

Reducing the rank of a matroid

Gwenaël Joret¹Adrian Vetta²¹ *Département d'Informatique, Université Libre de Bruxelles, Belgium*² *Department of Mathematics and Statistics, and School of Computer Science, McGill University, Canada**received 26th Nov. 2012, revised 24th Apr. 2015, accepted 23rd Aug. 2015.*

We consider the *rank reduction problem* for matroids: Given a matroid M and an integer k , find a minimum size subset of elements of M whose removal reduces the rank of M by at least k . When M is a graphical matroid this problem is the minimum k -cut problem, which admits a 2-approximation algorithm. In this paper we show that the rank reduction problem for transversal matroids is essentially at least as hard to approximate as the densest k -subgraph problem. We also prove that, while the problem is easily solvable in polynomial time for partition matroids, it is NP-hard when considering the intersection of two partition matroids. Our proof shows, in particular, that the maximum vertex cover problem is NP-hard on bipartite graphs, which answers an open problem of B. Simeone.

Keywords: matroid, maximum vertex cover, partial vertex cover

1 Introduction

Consider the well-known minimum k -cut problem: Given a graph G and an integer k , find a minimum size subset of edges whose removal increases the number of connected components by at least k . This problem is NP-hard, assuming k is part of the input, and several 2-approximation algorithms have been developed for it over the years [14, 17, 18, 20]. Notice that the minimum k -cut problem has a simple formulation in terms of matroids: Given a graph G and an integer k , find a minimum size subset of elements of the graphical matroid of G whose removal reduces its rank by at least k .

This observation motivates the study of the *rank reduction* problem in other classes of matroids. For example, is the rank reduction problem computationally hard and, if so, does it admit approximation algorithms with good approximation guarantees (as is the case for graphical matroids)? Moreover, as we will see, many fundamental problems can be formulated in this rank reduction framework.

In this paper, our focus is on the case of transversal matroids. First, we show that the rank reduction problem in transversal matroids is roughly at least as hard to approximate as the densest k -subgraph problem: Given a graph G and an integer k , find a subset of k vertices inducing a subgraph with a maximum number of edges. (Note that in our reduction the parameter k is not necessarily the same one as in the rank reduction problem.) The densest k -subgraph problem can be approximated to within a factor of $O(n^{\frac{1}{4}+\epsilon})$ due to a recent breakthrough result of [4]. Moreover, it is widely believed [5, 1, 10, 11] that its hardness is also close to this upper bound – indeed, Bhaskara et al. [5] present $n^{\Omega(1)}$ lower bounds for lift and project methods based upon the Sherali-Adams and the Lasserre hierarchies. We will show in particular that an $O(m^\epsilon)$ -approximation algorithm for the rank reduction problem on transversal matroids

(where m denotes the number of elements) implies an $O(n^{4\epsilon})$ -approximation algorithm for the densest k -subgraph (where n is the number of vertices).

Secondly, we prove that while the rank reduction problem is easily solvable in polynomial time for partition matroids (a special class of transversal matroids), it is NP-hard when considering the intersection of two partition matroids. Our proof shows in particular that the maximum vertex cover problem—also known as the partial vertex cover problem—is NP-hard on bipartite graphs. Here, one is given a graph and a positive integer k , and the goal is to find a set of k vertices hitting as many edges as possible. The problem is obviously NP-hard on arbitrary graphs since it contains the vertex cover problem as a special case. Whether it remained NP-hard on bipartite graphs was an open problem of B. Simeone (see [15]). We note that we learned after finishing this paper that Apollonio and Simeone [2] independently obtained a proof of this result.

2 Preliminaries

In this section we give the necessary definitions and notations. All graphs and matroids in this paper are finite, and “graph” will always mean an undirected simple graph. We use the shorthands $|G|$ and $\|G\|$ for the number of vertices and edges of a graph G , respectively. We denote by $\mu(G)$ the maximum size of a matching in G , which we call the *matching number* of G .

A *matroid* M is a pair (E, \mathcal{I}) where \mathcal{I} is a family of subsets, called the *independent sets*, of the *ground set* E satisfying the following three axioms:

- the empty set \emptyset is independent;
- every subset of an independent is again independent, and
- if X and Y are two independent sets with $|X| > |Y|$ then there exists $x \in X \setminus Y$ such that $Y \cup \{x\}$ is independent.

The inclusion-wise maximal independent sets are the *bases* of the matroid M ; as follows from the third axiom the bases all have the same cardinality. The *rank function* of M is the function $r : 2^E \rightarrow \mathbb{N}$ that assigns to each subset X of elements of E the maximum size $r(X)$ of an independent set contained in X , called the *rank* of X . In particular, $r(E)$ is the cardinality of a basis of M , which is called the *rank* of M .

The *rank reduction* problem for matroids is defined as follows: Given a positive integer k and a matroid on a set A of elements with rank function r , the goal is to find a minimum size subset $X \subseteq A$ such that $r(A \setminus X) \leq r(A) - k$.

For example consider the case of graphical matroids: Given a graph $G = (V, E)$, the *graphical matroid* of G is obtained by taking E as ground set, and letting a subset F of edges be independent if and only if the corresponding subgraph is acyclic. Here the rank reduction problem is the minimum k -cut problem.

As stated, here we study transversal matroids. A bipartite graph G with bipartition (A, B) induces a matroid M as follows: The matroid M has A as ground set, and $X \subseteq A$ is independent in M if and only if there exists a matching of G covering X . The fact that this is indeed a matroid is well-known; see for instance [19]. Any matroid M that can be obtained this way is called a *transversal matroid*, and the bipartite graph G is said to be a *model* for M . Observe that, letting r denote the rank function of M , the rank $r(X)$ of $X \subseteq A$ is equal to $\mu(G[X \cup B])$. Also note that being a transversal matroid is a hereditary property, in the sense that for each set $X \subseteq A$, taking the restriction $I \cap X$ of all independent sets I yields a transversal matroid on ground set X .

