

(In)approximability of Maximum Minimal FVS

Louis Dublois

Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, Paris, France
louis.dublois@gmail.com

Tesshu Hanaka 

Chuo University, Tokyo, Japan
hanaka.91t@g.chuo-u.ac.jp

Mehdi Khosravian Ghadikolaei

Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, Paris, France
m.khosravian@gmail.com

Michael Lampis 

Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, Paris, France
michail.lampis@dauphine.fr

Nikolaos Melissinos

Université Paris-Dauphine, PSL Research University, CNRS, UMR 7243, LAMSADE, Paris, France
nikolaos.melissinos@dauphine.eu

Abstract

We study the approximability of the NP-complete MAXIMUM MINIMAL FEEDBACK VERTEX SET problem. Informally, this natural problem seems to lie in an intermediate space between two more well-studied problems of this type: MAXIMUM MINIMAL VERTEX COVER, for which the best achievable approximation ratio is \sqrt{n} , and UPPER DOMINATING SET, which does not admit any $n^{1-\epsilon}$ approximation. We confirm and quantify this intuition by showing the first non-trivial polynomial time approximation for MAX MIN FVS with a ratio of $O(n^{2/3})$, as well as a matching hardness of approximation bound of $n^{2/3-\epsilon}$, improving the previous known hardness of $n^{1/2-\epsilon}$. Along the way, we also obtain an $O(\Delta)$ -approximation and show that this is asymptotically best possible, and we improve the bound for which the problem is NP-hard from $\Delta \geq 9$ to $\Delta \geq 6$.

Having settled the problem's approximability in polynomial time, we move to the context of super-polynomial time. We devise a generalization of our approximation algorithm which, for any desired approximation ratio r , produces an r -approximate solution in time $n^{O(n/r^{3/2})}$. This time-approximation trade-off is essentially tight: we show that under the ETH, for any ratio r and $\epsilon > 0$, no algorithm can r -approximate this problem in time $n^{O((n/r^{3/2})^{1-\epsilon})}$, hence we precisely characterize the approximability of the problem for the whole spectrum between polynomial and sub-exponential time, up to an arbitrarily small constant in the second exponent.

2012 ACM Subject Classification Mathematics of computing \rightarrow Graph algorithms; Theory of computation \rightarrow Approximation algorithms analysis

Keywords and phrases Approximation Algorithms, ETH, Inapproximability

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2020.3

Related Version A full version of the paper is available at <https://arxiv.org/abs/2009.09971>.

Funding This work is partially supported by PRC CNRS JSPS project PARAGA and by JSPS KAKENHI Grant Number JP19K21537.

1 Introduction

In a graph $G = (V, E)$, a set $S \subseteq V$ is called a *feedback vertex set* (fvs for short) if the subgraph induced by $V \setminus S$ is a forest. Typically, fvs is studied with a minimization objective: given a graph we are interested in finding the best (that is, smallest) fvs. In this paper we are interested in an objective which is, in a sense, the inverse: we seek an fvs S which is as *large* as possible, while still being minimal. We call this problem MAX MIN FVS.



© Louis Dublois, Tesshu Hanaka, Mehdi Khosravian Ghadikolaei, Michael Lampis, and Nikolaos Melissinos;
licensed under Creative Commons License CC-BY

31st International Symposium on Algorithms and Computation (ISAAC 2020).

Editors: Yixin Cao, Siu-Wing Cheng, and Minming Li; Article No. 3; pp. 3:1–3:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

MaxMin and MinMax versions of many famous optimization problems have recently attracted much interest in the literature (we give references below) and MAX MIN FVS can be seen as a member of this framework. Although the initial motivation for studying such problems was a desire to analyze the worst possible performance of a naive heuristic, these problems have gradually been revealed to possess a rich combinatorial structure that makes them interesting in their own right. Our goal in this paper is to show that MAX MIN FVS displays an interesting complexity behavior with respect to its approximability.

Our motivation for focusing on MAX MIN FVS is the contrast between two of its more well-studied cousins: the MAX MIN VERTEX COVER and UPPER DOMINATING SET problems, where the objective is to find the largest minimal vertex cover or dominating set respectively. At first glance, one would expect MAX MIN VC to be the easier of these two problems: both problems can be seen as trying to find the largest minimal hitting set of a hypergraph, but in the case of MAX MIN VC the hypergraph has a very restricted structure, while in UDS the hypergraph is essentially arbitrary. This intuition turns out to be correct: while UDS admits no $n^{1-\epsilon}$ -approximation [5], MAX MIN VC admits a \sqrt{n} -approximation (but no $n^{1/2-\epsilon}$ -approximation) [9].

This background leads us to the natural question of the approximability of MAX MIN FVS. Intuitively, one may be tempted to think that this problem should be harder than MAX MIN VC, since hitting cycles is more complex than hitting edges, but easier than UDS, since hitting cycles still offers us more structure than an arbitrary hypergraph. However, to the best of our knowledge, no $n^{1-\epsilon}$ -approximation algorithm is currently known for MAX MIN FVS (so the problem could be as hard as UDS), and the best hardness of approximation bound known is $n^{1/2-\epsilon}$ [33] (so the problem could be as easy as MAX MIN VC).

Our main contribution in this paper is to fully answer this question, confirming and precisely quantifying the intuition that MAX MIN FVS is a problem that lies “between” MAX MIN VC and UDS: We give a polynomial-time approximation algorithm with ratio $O(n^{2/3})$ and a hardness of approximation reduction which shows that (unless $P = NP$) no polynomial-time algorithm can obtain a ratio of $n^{2/3-\epsilon}$, for any $\epsilon > 0$. This completely settles the approximability of the problem in polynomial time. Along the way, we also give an approximation algorithm with ratio $O(\Delta)$, show that no algorithm can achieve ratio $\Delta^{1-\epsilon}$, for any $\epsilon > 0$, and improve the best known NP-completeness proof for MAX MIN FVS from $\Delta \geq 9$ [33] to $\Delta \geq 6$, where Δ is the maximum degree of the input graph.

One interesting aspect of our results is that they have an interpretation from extremal combinatorics which nicely mirrors the situation for MAX MIN VC. Recall that a corollary of the \sqrt{n} -approximation for MAX MIN VC [9] is that any graph without isolated vertices has a minimal vertex cover of size at least \sqrt{n} , and this is tight (see Remark 13). Hence, the algorithm only needs to trivially preprocess the graph (deleting isolated vertices) and then find this set, which is guaranteed to exist. Our algorithms can be seen in a similar light: we prove that if one applies two almost trivial pre-processing rules to a graph (deleting leaves and contracting edges between degree-two vertices), a minimal fvs of size at least $n^{1/3}$ (and $\Omega(n/\Delta)$) is always guaranteed to exist, and this is tight (Corollary 11 and Remark 12). Thus, the approximation ratio of $n^{2/3}$ is automatically guaranteed for any graph where we exhaustively apply these very simple rules and our algorithms only have to work to construct the promised set. This makes it somewhat remarkable that the ratio of $n^{2/3}$ turns out to be best possible.

