \sum_{u \in V} w(v, u)$.
 56 The h -hop neighborhood of v , denoted $\Gamma_h(v)$, is formally defined as the vertices in set $\Gamma_h(v) = \{u \mid \text{there is a path of}$
 $\Gamma_1(v)$, is the set of vertices adjacent to v ; the 2-hop neighborhood, $\Gamma_2(v)$, is the set of vertices 2-hops away from v , and so forth.

58 A similarity score $S(u, v)$ can be computed from a pair of vertices u and v . A fundamental structural similarity function between a pair of vertices is given by the number of common neighbors, defined as $S_{cn}(u, v) = |\Lambda(u, v)|$,

60 $\Lambda(u, v) = \{\Gamma_1(u) \cap \Gamma_1(v)\}$. Alternatively, a weighted common neighbors similarity function can be defined by equation 1, where the term $\log(1 + s(z))$ is used to prevent negative scores [35].

$$S_{wcn}(u, v) = \sum_{z \in \Lambda(u, v)} \frac{\omega(u, z) + \omega(v, z)}{\log(1 + s(z))} \quad s(u) = \sum_{z \in \Gamma_1(u)} \omega(u, z) \quad (1)$$

62 The local clustering coefficient of a vertex is given by the probability of its neighbors being connected [64, 39]. This statistics is closely related to transitivity, which measures the relative frequency of triangles in the vertex neighborhood. The clustering coefficient is defined by equation 2:

$$cc(v) = \frac{2tr(v)}{|\Gamma_1(v)|(|\Gamma_1(v)| - 1)}, \quad (2)$$

66 where $tr(u)$ denotes the number of edges $(v, z) \in E$, such that $v, z \in \Gamma_1(u)$ and cc relies on the enumeration of the triangles in the network. However, as triangles do not occur in bipartite networks this definition is not valid. An equivalent metrics for bipartite graphs was introduced by Latapy et al. [33], defined in equation 3, which captures the 68 overlap between vertex neighborhoods in the same layer.

$$cc_b(v, u) = \frac{|\Gamma_1(u) \cap \Gamma_1(v)|}{|\Gamma_1(u) \cup \Gamma_1(v)|} \quad cc_b(v) = \frac{\sum_{u \in \Gamma_2(v)} cc_b(v, u)}{|\Gamma_2(v)|}, \quad (3)$$

2.2. General-purpose multilevel optimization

70 A multilevel optimization is formally defined as a meta-heuristics that combines different heuristics to guide, 72 modify and possibly fix a solution obtained from a target algorithm (or operations of the subordinate heuristics, local search or global search) and refines this solution over multiple iterations. It operates in three phases, namely 74 *coarsening*, *solution finding* and *uncoarsening*. In the coarsening phase, the network size is successively reduced to 76 obtain coarser network representations; in the solution finding phase a starting solution is obtained applying the target algorithm in the coarsest representation; in the uncoarsening phase, the starting solution is successively projected back 78 to the intermediate networks and refined, until obtaining the final solution.

Figure 1 illustrates such a process, considering an initial network G_0 (in which the original problem instance is 78 defined), where G_L denotes the coarsest network obtained after L coarsening steps (levels), S_L denotes the starting solution obtained in G_L , and S_0 denotes the final refined solution obtained in G_0 .

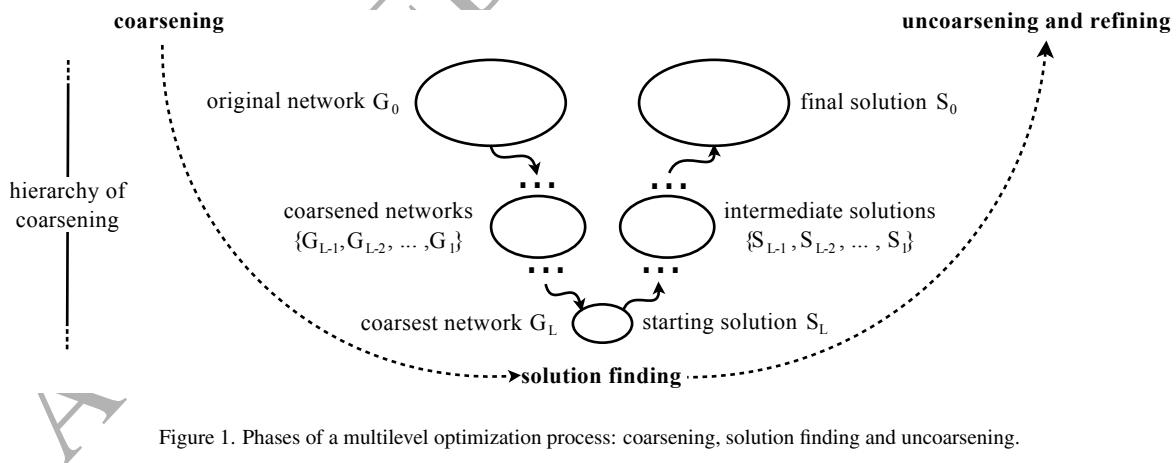


Figure 1. Phases of a multilevel optimization process: coarsening, solution finding and uncoarsening and refining.

80 Coarsening

82 The coarsening phase constructs a hierarchy of coarsened networks G_l from the initial network G_0 , yielding intermediate 84 network approximations on multiple levels-of-detail. The process requires two algorithms, namely *matching*, which defines which vertices will be merged, and *contracting*, which builds the reduced representation, given the matching. Let $G_l = (V_l, E_l, \sigma_l, \omega_l)$ be the network model coarsened at level l , with $|V_0| > |V_1| > \dots > |V_l|$.

Coarsening starts with the matching step. According to some given restriction, in general, pairs of vertices are selected for matching, producing a set of unordered pairs called *vertex matching*, *independent edge set* or simply *matching*. Formally, a matching M consists of a set of pairwise non-adjacent edges. Heavy-edge matching is a popular algorithm for this purpose, which attempts to find a matching of maximal weight [25].

At any level, the coarsening of a network must preserve its topological features, implying that vertex and edge weights of the reduced network must reflect the connectivity of its parent. This will be guaranteed by a proper choice of matching strategy, which is a key component of effective multilevel optimizations. A matching strategy inadequate to support the solution finding phase will impair the quality of the solution derived by a multilevel algorithm and its performance.

Once the matching is defined, a contracting algorithm constructs the coarsened network, by joining matched vertex pairs into a single *super-vertex* (sV). A child network G_{l+1} will inherit the non-joined vertices from its parent. In order for G_{l+1} to be a good proxy to its parent, given a super-vertex $sV = \{v, u\} \in V_{l+1}$ its weight $\sigma(sV)$ is computed as the sum of weights $\sigma(v)$ and $\sigma(u)$, $\{u, v\} \in V_l$. Furthermore, the edges incident to vertices $\{u, v\} \in V_l$ are joined to obtain the so-called super-edges incident to sV .

Solution finding

This phase employs the target algorithm to solve the problem on the smallest network G_l . Let S be the set of all possible solutions in the coarsest problem instance. Given an objective (cost) function $f : S \rightarrow \mathbb{R}$ (or \mathbb{N}) that assigns a cost to each solution in S , the aim is to find a state $s \in S$ with minimum (or maximum) cost. For instance, in the traveling salesman problem $f(s)$ expresses the length of tour s , whereas in a network community detection problem it denotes some measure of community quality. Since the coarsest network is possibly very small, it becomes feasible to employ computationally expensive target algorithms to find a starting solution [22].

Uncoarsening and refinement

The uncoarsening (also known as solution projection) phase successively transfers the solution available at a current level to the upper level in the hierarchy, i.e., the solution obtained in the coarsest network G_l is successively projected through intermediate networks $G_{l-1}, G_{l-2}, \dots, G_1$ up to the original network G_0 .

Solution S_l is constructed from S_{l+1} simply by assigning vertices $\{u, v\} \in V_l$ to the same set of their parent $sV \in V_{l+1}$. Although S_l is a local minima of f in G_l , this may not be the case of solution S_{l-1} , derived for the upper level G_{l-1} , with respect to G_l . Therefore, a refinement heuristics can be applied to avoid local minima and improve solution quality. Local operations can move the solution towards a lower cost neighboring solution in the search space; for instance, in a community detection problem vertices can be moved between adjacent communities to improve a target quality measure.

3. Related Work

Early studies of multilevel optimization were mostly designed to speed up the recursive bisection problem, as its high computational cost prevents wider applicability [6]. One of the first theoretical analysis was presented by Karypis and Kumar [22], who demonstrated multilevel approaches can find high-quality communities in a variety of networks. Later, Karypis and Kumar [25] introduced the now widely adopted matching algorithms HEM (Heavy Edge Matching), LEM (Light Edge Matching) and MCH (Modified Edge Matching). Other studies relevant for the development and expansion of the multilevel approach were conducted by Walshaw and Cross [60], who presented a theoretical multilevel formulation of the Kernighan-Lin method for mesh partitioning, and by Korosec et al. [29], who introduced a new multilevel colony optimization applicable to several optimization problems. Sciences [58] and Walshaw [59] provided compelling evidence the multilevel framework is an extremely useful addition to combinatorial optimization toolkits, although it can not be considered a panacea.

Previous work on multilevel optimization in complex networks can be broadly organized into four categories. A first category encompasses studies that explore features of network domains; e.g., Abou-Rjeili and Karypis [1] studied scale-free networks (with power-law degree distribution), Oliveira and Seok [43] designed a multilevel spectral approach that explores features of biological networks (protein complexes) and Valejo et al. [53] considered properties of social networks, such as high transitivity and assortativity. A second group comprises contributions focused on

scalability, e.g., investigating parallel and distributed paradigms that improve the performance of the coarsening and refinement phases [5, 23, 26, 24, 61, 47, 3, 52, 51, 16, 49, 62, 31, 32]. A third category of studies is concerned with the optimization of a target objective function, e.g., several contributions focused on improvements in modularity [40, 14, 48, 65, 41, 46, 15, 32]. Finally, there are significant contributions in applications, including (and not limited to) graph coloring [56], traveling salesman problem [57], graph drawing [55], biomedical feature selection [42], covering design [12], DNA sequencing [9], vehicle routing [44], semi-supervised learning [66], partitioning or community detection [8] and computation of centrality measures [11].

We are not aware of any previous effort concerned with the usage of multilevel strategies to solve optimization problems in bipartite networks, despite the relevance of this kind of network to real-world modeling problems. Current multilevel methods cannot be directly applied in bipartite models without adaptations, e.g., they do not consider vertex types, whereas the layers in a bipartite network usually represent distinct types of entities that must be handled independently. For the sake of illustration, suppose a text document collection modeled as a document-word bipartite network. For a start, matching vertices that represent words and vertices that represent documents (i.e., matching of vertices of different layers) would not be meaningful in most application scenarios. Moreover, as the number of words is typically much higher, coarsening the word layer may be sufficient to reduce the asymptotic convergence of a target algorithm.

4. A multilevel approach in bipartite networks

This section introduces a multilevel optimization framework designed to handle bipartite networks. Bearing in mind the previous discussion, we have defined two restrictions that establish the major distinction between current multilevel methods and the one proposed here:

1. Vertices are only allowed to match their set of two-hop neighbors.
2. The matching algorithm must operate on vertices of the same layer.

Such restrictions support cost-effective implementations of multilevel strategies in bipartite networks. From the definitions of bipartite networks and two-hop neighborhoods, the first restriction implies adjacent vertices are not matched. Furthermore, vertices can only match others in their two-hop neighborhood set, rather than any non-adjacent vertex. Considering, for instance, the network in Figure 2, vertex u_1 can match vertices u_2 and u_3 , which are non-adjacent and are in $\Gamma_2(u_1)$.

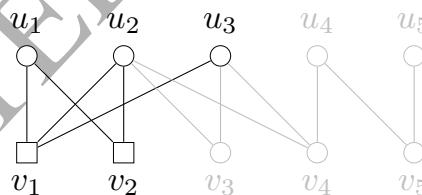


Figure 2. Example of a bipartite network with $V_1 = \{u_1, u_2, u_3, u_4, u_5\}$ and $V_2 = \{v_1, v_2, v_3, v_4\}$.

The first restriction alone does not enforce independent handling of layers, which is guaranteed by the second restriction, which states layers (either one of them or both) will be processed independently in the coarsening phase. Therefore, the coarsening will not match vertices of different kinds and distinct coarsening and refinement algorithms, or distinct parameterizations of the same algorithm may be adopted in processing each layer. Such a restriction also favors the adoption of distributed or parallel processing strategies.