A special case of transversal matroids are partition matroids. Here we are given a collection E_1, \dots, E_p of disjoint sets and integers d_1, \dots, d_p such that $0 \leq d_i \leq |E_i|$ for each $i \in \{1, \dots, p\}$. One can define a corresponding matroid with ground set $E := E_1 \cup \dots \cup E_p$ by letting $X \subseteq E$ be independent if and only if $|E_i \cap X| \leq d_i$ for each $i \in \{1, \dots, p\}$. Such a matroid is called a *partition matroid*, with model $((E_1, d_1), \dots, (E_p, d_p))$. This corresponds to a transversal matroid on a bipartite graph G with bipartition (E, B) , where B has d_i vertices that are adjacent to all vertices in E_i , and none other, for each set E_i . Notice that partition matroids are also hereditary.

Throughout, since we restrict ourselves to specific families of matroids, we assume that the matroid is given concisely and not given explicitly as a set system in input. Specifically, a corresponding model of the matroid is provided: a bipartite graph for a transversal matroid, a graph for a graphical matroid, etc.

More generally, the rank reduction problem can be considered on the intersection of matroids. Given two matroids $M_1 = (E, \mathcal{I}_1)$ and $M_2 = (E, \mathcal{I}_2)$ with common ground set E , the intersection $M_1 \cap M_2$ of M_1 and M_2 is the pair (E, \mathcal{I}) where \mathcal{I} is the family of sets $X \subseteq E$ that are independent in both M_1 and M_2 , which are said to be the independent sets of $M_1 \cap M_2$. While the independence system (E, \mathcal{I}) is not necessarily a matroid anymore, it enjoys several of the nice properties of matroids (see [19]). In particular, letting as before the *rank* $r(X)$ of $X \subseteq E$ be the maximum size of an independent set of $M_1 \cap M_2$ contained in X , the rank $r(E)$ of $M_1 \cap M_2$ can be computed in polynomial time given access to the rank functions r_1 and r_2 of M_1 and M_2 , respectively, by a classical result of Edmonds (see [19]). We examine the rank reduction problem for the intersection of two partition matroids in Section 4.

3 Transversal Matroids

We start our investigation of the rank reduction problem with an easy observation, namely that the problem can be solved in polynomial time if the input matroid is a partition matroid.

Theorem 3.1 *The rank reduction problem can be solved in polynomial time on partition matroids.*

Proof: Let M be a given partition matroid with model $((E_1, d_1), \dots, (E_p, d_p))$ and rank function r . Let $E := E_1 \cup \dots \cup E_p$ denote the ground set of M . Observe that $r(E) = \sum_{i=1}^p d_i$. Let k be the given parameter for the rank reduction problem on M . We may assume that $1 \leq k \leq r(E)$. Let $c_i := |E_i| - d_i$ for each $i \in \{1, \dots, p\}$. Given $X \subseteq E$, the rank $r(E \setminus X)$ of the set $E \setminus X$ is equal to $\sum_{i=1}^p \min(|E_i \setminus X|, d_i)$.

Let $X \subseteq E$ be such that $r(E \setminus X) \leq r(E) - k$. Moreover, assume X is inclusion-wise minimal with this property. Then, for each $i \in \{1, \dots, p\}$, either $|E_i \cap X| \geq c_i + 1$ or $E_i \cap X = \emptyset$. Moreover, letting J be the subset of indices $i \in \{1, \dots, p\}$ such that $E_i \cap X \neq \emptyset$, we have that $\sum_{i \in J} d_i \geq k$ and $|X| = k + \sum_{i \in J} c_i$.

Conversely, suppose $J' \subseteq \{1, \dots, p\}$ is such that $\sum_{i \in J'} d_i \geq k$. Then choosing arbitrarily c_i elements of E_i , for each $i \in J'$, plus k additional elements from $\bigcup_{i \in J'} E_i$ gives a set X' with $|X'| = k + \sum_{i \in J'} c_i$ such that $r(E \setminus X') \leq r(E) - k$.

Therefore, computing an optimal solution to the rank reduction problem reduces to the problem of finding a subset $J \subseteq \{1, \dots, p\}$ such that $\sum_{i \in J} d_i \geq k$ and $\sum_{i \in J} c_i$ is minimum. Thus we obtain a knapsack problem. Moreover, as c_i and d_i are at most $|E_i|$, they are of polynomial size when encoded in unary. Thus the knapsack problem can be solved easily in polynomial time using dynamic programming. \square

While the rank reduction problem admits a simple polynomial-time algorithm on partition matroids, the problem turns out to be more difficult on the broader class of transversal matroids. In fact, up to some degree, the problem can be viewed as a generalization of the *densest k -subgraph* problem. In the latter problem, one is given a graph G and a positive integer k , and the aim is to find a subgraph H of G with $|H| = k$ and $\|H\|$ maximum. Towards this goal, we consider a closely related problem, the *minimum t -edge subgraph* problem: Given a graph G and a positive integer t , the goal is to find a subgraph H of G with $\|H\| = t$ and $|H|$ minimum.

We start by drawing a connection between the rank reduction problem on transversal matroids and the minimum t -edge subgraph problem. Then we will extend the connection to the densest k -subgraph problem.

Lemma 3.2 *For each constant ε with $0 < \varepsilon < 1/2$, every $O(m^\varepsilon)$ -approximation algorithm for the rank reduction problem on transversal matroids with m elements can be turned into an $O(n^{2\varepsilon})$ -approximation algorithm for the minimum t -edge subgraph problem on graphs with n vertices.*

Proof: Let $G = (V, E)$ be an instance of the minimum t -edge subgraph problem. Let $n := |V|$. Let V_1, V_2, \dots, V_n be n disjoint copies of V . Let E' be a disjoint copy of E . Let H be the bipartite graph with bipartition (A, B) where $A = V_1 \cup V_2 \cup \dots \cup V_n \cup E'$ and $B = E$, and where $u \in A$ is adjacent to $v \in B$ if either u corresponds to a vertex of G that is incident to the edge corresponding to v in G , or if u and v correspond to the same edge of G .