Having settled the approximability of MAX MIN FVS in polynomial time, we consider the question of how much time needs to be invested if one wishes to guarantee an approximation ratio of r (which may depend on n) where $r < n^{2/3}$. This type of time-approximation trade-off was extensively studied by Bonnet et al. [8], who showed that MAX MIN VERTEX COVER admits an r -approximation in time $2^{O(n/r^2)}$ and this is optimal under the randomized ETH.

For MAX MIN FVS we cannot hope to obtain a trade-off with performance exponential in n/r^2 , as this implies a polynomial-time \sqrt{n} -approximation. It therefore seems more natural to aim for a running time exponential in $n/r^{3/2}$. Indeed, generalizing our polynomial-time approximation algorithm, we show that we can achieve an r -approximation in time $n^{O(n/r^{3/2})}$. Although this algorithm reuses some ingredients from our polynomial-time approximation, it is significantly more involved, as it is no longer sufficient to compare the size of our solution to n . We complement our result with a lower bound showing that our algorithm is essentially best possible under the randomized ETH for any r (not just for polynomial time), or more precisely that the exponent of the running time of our algorithm can only be improved by $n^{o(1)}$ factors.

Related work. To the best of our knowledge, MAX MIN FVS was first considered by Mishra and Sikdar [33], who showed that the problem does not admit an $n^{1/2-\epsilon}$ approximation (unless $P = NP$), and that it remains APX-hard for $\Delta \geq 9$. On the other hand, UDS and MAX MIN VC are well-studied problems, both in the context of approximation and in the context of parameterized complexity [1, 5, 9, 11, 13, 14, 19, 28, 30, 34, 36]. Many other classical optimization problems have recently been studied in the MaxMin or MinMax framework, such as MAX MIN SEPARATOR [25], MAX MIN CUT [21], MIN MAX KNAPSACK (also known as the LAZY BUREAUCRAT PROBLEM) [3, 23, 24], and MAX MIN EDGE COVER [32, 26]. Some problems in this area also arise naturally in other forms and have been extensively studied, such as MIN MAX MATCHING (also known as EDGE DOMINATING SET [29]) and GRUNDY COLORING, which can be seen as a Max Min version of COLORING [2, 6].

The idea of designing super-polynomial time approximation algorithms which obtain guarantees better than those possible in polynomial time has attracted much attention in the last decade [4, 10, 16, 18, 20, 22, 31]. As mentioned, the result closest to the time-approximation trade-off we give in this paper is the approximation algorithm for MAX MIN VC given by Bonnet et al. [8]. It is important to note that such trade-offs are only generally known to be tight up to poly-logarithmic factors in the exponent of the running time. As explained in [8], current lower bound techniques can rule out improvements in the running time that shave at least n^ϵ from the exponent, but not improvements which shave poly-logarithmic factors, due to the state of the art in quasi-linear PCP constructions. Indeed, such improvements are sometimes possible [4] and are conceivable for MAX MIN VC and MAX MIN FVS. Lower bounds for this type of algorithm rely on the (randomized) Exponential Time Hypothesis (ETH), which states that there is no (randomized) algorithm for 3-SAT running in time $2^{o(n)}$.

2 Preliminaries

We use standard graph-theoretic notation and only consider simple loop-less graphs. For a graph $G = (V, E)$ and $S \subseteq V$ we denote by $G[S]$ the graph induced by S . For $u \in V$, $G - u$ is the graph $G[V \setminus \{u\}]$. We write $N(u)$ to denote the set of neighbors of u and $d(u) = |N(u)|$ to denote its degree. For $S \subseteq V$, $N(S) = \cup_{u \in S} N(u) \setminus S$. We use $\Delta(G)$ (or simply Δ) to denote the maximum degree of G . For $uv \in E$ the graph G/uv is the graph obtained by contracting the edge uv , that is, replacing u, v by a new vertex connected to $N(u) \cup N(v)$. In this paper we will only apply this operation when $N(u) \cap N(v) = \emptyset$, so the result will always be a simple graph.

A forest is a graph that does not contain cycles. A feedback vertex set (fvs for short) is a set $S \subseteq V$ such that $G[V \setminus S]$ is a forest. An fvs S is minimal if no proper subset of S is an fvs. It is not hard to see that if S is minimal, then every $u \in S$ has a *private cycle*, that

is, there exists a cycle in $G[(V \setminus S) \cup \{u\}]$, which goes through u . A vertex u of a feedback vertex set S that does not have a private cycle (that is, $S \setminus \{u\}$ is also an fvs), is called *redundant*. For a given fvs S , we call the set $F := V \setminus S$ the corresponding induced forest. If S is minimal, then F is maximal.

The main problem we are interested in is MAX MIN FVS: given a graph $G = (V, E)$, find a minimal fvs of G of maximum size. Since this problem is NP-hard, we will be interested in approximation algorithms. An approximation algorithm with ratio $r \geq 1$ (which may depend on n , the order of the graph) is an algorithm which, given a graph G , returns a solution of size at least $\frac{\text{mmfvs}(G)}{r}$, where $\text{mmfvs}(G)$ is the size of the largest minimal fvs of G .

We make two basic observations about our problem: deleting vertices or contracting edges can only decrease the size of the optimal solution.

► **Lemma 1.** *Let $G = (V, E)$ be a graph and $u \in V$. Then, $\text{mmfvs}(G) \geq \text{mmfvs}(G - u)$. Furthermore, given any minimal feedback vertex set S of $G - u$, it is possible to construct in polynomial time a minimal feedback vertex set of G of the same or larger size.*

► **Lemma 2.** *Let $G = (V, E)$ be a graph, $u, v \in V$ with $N(u) \cap N(v) = \emptyset$ and $uv \in E$. Then $\text{mmfvs}(G) \geq \text{mmfvs}(G/uv)$. Furthermore, given any minimal feedback vertex set S of G/uv , it is possible to construct in polynomial time a minimal feedback vertex set of G of the same or larger size.*

Proof. Before we prove the Lemma we note that the contraction operation, under the condition that $N(u) \cap N(v) = \emptyset$, preserves acyclicity in a strong sense: G is acyclic if and only if G/uv is acyclic. Indeed, if we contract an edge that is part of a cycle, this cycle must have length at least 4, and will therefore give a cycle in G/uv . Of course, contractions never create cycles in acyclic graphs.

Let $G' := G/uv$, w be the vertex of G' which has replaced u, v , $V' = V(G')$, and S be a minimal fvs of G' . We have two cases: $w \in S$ or $w \notin S$.

In case $w \in S$, we start with the set $S' = (S \setminus \{w\}) \cup \{u, v\}$. It is not hard to see that S' is an fvs of G . Furthermore, no vertex of $S' \setminus \{u, v\}$ is redundant: for all $z \in S \setminus \{w\}$, there is a cycle in $G'[(V' \setminus S) \cup \{z\}]$, therefore there is also a cycle in $G[(V \setminus S') \cup \{z\}]$. Furthermore, we claim that $S' \setminus \{u, v\}$ is not a valid fvs. Indeed, there must be a cycle contained (due to minimality) in $G_1 = G'[(V' \setminus S) \cup \{w\}]$. Therefore, if there is no cycle in $G_2 = G[(V \setminus S') \cup \{u, v\}]$, we get a contradiction, as G_1 can be obtained by G_2 by contracting the edge uv and contracting edges preserves acyclicity. We conclude that even if S' is not minimal, if we remove vertices until it becomes minimal, we will remove at most one vertex, so the size of the fvs obtained is at least $|S|$.