Algorithm 4.1 summarizes the general *multilevel optimization framework for bipartite networks (MOB)*. Similarly to the standard multilevel approach, it comprises the phases of coarsening (lines 1–4), solution finding (line 6) and uncoarsening (lines 7–9). It takes as inputs the initial bipartite network $G = (V_1 \cup V_2, E, \sigma, \omega)$, and for each layer a maximum number of coarsening levels L and a layer reduction factor rf .

The coarsening is applied to each layer (line 1), level by level until the desired reduction factor is attained, by calls to a matching algorithm (line 3) and a contracting algorithm (line 4). The matching algorithm considers the above restrictions in selecting the vertex pairs to produce the list of independent edges (line 3). The reduction factor

($rf \in [0, 0.5]$) is multiplied by the number of vertices to determine the maximum matching number. If $rf = 0.5$ (the maximum reduction factor), each coarsening iteration will (potentially) reduce the number of vertices by a factor of two, yielding a logarithmic decrease in network size along the process. This is not guaranteed if the network is disconnected or highly sparse, since in this case there may be an insufficient number of edges. The call to the contracting algorithm in the next step (line 4) is responsible for creating the coarsened bipartite network, in which any matched vertex pairs have been merged into super-vertices.

Algorithm 4.1: MOb: multilevel optimization framework for bipartite networks.

Input:

bipartite network	: $G = (V, E, \sigma, \omega)$
maximal number of levels	: array $L = \{L_i \mid L_i \in [0, n] \subset \mathbb{Z}\}$ with $1 \leq L \leq 2$
reduction factor for each layer	: array $rf = \{rf_i \mid rf_i \in (0, 0.5] \subset \mathbb{R}\}$ with $1 \leq rf \leq 2$
layers to be coarsened	: array $layers = \{i \mid i \in \{1, 2\} \subset \mathbb{Z}\}$ with $1 \leq layers \leq 2$

Output:

solution	: S
----------	-------

```

1 for  $i \in layers$  do
2   while  $l \leq L_i$  or layer  $i$  is not as small as desired do
3      $M \leftarrow$  matching( $G_l, i, rf_i$ );
4      $G_{l+1} \leftarrow$  contracting( $G_l, M$ );
5     increment  $l$ ;
6    $S_l \leftarrow$  target_algorithm( $G_l$ );
7   while  $l \neq 0$  do
8      $S_{l-1} \leftarrow$  uncoarsening( $G_{l-1}, G_l, S_l$ );
9      $S_{l-1} \leftarrow$  refining( $G_{l-1}, S_{l-1}$ );
10    decrement  $l$ ;

```

Return: S

The target algorithm is then executed in the coarsest network G_l to obtain a starting solution S_l (line 6). Finally, in the subsequent uncoarsening phase this solution is projected back, up to G_0 , through the space of intermediate solutions $S_{l-1}, S_{l-2}, \dots, S_1, S_0$ (lines 7–10), possibly refining the solutions at each level (line 9). Although we do not investigate refinement algorithms in this work, the rationale is to apply a local search strategy to improve the current solution, i.e., algorithms should explore small regions of the solution space, in order to reduce impact on performance and scalability. We do introduce novel algorithms for matching, contracting and uncoarsening.

4.1. Matching in bipartite networks

We propose a straightforward matching algorithm called *random greedy matching for bipartite networks (RGMb)*, described in Algorithm 4.2, where $S(u, v)$ denotes a similarity function.

A vertex u is picked randomly (Line 4) at each iteration, and as long as it has not yet been matched, one of its unmatched two-hop neighbors v so that $S(u, v)$ is maximal is chosen (Line 5). Vertices u and v are marked as matched and removed from the list of unmatched vertices (Line 7). The process iterates until no more vertices can be eliminated from the list. For unweighted edges, a random neighbor is selected for matching, otherwise, the heaviest adjacent edge is selected.

One may consider different similarity functions, e.g., common neighbors similarity S_{cn} or weighted common neighbors similarity S_{wcn} . For illustration, consider a graph $G = (V_1 \cup V_2, E)$, shown in Figure 3, with $V_1 = \{u_1, u_2, u_3, u_4, u_5\}$ and $V_2 = \{v_1, v_2, v_3, v_4, v_5\}$, $\Gamma_2(u_2) = \{u_1, u_3, u_4\}$ and u_1, u_2, u_3 and u_4 unmatched. Suppose vertex u_2 has been randomly chosen for matching. Adopting S_{cn} as the similarity function, then $S_{cn}(u_2, u_1) = 2$, $S_{cn}(u_2, u_3) = 3$ and $S_{cn}(u_2, u_4) = 1$, therefore, pair $\{u_2, u_3\}$ would be included in the matching, as illustrated in Figure 3(a). Alternatively, for a choice of S_{wcn} as similarity function, $S_{wcn}(u_1, u_2) = 11$, $S_{wcn}(u_1, u_3) = 7$ and $S_{wcn}(u_2, u_4) = 2.5$ and pair $\{u_1, u_2\}$ would be included, as illustrated in Figure 3(b).

Algorithm 4.2: *RGMb*: random greedy matching for bipartite networks.**Input:**

bipartite network : $G = (V, E, \sigma, \omega)$
 selected layer : $l \in \{1, 2\}$
 reduction factor : $rf \in (0, 0.5] \subset \mathbb{R}$

Output:

matching : array of tuples $M = \{(u, v) \mid u, v \in V\}$

```

1 matching  $M \leftarrow \emptyset$ ;
2 merge count  $mc \leftarrow rf * |V|$ ;
3 while  $mc > 0$  do
4   randomly select a vertex  $u \in V_l$ ;
5   select  $v \in \Gamma_2(u)$  of maximal  $S(u, v)$ ;
6    $M \leftarrow M \cup (u, v)$ ;
7   remove  $u$  and  $v$  from  $V_l$ ;
8   decrement  $mc$ ;
```

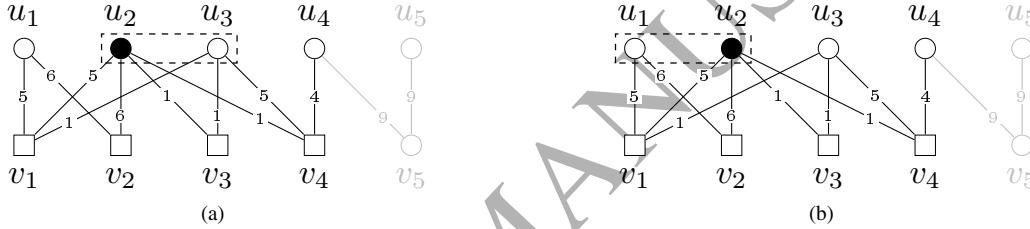
Return: M 

Figure 3. Bipartite network, such that $V_1 = \{u_1, u_2, u_3, u_4, u_5\}$ and $V_2 = \{v_1, v_2, v_3, v_4\}$; (a) pairs $\{u_2, u_3\}$ included in the matching; (b) pairs $\{u_1, u_2\}$ included in the matching.

198 Algorithm *RGMb* does not ensure an optimal choice over all possible matchings. An alternative strategy would
 choose from a list of vertex pairs sorted in decreasing order of similarity scores, which may be kept in an appropriate
 200 data structure, e.g., a heap or a priority queue. An alternative algorithm that implements this strategy is called *greedy*
 202 *sorted matching for bipartite networks (GMb)* (Algorithm 4.3): it selects the best possible match for a vertex from its
 two-hop neighborhood (line 4).

204 *GMb* constructs a priority queue of potential matches ranked by similarity (Lines 3–5) and uses it to retrieve
 206 optimal matching choices (Line 8). In case of a tie, it makes a random choice, in order to favor further exploration of
 208 the solution space over multiple iterations. If a vertex is selected that has already been matched, it is skipped. *GMb*,
 210 albeit slower than its random search counterpart *RGMb*, is more robust and yields better performance (see Section 5).

212 For illustration, consider the graph $G = (V_1 \cup V_2, E)$ depicted in Figure 4, with $V_1 = \{u_1, u_2, u_3, u_4, u_5\}$ and
 214 $V_2 = \{v_1, v_2, v_3, v_4, v_5\}$. Suppose the matching is being performed on layer V_1 . Adopting the S_{cn} similarity, pairs
 $\{u_2, u_3\}$ and $\{u_4, u_5\}$ would be included in the matching, as illustrated in Figure 4(a). Alternatively, for a choice of S_{wcn}
 216 similarity, pairs $\{u_1, u_2\}$ and $\{u_4, u_5\}$ would be included, as shown in Figure 4(b).

218 Variations of the general matching algorithm can be obtained depending on the combined choices of matching
 220 strategy and similarity function $S(u, v)$. We report four variants, namely $RGMb_{cn}$ and $RGMb_{wcn}$ for a choice of
 222 random matching with, respectively, S_{cn} and S_{wcn} similarity; GMb_{cn} and GMb_{wcn} for the equivalent combinations
 224 with greedy matching.

4.2. Contracting in bipartite networks

216 Algorithm 4.4 (*Cb, contracting of a bipartite network from a matching*) describes how to create a coarsened
 bipartite network from a matching M . It takes as input a bipartite network G_l and a corresponding matching M . First,

Algorithm 4.3: GMb: greedy matching for bipartite networks.**Input:**

bipartite network : $\mathbf{G} = (V, E, \sigma, \omega)$
 selected layer : $l \in \{1, 2\}$
 reduction factor : $rf \in (0, 0.5] \subset \mathbb{R}$

Output:

matching : array of tuples $M = \{(u, v) \mid u, v \in V\}$

```

1 matching M ← ∅;
2 Q ← max priority queue;
3 forall the  $u \in V_l$  do
4   select  $v \in \Gamma_2(u)$  of maximal  $\mathbf{S}(u, v)$ ;
5   insert  $(u, v)$  into Q with priority  $\mathbf{S}(u, v)$ ;
6 merge count mc ←  $rf * |V|$ ;
7 while  $mc > 0$  do
8    $(u, v) \leftarrow$  remove from Q;
9   if  $u$  and  $v$  unmatched then
10    M ← M ∪  $(u, v)$ ;
11   decrement mc;

```

Return: m

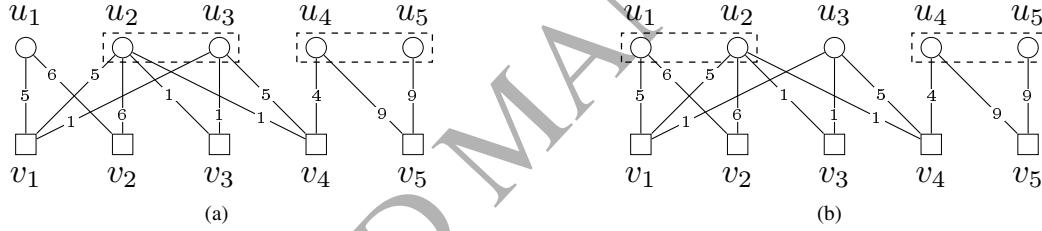


Figure 4. Bipartite network, such that $V_1 = \{u_1, u_2, u_3, u_4, u_5\}$ and $V_2 = \{v_1, v_2, v_3, v_4\}$; (a) pairs $\{u_2, u_3\}$ and $\{u_4, u_5\}$ included in the matching; (b) pairs $\{u_1, u_2\}$ and $\{u_4, u_5\}$ included in the matching.

218 vertex pairs $\{u, v\} \in M$ are mapped into a *successor vector* (lines 5–6) and joined into a single super-vertex $sV \in G_{l+1}$ (lines 4–9). The *successor vector* will be accessed in the uncoarsening phase to assign pair $\{u, v\} \in G_{l+1}$ to the same 220 subset of its super-vertex sV . Any vertices not included in M are inherited by G_{l+1} (line 10). After the mapping, 222 adjacent edges in E_l are joined and added to E_{l+1} (lines 11–18), i.e., each pair $(u, v) \in E_l$ holds its successors (w, z) and if edge (w, z) already exists in E_{l+1} its weight is increased by $\omega_l(u, v)$ (lines 15–16); otherwise, a new edge (w, z) with weight $\omega_l(u, v)$ is inserted into E_{l+1} (lines 17–18).

224 Figure 5 illustrates the contracting process. Considering, for instance, the previous example, in which pairs $\{u_2, u_3\}$ and $\{u_4, u_5\}$ have been included in the matching, as illustrated in Figure 5(a), the resulting coarsened network 226 is depicted in Figure 5(b).