Let r denote the rank function of the transversal matroid induced by H on A ; thus for $X \subseteq A$, $r(X)$ is the maximum size of a matching in $H[X \cup B]$. Obviously, $r(A) = |E|$, since every $v \in B$ can be matched to its copy in A . Let $m := |A|$ denote the number of elements of the transversal matroid. Now consider the rank reduction problem on this matroid with $k = t$. Recall that a feasible solution is a subset $X \subseteq A$ such that $r(A \setminus X) \leq r(A) - t = |E| - t$.

As is well known, we have that $r(A \setminus X) \leq r(A) - t$ for $X \subseteq A$ if and only if there exists $Y \subseteq B$ such that $|N_H(Y) \setminus X| \leq |Y| - t$, where $N_H(Y)$ denotes the set of vertices of H that have a neighbor in Y .⁽ⁱ⁾ Such a set Y is said to be a *witness* for X . The set Y defines in turn a corresponding subgraph G_Y of G consisting of all the edges of G included in Y , and the vertices of G incident to those edges. By definition of H , the set $N_H(Y)$ consists of the n copies of each vertex of G_Y , along with the copies in E' of each edge of G_Y . Observe that any set X' obtained by taking the n copies in A of each vertex of G_Y and t arbitrarily chosen edges of G_Y in A is such that $|N_H(Y) \setminus X'| = |Y| - t$. Moreover, since $|N_H(Y) \setminus X| \leq |Y| - t$ and $X' \subseteq N_H(Y)$, it follows that $|X'| \leq |X|$, that is, X' is a solution of size no greater than X and having the same witness Y . Such a pair (X', Y) is called a *canonical pair*.

Now, if a canonical pair (X, Y) is such that $|Y| > t$, then $N_H(Y) \setminus X$ consists of exactly $|Y| - t > 0$ edges of G_Y (or more precisely, their copies in E'). For each $u \in N_H(Y) \setminus X$ with corresponding copy

⁽ⁱ⁾ This is a consequence of Hall's Marriage Theorem, as we now explain for completeness. Add $t - 1$ new vertices to the set A , yielding a set A' , and make each of them adjacent to every vertex in B . Let H' be the bipartite graph obtained from H in this manner. Let $X \subseteq A$. Then every matching M in $H \setminus X$ with $|M| \leq |B| - (t - 1)$ can be extended to a matching M' of $H' \setminus X$ with $|M'| = |M| + t - 1$. Conversely, every matching M' in $H' \setminus X$ with $|M'| \geq t - 1$ yields a matching M of $H \setminus X$ with $|M| \geq |M'| - (t - 1)$ by discarding the at most $t - 1$ edges of M' incident to the vertices in $A' \setminus A$. Hence, $H \setminus X$ has a matching of size $|B| - (t - 1)$ —or equivalently, $r(A \setminus X) \geq r(A) - (t - 1)$ —if and only if B can be completely matched in $H' \setminus X$. By Hall's theorem, the latter happens if and only if $|N_{H' \setminus X}(Y)| \geq |Y|$ for every $Y \subseteq B$, which is equivalent to $|N_H(Y) \setminus X| \geq |Y| - (t - 1)$ for every $Y \subseteq B$. Therefore, $r(A \setminus X) \leq r(A) - t$ if and only if there exists $Y \subseteq B$ such that $|N_H(Y) \setminus X| \leq |Y| - t$.

v in B , we have that $X \setminus \{u\}$ is again a solution to the rank reduction problem, with witness $Y \setminus \{v\}$, and of size smaller than X .

To summarize the above discussion, given an arbitrary set $X' \subseteq A$ such that $r(A \setminus X') \leq r(A) - t$, one can in polynomial time compute a canonical pair (X, Y) with $|Y| = ||G_Y|| = t$ and $|X| = n|G_Y| + t \leq |X'|$.

Conversely, for every subgraph $G' \subseteq G$ with $||G'|| = t$, there is a natural corresponding canonical pair (X, Y) , where Y contains the copies in B of the t edges of G' , and where $X = N_H(Y)$. Letting x^* and j^* denote the size of an optimal solution for the rank reduction and minimum t -edge subgraph problems, respectively, it follows that $x^* = nj^* + t$.

Now suppose that the rank reduction problem admits a cm^ε -approximation algorithm, where $0 < \varepsilon < 1/2$ and $c > 0$ are absolute constants. Letting (X, Y) be a canonical pair with $|Y| = t$ obtained using this algorithm, we have

$$n|G_Y| + t = |X| \leq cm^\varepsilon x^* = c(n^2 + |E|)^\varepsilon (nj^* + t) \leq c(2n^2)^\varepsilon (nj^* + t)$$

and hence

$$|G_Y| \leq \frac{c(2n^2)^\varepsilon (nj^* + t) - t}{n} \leq 2cn^{2\varepsilon} \left(j^* + \frac{t}{n} \right) \leq 2cn^{2\varepsilon} (j^* + j^*) = 4cn^{2\varepsilon} j^*.$$

(In the last inequality we used the fact that $(j^*)^2 \geq t$, and thus $\frac{t}{n} \leq \frac{(j^*)^2}{n} \leq \frac{j^* n}{n} = j^*$.) Therefore, G_Y is a t -edge subgraph whose order is within a $4cn^{2\varepsilon}$ -factor of optimal. \square

As pointed out to us by an anonymous referee, the following lemma is implicit in the recent work of Chlamtac, Dinitz, and Krauthgamer [8] (in [8], the minimum t -edge subgraph problem is called the *smallest m -edge subgraph problem*). We include a proof nevertheless, for completeness.