In case $w \notin S$, we will return the same set S . Let $F = V \setminus S$, $F' = V' \setminus S$. By definition, $G'[F']$ is acyclic. To see that $G[F]$ is also a forest, we note that $G'[F']$ is obtained from $G[F]$ by contracting uv , and as we noted in the beginning, the contractions we use strongly preserve acyclicity. To see that S is minimal, take $z \in S$ and consider the graphs $G_1 = G[(V \setminus S) \cup \{z\}]$ and $G_2 = G'[(V' \setminus S) \cup \{z\}]$. We see that G_2 can be obtained from G_1 by contracting uv . But G_2 must have a cycle, by the minimality of S , so G_1 also has a cycle. Thus, S is minimal in G . ◀

3 Polynomial Time Approximation Algorithm

In this section we present a polynomial-time algorithm which guarantees an approximation ratio of $n^{2/3}$. As we show in Theorem 17, this ratio is the best that can be hoped for in polynomial time. Later (Theorem 15) we show how to generalize the ideas presented here to obtain an algorithm that achieves a trade-off between the approximation ratio and the (sub-exponential) running time, and show that this trade-off is essentially optimal.

On a high level, our algorithm proceeds as follows: first we identify some easy cases in which applying Lemma 1 or Lemma 2 is *safe*, that is, the value of the optimal is guaranteed to stay constant, namely deleting vertices of degree at most 1, and contracting edges between vertices of degree 2. After we apply these reduction rules exhaustively, we compute a minimal fvs S in an arbitrary way. If S is large enough (larger than $n^{1/3}$), we simply return this set.

If not, we apply some counting arguments to show that a vertex $u \in S$ with high degree ($\geq n^{2/3}$) must exist. We then have two cases: either we are able to construct a large minimal fvs just by looking at the neighborhood of u in the forest (and ignoring $S \setminus \{u\}$), or u must share many neighbors with another vertex $v \in S$, in which case we construct a large minimal fvs in the common neighborhood of u, v .

Because our algorithm is constructive (and runs in polynomial time), we find it interesting to remark an interpretation from the point of view of extremal combinatorics, given in Corollary 11.

3.1 Basic Reduction Rules and Combinatorial Tools

We begin by showing two *safe* versions of Lemmas 1, 2.

► **Lemma 3.** *Let G, u be as in Lemma 1 with $d(u) \leq 1$. Then $\text{mmfvs}(G - u) = \text{mmfvs}(G)$.*

► **Lemma 4.** *Let G, u, v be as in Lemma 2 with $d(u) = d(v) = 2$. Then $\text{mmfvs}(G/uv) = \text{mmfvs}(G)$.*

Proof. Let $G' = G/uv$, w be the vertex that replaced u, v in G' , and $V' = V(G')$.

We only need to show that $\text{mmfvs}(G) \leq \text{mmfvs}(G')$, as the other direction is given by Lemma 2. Let S be a minimal fvs of G . We consider two cases:

If $u, v \notin S$, then we claim that S is also a minimal fvs of G' . Indeed, $G'[V' \setminus S]$ is obtained from $G[V \setminus S]$ by contracting uv , so both are acyclic. Furthermore, for all $z \in S$, $G'[(V' \setminus S) \cup \{z\}]$ is obtained from $G[(V \setminus S) \cup \{z\}]$ by contracting uv , therefore both have a cycle, hence no vertex of S is redundant in G' .

If $\{u, v\} \cap S \neq \emptyset$, we claim that exactly one of u, v is in S . Indeed, if $u, v \in S$, then $G'[(V' \setminus S) \cup \{u\}]$ does not contain a cycle going through u , as u has degree 1 in this graph. Without loss of generality, let $u \in S$, $v \notin S$. We set $S' := (S \setminus \{u\}) \cup \{w\}$ and claim that S' is a minimal fvs of G' . It is not hard to see that S' is an fvs of G' , since it corresponds to deleting $S \cup \{v\}$ from G . To see that it is minimal, for all $z \in S' \setminus \{w\}$ we observe that $G'[(V' \setminus S') \cup \{z\}]$ obtained from $G'[(V \setminus S) \cup \{z\}]$ by deleting v , which has degree 1. Therefore, this deletion strongly preserves acyclicity. Finally, to see that w is not redundant for S' we observe that $G'[(V' \setminus S) \cup \{u\}]$ has a cycle, and this cycle must be present in $G'[(V' \setminus S') \cup \{w\}]$, which is obtained from the former graph by contracting uv . ◀

► **Definition 5.** *For a graph $G = (V, E)$ we say that G is reduced if it is not possible to apply Lemma 3 or Lemma 4 to G .*

We now present a counting argument which will be useful in our algorithm and states, roughly, that if in a reduced graph we find a (not necessarily minimal) fvs, that fvs must have many neighbors in the corresponding forest.

► **Lemma 6.** *Let $G = (V, E)$ be a reduced graph and $S \subseteq V$ a feedback vertex set of G . Let $F = V \setminus S$. Then, $|N(S) \cap F| \geq \frac{|F|}{4}$.*

We note that Lemma 6 immediately gives an approximation algorithm with ratio $O(\Delta)$.

► **Lemma 7.** *In a reduced graph G with n vertices and maximum degree Δ , every feedback vertex set has size at least $\frac{n}{5\Delta}$.*

► Remark 8. Lemma 6 is tight.

3.2 Polynomial Time Approximation and Extremal Results

We begin with a final intermediate lemma that allows us to construct a large minimal fvs in any reduced graph that is a forest plus one vertex.

► **Lemma 9.** *Let $G = (V, E)$ be a reduced graph and $u \in V$ such that $G - u$ is acyclic. Then it is possible to construct in polynomial time a minimal feedback vertex set S of G with $|S| \geq d(u)/2$.*

► **Theorem 10.** *There is a polynomial time approximation algorithm for MAX MIN FVS with ratio $O(n^{2/3})$.*

Proof. We are given a graph $G = (V, E)$. We begin by applying Lemmas 3,4 exhaustively in order to obtain a reduced graph $G' = (V', E')$. Clearly, if we obtain a $|V'|^{1/3}$ approximation in G' , since the reductions we applied do not change the optimal, and we can construct a solution of the same size in G , we get a $|V'|^{2/3} \leq |V|^{2/3}$ approximation ratio in G . So, in the remainder, to ease presentation, we assume G is already reduced and has n vertices.

Our algorithm begins with an arbitrary minimal fvs S . This can be constructed, for example, by starting with $S = V$ and removing vertices from S until it becomes minimal. If $|S| \geq n^{1/3}$ then we return S . Since the optimal solution cannot have size more than n , we achieve the claimed ratio.

Suppose then that $|S| < n^{1/3}$. Let F be the corresponding forest. We have $|F| > n - n^{1/3} > n/2$ for sufficiently large n . By Lemma 6, $|N(S) \cap F| \geq n/8$. Since $|S| < n^{1/3}$ there must exist $u \in S$ such that u has at least $\frac{n^{2/3}}{8}$ neighbors in F .