228 The capability of handling layers with distinct parameter settings (rf and L) and coarsening algorithms is a compelling 229 feature of the proposed framework. This is advantageous in many situations, since layers in many real network 230 models are highly unbalanced in size, as in the already mentioned document-word networks. Coarsening just one of 231 the layers is also a useful feature, for example, in dimensionality reduction problems, e.g., to reduce the space of 232 words in document-word networks, as illustrated in Section 5.3.

Algorithm 4.4: Cb : Contracting of a bipartite network from a matching**Input:**

bipartite network : $G_l = (V_l, E_l, \sigma_l, \omega_l)$
 matching : array of tuples $M = \{(u, v) \mid u, v \in V\}$

Output:

coarsened bipartite network : $G_{l+1} = (V_{l+1}, E_{l+1}, \sigma_{l+1}, \omega_{l+1})$

```

232
1   $n \leftarrow 1$ ;
2  sucessor  $\leftarrow \emptyset$ ;
3   $G_{l+1} = (V_{l+1}, E_{l+1}, \sigma_{l+1}, \omega_{l+1}) \leftarrow$  empty bipartite network;
4  forall the  $(u, v) \in M$  do
5    sucessor[u]  $\leftarrow n$ ;
6    sucessor[v]  $\leftarrow n$ ;
7    add new super-vertex  $sV = \{u, v\}$  to  $V_{l+1}$  with id  $n$ ;
8     $\sigma_{l+1}(v) \leftarrow \sigma_l(u) + \sigma_l(v)$ ;
9    increment  $n$ ;
10 Vertices not merged are inherited by  $G_{l+1}$ ;
11 forall the  $(v, u) \in E_l$  do
12    $w \leftarrow$  sucessor[u];
13    $z \leftarrow$  sucessor[v];
14   if  $w \neq z$  or they are not super-vertices in  $E_{l+1}$  then
15     if  $(w, z) \in E_{l+1}$  then
16        $E_{l+1} \leftarrow \omega_{l+1}(w, z) + \omega_l(u, v)$ ;
17     else
18        $E_{l+1} \leftarrow$  new edge  $(w, z)$  with weight  $\omega_l(u, v)$ ;

```

Return: $G_{l+1} = (V_{l+1}, E_{l+1})$

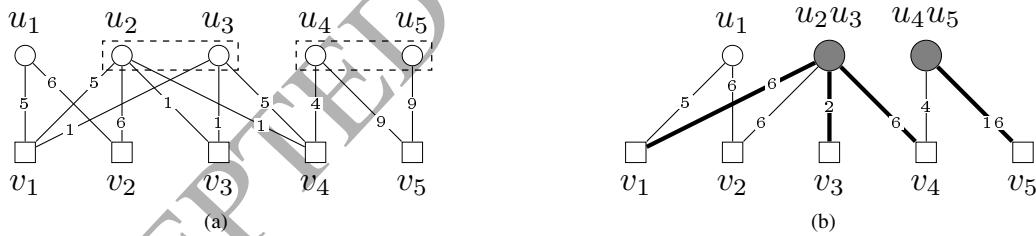


Figure 5. Contracting of a layer with the GMb_{cn} matching algorithm: (a) selected matching in original bipartite network; (b) contracting process with newly formed super-vertices and modified edges highlighted.

4.3. Uncoarsening algorithm

The solution finding phase executes the target algorithm to obtain a starting solution S_L . The representation of S_L depends on the problem being handled. In a community detection problem, it is described as a partitioning of the vertex set into non-empty partitions P_k with $\cup P_k = S_L$, $P_k \subseteq V_L$. Each iteration of the uncoarsening phase projects the current solution, obtained in network G_l , to network G_{l-1} as described in Algorithm 4.5 Ub (*uncoarsening of bipartite networks*). The rationale is, for each vertex $u \in V_{l-1}$ in network G_{l-1} (line 2), to obtain its corresponding successor vertex $w \in V_l$ (line 3) and assign vertex u to its partition (lines 4–5). An implementation can keep the successor vector as a global variable, or it may be provided as a network attribute.

As discussed later (Section 6), the algorithm can be adapted with little effort to handle other types of problems and representations, e.g., in link prediction, S would be a set of predicted edges $E_p \notin E$. In some problems the solution

finding and uncoarsening phases are omitted, e.g., in dimension reduction the coarsest bipartite network is itself the final solution (see Section 5.3).

Algorithm 4.5: *Ub*: Uncoarsening of bipartite networks.

Input:

bipartite network	: $G_{l-1} = (V_{l-1}, E_{l-1}, \sigma_{l-1}, \omega_{l-1})$
bipartite network	: $G_l = (V_l, E_l, \sigma_l, \omega_l)$
solution S_L	: $S_L = \{P_1, P_2, \dots, P_k\}$

Output:

projected solution	: S_{l-1}
--------------------	-------------

```

1 new solution  $S_{l-1} \leftarrow \emptyset$ ;
2 forall the  $u \in V_{l-1}$  do
3    $w \leftarrow \text{successor}[u]$ ;
4    $P_i \leftarrow \text{get\_partition}(w)$  with  $i \in \{1, \dots, k\}$ ;
5   assign  $u$  to  $P_i$ ;
```

Return: S_{l-1}

4.4. Computational complexity

Consider the multilevel strategy applied to both layers of a network. Computing a matching requires, for each selected vertex u , to find its h -hop neighborhood $\Gamma_h(u)$, which takes $O(n\langle\kappa\rangle h)$ time, where $\langle\kappa\rangle$ is the average network degree and h specifies the number of distance hops. Most real networks are sparse, i.e., $\langle\kappa\rangle \ll n$ and $m \approx n$. This study only considers two-hop neighborhoods ($h = 2$), which leads to notable computational savings, since Γ_2 can be reached in nearly linear time. Appropriate data structures and parallelism can be employed to maximize the efficiency of the matching step [21, 27, 28]. The subsequent contracting step relies directly on the matching and can be very fast in sparse networks. A coarsened network can be mapped and built in time $O(|E|)$, since the process is iterated over all edges of its parent network.

If the coarsening parameter is set to its maximum value ($rf = 0.5$), at each iteration the current network size will be reduced at most by a factor of two relative to its parent. In this case, given $1 < L < n$, the coarsening steps would produce networks of sizes $(n, n/2, \dots, n/2^L)$, yielding a reduction factor $O(\log_2(n))$ over the L iterations. For instance, departing from a network with $n = 100$ vertices, $rf = 0.5$ and $L = 2$, potentially 50 vertex pairs will be merged in the first iteration (from G_0 to G_1) and 25 in the second (from G_1 to G_2), a reduction factor of 4. In most real-world problems, such a reduction factor would suffice to produce a manageable model to run the optimization algorithm. As the cost of coarsening decreases drastically at each level, the number of levels L may be neglected in the estimation of the computational cost.

The complexity of the solution finding phase is determined by the target algorithm. If T_{tg} is the cost of the target algorithm on a bipartite network with n vertices, and T_{ss} is the cost to obtain the starting solution, it is not too difficult to show that $T_{ss} = T_{tg}/(2^L)$. Insofar as the hierarchy of approximations produced by coarsening, $L = 1$ yields $(n, n/2)$, $L = 2$ yields $(n, n/2, n/4)$, $L = 3$ yields $(n, n/2, n/4, n/8)$ and so on. Hence, the number of vertices in the coarsest network is approximately $n/(2^L)$ and therefore $T_{ss} = T_{tg}/(2^L)$. Parameter L can thus be chosen to control the cost deemed acceptable to obtain an initial solution.

As an illustration, consider a network with $n = 100$ vertices and two target algorithms, X and Y with $T_X = O(n)$ and $T_Y = O(n^2)$, hence, $T_X = O(100)$, $T_Y = O(10,000)$ and $T_X = \sqrt{T_Y}$. Consequently, if one wishes to execute algorithm Y on a network with $n = 10,000$ with a runtime similar to that of algorithm X , the network size should be reduced to $\sqrt{n} = 100$, which requires setting $L \approx 7$ and $rf = 0.5$. Of course, the asymptotic complexity of algorithm Y remains quadratic, but the coarsening allows its execution with a runtime equivalent to the linear one.

Analogous to the solution finding phase, the complexity of the uncoarsening phase depends on the target problem (target algorithm), and also on the choice of refinement algorithm, if one is used. In problems such as community detection and classification, it is not necessary to project the solution along the intermediate levels, rather it can be projected directly to the original network, since successors have been stored in an array (or other data structure). In

these cases, this phase as implemented in Algorithm 4.5 would have linear cost in the number of vertices, $O(n)$. As we do not investigate refinement algorithms, an analysis of their asymptotic complexity is suppressed.

5. Experimental Results and Analysis

We assessed the proposed framework guided by the following research questions:

1. How does coarsening affect a network's topological features at each level?
2. Can the framework improve the efficiency of optimization algorithms without incurring significant losses in the solution quality?
3. How does the framework impact on the scalability of the local search strategy?
4. To which extent can it enable the running of high-cost algorithms on large networks?
5. Is it sufficiently general for handling different combinatorial optimization problems?

We conducted three experimental studies to find answers to the questions. A first study investigated the impact of the coarsening process on the topological properties of a well-known author-paper network. In a second study, the framework was employed in the context of community detection, which is a prototypical application of multilevel strategies. Finally, in a third study, it was employed in dimension reduction in a text classification scenario.

The experiments were executed in a Linux machine with 8-core processor with 3.7 GHz CPU and 64 GB main memory. The framework¹ was implemented in Python with *igraph* library². We report average values obtained from 30 executions for algorithms that rely on random strategies.

5.1. Analysis of topological properties

We considered the scientific collaboration network Cond-Mat³, which describes co-authorships of preprints posted from 1995 to 1999 in the Condensed Matter section of the arXiv repository, to address the first research question. It has 38,742 vertices (representing authors and papers) and 58,595 edges (co-authorship relations) (additional information can be found elsewhere [36, 37]). Our interest was to observe how a progressive coarsening affects the intrinsic topological properties of a network.

The Cond-Mat network has been extensively analyzed and is known to have characteristic features regarding degree distribution and clustering coefficients, as depicted in Figure 6. According to Figure 6(a), the degree distribution follows a power-law relationship characteristic of scale-free networks, with a vast majority of low-degree vertices and a few vertices of very high-degree, i.e., the so-called “hubs”.

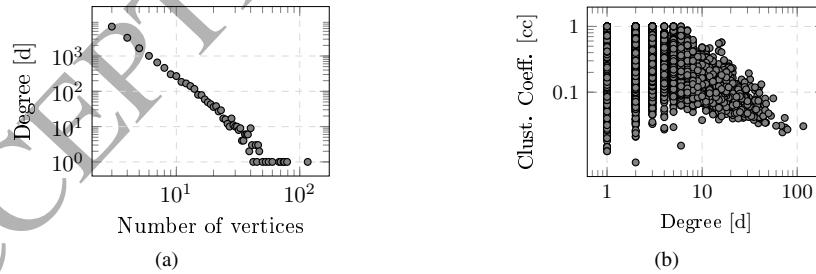


Figure 6. Scale-free properties of the Cond-Mat bipartite network. (a): histogram of degree distribution; (b): histogram of two-mode local clustering coefficient.

An inverse relation between vertex clustering coefficients and vertex degrees is also evident in Figure 6(b). This is a particular feature of this network: the hubs, i.e., authors with many collaborators, have low-clustered neighborhoods

¹available from <https://github.com/alanvalejo/mob>

²available from <http://igraph.org/python/>

³available from <https://toreopsahl.com/datasets/#newman2001>

306 and they tend to interact with collaborators from distinct groups, who are not usually collaborators themselves. In
 307 contrast, authors with few collaborators have highly-clustered neighborhoods, i.e., they often interact within smaller
 308 and more restrict research groups whose members also collaborate with each other.

310 We investigated how the scale-free properties of degree distribution and two-mode clustering coefficient are af-
 311 fected as the network is progressively coarsened by matching algorithms $RGMb$ and GMb , with input parameters set
 312 as $L = 10$ and $rf = 0.5$. Figure 7 shows the reduction factor on $|V| = n$ and $|E| = m$ at each level of the coarsening
 313 hierarchy, i.e., starting from the original network (level 0) to the coarsest one (level 10). The curves depict the
 314 percentages of remaining vertices (left) and edges (right) relative to the initial numbers.