Lemma 3.3 *For each constant ε with $0 < \varepsilon < 1/2$, every $O(n^\varepsilon)$ -approximation algorithm for the minimum t -edge subgraph problem can be turned into an $O(n^{2\varepsilon})$ -approximation algorithm for the densest k -subgraph problem.*

Proof: Suppose that the minimum t -edge subgraph problem admits a $c'n^\varepsilon$ -approximation algorithm, which we denote \mathcal{A} , where $0 < \varepsilon < 1$ and $c' \geq 1$ are absolute constants. Let G be an instance of the densest k -subgraph problem. As before, we let n and m denote the number of vertices and edges of G , respectively. We may assume $n \geq k \geq 2$. Since $c' \geq 1$ and $n^\varepsilon \geq 1$, there exists c with $c' \leq c \leq 2c'$ such that cn^ε is an integer. We will consider the approximation factor of \mathcal{A} to be cn^ε in what follows, to avoid cumbersome floors and ceilings in the calculations.

Run algorithm \mathcal{A} on G with $t = 1, 2, \dots, m$. Let H_i be the i th subgraph returned by the algorithm. Clearly, we may suppose that $|H_i| \leq |H_{i+1}|$ for each $i \in \{1, \dots, m-1\}$.

Let z^* denote the number of edges in an optimal solution to the densest k -subgraph problem on G . If $|H_m| \leq kcn^\varepsilon$ then let $t' := m$, otherwise let t' be the index in $\{1, \dots, m-1\}$ such that $|H_{t'}| \leq kcn^\varepsilon$ and $|H_{t'+1}| > kcn^\varepsilon$. Since algorithm \mathcal{A} is a cn^ε -approximation algorithm, and since either $t' = m$ or $|H_{t'+1}| > kcn^\varepsilon$, it follows that every subgraph of G with exactly k vertices has at most t' edges, that is, $z^* \leq t'$.

Let $q := \lceil |H_{t'}| / \lfloor k/2 \rfloor \rceil$. Observe that $2 \leq q \leq 3cn^\varepsilon$. Let V_1, V_2, \dots, V_q be a partition of the vertex set of $H_{t'}$ into q subsets with $|V_1| = |V_2| = \dots = |V_{q-1}| = \lfloor k/2 \rfloor$ and $|V_q| = |H_{t'}| - (q-1)\lfloor k/2 \rfloor \leq \lfloor k/2 \rfloor$. Let (V_i, V_j) be a pair with $i \neq j$ such that $\|H_{t'}[V_i \cup V_j]\|$ is maximized. By the pigeonhole principle,

$$\|H_{t'}[V_i \cup V_j]\| \geq \frac{\|H_{t'}\|}{\binom{q}{2}} = \frac{t'}{\binom{q}{2}} \geq \frac{t'}{(3cn^\varepsilon)^2} \geq \frac{z^*}{9c^2n^{2\varepsilon}}.$$

If $|V_i \cup V_j| = k$, then let $H := H_{t'}[V_i \cup V_j]$. If, on the other hand, $|V_i \cup V_j| < k$, then let $H := H_{t'}[V_i \cup V_j \cup X]$ where X is an arbitrary subset of $V(H_{t'}) \setminus (V_i \cup V_j)$ of size $k - |V_i \cup V_j|$. Thus in both cases $|H| = k$ and $\|H\| \geq z^*/9c^2n^{2\varepsilon}$. Hence, H is a solution to the densest k -subgraph problem on G whose number of edges is within a $9c^2n^{2\varepsilon}$ -factor of the optimum. \square

Combining Lemma 3.2 and 3.3 gives:

Theorem 3.4 *For each constant ε with $0 < \varepsilon < 1/4$, every $O(m^\varepsilon)$ -approximation algorithm for the rank reduction problem on transversal matroids with m elements can be turned into an $O(n^{4\varepsilon})$ -approximation algorithm for the densest k -subgraph problem on graphs with n vertices.*

As discussed in the introduction, the best approximation algorithm for the densest k -subgraph problem currently known has an approximation ratio of $O(n^{1/4+\delta})$ for any fixed $\delta > 0$ [4] and it is conjectured that the inapproximability of the problem is of a similar magnitude. It would be nice to obtain strong inapproximability bounds for the rank reduction problem that do not rely on this conjecture. One approach may be to analyze hypergraphs as the rank reduction problem in transversal matroids incorporates the hypergraph version of the minimum t -edge subgraph problem. That is, we wish to select as few vertices as possible that induce at least t hyperedges. Perhaps surprisingly, little is known about this problem. As far as we are aware, the only specific hardness result is NP-hardness due Vinterbo [21] who studied the problem in the context of making medical databases anonymous.

We conclude this section with a remark about the approximability of the minimum t -edge subgraph problem itself. Given the existence of a $O(n^{1/4+\delta})$ -approximation algorithm for the densest k -subgraph problem, in view of Lemma 3.3 it is perhaps natural to wonder whether one could achieve a $O(n^{1/8+\delta})$ -approximation for the former problem. While this is still open as far as we know, Chlamtac *et al.* [8] recently made progress in that direction by describing an algorithm for the minimum t -edge subgraph problem with an approximation ratio of $O(n^{3-2\sqrt{2}+\delta}) = O(n^{0.1716+\delta})$ for fixed $\delta > 0$.

4 The Maximum Vertex Cover Problem in Bipartite Graphs

As we have seen, the rank reduction problem admits a fairly simple polynomial-time algorithm on partition matroids but becomes much harder on transversal matroids, in the sense that approximation algorithms offering good guarantees seem unlikely to exist. Another interesting generalization of the case of partition matroids is to consider the intersection of two partition matroids.