Let $w \in F \cap N(u)$. We say that w is a *good* neighbor of u if there exists $w' \in F \cap N(u)$ with $w' \neq w$ and w' is in the same tree of $G[F]$ as w . Otherwise w is a *bad* neighbor of u . By extension, a tree of $G[F]$ that contains a good (resp. bad) neighbor of u will be called a good (resp. bad) tree. Every vertex of $N(u) \cap F$ is either good or bad.

We have argued that $|N(u) \cap F| \geq \frac{n^{2/3}}{8}$. We distinguish two cases: either u has at least $\frac{n^{2/3}}{16}$ good neighbors in F , or it has at least that many bad neighbors in F .

In the former case, we delete from the graph the set $S \setminus \{u\}$ and apply Lemmas 3, 4 exhaustively again. We claim that the number of good neighbors of u does not decrease in this process. Indeed, two good neighbors of u cannot be contracted using Lemma 4, since they have a common neighbor (namely u). Furthermore, suppose w is the first good neighbor of u to be deleted using Lemma 3. This would mean that w currently has no other neighbor except u . However, since w is good, initially there was a $w' \in N(u)$ in the same tree of $G[F]$ as w . The vertex w' has not been deleted (since we assumed w is the first good neighbor to be deleted). Furthermore, Lemmas 3, 4 cannot disconnect two vertices which are in the same component, so we get a contradiction. We therefore have a reduced graph, where $\{u\}$ is an fvs, and $d(u) \geq \frac{n^{2/3}}{16}$. By Lemma 9 we obtain a minimal fvs of size at least $\frac{n^{2/3}}{32}$, which is an $O(n^{1/3})$ approximation.

In the latter case, u has at least $\frac{n^{2/3}}{16}$ bad neighbors in F . Consider a bad tree T . We claim that T must have a neighbor in $S \setminus \{u\}$, because T has at least two leaves, at most one of which is a neighbor of u (since T is bad). If the second leaf is not connected to S , it will be deleted by Lemma 3. Furthermore, since u is connected to one vertex in each bad tree, u is connected to at least $\frac{n^{2/3}}{16}$ bad trees.

We now find the vertex $v \in S \setminus \{u\}$ such that v is connected to the maximum number of bad trees connected to u . Since $|S| \leq n^{1/3}$, v must be connected to at least $\frac{n^{1/3}}{16}$ bad trees connected to u . We now delete from the graph the set $S \setminus \{u, v\}$ as well as all trees of $G[F]$, except the bad trees connected to u, v . Furthermore, in each bad tree T connected to both u, v let $u' \in T \cap N(u)$ and $v' \in T \cap N(v)$ such that u', v' are as close as possible in T (note that perhaps $v' = u'$). We delete all vertices of the tree T except those on the path from v' to u' . Then, we contract all internal edges of this path (note that internal vertices of the path are not connected to $\{u, v\}$ by the selection of u', v'). It is not hard to verify that, by using Lemmas 1, 2, if we are able to produce a large minimal fvs in the resulting graph, we obtain a solution for G . Furthermore, in the resulting graph, every bad tree T connected to u, v has been reduced to a single vertex connected to u, v . So the graph is now either a $K_{2,s}$, with $s \geq \frac{n^{1/3}}{16}$, or the same graph with the addition of the edge uv . In either case, it is not hard to see that starting with the fvs that contains all vertices except $\{u, v\}$, and making it minimal, we obtain a solution of size at least $s - 1$ which gives an approximation ratio of $O(n^{2/3})$. \blacktriangleleft

► **Corollary 11.** *For any reduced graph G on n vertices we have $\text{mmfvs}(G) = \Omega(n^{1/3})$.*

► **Remark 12.** Corollary 11 is tight.

Finally, we remark that a similar combinatorial point of view can be taken for the related problem of MAX MIN VC, giving another intuitive explanation for the difference in approximability between the two problems.

► **Remark 13.** Any graph $G = (V, E)$ without isolated vertices, has a minimal vertex cover of size at least $\sqrt{|V|}$, and this is asymptotically tight.

4 Sub-exponential Time Approximation

In this section we give an approximation algorithm that generalizes our $n^{2/3}$ -approximation and is able to guarantee any desired performance, at the cost of increased running time. On a high level, our initial approach again constructs an arbitrary minimal fvs S and if S is clearly large enough, returns it. However, things become more complicated from then on, as it is no longer sufficient to consider vertices of S individually or in pairs. We therefore need several new ideas, one of which is given in the following lemma, which states that we can find a constant factor approximation in time exponential in the size of a given fvs. This will be useful as we will use the assumption that S is “small” and then cut it up into even smaller pieces to allow us to use Lemma 14.

► **Lemma 14.** *Given a graph $G = (V, E)$ on n vertices and a feedback vertex set $S \subseteq V$ of size k , it is possible to produce a minimal fvs S' of G of size $|S'| \geq \frac{\text{mmfvs}(G)}{3}$ in time $n^{O(k)}$.*

► **Theorem 15.** *There is an algorithm which, given an n -vertex graph $G = (V, E)$ and a value r , produces an r -approximation for MAX MIN FVS in G in time $n^{O(n/r^{3/2})}$.*

Proof. First, let us note that we may assume that r is $\omega(1)$, because if r is bounded by a constant, then we can solve the problem exactly in the given time. To ease presentation, we will give an algorithm with approximation ratio $O(r)$. A ratio of exactly r can be obtained by multiplying r with an appropriate (small) constant.

Our algorithm borrows several of the basic ideas from Theorem 10, but requires some new ingredients (including Lemma 14). The first step is, again, to construct a minimal fvs S in some arbitrary way, for example by setting $S = V$ and then removing vertices from S until

it becomes minimal. If $|S| \geq n/r$ we are done, as we already have an r -approximation, so we simply return S . From this point, this algorithm departs from the algorithm of Theorem 10, because it is no longer sufficient to compare the size of the returned solution with a function of n (we need to compare it to the actual optimal in order to obtain a ratio of r), and because we need to partition S into non-trivial parts that contain more than one vertex. The algorithm proceeds as follows:

Let $k = \lceil \sqrt{r} \rceil$ and partition S into k parts of (almost) equal size S_1, \dots, S_k . Our algorithm proceeds as follows: for each $i, j \in \{1, \dots, k\}$ (not necessarily distinct) consider the graph $G_{i,j}$ obtained by deleting all vertices of $S \setminus (S_i \cup S_j)$. Compute, using Lemma 14 a solution for $G_{i,j}$, taking into account that $S_i \cup S_j$ is a feedback vertex set of this graph. Output the largest of the solutions found, using Lemma 1 to transform them into solutions of G (or output S if it is larger than all solutions).

The algorithm clearly runs in the promised time: $|S_i \cup S_j| \leq \frac{2n}{rk}$, so the algorithm of Lemma 14 takes time $n^{O(n/r^{3/2})}$ and is executed a polynomial number of times.

Let us now analyze the approximation ratio of the produced solution. Let S_{OPT} be an optimal solution and let $F := V \setminus S$ and $F_{OPT} = V \setminus S_{OPT}$ be the induced forests corresponding to S and to the optimal solution. We would like to argue that one of the considered subproblems contains at least a $\frac{1}{r}$ fraction of S_{OPT} and that most (though not all) of these vertices form part of a minimal fvs of that subgraph.