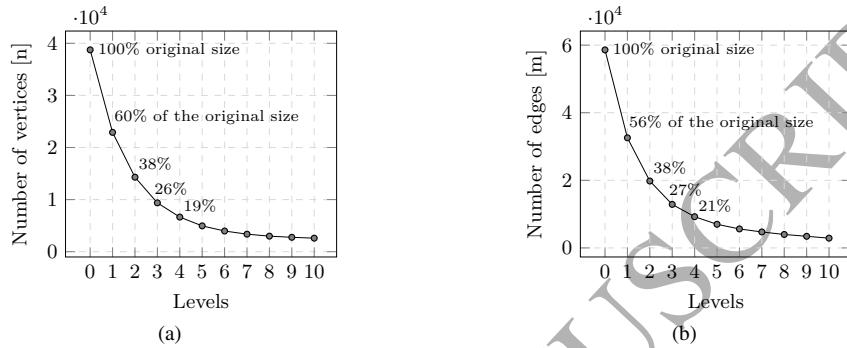


Figure 7. Percentages of remaining vertices (left) and edges (right) in the networks along the coarsening hierarchy.

314 Figures 8 and 9 show curves for degree distribution and clustering coefficient, respectively, in each coarsened
 315 network, from levels 1 to 10 (the top two rows refer to $RGMb$ and the bottom two rows refer to GMb). Both algorithms
 316 yielded very similar distributions of degree and clustering coefficient, which suggests the random exploration of
 317 the solution space adopted by $RGMb$ has no strong impact on the topological features of intermediate networks, in
 318 comparison with GMb .

320 The characteristic behavior of degree distribution observed in the original network is reasonably preserved in the
 321 coarsened models down to level 3, i.e., coarsened networks at levels 1, 2 and 3 still contain few hubs and many
 322 low-degree vertices, a behavior that gradually changes from level 4 onwards.

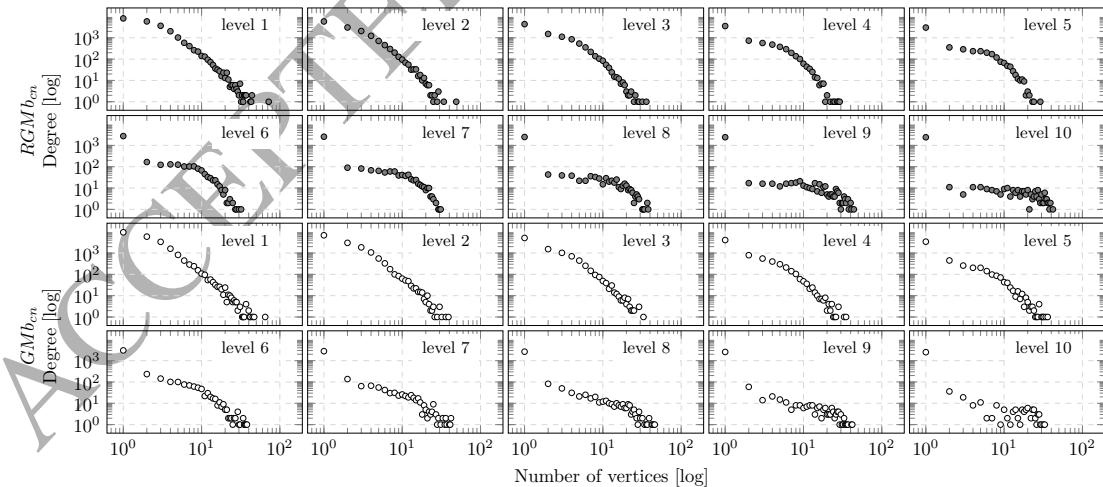


Figure 8. Degree distribution of the networks at the 10 levels of the coarsening hierarchy ($rf = 0.5$). Top graphs refer to coarsening with $RGMb$ (random strategy) and bottom graphs refer to coarsening with GMb .

322 Particularly from level 8, vertex degrees become more homogeneous, until the original topological features have
 323 been completely lost in the final network at level 10. A similar pattern is observed in the clustering coefficient, i.e.,

324 again, the original behavior of neighborhoods of hub authors is preserved up to level 3. From level 5 onwards, most
 325 vertices converge to lower clustering coefficients and the network's characteristic behavior is completely lost from
 326 level 8.

327 At the initial coarsening levels (1, 2 and 3) only vertices with many two-hop common neighbors can compose a
 328 matching. The newly formed super-vertices thus preserve the neighborhood properties of their predecessors, hence,
 329 the dominant topological properties of the parent network. In other words, it is likely that authors matched at the
 330 early coarsening levels indeed have many common collaborators. However, at later levels, authors with few common
 331 collaborators may be forced to join, which corrupts the original topological relations in the network. Notice the sizes
 332 of networks at levels 9 and 10 were reduced to nearly 10% of the original network (see Figure 7), which implies
 333 the coarsest networks are mostly formed by heavy-weight super-vertices ($\gg \sigma(sV)$) sparsely connected. Such super-
 334 vertices of low-degree form clustered structures with few triangles.

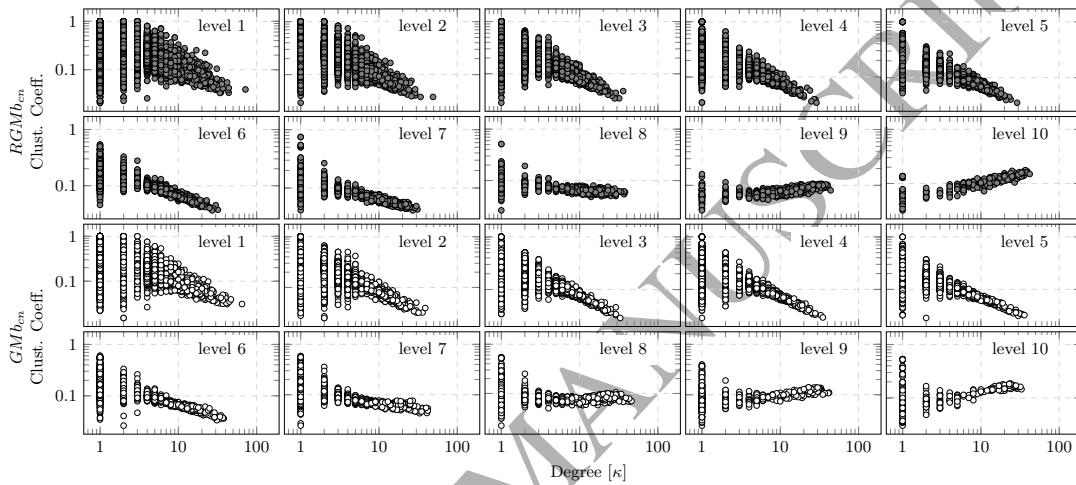


Figure 9. Two-mode local clustering coefficient of the networks at the 10 levels of the coarsening hierarchy ($rf = 0.5$). Top graphs refer to coarsening with $RGMB$ (random strategy) and bottom graphs refer to coarsening with GMb .

335 Results from this analysis indicate a limited reduction of an initial network, up to two or three levels, can preserve
 336 its relevant topological features. Furthermore, it is evident that choosing the appropriate coarsening level is critical,
 337 and depends on properties of the target application and dataset. Establishing a suitable trade-off between accuracy
 338 and runtime may require empirical verification in each case. This inherent limitation can be associated to the well-
 339 known overfitting/underfitting problem; however, the coarsening phase in our multilevel process “generalizes” the data
 340 instead of an objective function. This may explain why, in some cases, better solutions were obtained on the coarsened
 341 networks. In general, extensive coarsening reduces execution times of a target algorithm, but it can lead to excessive
 342 generalization of the data with significant degradation of topological features, and possibly algorithm accuracy. In
 343 contrast, limited coarsening preserves topological features and accuracy, at the expense of higher execution times.

344 5.2. Performance in benchmark networks for community detection

345 Research efforts on multilevel optimization have been strongly motivated by community detection (or graph parti-
 346 tioning) problems, which makes this a benchmark problem. Algorithms for community detection in networks split the
 347 vertices into disjoint groups (or communities), so as to minimize the number of edges between distinct communities
 348 [17]. Barber's modularity optimization [4] is often employed to identify community structures in bipartite networks.
 349 Formally, it quantifies the extent of communities formed in both layers relative to a null bipartite network model.
 350 Beckett [7] introduced the *LPAwb+* algorithm⁴, which maximizes Barber's modularity through label propagation
 351 in weighted bipartite networks and showed it has competitive performance compared with state-of-the-art methods.
 352 However, it is a computationally costly algorithm and becomes unfeasible in large-scale networks.

⁴available in <https://github.com/sjbeckett/weighted-modularity-LPAwbPLUS>.

In order to address the second research question, the *MOb* framework was tested considering Beckett's algorithm *LPAwb+* as a target algorithm. Therefore, *MOb* (Algorithm 4.1) performs the coarsening, runs *LPAwb+* to find the community structure in the coarsest network, and projects the solution to obtain the community structure in the original network. Our goal was to investigate whether *MOb* can yield solutions statistically equivalent in quality to the standard *LPAwb+*, whilst increasing its scalability. Results are compared with those obtained with the standard *LPAwb+*, used as baseline.

We investigated the performance of the four instances of *MOb*, listed in Table 1. They were executed with parameter settings $rf = 0.5$ and $L = [1, 2, 3]$ in a set of 15 synthetic weighted bipartite networks, identified as R1-R15 (hereafter each *MOb* instance is referred to by its name).

Algorithm	Coarsening		Solution finding	Uncoarsening
	Matching	Contracting		
<i>MOb-RGMb_{cn}</i>	<i>RGMb</i> with S_{cn}	<i>Cb</i>	<i>LPAwb+</i>	<i>Ub</i>
<i>MOb-RGMb_{wcn}</i>	<i>RGMb</i> with S_{wcn}			
<i>MOb-GMb_{cn}</i>	<i>GMb</i> with S_{cn}			
<i>MOb-GMb_{wcn}</i>	<i>GMb</i> with S_{wcn}			

Table 1. *MOb* applied to target algorithm *LPAwb+* (for the community detection problem): there are two choices of matching algorithm (*RGMb* and *GMb*) employed with two choices of similarity function (S_{cn} and S_{wcn}). All *MOb* algorithms were executed with $rf = 0.5$ and $L = [1, 2, 3]$.

Synthetic networks were obtained with the community model described in [7], which creates unbalanced and randomly positioned community structures. Networks of sizes $n = |V_1 + V_2|$ were generated within the range [1,000; 15,000] at increments of 1,000 and the number of communities was set to $0.01 * n$. Edge weights were randomly assigned from a skewed negative binomial distribution and noise was introduced in the connection patterns by reconnecting a percentage of the edges between and within communities.

Performance was measured in terms of accuracy, by means of the NMI (normalized mutual information), which compares the solution found by a selected algorithm with the baseline [30], we also measured execution times. Table 2 shows the NMI accuracy values in the 15 networks. The highest values are in bold and values equal to or higher than baseline *LPAwb+* are highlighted with a gray background. The best performances were achieved by *MOb-GMb_{cn}* with one level of coarsening ($L = 1$) on 11 out of the 15 networks. Baseline *LPAwb+* yielded the best performance in three networks, whereas *MOb-RGMb_{cn}* with $L = 3$ yielded the worst results.

Algorithm		Dataset														
Name	Levels [L]	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15
<i>LPAwb+</i>	0	0.918	0.926	0.983	0.972	0.964	0.990	0.984	0.999	0.999	0.985	0.989	0.996	0.995	0.987	0.992
<i>MOb-RGMb_{cn}</i>	1	0.984	0.985	0.983	0.988	0.990	0.991	0.992	0.991	0.991	0.992	0.992	0.992	0.991	0.991	0.992
<i>MOb-RGMb_{cn}</i>	2	0.952	0.968	0.963	0.977	0.977	0.976	0.978	0.978	0.978	0.977	0.979	0.981	0.982	0.982	0.980
<i>MOb-RGMb_{cn}</i>	3	0.866	0.910	0.923	0.936	0.943	0.944	0.952	0.954	0.950	0.957	0.954	0.957	0.956	0.955	0.957
<i>MOb-RGMb_{wcn}</i>	1	0.982	0.988	0.986	0.988	0.990	0.990	0.991	0.987	0.992	0.992	0.991	0.991	0.992	0.992	0.992
<i>MOb-RGMb_{wcn}</i>	2	0.942	0.960	0.966	0.968	0.974	0.973	0.978	0.978	0.977	0.978	0.977	0.977	0.979	0.979	0.978
<i>MOb-RGMb_{wcn}</i>	3	0.905	0.922	0.948	0.954	0.954	0.951	0.953	0.951	0.960	0.959	0.958	0.961	0.959	0.961	0.960
<i>MOb-GMb_{cn}</i>	1	0.994	0.993	0.993	0.995	0.996	0.996	0.995	0.993	0.994	0.995	0.995	0.995	0.996	0.995	0.996
<i>MOb-GMb_{cn}</i>	2	0.981	0.986	0.982	0.983	0.988	0.987	0.988	0.987	0.988	0.988	0.988	0.988	0.988	0.990	0.990
<i>MOb-GMb_{cn}</i>	3	0.901	0.934	0.963	0.966	0.971	0.968	0.973	0.974	0.974	0.975	0.972	0.975	0.976	0.975	0.977
<i>MOb-GMb_{wcn}</i>	1	0.990	0.989	0.992	0.992	0.995	0.993	0.992	0.992	0.993	0.994	0.995	0.995	0.995	0.994	0.994
<i>MOb-GMb_{wcn}</i>	2	0.969	0.985	0.985	0.979	0.988	0.985	0.987	0.985	0.987	0.987	0.989	0.988	0.988	0.989	0.989
<i>MOb-GMb_{wcn}</i>	3	0.965	0.973	0.973	0.972	0.975	0.979	0.977	0.977	0.978	0.977	0.978	0.980	0.978	0.979	0.979

Table 2. NMI accuracy values of all *MOb* instances and baseline *LPAwb+* in 15 synthetic networks. The values are average results over 30 executions. The highest accuracy values are in bold and values equal to or higher than the baseline (*LPAwb+*) are highlighted (gray background).