As is well-known, the set of matchings of a bipartite graph $G = (V, E)$ with bipartition (A, B) can be modeled as the family of common independent sets of two partition matroids M_1 and M_2 defined on E : Take M_1 to be the partition matroid with model $((E(u_1), 1), \dots, (E(u_a), 1))$ and M_2 the partition matroid with model $((E(v_1), 1), \dots, (E(v_b), 1))$, where $A = \{u_1, \dots, u_a\}$, $B = \{v_1, \dots, v_b\}$, and for $w \in V$ the set $E(w)$ denotes the set of edges incident to w . Hence, in this specific case the rank reduction problem on $M_1 \cap M_2$ amounts to finding a subset F of edges of G of minimum size such that

$\mu(G - F) \leq \mu(G) - k$. In this section we show that this problem is NP-hard. More accurately, we show that a problem polynomially equivalent to it, the maximum vertex cover problem on bipartite graphs, is NP-hard; see Theorem 4.1. This solves an open problem of B. Simeone (see [15]).

The *maximum vertex cover problem* (also known as the *partial vertex cover problem*) is defined as follows: Given a graph $G = (V, E)$ and a positive integer $k \leq |G|$, find a subset X of vertices of G with $|X| = k$ such that the number of edges covered by X is maximized. (An edge e of G is covered by X if e has at least one endpoint in X .)

Now, if G is bipartite, t is a positive integer with $t \leq \mu(G)$, and F is a subset of edges of G such that $\mu(G - F) \leq \mu(G) - t =: k$, then by König's theorem $G - F$ has a vertex cover X of size k , and hence X covers at least $|E| - |F|$ edges of G . (We remark that X could cover some edges of F too, and that X can be computed in polynomial time given F). Conversely, for every set $X \subseteq V$ with $|X| = k$, the set F of edges of G not covered by X is such that $\mu(G - F) \leq k$. Therefore, for bipartite graphs, the maximum vertex cover problem is polynomially equivalent to that of finding a minimum-size set of edges decreasing the matching number by a prescribed amount.

It should be noted that two recent works [9, 6] with an overlapping set of authors claim that the NP-hardness of the maximum vertex cover problem on bipartite graphs can be derived directly from the reduction of Corneil and Perl [7] showing that the densest k -subgraph problem is NP-hard on bipartite graphs. However, the argument relating the latter reduction to the maximum vertex cover problem, described explicitly in [9, Lemma 4], is flawed.⁽ⁱⁱ⁾ We also mention that the proof of Theorem 1 in [6], showing that a related problem called the maximum quasi-independent set problem is NP-hard on bipartite graphs, relies on the assumption that the maximum vertex cover problem is NP-hard on bipartite graphs. Thus our result also fills a gap in that proof.

Theorem 4.1 *The maximum vertex cover problem is NP-hard on bipartite graphs.*

Before proving Theorem 4.1, we need to introduce a technical lemma.

Lemma 4.2 *Let ℓ be an integer with $\ell \geq 6$. Then the integer program*

$$\begin{aligned} \text{minimize} \quad & x + 2y + 3z \\ \text{s. t.} \quad & x + y + z - s = \binom{\ell}{2} - \ell \\ & x \leq \binom{s}{2} \\ & x, y, z, s \in \mathbb{N} \end{aligned}$$

has a unique optimal solution given by $x = \binom{\ell}{2}, y = 0, z = 0, s = \ell$.

⁽ⁱⁱ⁾ As mentioned in [9, Lemma 3], the maximum vertex cover problem in bipartite graphs is polynomially equivalent to the densest k -subgraph problem in complements of bipartite graphs. Thus one may equivalently consider the complexity of the latter problem. In the proof of Lemma 4 in [9], the authors point out that the reduction of [7] implies that the problem of finding a densest k -subgraph in the complement of a bipartite graph with bipartition (A, B) with $k = q + \binom{q}{2}$ and with the extra requirement that it contains exactly q vertices from A and $\binom{q}{2}$ vertices from B is NP-hard. From this they wrongly conclude that the densest k -subgraph problem, without this extra constraint, is also NP-hard on complements of a bipartite graphs. (In fact, the instances obtained via the reduction in [7] satisfy $|B| \geq k$, and thus a densest k -subgraph is trivially obtained by taking k vertices in the clique B .)

Proof: The proof is a straightforward case analysis. Consider an optimal solution x, y, z, s to the integer program and, arguing by contradiction, assume it differs from the solution described above. Let $f(x, y, z) := x + 2y + 3z$.

Case 1: $s \leq 2$. We have $x \leq s$ and thus

$$y + z = \binom{\ell}{2} - \ell + s - x \geq \binom{\ell}{2} - \ell.$$

It follows that $f(x, y, z) \geq 2(y + z) \geq 2\binom{\ell}{2} - 2\ell$. But $2\binom{\ell}{2} - 2\ell > \binom{\ell}{2}$ since $\ell \geq 6$, contradicting the optimality of the solution.

Case 2: $3 \leq s < \ell$.

$$y + z = \binom{\ell}{2} - \ell + s - x \geq \binom{\ell}{2} - \ell + s - \binom{s}{2} = \left(\binom{\ell}{2} - \binom{s}{2} \right) - (\ell - s) \geq 2.$$

Here, the last inequality follows from the fact that $\binom{s+1}{2} - \binom{s}{2} \geq 3$.

Now, increment s by 1, x by 3, and decrement y and z in such a way that they remain non-negative integers and that the sum $y + z$ decreases by exactly 2. The modified solution is still feasible and $f(x, y, z)$ decreases by at least 1, a contradiction.

Case 3: $s = \ell$. Then $y + z \geq 1$, since otherwise we would have the solution described in the lemma statement. It follows that $x \leq \binom{\ell}{2} - 1$. Thus we can increment x by 1 and decrement by 1 a positive variable among y, z . This strictly decreases $f(x, y, z)$, a contradiction.