To be more precise, we will define the notion of “type” for each $u \in S_{OPT} \cap F$. For each such u there must exist a cycle in the graph $G[F_{OPT} \cup \{u\}]$ (if not, this would contradict the minimality of S_{OPT}). Call this cycle $c(u)$ (select one such cycle arbitrarily if several exist). The cycle $c(u)$ must intersect S , as S is an fvs. Let v be the vertex of $c(u) \cap S$ closest to u on the cycle. Let v' be the vertex of $c(u) \cap S$ that is closest to u if we traverse the cycle in the opposite direction (note that v, v' are not necessarily distinct). Suppose that $v \in S_i, v' \in S_j$ and without loss of generality $i \leq j$. We then say that $u \in S_{OPT} \cap F$ has type (i, j) . In this way, we define a type for each $u \in S_{OPT} \cap F$. Note that according to our definition, all internal vertices of the path in $c(u)$ from u to v (and also from u to v') belong in $F_{OPT} \cap F$.

According to the definition of the previous paragraph, there are $k(k+1)/2 \leq r$ possible types of vertices in $S_{OPT} \cap F$. Therefore, there must be a type (i, j) such that at least $\frac{|S_{OPT} \cap F|}{r}$ vertices have this type. We now concentrate on the graph $G_{i,j}$, for the type (i, j) which satisfies this condition. Our algorithm constructed $G_{i,j}$ by deleting all of S except $S_i \cup S_j$. We would like to claim that this graph has a minimal feedback vertex set of size comparable to $\frac{|S_{OPT} \cap F|}{r}$.

For the sake of the analysis, construct a minimal feedback vertex set S^* of $G_{i,j}$ as follows: we begin with the fvs $S^* = S_{OPT} \cap (F \cup S_i \cup S_j)$ and the corresponding induced forest $F^* = F_{OPT} \cap (F \cup S_i \cup S_j)$. The set S^* is a feedback vertex set as it contains all vertices of S_{OPT} found in $G_{i,j}$ and S_{OPT} is a feasible feedback vertex set of all of G . We then make S^* minimal by arbitrarily removing redundant vertices. Call the resulting set $S^{**} \subseteq S^*$ and the corresponding induced forest $F^{**} \supseteq F^*$.

Our main claim now is that the number of vertices of $S^* \cap F$ of type (i, j) which were “lost” in the process of making S^* minimal, is upper-bounded by $|S_i \cup S_j|$. Formally, we claim that $|\{u \in (S^* \cap F) \setminus S^{**} \mid u \text{ has type } (i, j)\}| \leq |S_i \cup S_j|$. Indeed, consider such a vertex $u \in (S^* \cap F) \setminus S^{**}$ of type (i, j) , let $c(u)$ be the cycle that defines its type and v, v' the vertices of $S_i \cup S_j$ which are closest to u on the cycle in either direction. All vertices of $c(u)$ in the paths from u to v and from u to v' belong to $F_{OPT} \cap F$, therefore also to F^* . If u was removed as redundant, this means that v, v' must have been in distinct connected

components at the moment u was removed from the feedback vertex set (and also that v, v' are distinct). However, the addition of u to the induced forest creates a path from v to v' in the induced forest and hence decreases the number of connected components (that is, trees in the induced forest) containing vertices of $S_i \cup S_j$. The number of such connected components cannot decrease more than $|S_i \cup S_j|$ times, therefore, during the process of making S^* minimal we may have removed at most $|S_i \cup S_j|$ vertices of type (i, j) from $S^* \cap F$.

Using the above analysis and the assumption that S^* contains at least $\frac{|S_{OPT} \cap F|}{r}$ vertices of type (i, j) , we conclude that $\text{mmfvs}(G_{i,j}) \geq |S^{**}| \geq \frac{|S_{OPT} \cap F|}{r} - |S_i \cup S_j|$. We now note that if $|S_{OPT} \cap S| \geq \frac{|S_{OPT}|}{r}$, then S is already an r -approximation, so it is safe to assume $|S_{OPT} \cap F| \geq \frac{(r-1)|S_{OPT}|}{r}$. Furthermore, $|S_i \cup S_j| \leq \frac{2|S|}{\sqrt{r}} \leq \frac{2|S_{OPT}|}{r\sqrt{r}}$, where again we are assuming that S is not already an r -approximation. Putting things together we get $\text{mmfvs}(G_{i,j}) \geq \frac{(r-1)|S_{OPT}|}{r^2} - \frac{2|S_{OPT}|}{r\sqrt{r}} \geq \frac{|S_{OPT}|}{2r}$, for sufficiently large r . Hence, since the algorithm will return a solution that is at least as large as $\frac{\text{mmfvs}(G_{i,j})}{3}$, we obtain an $O(r)$ -approximation. \blacktriangleleft

5 Hardness of Approximation and NP-hardness

In this section we establish lower bound results showing that the approximation algorithms given in Theorems 10 and 15 are essentially optimal, under standard complexity assumptions.

5.1 Hardness of Approximation in Polynomial Time

We begin by showing that the best approximation ratio achievable in polynomial time is indeed (essentially) $n^{2/3}$. For this, we rely on the celebrated result of Håstad on the hardness of approximating MAX INDEPENDENT SET, which was later derandomized by Zuckerman, cited below.

► **Theorem 16** ([27, 35]). *For any $\epsilon > 0$, there is no polynomial time algorithm which approximates MAX INDEPENDENT SET with a ratio of $n^{1-\epsilon}$, unless $P = NP$.*

Starting from this result, we present a reduction to MAX MIN FVS.

► **Theorem 17.** *For any $\epsilon > 0$, MAX MIN FVS is inapproximable within a factor of $n^{2/3-\epsilon}$ unless $P = NP$.*

Proof. We give a gap-preserving reduction from MAX INDEPENDENT SET, which cannot be approximated within a factor of $n^{1-\epsilon}$, unless $P = NP$. We are given a graph $G = (V, E)$ on n vertices as an instance of MAX INDEPENDENT SET. Recall that $\alpha(G)$ denotes the size of the maximum independent set of G .

We transform G into an instance of MAX MIN FVS as follows: For every pair of $u, v \in V$, we add n vertices such that they are adjacent only to u and v . We denote by I_{uv} the set of such vertices. Then I_{uv} is an independent set. Let $G' = (V', E')$ be the constructed graph.