Limited coarsening levels (mainly $L = 1$) yielded higher accuracy values, reinforcing that a controlled coarsening can filter the solution space by joining promising vertex pairs and removing irrelevant high-cost solutions, while

preserving important topological features. In contrast, accuracy deteriorated with more extensive coarsening ($L = 3$), which does not necessarily preserve the original topological properties and is likely to blur the boundaries between adjacent communities. The effect of parameter L depends on network size, i.e.; differences in algorithm accuracy are likely to decrease as the network sizes increase, hinting that higher values of L might be successfully adopted in handling larger networks.

Although *LPAwb+* achieved the best performance in three networks out of the fifteen, the corresponding accuracy values attained by the *MOb* instances are very close for these networks. Indeed, in these specific cases accuracy values differ up to 0.006 in R8, up to 0.005 in R9 and up to 0.001 in R12. Interestingly, all *MOb* instances yielded similar accuracy values and more stable results than the standard *LPAwb+*. The accuracy values obtained with *MOb-GMb_{cn}* and $L = 1$, for instance, are within the range [0.994, 0.996], whereas for the standard *LPAwb+* they are within the range [0.918, 0.999], as shown in Figures 10 and 11.

Figure 10 depicts the averages and standard deviations of the accuracy values, whereas Figure 11 shows the dispersion of their distribution and outliers, considering in both cases the alternative settings of parameter L . The bar plots in Figure 10(a) reveal superior performance and stability of the four *MOb* instances when $L = 1$, confirmed by their higher average accuracies and narrower standard deviations. When $L = 2$, only *MOb-GMb_{cn}* and *MOb-GMb_{wcn}* yielded a slightly superior solution in comparison to *LPAwb+*. For $L = 3$, *LPAwb+* yielded better results than any *MOb* instance, a consequence of the extensive network reduction.

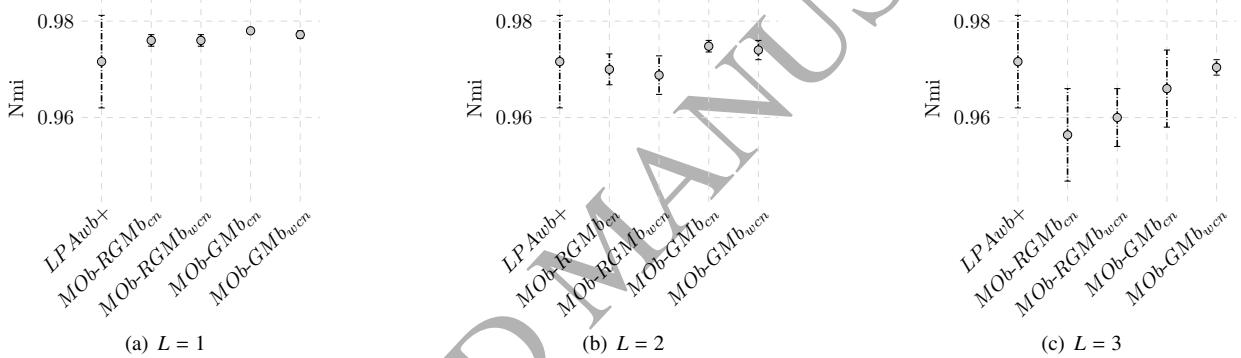


Figure 10. Averages and standard deviations of the NMI accuracy values obtained with *LPAwb+* and four *MOb* instances in three settings of parameter L (number of levels). (a) $L = 1$, 10(b) $L = 2$ and 10(c) $L = 3$.

The box plots in Figure 11(a) reveal that for $L = 1$ all *MOb* instances yielded accuracy values within a narrower distribution and higher averages than *LPAwb+*.

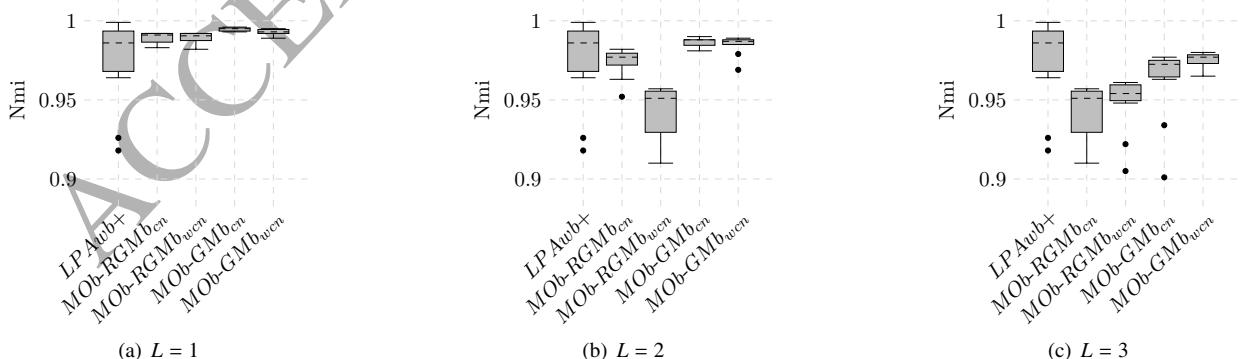


Figure 11. Shape distribution, variability, and median of the accuracy values yielded by *LPAwb+* and the four instances of *MOb* considering three settings of parameter L . (a) $L = 1$, 10(b) $L = 2$ and 10(c) $L = 3$.

394 We can conclude *MOb* instances yielded, in general, higher accuracy and improved stability in comparison to
 395 *LPAwb+*. In summary, the experimental evidence regarding solution quality (average, standard deviation and dis-
 396 persion of the accuracy values) suggests the multilevel framework stabilizes and improves the performance of the
 397 algorithm.

398 A Nemenyi post-hoc test [13] was applied to the results in Table 2 to verify statistical differences in the algorithms
 399 performances and the results are shown in Figure 12 for (a) $L = 1$, (b) $L = 2$ and (c) $L = 3$. The critical difference
 400 (CD) is indicated at the top of each diagram and the algorithms' average ranks are placed on the horizontal axes,
 401 with the best ranked algorithms to the left. A black line connects algorithms if no significant difference has been
 402 detected among them. According to the the Nemenyi statistics, the critical value for comparing the mean-ranking of
 403 two different algorithms at 95 percentile is 1.58.

404 Let us consider the outcome of the post-hoc test for $L = 1$, i.e., when the number of vertices n is reduced by a
 405 factor of two, shown in Figure 12(a). *MOb-GMB_{cn}* was ranked best, followed by *MOb-GMB_{wcn}* and *MOb-RGMb_{wcn}*
 406 and then *LPAwb+*. Furthermore, *MOb-GMB_{cn}*, *MOb-GMB_{wcn}* and *MOb-RGMb_{wcn}* presented statistically significant
 407 differences compared with standard *LPAwb+*. Interestingly, for $L = 2$ (Figure 12(b)), *MOb-GMB_{cn}* and *MOb-*
 408 *GMB_{wcn}* remain ranked first, however, no statistically significant difference was observed in relation to *LPAwb+*.
 409 Finally, for $L = 3$ (Figure 12(c)), *LPAwb+* was ranked first, with no statistically significant difference observed in
 410 relation to *MOb-GMB_{cn}* or *MOb-GMB_{wcn}*. However, parameter settings $L = 3$ and $rf = 0.5$ implied, in this case, in
 411 reducing the original size by a factor of 75%, which explains the poor performance of all *MOb* instances.

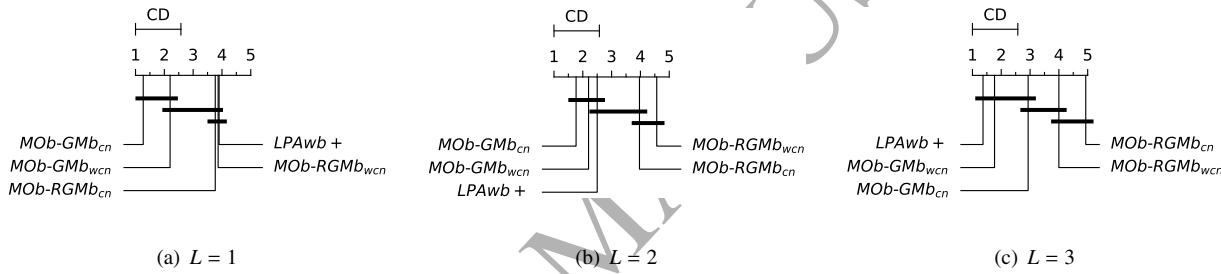


Figure 12. Nemenyi post-hoc test for *LPAwb+* and the four *MOb* instances.

412 We also assessed the scalability of the *MOb* instances to investigate the third and fourth research questions. Their
 413 performance was analyzed considering each individual network and the total time spent in the experiments. Table 3
 414 shows the absolute execution times (in seconds) on each network - values refer to the average times from 30 executions
 415 of each scenario.

Algorithm		Dataset															sum
Name	Levels [L]	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	
<i>LPAwb+</i>	0	14	96	308	904	2,782	2,800	7,146	15,925	39,197	56,119	66,729	75,990	97,392	224,032	302,442	891,875
<i>MOb-RGMb_{cn}</i>	1	2	10	31	77	160	328	616	1,106	1,787	4,148	5,386	8,108	10,242	15,235	23,174	70,410
<i>MOb-RGMb_{cn}</i>	2	1	3	8	19	42	77	154	267	428	1,081	1,296	1,989	2,372	3,559	5,283	16,577
<i>MOb-RGMb_{cn}</i>	3	1	1	4	7	14	25	49	86	142	304	389	586	694	940	1,310	4,552
<i>MOb-RGMb_{wcn}</i>	1	3	11	34	79	174	322	643	1,424	1,715	4,192	5,335	8,582	9,355	15,121	22,535	69,525
<i>MOb-RGMb_{wcn}</i>	2	1	4	9	24	44	101	165	282	460	1,049	1,399	1,911	2,245	3,281	4,866	15,840
<i>MOb-RGMb_{wcn}</i>	3	1	3	6	11	18	33	59	91	139	323	396	522	713	1,056	1,381	4,751
<i>MOb-GMB_{cn}</i>	1	2	13	33	78	165	329	623	1,111	2,680	4,145	5,561	6,705	9,091	18,136	24,732	73,402
<i>MOb-GMB_{cn}</i>	2	1	5	11	24	46	84	168	298	721	1,088	1,492	1,657	2,455	4,013	5,555	17,619
<i>MOb-GMB_{cn}</i>	3	1	3	6	13	22	34	62	90	239	332	471	481	672	1,180	1,596	5,204
<i>MOb-GMB_{wcn}</i>	1	3	13	39	83	180	338	688	1,135	2,864	4,392	5,819	6,317	8,932	19,016	25,674	75,493
<i>MOb-GMB_{wcn}</i>	2	3	8	15	32	55	96	187	326	744	1,100	1,447	1,719	2,345	4,050	4,574	16,700
<i>MOb-GMB_{wcn}</i>	3	3	7	12	20	31	45	80	117	287	392	469	536	745	1,233	1,380	5,357
sum		36	178	514	1,371	3,733	4,613	10,639	22,257	51,403	78,664	96,189	115,104	147,253	310,853	424,502	1,267,307

Table 3. Absolute runtime (seconds) of *LPAwb+* and four *MOb* instances on each network.