Case 4: $s > \ell$. Then $x + y + z = \binom{\ell}{2} + s - l > \binom{\ell}{2}$. But $x \leq \binom{\ell}{2} < \binom{s}{2}$, otherwise the solution cannot be minimum. Therefore $y + z \geq 1$. Thus we improve the solution by incrementing x by 1 and decrementing by 1 a positive variable among y, z . \square

Now we may turn to the proof of Theorem 4.1.

Proof of Theorem 4.1.: The reduction is from the NP-complete problem *Clique*: Given a graph H and an integer ℓ , decide whether H contains a clique on ℓ vertices or not. We may assume $\ell \geq 6$ (otherwise, we simply check the existence of an ℓ -clique by brute force). We may also suppose that H has minimum degree at least 2. Indeed, a vertex with degree at most 1 cannot be part of an ℓ -clique, and thus those vertices can iteratively be removed from the graph. Finally, we assume that $\|H\| \geq |H| + \binom{\ell}{2}$. This last assumption can also be made without loss of generality. Indeed, if $\|H\|$ is too small then one can simply consider the disjoint union of H with a large enough 3-regular graph; since $\ell \geq 6$ no vertex from this new 3-regular component can be part of an ℓ -clique.

We build an instance (G, k) of the maximum vertex cover problem as follows. First, create two adjacent vertices a_v and b_v for every vertex $v \in V(H)$, and similarly two adjacent vertices a_e and b_e for every edge $e \in E(H)$. Next, for every edge $uv \in E(H)$, add the edges $a_{uv}b_u, b_{uv}a_u$ and $a_{uv}b_v, b_{uv}a_v$. Finally, let

$$k := |H| + \|H\| - \binom{\ell}{2} + \ell.$$

Observe that G is bipartite with bipartition

$$(\{a_x : x \in V(H) \cup E(H)\}, \{b_x : x \in V(H) \cup E(H)\}).$$

See Figure 1 for an illustration of the construction.




Fig. 1: Illustration of the construction of the bipartite graph G . (We note that this graph H cannot be a valid instance of the problem because $\ell \geq 6$ and thus H cannot satisfy $||H|| \geq |H| + \binom{\ell}{2}$; our aim here is only to illustrate how G is obtained from H on a small example.)

A feasible solution for this instance of the maximum vertex cover problem is a subset X of vertices of G with $|X| = k$, which we call a *partial vertex cover* for short. We let $c(X)$ denote the number of edges covered by such a set X . Let OPT denote the maximum of $c(X)$ over every partial vertex cover X of G .

A partial vertex cover X of G is *nice* if

$$X \cap \{a_u, b_u\} \in \{\{a_u\}, \{a_u, b_u\}\}$$

for every $u \in V(H)$ and

$$X \cap \{a_e, b_e\} \in \{\emptyset, \{a_e\}\}$$

for every $e \in E(H)$.

Claim 4.3 Given a partial vertex cover X of G , one can find a nice partial vertex cover X' of G with $c(X') \geq c(X)$.

Proof: First we define a partial vertex cover \tilde{X} based on X which is close to being nice: Let

$$\begin{aligned} \tilde{X} := & \{a_u, b_u : u \in V(H), a_u, b_u \in X\} \cup \{a_u : u \in V(H), |\{a_u, b_u\} \cap X| = 1\} \\ & \cup \{a_e, b_e : e \in E(H), a_e, b_e \in X\} \cup \{a_e : e \in E(H), |\{a_e, b_e\} \cap X| = 1\}. \end{aligned}$$

By construction $|\tilde{X}| = |X|$. Clearly, an edge $a_x b_x$ with $x \in V(H) \cup E(H)$ is covered by \tilde{X} if and only if it is covered by X . Also, given a pair (u, e) of vertex $u \in V(H)$ and edge $e \in E(H)$ such that u is incident to e in H , the set \tilde{X} covers at least as many edges in $\{a_e b_u, b_e a_u\}$ as X (though not necessarily the same ones). It follows that $c(\tilde{X}) \geq c(X)$.

A useful property of the set \tilde{X} is that if $b_x \in \tilde{X}$ for some $x \in V(H) \cup E(H)$ then necessarily $a_x \in \tilde{X}$. For simplicity we call this property the *a-property* of \tilde{X} .

We need to introduce an additional definition. An element $x \in V(H) \cup E(H)$ is said to be *bad* in a partial vertex cover Y of G if either $x \in V(H)$ and $a_x, b_x \notin Y$ (x is a bad vertex), or $x \in E(H)$ and $a_x, b_x \in Y$ (x is a bad edge). Observe that Y is nice if and only if Y has the *a-property* and there is no bad element.

Suppose $e = uv$ is an edge of H which is bad in \tilde{X} . If u or v is also bad in \tilde{X} , say u , then let

$$\tilde{X}' := (\tilde{X} - \{b_e\}) \cup \{a_u\}.$$

We have $a_e \in \tilde{X}$, thus the edge $a_e b_e$ is still covered by \tilde{X}' . Since \tilde{X}' covers also $b_e a_u$, there is at most one edge incident to b_e in G (namely, $b_e a_v$) which is not covered by \tilde{X}' . On the other hand, \tilde{X}' covers the previously uncovered edge $a_u b_u$. Hence, \tilde{X}' is a partial vertex cover with $c(\tilde{X}') \geq c(\tilde{X})$. Observe that \tilde{X}' still has the *a-property*, and the edge e is no longer bad in \tilde{X}' .