We now make the following two claims:

▷ **Claim 18.** $\text{mmfvs}(G') \geq (n-1)\binom{\alpha(G)}{2}$

Proof. We construct a minimal fvs of G' as follows: let C be a minimum vertex cover of G . Then we begin with the set that contains C and the union of all I_{uv} (which is clearly an fvs) and remove vertices from it until it becomes minimal. Let S be the final minimal fvs. We observe that for all $u, v \in V \setminus C$, S contains at least $n-1$ of the vertices of I_{uv} . Since C is a minimum vertex cover of G , there are $\binom{\alpha(G)}{2}$ pairs $u, v \in V \setminus C$. \blacktriangleleft

▷ **Claim 19.** $\text{mmfvs}(G') \leq n \binom{2\alpha(G)}{2} + n$

Proof. Let S be a minimal fvs of G' and F be the corresponding forest. It suffices to show that $|S \setminus V| \leq n \binom{2\alpha(G)}{2}$, since $|S \cap V| \leq n$. Consider now a set I_{uv} . If $u \in S$ or $v \in S$, then $I_{uv} \cap S = \emptyset$, because all vertices of I_{uv} have at most one neighbor in F , and are therefore redundant. So, I_{uv} contains (at most n) vertices of S only if $u, v \in F$. However, $|F \cap V| \leq 2\alpha(G)$, because F is bipartite, so $F \cap V$ induces two independent sets, both of which must be at most equal to the maximum independent set of G . So the number of pairs $u, v \in F \cap V$ is at most $\binom{2\alpha(G)}{2}$ and since each corresponding I_{uv} has size n , we get the promised bound. ◁

The two claims together imply that there exist constants c_1, c_2 such that (for sufficiently large n) we have $c_1 n (\alpha(G))^2 \leq \text{mmfvs}(G') \leq c_2 n (\alpha(G))^2$. That is, $\text{mmfvs}(G') = \Theta(n(\alpha(G))^2)$.

Suppose now that there exists a polynomial-time approximation algorithm which, given a graph G' , produces a minimal fvs S with the property $\frac{\text{mmfvs}(G')}{r} \leq |S| \leq \text{mmfvs}(G')$, that is, there exists an r -approximation for MAX MIN FVS. Running this algorithm on the instance we constructed, we obtain that $\frac{c_1 n (\alpha(G))^2}{r} \leq |S| \leq c_2 n (\alpha(G))^2$. Therefore, $\frac{\alpha(G)}{\sqrt{rc_2/c_1}} \leq \sqrt{\frac{|S|}{c_2 n}} \leq \alpha(G)$. As a result, we obtain an $O(\sqrt{r})$ approximation for the value of $\alpha(G)$. We therefore conclude that, unless $P = NP$, any such algorithm must have $\sqrt{r} > n^{1-\epsilon}$, for any $\epsilon > 0$, hence, $r > n^{2-\epsilon}$, for any $\epsilon > 0$. Since the graph G' has $N = \Theta(n^3)$ vertices, we get that no approximation algorithm can achieve a ratio of $N^{2/3-\epsilon}$. ◀

We notice that in the construction of the previous theorem, the maximum degree of the graph is approximately equal to the approximation gap. Thus, the following corollary also holds.

► **Corollary 20.** *For any positive constant ϵ , MAX MIN FVS is inapproximable within a factor of $\Delta^{1-\epsilon}$ unless $P = NP$.*


5.2 Hardness of Approximation in Sub-Exponential Time

In this section we extend Theorem 17 to the realm of sub-exponential time algorithms. We recall the following result of Chalermsook et al.

► **Theorem 21** ([12]). *For any $\epsilon > 0$ and any sufficiently large r , if there exists an r -approximation algorithm for MAX INDEPENDENT SET running in $2^{(n/r)^{1-\epsilon}}$, then the randomized ETH is false.*

We remark that Theorem 21, which gives an almost tight running time lower bound for MAX INDEPENDENT SET, has already been used as a starting point to derive a similarly tight bound for the running time of any sub-exponential time approximation for MAX MIN VC. Here, we modify the proof of Theorem 17 to obtain a similarly tight result for MAX MIN FVS. Nevertheless, the reduction for MAX MIN FVS is significantly more challenging, because the ideas used in Theorem 17 involve an inherent quadratic (in n) blow-up of the size of the instance. As a result, in addition to executing an appropriately modified version of the reduction of Theorem 17, we are forced to add an extra “sparsification” step, and use a probabilistic analysis with Chernoff bounds to argue that this step does not destroy the inapproximability gap.

► **Theorem 22.** *For any $\epsilon > 0$ and any sufficiently large r , if there exists an r -approximation algorithm for MAX MIN FVS running in $2^{(n/r^{3/2})^{1-\epsilon}}$, then the randomized ETH is false.*



■ **Figure 1** The edge gadget of $e = (u, v)$ in the constructed graph G .

5.3 NP-hardness for $\Delta = 6$

► **Theorem 23.** *MAX MIN FVS is NP-hard on planar bipartite graphs with $\Delta = 6$.*

Proof. We give a reduction from MAX MIN VC, which is NP-hard on planar bipartite graphs of maximum degree 3 [36]. Note that the NP-hardness in [36] is stated for MINIMUM INDEPENDENT DOMINATING SET, but any independent dominating set is also a maximal independent set (and vice-versa) and the complement of the minimum maximal independent set of any graph is a maximum minimal vertex cover. Thus, we also obtain NP-hardness for MAX MIN VC on the same instances.

We are given a graph $G = (V, E)$. For each edge $e = (u, v) \in E$, we add a path of length three from u to v going through two new vertices $e^{(1)}, e^{(2)}$ (see Figure 1). Note that $u, e^{(1)}, e^{(2)}, v$ form a cycle of length 4. Then we add two cycles of length 4, $e^{(i)}, c_{e1}^{(i)}, c_{e2}^{(i)}, c_{e3}^{(i)}$ and $e^{(i)}, c_{e4}^{(i)}, c_{e5}^{(i)}, c_{e6}^{(i)}$ for $i \in \{1, 2\}$. Let $G' = (V', E')$ be the constructed graph. Because $\Delta(G) = 3$, we have $\Delta(G') = 6$. Moreover, since G is planar and bipartite, G' is also planar and bipartite. We will show that there is a minimal vertex cover of size at least k in G if and only if there is a minimal feedback vertex set of size at least $k + 4|E|$ in G' .

Given a minimal vertex cover S of size at least k in G , we construct the set $S' = S \cup \bigcup_{e \in E} \{c_{e1}^{(1)}, c_{e4}^{(1)}, c_{e1}^{(2)}, c_{e4}^{(2)}\}$. Then $|S'| \geq k + 4|E|$. Let us first argue that S' is an fvs of G' . For each $e = (u, v) \in E$ we have at least one of $u, v \in S$, without loss of generality let $u \in S$. Now in $G'[V' \setminus S']$ the edges $(e^{(1)}, e^{(2)})$ and $(e^{(2)}, v)$ are bridges and therefore cannot be part of any cycle. The remaining cycles going through $e^{(1)}, e^{(2)}$ are handled by $\{c_{e1}^{(1)}, c_{e4}^{(1)}, c_{e1}^{(2)}, c_{e4}^{(2)}\}$. Furthermore, since $G'[V' \setminus S]$ is an independent set, it is also acyclic. To see that S' is a minimal fvs, we remark that for each $c_{e1}^{(i)}, c_{e4}^{(i)}$ contained in S' there is a private cycle in $G'[V' \setminus S']$. We also note that since S is a minimal vertex cover of G , for each $u \in S$, there exists $v \notin S$ with $e = (u, v) \in E$. This means that u has the private cycle formed by $\{u, v, e^{(1)}, e^{(2)}\}$ in $G'[V' \setminus S']$. Therefore, S' is a minimal fvs.