416 The longest execution time of standard *LPAwb+* was 302,442 seconds (time to process the largest network) and
 418 the shortest was 14 seconds (time to process the smallest one). *MOb-GMb_{wcn}* was the most expensive *MOb* instance,
 420 consuming ($L = 1$) 25,674 seconds on the largest network and 3 seconds on the smallest one. Therefore, regarding its
 422 maximum and minimum execution times, respectively, *MOb-GMb_{wcn}* run 11.8 to 4.6 times faster than the standard
 424 *LAPwb+*. The maximum and minimum running times of the least expensive *MOb* instance, *MOb-RGMb_{cn}* ($L = 3$),
 426 were 1,310 seconds and 1 second, respectively. Therefore, *MOb-RGMb_{cn}* run 230 to 14 times faster than *LAPwb+*.

422 The analysis in Section 5.1 revealed that a network coarsened at level 3 has roughly 25% of its original size. Let
 424 us consider, for example, network R15 ($n = 15,000$): we know it has been reduced to 3,750 vertices at level 3 and
 426 algorithm *MOb-RGMb_{cn}* processed it in 1,310 seconds (see Table 3). This is close to the execution time of standard
 428 *LPAwb+* on network R4 (of size $n = 4,000$, similar to the size of R15 coarsened at level 3), i.e., 904 seconds. As
 430 *MOb-RGMb_{cn}* executes the coarsening/uncoarsening steps, the actual time spent running *LPAwb+* to find the solution
 432 is roughly similar in both cases, but *MOb-RGMb_{cn}* is handling a network nearly four times larger.

428 The total time spent running the experiments was 1,267.307 seconds, or nearly 352 hours. Figure 13 shows the
 430 contribution of each algorithm to the total time, considering both absolute values (seconds) (a) and relative values
 432 (percentages) (b).

434 In the best case, the *MOb* instances reduced execution time from 891,875 seconds (nearly 208.8 hours) required by
 432 the standard *LPAwb+* to 4,552.4 seconds (1.26 hours), which implies *LPAwb+* was nearly 195 times slower than its
 434 *MOb* instantiations. Executing *LPAwb+* consumed over 70% of the time spent in the experiments, whereas roughly
 6% of the time was spent running the *MOb* instances.

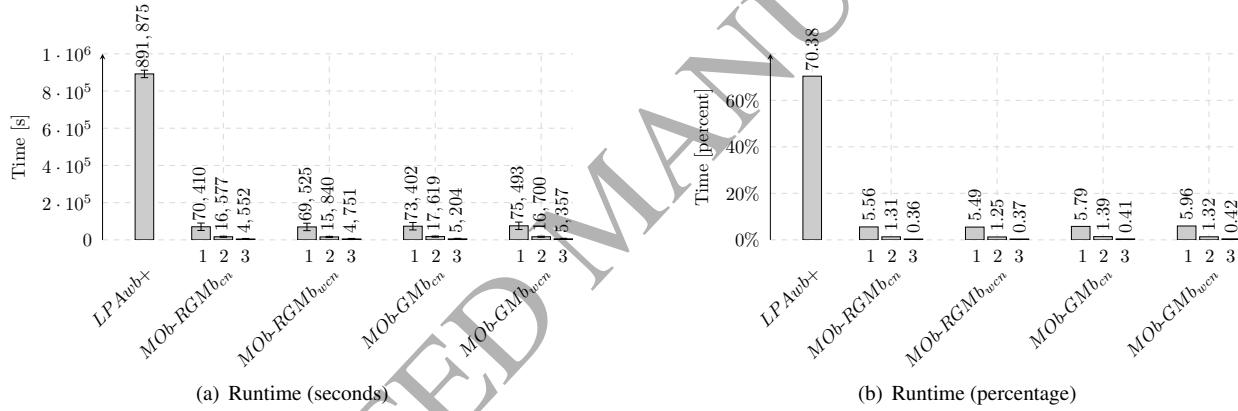


Figure 13. Contribution of each algorithm to the total time of the experiments: absolute values (seconds) (a) and percentages (b), 100% runtime corresponds to 1,127.300 seconds. Each bar refers to a *MOb* instance executed with a particular setting of parameter $L = [1, 2, 3]$, with the specific execution time and standard deviations shown on top in (a) and relative percent times shown on top in (b).

436 Finally, we assessed the impact of each phase (coarsening, target algorithm on coarsest network and uncoarsening)
 438 on the execution time of the multilevel process; we analyzed algorithm behavior separately on each network and then
 440 in relation to the total time of the experiments.

442 The relative contributions of each multilevel phase for each network are shown in Figure 14 (for legibility, we
 444 show bars for 12 out of the 15 networks). On the smallest network ($n = 1,000$) the coarsening phase consumed nearly
 446 52% of the total execution time and the local search step consumed nearly 45%. On the other hand, as the networks
 448 increase, the time spent on the coarsening relative to the solution finding gradually decreases. On the larger networks,
 the coarsening phase consumed roughly 1% of the total execution time, in contrast to roughly 99% of the solution
 finding phase. The time spent on the uncoarsening phase was negligible.

444 Figure 15 shows the relative contribution of each multilevel phase to the total time of the experiments. In general,
 446 the coarsening phase consumed less than 1% of the total execution time, the solution finding phase (executing
 448 *LPAwb+*) consumed over 99%, and the time spent in the uncoarsening phase was negligible. These results indicate
 coarsening and uncoarsening exerted no significant influence on the scalability of the multilevel process and provide
 empirical evidence the multilevel strategy is a promising approach to scale network optimization algorithms.

448 From this empirical investigation we conclude: (i) the proposed *MOb* approach yielded more accurate and stable

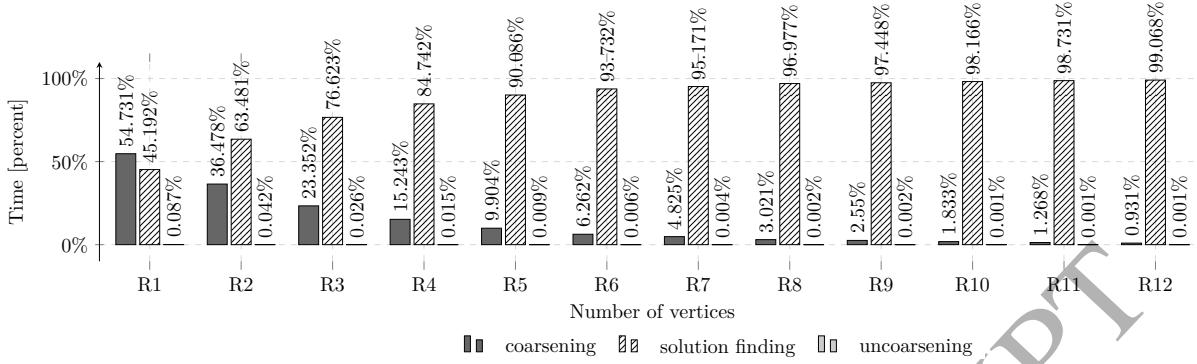


Figure 14. Relative contribution (percentage) of each multilevel phase to the execution time, on different network sizes n (from 1,000 for R1 to 12,000 for R12).

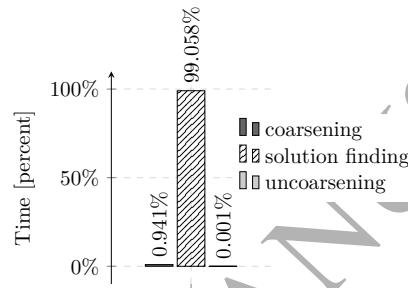


Figure 15. Average relative contribution (percentage of execution time) of each stage of MOB in all 15 networks.

450 results compared to the standard *LPAwb+*; (ii) although solution quality degrades as the network is progressively
451 coarsened, runtime drops drastically at each additional coarsening level; hence, a successful solution requires estab-
452 lishing a suitable trade-off between accuracy and execution time.

5.3. A test case on dimensionality reduction

454 We illustrate how *MOB* framework can be adapted to perform dimensionality reduction in the context of text clas-
455 sification, having the k -Nearest Neighbor classifier (kNN) as the target algorithm, in order to exemplify its application
456 in a different kind of optimization problem.

458 Documents are often represented, in text classification tasks, as multidimensional feature vectors, in which each
459 dimension maps a particular term. As the dimensionality of the representation space has strong impact in classification
460 performance, such tasks are often preceded by a dimension reduction step. Specifically, a kNN classifier that employs
461 a naïve search strategy has time complexity $O(ndk)$ for a fixed k , where n is the cardinality of the training set and d
462 denotes the dimensionality of the document representation.

463 Alternatively, a document corpus can be represented as a bipartite network $G = ((V_1 \cup V_2), E, \sigma, \omega)$, where $V_1 = \{d_1, \dots, d_r\}$ is the set of documents and $V_2 = \{t_1, \dots, t_s\}$ is the set of terms. An edge (u, v) exists if term t_u occurs
464 in document d_v , and the term's frequency determines the corresponding edge weight $w(u, v)$. Such a network is
465 represented as a bi-adjacency matrix $A_{r \times s}$, where $r = |V_1|$, $s = |V_2|$ and $A_{u,v} = w(u, v)$ if $(u, v) \in E$. Dimensionality
466 reduction is aimed at obtaining a lower dimensional matrix $A'_{r' \times s'}$, with $r' = r$ and $s' \ll s$.

467 The proposed solution is described in Algorithm 5.1, **multilevel dimensionality reduction** (*Mdr*), which only re-
468 quires a coarsening phase (lines 1-3). It takes as inputs the initial bipartite network G , the term layer t_l to be coarsened,
469 the desired maximum number of levels L and reduction factor rf , and returns the bi-adjacency matrix of the coarsest
470 network (line 4).

471 We adopted *RGMb_{cn}* as the contracting algorithm in the *Mdr* implementation, and used *Mdr* in combination with
472 a kNN classifier, assuming $k = 3$ and the Euclidean distance as similarity measure, hereafter referred to as *Mdr-kNN*.

The results obtained considering a reduction factor $rf = 0.5$ and L varying in the range [1, 2, 3, 4, 5, 6, 7, 8, 9, 10] were compared to an equivalent kNN setting that employed PCA (Principal Components Analysis) [20] for dimensionality reduction ($PCA-kNN$). The complexity of PCA is $O(d^2n + d^3)$, where n denotes the number of samples and d the data dimensionality.

Algorithm 5.1: *Mdr: multilevel dimensionality reduction.*

Input:

bipartite network : $G = (V, E, \sigma, \omega)$
 term layer : $t_l \in \{1, 2\}$
 maximal number of levels for term layer : $L \in [0, n] \subset \mathbb{Z}$
 reduction factor for term layer : $rf \in (0, 0.5] \subset \mathbb{R}$

Output:

low-dimensional bi-adjacency matrix : A'

1 **while** $l \leq L$ **or** layer l is as small as desired **do**
 2 $M \leftarrow \text{matching}(G_l, t_l, rf);$
 3 $G_{l+1} \leftarrow \text{contracting}(G_l, M);$
 4 matrix $A' \leftarrow \text{bi-adjacency matrix of } G_l;$

Return: A'

Implementations $Mdr-kNN$ and $PCA-kNN$ were used to classify thirteen real document-term networks available from the literature, described in Table 4, considering a cross-validation with permutation testing and ten-fold cross-validation to estimate classification error. The training set was randomly split into ten equal-sized subsets, so the classification model was trained in nine subsets and tested on the remaining one.

	cstr	oh0	oh5	0h10	syskillweber	opinosis	classic4	hitech	industry	irish	la1s	enron	acm
Documents	299	1,003	918	1,050	334	6,457	7,095	2,301	8,817	1,660	3,204	13,199	3,493
Terms	1,726	3,183	3,013	3,239	4,340	2,693	7,749	10,000	10,000	8,659	10,000	10,000	10,000
Domain	Scientific	Medical	Medical	Medical	Web pages	Sentiment	Scientific	News	Web page	Sentiment	News	E-mail	Scientific

Table 4. Properties of text collections considered.

Figure 16 shows the performance of the $Mdr-kNN$ classifier as a function of the number of levels in the coarsening hierarchy, where $L = 0$ refers to classification with no dimensionality reduction. Figures 16(a), (b) and (c) show accuracy, dimensionality of the representation (number of terms) and execution times, respectively.