If, on the other hand, none of u, v is bad in \tilde{X} , then $a_u, a_v \in \tilde{X}$ by the *a-property*. Since $a_e \in \tilde{X}$, it follows that $c(\tilde{X} - \{b_e\}) = c(\tilde{X})$. There exists an element $x \in V(H) \cup E(H)$ such that $a_x \notin \tilde{X}$, because $|\tilde{X}| = k = |H| + ||H|| - \binom{\ell}{2} + \ell < |H| + ||H||$ (since $\ell \geq 6$). Let then

$$\tilde{X}' := (\tilde{X} - \{b_e\}) \cup \{a_x\}.$$

The set \tilde{X}' is a partial vertex cover with the *a-property* and with $c(\tilde{X}') \geq c(\tilde{X})$. Moreover, the edge e is no longer bad in \tilde{X}' .

Now apply iteratively the above modifications on \tilde{X} as long as there exists a bad edge. This results in a partial vertex cover \hat{X} with the *a-property*, without bad edges, and with $c(\hat{X}) \geq c(\tilde{X})$.

Next we deal with bad vertices in \hat{X} . Suppose $u \in V(H)$ is such a vertex, that is, $a_u, b_u \notin \hat{X}$. Consider two edges e, f incident to u in H . (Recall that H has minimum degree at least 2.) Since $|\hat{X}| = k = |H| + ||H|| - \binom{\ell}{2} + \ell$ and $||H|| \geq |H| + \binom{\ell}{2}$ by our assumption on H , we have $|\hat{X}| > 2|H|$. Together with the *a-property* of \hat{X} , it follows that $a_{e'} \in \hat{X}$ for some edge $e' \in E(H)$ (possibly $e' = e$ or $e' = f$). Note that $b_{e'} \notin \hat{X}$, because otherwise e' would be a bad edge for \hat{X} . Let

$$\hat{X}' := (\hat{X} - \{a_{e'}\}) \cup \{a_u\}.$$

Since $a_{e'}$ has degree 3 in G we have $c(\hat{X} - \{a_{e'}\}) \geq c(\hat{X}) - 3$. Furthermore, $b_e, b_f \notin \hat{X}$ as there are no bad edges in \hat{X} . Thus, the three edges $a_u b_u, b_e a_u, b_f a_u$ of G were not covered by \hat{X} but are covered by \hat{X}' , so we have $c(\hat{X}') \geq c(\hat{X})$. Similarly as before, the partial vertex cover \hat{X}' has the *a-property* and one less bad vertex than \hat{X} . Therefore, iterating this procedure as long as there is a bad vertex, we eventually obtain a partial vertex cover X' with $c(X') \geq c(\hat{X}) \geq c(X)$ having the *a-property* and no bad element, as desired. \square

Consider a nice partial vertex cover X of G . Let $S(X)$ be the set of vertices $u \in V(H)$ such that $a_u, b_u \in X$, and let $s(X) := |S(X)|$. An edge $e = uv$ of H satisfies exactly one of the following three conditions:

1. $u, v \in S(X)$;
2. exactly one of u, v belongs to $S(X)$,
3. $u, v \notin S(X)$.

We say that edge e is of *type* i ($i \in \{1, 2, 3\}$) if e satisfies the i th condition above *and* moreover $a_e, b_e \notin X$. (We will focus on edges e of H such that $a_e, b_e \notin X$ in what follows, which is why the other ones do not get assigned a type.) Let $E_i(X)$ be the set of edges of H with type i , and let $e_i(X) := |E_i(X)|$.

Claim 4.4 *Let X be a nice partial vertex cover. Then*

$$c(X) = \|G\| - e_1(X) - 2e_2(X) - 3e_3(X).$$

Proof: As X is nice, $a_v \in X$ for all $v \in V$. Therefore every edge of the form $a_v b_v$ or $a_v b_e$ is covered. Also, for each edge $e = uv \in E(H) \setminus \cup_{1 \leq i \leq 3} E_i(X)$ we have that $a_e b_e, a_e b_u, a_e b_v$ are all covered. Thus the only uncovered edges are incident to vertices a_e where $e \in \cup_{1 \leq i \leq 3} E_i(X)$. Suppose $e \in \cup_{1 \leq i \leq 3} E_i(X)$ and let us consider which edges among the three edges $a_e b_e, a_e b_u, a_e b_v$ are covered by X . If $e \in E_1(X)$, then X covers $a_e b_u$ and $a_e b_v$ but not $a_e b_e$. If $e \in E_2(X)$, then X covers exactly one of $a_e b_u, a_e b_v$, and avoids $a_e b_e$. If $e \in E_3(X)$, then X covers none of the three edges. Hence, the total number of edges not covered by X is exactly $e_1(X) + 2e_2(X) + 3e_3(X)$. \square

Claim 4.5 *Let X be a nice partial vertex cover. Then*

$$c(X) \leq \|G\| - \binom{\ell}{2},$$

with equality if and only if $s(X) = \ell$, $e_1(X) = \binom{\ell}{2}$, $e_2(X) = e_3(X) = 0$.

Proof: Let $x := e_1(X)$, $y := e_2(X)$, $z := e_3(X)$ and $s := s(X)$. Then $c(X) = \|G\| - f(x, y, z)$, where $f(x, y, z) = x + 2y + 3z$ by the previous claim.

Every edge in $E_1(X)$ has its two endpoints in $S(X)$; hence,

$$\binom{s}{2} \geq x. \tag{1}$$

Also,

$$|H| + \|H\| - (x + y + z - s) = |X| = k = |H| + \|H\| - \left(\binom{\ell}{2} - \ell \right),$$

and thus

$$x + y + z - s = \binom{\ell}{2} - \ell. \tag{2}$$

Since $\ell \geq 6$ and x, y, z, s, ℓ are non-negative integers satisfying (1) and (2), by Lemma 4.2 we have $f(x, y, z) \geq \binom{\ell}{2}$, or equivalently,

$$c(X) \leq \|G\| - \binom{\ell}{2}.$$

Moreover, equality holds if and only if $x = \binom{\ell}{2}, y = 0, z = 0, s = \ell$ by the same lemma. \square

It follows from Claims 4.3 and 4.5 that

$$OPT \leq \|G\| - \binom{\ell}{2}.$$

If H has an ℓ -clique $K \subseteq H$, then the subset $X \subseteq V(G)$ defined by

$$\begin{aligned} X := & \{a_u, b_u : u \in V(K)\} \cup \{a_u : u \in V(H) - V(K)\} \\ & \cup \{a_e : e \in E(H) - E(K)\} \end{aligned}$$

is a partial vertex cover of G with $c(X) = \|G\| - \binom{\ell}{2}$, implying $OPT = \|G\| - \binom{\ell}{2}$.