Conversely, suppose we are given a minimal fvs S' of G' with $|S'| \geq k + 4|E|$. We will edit S' so that it contains only vertices in $V' \setminus \bigcup_{e \in E} \{e^{(1)}, e^{(2)}\}$, without decreasing its size.

First, suppose $e^{(1)}, e^{(2)} \in S'$, for some $e \in E$. We construct a new minimal fvs $S'' = S' \setminus \{e^{(2)}\} \cup \{c_{e1}^{(2)}, c_{e4}^{(2)}\}$ which is larger than S' , since by minimality we have $c_{ei}^{(2)} \notin S'$ for $i \in \{1, \dots, 6\}$. It is not hard to see that S'' is indeed an fvs, as no cycle can go through $e^{(2)}$ in $G'[V' \setminus S'']$. The two vertices we added have a private cycle, while all vertices of $S' \cap S''$ retain their private cycles, so S'' is a minimal fvs. As a result in the remainder we assume that S' contains at most one of $\{e^{(1)}, e^{(2)}\}$ for all $e \in E$.

Suppose now that for some $e = (u, v) \in E$, we have $S' \cap \{u, v\} \neq \emptyset$ and $S' \cap \{e^{(1)}, e^{(2)}\} \neq \emptyset$. Without loss of generality, let $e^{(1)} \in S'$. We set $S'' = S' \setminus \{e^{(1)}\} \cup \{c_{e1}^{(1)}, c_{e4}^{(1)}\}$ and claim that S'' is a larger minimal fvs than S . Indeed, no cycle goes through $e^{(1)}$ in $G'[V' \setminus S'']$, the new

vertices we added to S' have private cycles, and all vertices of $S' \cap S''$ retain their private cycles in $G'[V' \setminus S'']$. Therefore, we can now assume that if for some $e = (u, v) \in E$ we have $S' \cap \{e^{(1)}, e^{(2)}\} \neq \emptyset$ then $u, v \notin S'$.

For the remaining case, suppose that for some $e = (u, v) \in E$ we have $u, v \notin S'$ and (without loss of generality) $e^{(1)} \in S'$. We construct the set $S'' = S' \setminus \{e^{(1)}\} \cup \{c_{e1}^{(1)}, c_{e4}^{(2)}, u\}$. Note that $|S''| \geq |S'| + 2$. It is not hard to see that S'' is an fvs, since by adding $c_{e1}^{(1)}, c_{e4}^{(1)}, v$ to our set we have hit all cycles containing $e^{(1)}$ in G' . The problem now is that S'' is not necessarily minimal. We greedily delete vertices from S'' to obtain a minimal fvs S^* . We claim that in this process we cannot delete more than two vertices, that is $|S^* \setminus S''| \leq 2$. To see this, we first note that $c_{e1}^{(1)}, c_{e4}^{(2)}, u$ cannot be removed from S'' as they have private cycles in $G'[V' \setminus S'']$. Suppose now that $w_1 \in S'' \setminus S^*$ is the first vertex we removed from S'' , so $G'[(V' \setminus S'') \cup \{w_1\}]$ is acyclic. This vertex must have had a private cycle in $G'[V' \setminus S']$, which was necessarily going through u . Therefore, $G'[(V' \setminus S'') \cup \{w_1\}]$ has a path connecting two neighbors of u and this path does not exist in $G'[(V' \setminus S'')]$. With a similar reasoning, removing another vertex $w_2 \in S''$ from the fvs will create a second path between neighbors of u in the induced forest. We conclude that this cannot happen a third time, since $|N(u)| \leq 3$, and if we create three paths between neighbors of u , this will create a cycle. As a result, $|S^*| \geq |S'|$. We assume in the remainder that S' does not contain $e^{(1)}, e^{(2)}$ for any $e \in E$.

Now, given a minimal fvs S' of G' with $|S'| \geq k + 4|E|$ and $S' \cap (\cup_{e \in E} \{e^{(1)}, e^{(2)}\}) = \emptyset$ we set $S = S' \cap V$ and claim that S is a minimal vertex cover of G with $|S| \geq k$. Indeed S is a vertex cover, as for each $e = (u, v) \in E$, if $u, v \notin S'$ then we would get the cycle formed by $\{u, v, e^{(1)}, e^{(2)}\}$. To see that S is minimal, suppose $N_G[u] \subseteq S'$. We claim that in that case u has no private cycle in $G'[V' \setminus S']$ (this can be seen by deleting all bridges in $G'[V' \setminus S']$, which leaves u isolated). This contradicts the minimality of S' as an fvs of G' . Finally, we argue that $|S' \setminus V| \leq 4|E|$, which gives the desired bound on $|S|$. Consider an $e = (u, v) \in E$. S' cannot contain more than one vertex among $c_{e1}^{(1)}, c_{e2}^{(1)}, c_{e3}^{(1)}$, since any of these vertices hits the cycle that goes through the others. With similar reasoning for the three other length-four cycles we conclude that S' contains at most 4 vertices for each edge $e \in E$. ◀

6 Conclusions

We have essentially settled the approximability of MAX MIN FVS for polynomial and sub-exponential time, up to sub-polynomial factors in the exponent of the running time. It would be interesting to see if the running time of our sub-exponential approximation algorithm can be improved by poly-logarithmic factors in the exponent, as in [4]. In particular, improving the running time to $2^{O(n/r^{3/2})}$ seems feasible, but would likely require a version of Lemma 14 which uses more sophisticated techniques, such as Cut&Count [7, 15, 17].

Another problem of similar spirit which deserves to be studied is MAX MIN OCT, where an odd cycle transversal (OCT) is a set of vertices whose removal makes the graph bipartite. This problem could also potentially be “between” MAX MIN VC and UDS, but obtaining a $n^{1-\epsilon}$ approximation for it seems much more challenging than for MAX MIN FVS.

References

- 1 Hassan AbouEisha, Shahid Hussain, Vadim V. Lozin, Jérôme Monnot, Bernard Ries, and Viktor Zamaraev. Upper domination: Towards a dichotomy through boundary properties. *Algorithmica*, 80(10):2799–2817, 2018. doi:10.1007/s00453-017-0346-9.
- 2 Pierre Aboulker, Édouard Bonnet, Eun Jung Kim, and Florian Sikora. Grundy coloring & friends, half-graphs, bicliques. In *STACS*, volume 154 of *LIPIcs*, pages 58:1–58:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.STACS.2020.58.