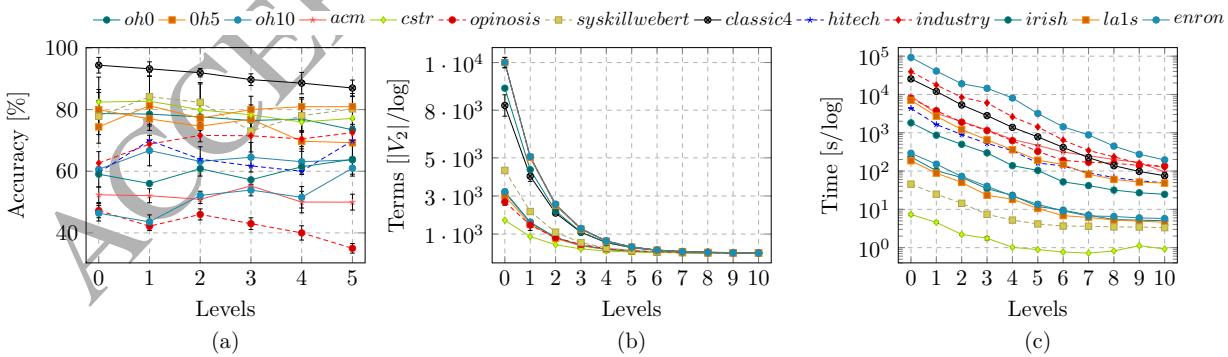


Figure 16. Performance (accuracy values) of the $Mdr-kNN$ classifier. $L = 0$ corresponds to classification with no dimension reduction applied.

A moderate decrease in accuracy values is observed up to level $L = 5$. In general, accuracy is stable at the early coarsening levels, which suggests the first coarsening iterations have limited impact in solution quality, as it is more

likely that highly correlated terms are being matched. Figure 16(a) shows accuracies up to level 5. Moreover, since from level 5 onwards the coarsening yields no significant reduction in the number of terms it may be interrupted at this point. Also, a sharp reduction in execution times is observed only up to this level. We remind the analysis presented in Section 5.1, which showed a network coarsened at level 5 has roughly 15% its original size.

We compared the classification accuracies of kNN without dimension reduction and $Mdr-kNN$ (with coarsening at $L = [1, 3, 5]$) and those of $PCA-kNN$ (with equivalent dimensionality reduction). The results are shown in Table 5. The highest accuracy values are shown in bold and values equivalent or superior to baseline kNN are shown with a gray background. $Mdr-kNN$ achieved the best performance in five out of the thirteen networks; $PCA-kNN$ performed best in seven networks and standard kNN performed best in only one of them.

Dataset	kNN	$L = 1$		$L = 3$		$L = 5$	
		$MdrkNN$	$PCA-kNN$	$Mdr-kNN$	$PCA-kNN$	$Mdr-kNN$	$PCA-kNN$
cstr	82.44 \pm 8.09	82.71 \pm 8.04	75.86 \pm 7.15	80.10 \pm 6.10	79.31 \pm 2.02	78.30 \pm 3.20	86.21 \pm 7.10
oh0	78.72 \pm 6.84	81.15 \pm 3.31	85.00 \pm 3.98	78.95 \pm 2.02	76.00 \pm 4.31	75.90 \pm 3.32	76.00 \pm 6.20
oh5	80.27 \pm 5.09	78.83 \pm 3.73	76.92 \pm 4.31	76.88 \pm 4.17	81.32 \pm 3.11	69.27 \pm 3.89	74.36 \pm 5.35
oh10	60.62 \pm 5.75	66.92 \pm 4.95	75.24 \pm 4.21	64.54 \pm 5.04	72.38 \pm 3.29	64.38 \pm 2.16	73.33 \pm 4.23
opinosis	47.25 \pm 2.32	42.18 \pm 1.42	43.57 \pm 1.20	45.90 \pm 4.92	55.19 \pm 8.15	35.98 \pm 3.06	49.15 \pm 5.78
syskillwebert	77.80 \pm 8.70	87.91 \pm 6.57	87.88 \pm 5.14	75.91 \pm 3.65	75.06 \pm 6.45	80.23 \pm 2.12	78.00 \pm 5.65
acm	52.37 \pm 2.54	52.96 \pm 2.26	54.44 \pm 2.62	55.23 \pm 3.15	53.00 \pm 2.98	49.98 \pm 1.87	49.10 \pm 4.87
hitech	59.57 \pm 3.55	70.00 \pm 2.59	64.35 \pm 3.24	61.74 \pm 2.34	66.10 \pm 3.10	69.90 \pm 2.55	66.90 \pm 3.50
industry	62.66 \pm 5.60	71.62 \pm 3.05	66.86 \pm 1.98	71.50 \pm 6.12	78.30 \pm 4.12	72.89 \pm 3.34	75.20 \pm 4.32
irish	59.04 \pm 4.28	63.86 \pm 2.34	57.83 \pm 4.32	58.23 \pm 3.87	56.90 \pm 4.07	63.91 \pm 5.12	63.10 \pm 5.98
la1s	74.38 \pm 6.74	81.25 \pm 4.10	83.75 \pm 5.20	80.00 \pm 4.12	82.12 \pm 2.23	80.94 \pm 4.80	82.30 \pm 4.76
classic4	94.39 \pm 6.04	93.82 \pm 2.55	93.51 \pm 3.15	89.72 \pm 4.01	75.87 \pm 2.68	87.00 \pm 3.12	91.90 \pm 3.28
enron	46.40 \pm 5.18	52.16 \pm 5.09	47.76 \pm 2.98	53.89 \pm 2.90	55.00 \pm 3.77	61.00 \pm 2.14	60.10 \pm 6.12

Table 5. Accuracy values obtained with standard kNN , $Mdr-kNN$ and $PCA-kNN$. The highest accuracy values are shown in bold and values equivalent or superior to the baseline kNN are highlighted in gray.

A Nemenyi post-hoc test was applied to the results in Table 5 in order to detect statistical differences among the algorithms. Demsar post-hoc test requires each algorithm and dataset to be independent, therefore, we perform the dimensionality reduction at each level separately. The results are shown in Figure 17 for (a) $L = 1$, (b) $L = 2$ and (c) $L = 3$. According to Nemenyi statistics, in all diagrams, the critical value for comparing the mean-ranking of two different algorithms at 95 percentile is 0.92. No significant difference was observed between the algorithms, therefore, they are connected by a bold line in each diagram. Albeit differences are not significant, Figure 17(a) shows $Mdr-kNN$ was ranked best for $L = 1$, whereas for $L = 2$ and $L = 3$ $PCA-kNN$ was ranked best (Figures 17(b) and 17(c)). Therefore, $Mdr-kNN$ proved competitive in terms of accuracy, in comparison to $PCA-kNN$ and baseline kNN .

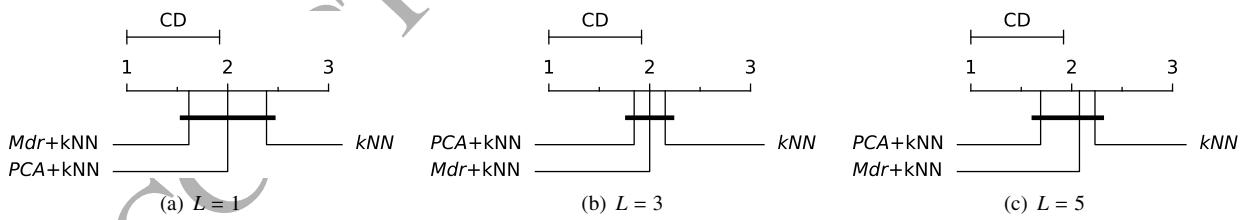


Figure 17. Nemenyi post-hoc test for kNN and its variant with Mdr and PCA for dimensionality reduction in three settings of parameter L . (a) $L = 1$, (b) $L = 3$ and (c) $L = 5$.

This case study has been presented as a preliminary investigation on the feasibility of extending the proposed multilevel framework to other combinatorial problems beyond community detection. The Mdr algorithm deserves further consideration and could incorporate additional capabilities, e.g., it would be convenient to reduce the feature space to a target dimensionality, rather than by a given reduction factor. Furthermore, the coarsening algorithm could take into account intrinsic characteristics of specific kinds of document-term networks by means of customized matching algorithms. Moreover, Mdr can be employed in connection with other classification algorithms.

6. The application of *MOb* to other combinatorial optimization problems

510 It is relatively straightforward to apply *MOb* to several combinatorial optimization problems beyond community
 511 detection. For instance, in Section 5.3 we illustrated its application to handle a dimensionality reduction problem over
 512 a bipartite network in which the two vertex layers represent objects and features, respectively. In this context, only
 513 the feature layer was coarsened and the reduced feature space is given directly by the adjacency matrix of the coarsest
 514 network. Likewise, it can be easily instantiated to handle overlapping or fuzzy community detection or classification
 515 problems.

516 In overlapping community detection, for each decomposed $sV \in V_{l+1}$, its original vertices $\{u, v\} \in V_l$ are assigned
 517 to the same set of communities as the corresponding super-vertex. If the coarsened network has a fuzzy structure, the
 518 strength of a vertex' pertinence to a community will be equal to that of its super-vertex. Similarly, in a classification
 519 problem original vertices $\{u, v\} \in V_l$ should be assigned to the same class of their super-vertex. Therefore, instantiating
 520 the framework to handle either problem would require just minor modifications in the projection algorithm (Algorithm
 521 4.5).

522 *MOb* might also be useful to support interactive visualization of large-scale bipartite networks, by means of
 523 navigation over a hierarchy of coarsened networks, which would demand a data structure to keep these intermediate
 524 networks.

525 Instantiation to other scenarios is not necessarily as straightforward, and may require further modifications in the
 526 proposed algorithms. For example, in the edge clustering problem (also called link communities) the contracting
 527 algorithm used in coarsening phase would require adjustments. Whereas in community detection a vertex inherits the
 528 same group assignment of its super-vertex, here each edge must inherit the connections from its super-edge. Therefore,
 529 an edge $e \in E_{l+1}$ incident to $sV \in V_{l+1}$ must refer to the edges incident to vertices $\{u, v\} \in V_l$. Algorithm 4.4 does not
 530 implement this function; however, it could be done by keeping an additional data structure similar to the successor
 531 vector.

532 Another possible application is in link prediction or recommendation problems, albeit this poses a more complex
 533 scenario for generalization. Link prediction methods rely on similarity between vertices, since similar vertices are
 534 likely to share common links. However, such information is not explicitly given for the super-vertices in a coarsened
 535 bipartite network. Nonetheless, the framework could be employed to reduce the number of required operations, i.e.
 536 super-vertices might be created grouping vertices with shared h -hop neighbors, thus filtering the search space. Link
 537 prediction could be performed in the uncoarsening phase through the decomposed super-vertices, and the solution
 538 finding and uncoarsening phases would be executed simultaneously.

539 As a final consideration, the matching algorithm used for coarsening should be carefully designed to incorporate
 540 the specific characteristics of each problem and context.

7. Conclusion and further research

541 Inspired by the potential of general-purpose multilevel strategies to scale optimization algorithms we have in-
 542 troduced algorithms of a novel multilevel optimization framework (*MOb*) for bipartite networks, and illustrated its
 543 application on two combinatorial optimization problems. Our framework accounts for the specificities of bipartite
 544 networks and provides a powerful tool for handling a variety of problems.

545 We investigated three empirical scenarios to illustrate strengths and limitations of the proposed *MOb* framework.
 546 A first study has shown that a controlled coarsening preserves relevant topological features of a network. A second
 547 study described an application in community detection, showing that *MOb* combined with a proper local search
 548 strategy can drastically improve speedup of a classic community detection algorithm while preserving solution quality.
 549 Finally, in a third study we considered text classification to illustrate how the general framework can be instantiated
 550 to handle different combinatorial optimization problems.

551 Our results provide compelling evidence that *MOb* offers a competitive approach to scale existing methods while
 552 preserving solution quality, and reinforce its usefulness for handling combinatorial optimization problems in large
 553 bipartite networks. Furthermore, the framework is flexible and can be adjusted to incorporate alternative and novel
 554 coarsening methods targeted at specific applications. We also discuss some general guidelines for future applications
 555 of combinatorial problems, such as link prediction, edge clustering, and interactive graph visualization.

Identifying the level of coarsening that will yield a suitable trade-off between accuracy and execution times is a critical issue in applying the proposed multilevel strategy. Currently, this is done by means of empirical investigation in each application problem and dataset, but it certainly deserves further investigation.

We also plan as future work to extend the framework to handle problems defined in heterogeneous networks, where edges connect vertices of multiple types. It would be applicable, e.g., to document-word networks indicating associations of the type document-word, word-word, and document-document; or networks describing relations between words, documents and authors. We are also interested in investigating distributed or parallel paradigms, as well as in application of *MOB* to supervised and unsupervised classification tasks. An implementation of the general framework is currently available and can be downloaded from <https://github.com/alanvalejo/mob>.