Conversely, if $OPT = \|G\| - \binom{\ell}{2}$, then there exists a partial vertex cover X of G with $c(X) = \|G\| - \binom{\ell}{2}$, and by Claim 4.3 we may assume that X is nice. From Claim 4.5 we then have $s(X) = \ell$, $e_1(X) = \binom{\ell}{2}$, $e_2(X) = e_3(X) = 0$, implying that $S(X)$ induces an ℓ -clique in H .

Therefore, we can decide in polynomial time if H has an ℓ -clique by checking if $OPT = \|G\| - \binom{\ell}{2}$. This concludes the proof. \square

The NP-hardness of the maximum vertex cover problem on bipartite graphs motivates the search for non-trivial approximation algorithms for this class of graphs. A recent result in this direction is due to Apollonio and Simeone [3], who gave an LP-based $(4/5)$ -approximation algorithm for bipartite graphs.

Acknowledgments

We are grateful to Attila Bernath and Tamás Kiraly for their comments on the proof of Theorem 4.1, and to the two anonymous referees for their helpful remarks and suggestions. We also thank Mohit Singh, Bill Cunningham and Anupam Gupta for interesting discussions, and Nicola Apollonio for providing us with a preliminary version of [2].

References

- [1] R. Andersen and K. Chellapilla, “Finding dense subgraphs with size bounds”, *Proceedings of the 6th International Workshop on Algorithms and Models for the Web-Graph*, pp. 25–37, 2009.
- [2] N. Apollonio and B. Simeone, “The maximum vertex coverage problem on bipartite graphs”. *Discrete Applied Mathematics*, **165**, pp. 37–48, 2014.
- [3] N. Apollonio and B. Simeone, “Improved Approximation of Maximum Vertex Coverage Problem on Bipartite Graphs”. *SIAM J. Discrete Math.*, **28(3)**, pp. 1137–1151, 2014.

- [4] A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan, “Detecting high log-densities: an $O(n^{\frac{1}{2}})$ approximation for densest k -subgraph”, *Proceedings of 42nd Symposium on Theory of Computing (STOC)*, pp. 201–210, 2010.
- [5] A. Bhaskara, M. Charikar, V. Guruswami, A. Vijayaraghavan, and Y. Zhou “Polynomial Integrality gaps for Strong relaxations of Densest k -subgraph”, *Proceedings of 23rd Symposium on Discrete Algorithms (SODA)*, pp. 388–405, 2012.
- [6] N. Bourgeois, A. Giannakos, G. Lucarelli, I. Milis, V. Paschos, and O. Pottie, “The max quasi-independent set problem”, *Journal of Combinatorial Optimization*, **23(1)**, pp. 94–117, Jan. 2012.
- [7] D. Corneil and Y. Perl. “Clustering and domination in perfect graphs”. *Discrete Applied Mathematics*, **9(1)**, pp. 27–39, 1984.
- [8] E. Chlamtac, M. Dinitz, and R. Krauthgamer. “Everywhere-Sparse Spanners via Dense Sub-graphsh”, *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pp. 758–767, 2012.
- [9] F. Della Croce and V. Paschos, “On the max k -vertex cover problem”, *Cahiers du LAMSADE*, **307**, Université Paris-Dauphine, 2011.
- [10] U. Feige, D. Peleg, and G. Kortsarz, “The dense k -subgraph problem”, *Algorithmica*, **29(3)**, pp. 410–421, 2001.
- [11] U. Feige and M. Seltser, “On the densest k -subgraph problem”, Technical Report CS97-16, Department of Mathematics and Computer Science, Weizmann Institute Of Science, 1997.
- [12] O. Goldschmidt and D. Hochbaum, “A polynomial time algorithms for the k -cut problem for fixed k ”, *Mathematics of Operations Research*, **19(1)**, pp. 24–37, 1994.
- [13] S. Khot, “Ruling out PTAS for graph min-bisection, densest subgraph and bipartite clique”, *Proceedings of the 44th Symposium on the Foundations of Computer Science (FOCS)*, pp. 136–145, 2004.
- [14] J. Naor and Y. Rabani, “Tree packing and approximating k -cuts”, *Proceedings of the 12th Symposium on Discrete Algorithms (SODA)*, pp. 26–27, 2001.
- [15] Open Problem Forum of the Egerváry Research Group, http://lemon.cs.elte.hu/egres/open/Partial_vertex_cover_in_bipartite_graphs.
- [16] J. Oxley, “What is a matroid?”, *Cubo*, **5**, pp. 179–218, 2003.
- [17] R. Ravi and A. Sinha, “Approximating k -cuts via network strength”, *Proceedings of the 13th Symposium on Discrete Algorithms (SODA)*, pp. 621–622, 2002.
- [18] H. Saran and V. Vazirani, “Finding k -cuts within twice the optimal”, *SIAM Journal on Computing*, **24**, pp. 101–108, 1995.
- [19] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, Springer, 2003.

- [20] L. Zhao, H. Nagamochi and T. Ibaraki, “Approximating the minimum k -way cut in a graph via minimum 3-way cuts”, *Journal of Combinatorial Optimization*, **5**, pp. 397–410, 2001.
- [21] S. Vinterbo, “A note on the hardness of the k -ambiguity problem”, *Technical Report*, 2002.