- 3 Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Steven Skiena. The lazy bureaucrat scheduling problem. *Inf. Comput.*, 184(1):129–146, 2003. doi:10.1016/S0890-5401(03)00060-9.
- 4 Nikhil Bansal, Parinya Chalermsook, Bundit Laekhanukit, Danupon Nanongkai, and Jesper Nederlof. New tools and connections for exponential-time approximation. *Algorithmica*, 81(10):3993–4009, 2019. doi:10.1007/s00453-018-0512-8.
- 5 Cristina Bazgan, Ljiljana Brankovic, Katrin Casel, Henning Fernau, Klaus Jansen, Kim-Manuel Klein, Michael Lampis, Mathieu Liedloff, Jérôme Monnot, and Vangelis Th. Paschos. The many facets of upper domination. *Theoretical Computer Science*, 717:2–25, 2018. doi:10.1016/j.tcs.2017.05.042.
- 6 Rémy Belmonte, Eun Jung Kim, Michael Lampis, Valia Mitsou, and Yota Otachi. Grundy distinguishes treewidth from pathwidth. *CoRR*, abs/2008.07425, 2020. doi:10.4230/LIPIcs.ESA.2020.14.
- 7 Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Inf. Comput.*, 243:86–111, 2015. doi:10.1016/j.ic.2014.12.008.
- 8 Édouard Bonnet, Michael Lampis, and Vangelis Th. Paschos. Time-approximation trade-offs for inapproximable problems. *Journal of Computer and System Sciences*, 92:171–180, 2018. doi:10.1016/j.jcss.2017.09.009.
- 9 Nicolas Boria, Federico Della Croce, and Vangelis Th. Paschos. On the max min vertex cover problem. *Discrete Applied Mathematics*, 196:62–71, 2015. doi:10.1016/j.dam.2014.06.001.
- 10 Nicolas Bourgeois, Bruno Escoffier, and Vangelis Th. Paschos. Approximation of min coloring by moderately exponential algorithms. *Inf. Process. Lett.*, 109(16):950–954, 2009. doi:10.1016/j.ipl.2009.05.002.
- 11 Arman Boyaci and Jérôme Monnot. Weighted upper domination number. *Electron. Notes Discret. Math.*, 62:171–176, 2017. doi:10.1016/j.endm.2017.10.030.
- 12 Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *FOCS*, pages 370–379. IEEE Computer Society, 2013. doi:10.1109/FOCS.2013.47.
- 13 Grant A. Cheston, Gerd Fricke, Stephen T. Hedetniemi, and David Pokrass Jacobs. On the computational complexity of upper fractional domination. *Discret. Appl. Math.*, 27(3):195–207, 1990. doi:10.1016/0166-218X(90)90065-K.
- 14 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000. doi:10.1007/s002249910009.
- 15 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer International Publishing, 2015. doi:10.1007/978-3-319-21275-3.
- 16 Marek Cygan, Lukasz Kowalik, and Mateusz Wykurz. Exponential-time approximation of weighted set cover. *Inf. Process. Lett.*, 109(16):957–961, 2009. doi:10.1016/j.ipl.2009.05.003.
- 17 Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, Michal Pilipczuk, Johan M. M. van Rooij, and Jakub Onufry Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. In *FOCS*, pages 150–159. IEEE Computer Society, 2011. doi:10.1109/FOCS.2011.23.
- 18 Marek Cygan and Marcin Pilipczuk. Exact and approximate bandwidth. *Theor. Comput. Sci.*, 411(40-42):3701–3713, 2010. doi:10.1016/j.tcs.2010.06.018.
- 19 Marc Demange. A note on the approximation of a minimum-weight maximal independent set. *Computational Optimization and Applications*, 14(1):157–169, 1999. doi:10.1023/A:1008765214400.

- 20 Bruno Escoffier, Vangelis Th. Paschos, and Emeric Tourniaire. Approximating MAX SAT by moderately exponential and parameterized algorithms. *Theor. Comput. Sci.*, 560:147–157, 2014. doi:10.1016/j.tcs.2014.10.039.
- 21 Hiroshi Eto, Tesshu Hanaka, Yasuaki Kobayashi, and Yusuke Kobayashi. Parameterized Algorithms for Maximum Cut with Connectivity Constraints. In *IPEC 2019*, pages 13:1–13:15, 2019. doi:10.4230/LIPIcs.IPEC.2019.13.
- 22 Dimitris Fotakis, Michael Lampis, and Vangelis Th. Paschos. Sub-exponential approximation schemes for csps: From dense to almost sparse. In *STACS*, volume 47 of *LIPIcs*, pages 37:1–37:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.STACS.2016.37.
- 23 Fabio Furini, Ivana Ljubic, and Markus Sinnl. An effective dynamic programming algorithm for the minimum-cost maximal knapsack packing problem. *European Journal of Operational Research*, 262(2):438–448, 2017. doi:10.1016/j.ejor.2017.03.061.
- 24 Laurent Gourvès, Jérôme Monnot, and Aris Pagourtzis. The lazy bureaucrat problem with common arrivals and deadlines: Approximation and mechanism design. In *FCT*, volume 8070 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2013. doi:10.1007/978-3-642-40164-0_18.
- 25 Tesshu Hanaka, Hans L. Bodlaender, Tom C. van der Zanden, and Hirotaka Ono. On the maximum weight minimal separator. *Theoretical Computer Science*, 796:294–308, 2019. doi:10.1016/j.tcs.2019.09.025.
- 26 Ararat Harutyunyan, Mehdi Khosravian Ghadikolaei, Nikolaos Melissinos, Jérôme Monnot, and Aris Pagourtzis. On the complexity of the upper r -tolerant edge cover problem. In Luís Soares Barbosa and Mohammad Ali Abam, editors, *Topics in Theoretical Computer Science - Third IFIP WG 1.8 International Conference, TTCS 2020, Tehran, Iran, July 1-2, 2020, Proceedings*, volume 12281 of *Lecture Notes in Computer Science*, pages 32–47. Springer, 2020. doi:10.1007/978-3-030-57852-7_3.
- 27 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math*, 182:105–142, 1999.
- 28 Michael A. Henning and Dinabandhu Pradhan. Algorithmic aspects of upper paired-domination in graphs. *Theor. Comput. Sci.*, 804:98–114, 2020. doi:10.1016/j.tcs.2019.10.045.
- 29 Ken Iwaida and Hiroshi Nagamochi. An improved algorithm for parameterized edge dominating set problem. *J. Graph Algorithms Appl.*, 20(1):23–58, 2016. doi:10.7155/jgaa.00383.
- 30 Michael S. Jacobson and Kenneth Peters. Chordal graphs and upper irredundance, upper domination and independence. *Discret. Math.*, 86(1-3):59–69, 1990. doi:10.1016/0012-365X(90)90349-M.
- 31 Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Improved (in-)approximability bounds for d -scattered set. In *WAOA*, volume 11926 of *Lecture Notes in Computer Science*, pages 202–216. Springer, 2019. doi:10.1007/978-3-030-39479-0_14.
- 32 Kaveh Khoshkhah, Mehdi Khosravian Ghadikolaei, Jérôme Monnot, and Florian Sikora. Weighted upper edge cover: Complexity and approximability. *J. Graph Algorithms Appl.*, 24(2):65–88, 2020. doi:10.7155/jgaa.00519.
- 33 Sounaka Mishra and Kripasindhu Sikdar. On the hardness of approximating some NP-optimization problems related to minimum linear ordering problem. *RAIRO Theor. Informatics Appl.*, 35(3):287–309, 2001. doi:10.1051/ita:2001121.
- 34 Meirav Zehavi. Maximum minimal vertex cover parameterized by vertex cover. *SIAM Journal on Discrete Mathematics*, 31(4):2440–2456, 2017. doi:10.1137/16M109017X.
- 35 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007. doi:10.4086/toc.2007.v003a006.
- 36 Igor E. Zverovich and Vadim E. Zverovich. An induced subgraph characterization of domination perfect graphs. *Journal of Graph Theory*, 20(3):375–395, 1995. doi:10.1002/jgt.3190200313.