566 Acknowledgments

Author A. Valejo is supported by a scholarship from the Brazilian Federal Agency for Support and Evaluation of Graduate Education (CAPES). This work has been partially supported by the State of São Paulo Research Foundation (FAPESP) grants 15/14228-9 and 17/05838-3; and the Brazilian Federal Research Council (CNPq) grants 302645/2015-2 and 305696/2013-0.

References

- [1] Abou-Rjeili, A., Karypis, G., 2006. Multilevel algorithms for partitioning power-law graphs. In: Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS). pp. 124–124.
- [2] Asratian, A. S., Denley, T. M. J., Häggkvist, R., 1998. Bipartite graphs and their applications. Cambridge University Press, New York, NY, USA.
- [3] Baños, R., Gil, C., Ortega, J., Montoya, F. G., 2004. A parallel multilevel metaheuristic for graph partitioning. *Journal of Heuristics* 10 (3), 315–336.
- [4] Barber, M. J., 2007. Modularity and community detection in bipartite networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 76 (6), 1–11.
- [5] Barnard, S. T., 1995. Pmrsb: Parallel multilevel recursive spectral bisection. In: Proceedings of the ACM/IEEE conference on Supercomputing (CDROM). pp. 1–27.
- [6] Barnard, S. T., Simon, H. D., 1994. Fast multilevel implementation of recursive spectral bisection for partitioning unstructured problems. *Concurrency: Practice and Experience* 6 (2), 101–117.
- [7] Beckett, S. J., 2016. Improved community detection in weighted bipartite networks. *Royal Society open science* 3 (1), 140536.
- [8] Bichot, C.-E., 2014. Graph Partitioning. ISTE - Wiley, Ch. A Partitioning requiring rapidity and quality: The multilevel method and partitions refinement algorithms, pp. 27–63.
- [9] Blum, C., Valles, M. Y., 2006. Multi-level ant colony optimization for dna sequencing by hybridization. In: Proceedings of the International Workshop on Hybrid Metaheuristics. Springer, pp. 94–109.
- [10] Brandt, A., 1988. Multilevel computations: Review and recent developments. In: Proceedings of the Multigrid Methods: theory, Applications and Supercomputing. Vol. 110. pp. 35–62.
- [11] Chernoskutov, M., Ineichen, Y., Bekas, C., 2015. Heuristic algorithm for approximation betweenness centrality using graph coarsening. *Procedia Computer Science* 66, 83–92.
- [12] Dai, C., Li, P. C., Toulouse, M., 2006. A cooperative multilevel tabu search algorithm for the covering design problem. In: Proceedings of the International Conference on Artificial Evolution. Vol. 3871 LNCS. Springer Berlin Heidelberg, pp. 119–130.
- [13] Demšar, J., 2006. Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research* 7, 1–30.
- [14] Djidjev, H. N., 2008. A scalable multilevel algorithm for graph clustering and community structure detection. In: Proceedings of the Algorithms and Models for the Web-Graph: Fourth International Workshop. pp. 117–128.
- [15] Djidjev, H. N., Onus, M., 2013. Scalable and accurate graph clustering and community structure detection. *IEEE Transactions on Parallel and Distributed Systems* 24 (5), 1022–1029.
- [16] Erciye, K., Alp, A., Marshall, G., 2005. Serial and parallel multilevel graph partitioning using fixed centers. In: Proceedings of the Current Trends in Theory and Practice of Computer Science. pp. 127–136.
- [17] Fortunato, S., 2010. Community detection in graphs. *Physics Reports* 486 (3-5), 75–174.
- [18] Han, J., Lu, J., Hu, Y., Zhang, G., 2015. Tri-level decision-making with multiple followers: Model, algorithm and case study. *Information Sciences* 311, 182 – 204.
- [19] Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N., Barabasi, A. L., 2000. The large-scale organization of metabolic networks. *Nature* 407 (6804), 651–654.
- [20] Jolliffe, I., 2002. Principal component analysis. Springer Series in Statistics. Springer.
- [21] Karpinski, M., Rytter, W., 1998. Fast parallel algorithms for graph matching problems. No. 9. Oxford University Press.
- [22] Karypis, G., Kumar, V., 1995. Analysis of multilevel graph partitioning. In: Proceedings of the ACM/IEEE Conference on Supercomputing. pp. 1–19.
- [23] Karypis, G., Kumar, V., 1996. Parallel multilevel graph partitioning. In: Proceedings of the International Conference on Parallel Processing. pp. 314–319.

[24] Karypis, G., Kumar, V., 1996. Parallel multilevel k-way partitioning scheme for irregular graphs. In: Proceedings of the ACM/IEEE Conference on Supercomputing.

[25] Karypis, G., Kumar, V., 1998. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on Scientific Computing (SISC)* 20 (1), 359–392.

[26] Karypis, G., Kumar, V., 1998. A parallel algorithm for multilevel graph partitioning and sparse matrix ordering. *Journal of Parallel and Distributed Computing* 48 (1), 71–95.

[27] Khan, A., Li, N., Yan, X., Guan, Z., Chakraborty, S., Tao, S., 2011. Neighborhood based fast graph search in large networks. In: Proceedings of the International Conference on Management of Data (SIGMOD). SIGMOD '11. ACM, New York, NY, USA, pp. 901–912.

[28] Kolllias, G., Sathe, M., Schenk, O., Grama, A., 2014. Fast parallel algorithms for graph similarity and matching. *Journal of Parallel and Distributed Computing* 74 (5), 2400 – 2410.

[29] Korosec, P., Sile, J., Robic, B., 2003. Mesh partitioning: A multilevel ant-colony-optimization algorithm. In: Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS). Vol. 00. pp. 1–8.

[30] Labatut, V., 2015. Generalized measures for the evaluation of community detection methods. *International Journal of Social Network Analysis and Mining (SNAM)* 2 (1), 44–63.

[31] Lasalle, D., Karypis, G., 2013. Multi-threaded graph partitioning. In: Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS). pp. 225–236.

[32] Lasalle, D., Karypis, G., 2015. Multi-threaded modularity based graph clustering using the multilevel paradigm. *Journal of Parallel and Distributed Computing* 76, 66–80.

[33] Latapy, M., Magnien, C., Vecchio, N. D., 2008. Basic notions for the analysis of large two-mode networks. *Social Networks* 30 (1), 31–48.

[34] Lu, J., Han, J., Hu, Y., Zhang, G., 2016. Multilevel decision-making: A survey. *Information Sciences* 346–347, 463 – 487.

[35] Lü, L., Zhou, T., 2010. Link prediction in weighted networks : The role of weak ties. *Europhysics Letters (EPL)* 89 (1), 18001.

[36] Newman, M. E. J., 2001. Scientific collaboration networks: I. network construction and fundamental results. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 64 (1), 1–8.

[37] Newman, M. E. J., 2001. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 64 (1), 016132.

[38] Newman, M. E. J., 2001. The structure of scientific collaboration networks. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 98 (2), 404–409.

[39] Newman, M. E. J., 2003. Mixing patterns in networks. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 67 (2), 026126.

[40] Newman, M. E. J., Jun, 2006. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)* 103 (23), 8577–8582.

[41] Noack, A., Rotta, R., 2009. Multi-level algorithms for modularity clustering. In: Vahrenhold, J. (Ed.), *Proceedings of the Experimental Algorithms: 8th International Symposium SEA*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 257–268.

[42] Oduntan, I. O., 2005. A multilevel search algorithm for feature selection in biomedical data. Ph.D. thesis, Dept. Computer Science, University of manitoba.

[43] Oliveira, S., Seok, S. C., 2006. A multilevel approach to identify functional modules in a yeast protein-protein interaction network. In: *Proceedings of the Computational Science (ICCS): 6th International Conference*, Reading. pp. 726–733.

[44] Rodney, D., Soper, A., Walshaw, C., 2007. The application of multilevel refinement to the vehicle routing problem. In: *Proceedings of the IEEE Symposium on Computational Intelligence in Scheduling (CISched)*. IEEE, pp. 212–219.

[45] Rossi, R. G., de Andrade Lopes, A., Rezende, S. O., 2016. Optimization and label propagation in bipartite heterogeneous networks to improve transductive classification of texts. *Information Processing & Management* 52 (2), 217 – 257.

[46] Rotta, R., Noack, A., 2011. Multilevel local search algorithms for modularity clustering. *Journal of Experimental Algorithms* 16 (2), 2.1.

[47] Schloegel, K., Karypis, G., Kumar, V., 2000. Parallel multilevel algorithms for multi-constraint graph partitioning. In: *Proceedings of the Euro-Par Parallel Processing, Proceedings*. Vol. 1900. pp. 296–310.

[48] Schuetz, P., Cafisch, A., 2008. Efficient modularity optimization by multistep greedy algorithm and vertex mover refinement. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 77 (4), 1–7.

[49] Schweitz, E. A., Agrawal, D. P., 2007. A parallelization domain oriented multilevel graph partitioner. *IEEE Transactions on Computers* 51 (12), 1435 – 1441.

[50] Teng, S.-H., 1999. *Algorithms for Parallel Processing*. Springer New York, Ch. Coarsening, Sampling, and Smoothing: Elements of the Multilevel Method, pp. 247–276.

[51] Trifunovic, A., Knottenbelt, W. J., 2004. A parallel algorithm for multilevel k-way hypergraph partitioning. In: *Proceedings of the Parallel and Distributed Computing*. IEEE, pp. 114–121.

[52] Trifunovic, A., Knottenbelt, W. J., 2004. Parkway 2.0: A parallel multilevel hypergraph partitioning tool. In: *Proceedings of the Computer and Information Sciences (ISCIS)*. pp. 789–800.

[53] Valejo, A., Valverde-Rebaza, J., Drury, B., Lopes, A. A., 2014. Multilevel refinement based on neighborhood similarity. In: *Proceedings of the International Database Engineering & Applications Symposium (IDEAS)*. IDEAS '14. ACM, New York, NY, USA, pp. 67–76.

[54] Valejo, A., Valverde-Rebaza, J., Lopes, A. A., Oct 2014. A multilevel approach for overlapping community detection. In: *Proceedings of the Brazilian Conference on Intelligent Systems (BRACIS)*. pp. 390–395.

[55] Walshaw, C., 2001. A multilevel algorithm for force-directed graph drawing. In: *Proceedings of the Graph Drawing: 8th International Symposium*, GD. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 171–182.

[56] Walshaw, C., May 2001. A multilevel approach to the graph colouring problem. Tech. Rep. 01/IM/69, School of Computing and Mathematical Science, Univeristy of Greenwich, London, UK.

[57] Walshaw, C., Sep. 2002. A multilevel approach to the travelling salesman problem. *Annals of Operations Research* 50 (5), 862–877.

[58] Walshaw, C., 2004. Multilevel refinement for combinatorial optimisation problems. *Annals of Operations Research* 131 (1), 325–372.

[59] Walshaw, C., 2008. *Hybrid Metaheuristics: An Emerging Approach to Optimization*. Vol. 114. Springer Berlin Heidelberg, Ch. Multilevel refinement for combinatorial optimisation: Boosting metaheuristic performance, pp. 261–289.

678 [60] Walshaw, C., Cross, M., 1998. Mesh partitioning: A multilevel balancing and refinement algorithm. SIAM Journal on Scientific Computing (SISC) 22 (1), 63–80.

680 [61] Walshaw, C., Cross, M., 2000. Parallel optimisation algorithms for multilevel mesh partitioning. Parallel Computing 26 (1), 1635–1660.

682 [62] Walshaw, C., Cross, M., 2007. Mesh partitioning techniques and domain decomposition methods. Saxe-Coburg Publications, Ch. JOSTLE: Parallel multilevel graph-partitioning software - An overview, pp. 27–58.

684 [63] Wang, Z., Zhang, D., Zhou, X., Yang, D., Yu, Z., Yu, Z., 2014. Discovering and profiling overlapping communities in location-based social networks. IEEE Transactions on Systems, Man, and Cybernetics: Systems 44 (4), 499–509.

686 [64] Watts, D. J., Strogatz, S. H., 1998. Collective dynamics of 'small-world' networks. Nature 393 (6684), 440–2.

688 [65] Ye, Z., Hu, S., Yu, J., 2008. Adaptive clustering algorithm for community detection in complex networks. Physical Review E 78, 046115.

[66] Zhang, C., Wang, F., 2010. A multilevel approach for learning from labeled and unlabeled data on graphs. Pattern Recognition 43 (6), 2301–2314.

[67] Zhang, G., Lu, J., Gao, Y., 2015. Multi-Level Decision Making: Models, Methods and Applications. Springer-Verlag Berlin, Heidelberg.