

Tight Running Time Lower Bounds for Vertex Deletion Problems

CHRISTIAN KOMUSIEWICZ, Philipps-Universität Marburg, Ger

For a graph class Π , the Π -VERTEX DELETION problem has as input an undirected graph $G = (V, E)$ and an integer k and asks whether there is a set of at most k vertices that can be deleted from G such that the resulting graph is a member of Π . By a classic result of Lewis and Yannakakis [J. Comput. Syst. Sci. '80], Π -VERTEX DELETION is NP-hard for all hereditary properties Π . We adapt the original NP-hardness construction to show that under the Exponential Time Hypothesis (ETH) tight complexity results can be obtained. We show that Π -VERTEX DELETION does not admit a $2^{o(n)}$ -time algorithm where n is the number of vertices in G . We also obtain a dichotomy for running time bounds that include the number m of edges in the input graph: On the one hand, if Π contains all edgeless graphs, then there is no $2^{o(n+m)}$ -time algorithm for Π -VERTEX DELETION. On the other hand, if there is a fixed edgeless graph that is not contained in Π and containment in Π can be determined in $2^{O(n)}$ time or $2^{O(m)}$ time, then Π -VERTEX DELETION can be solved in $2^{O(\sqrt{m})} + O(n)$ or $2^{O(m)} + O(n)$ time, respectively. We also consider restrictions on the domain of the input graph G . For example, we obtain that Π -VERTEX DELETION cannot be solved in $2^{o(\sqrt{n})}$ time if G is planar and Π is hereditary and contains and excludes infinitely many planar graphs. Finally, we provide similar results for the problem variant where the deleted vertex set has to induce a connected graph.

CCS Concepts: • **Theory of computation** → **Problems, reductions and completeness**; *Graph algorithms analysis*; • **Mathematics of computing** → *Graph algorithms*;

Additional Key Words and Phrases: Fine-grained complexity analysis, Exponential Time Hypothesis, graph properties

ACM Reference Format:

Christian Komusiewicz. 2018. Tight Running Time Lower Bounds for Vertex Deletion Problems. *ACM Trans. Comput. Theory* 1, 1 (January 2018), 18 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

In graph modification problems, the aim is to modify a given graph such that it fulfills a certain property Π , for example being acyclic, bipartite, having bounded diameter, or being edgeless. Formally, a graph property is simply a set of graphs. The focus is usually on nontrivial graph properties, that is, properties such that the set of graphs fulfilling the property and the set of graphs *not* fulfilling the property are infinite. Throughout this work, we consider only nontrivial graph properties.

The most common operations to modify the graph are vertex deletion, edge deletion and edge addition. Graph modification problems have many applications, for example in data anonymization [18, 23], data clustering [3], data visualization [4], in the modeling of complex networks via differential equations [8], and in social network analysis [2, 16].

Author's address: Christian Komusiewicz, Philipps-Universität Marburg, Hans-Meerwein-Str. 6, Marburg, 35032, Ger, komusiewicz@informatik.uni-marburg.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

1942-3454/2018/1-ART \$15.00

<https://doi.org/0000001.0000001>

In this work, we consider graph modification problems where only vertex deletions are allowed. Formally, we study the following problem:

Π -VERTEX DELETION

Input: An undirected graph $G = (V, E)$ and an integer k .

Question: Is there a set $S \subseteq V$ such that $|S| \leq k$ and $G[V \setminus S]$ is contained in Π ?

A natural class of graph properties occurring in many applications are the hereditary graph properties: A graph property is *hereditary* if it is closed under vertex deletion. Equivalently, if a graph G fulfills a hereditary property Π , then every *induced subgraph* of G fulfills Π . The family of hereditary graph properties includes all monotone graph properties (which are the properties closed under vertex and edge deletions) and all minor-closed graph properties (which are the properties closed under vertex and edge deletions and edge contractions).

Hereditary graph properties can be characterized by forbidden induced subgraphs. That is, there is a (possibly infinite) set of graphs \mathcal{F} such that G has property Π if and only if G does not contain a graph from \mathcal{F} as an induced subgraph. For example, a graph is a forest if and only if it does not contain any cycle as induced subgraph. For hereditary graph properties, there is the following classic hardness result.

THEOREM 1.1 ([17]). *Let Π be a hereditary nontrivial graph property. Then Π -VERTEX DELETION is NP-hard.*

While this result rules out the possibility of polynomial-time algorithms, it leaves open the existence of subexponential-time algorithms. Our aim is to rule out the existence of subexponential-time algorithms for these problems by assuming the Exponential Time Hypothesis (ETH) [14]. To this end, let us inspect the reduction that Yannakakis¹ devised to prove Theorem 1.1 more closely.

The source problem in this reduction is VERTEX COVER.

VERTEX COVER

Input: An undirected graph $G = (V, E)$ and an integer k .

Question: Is there a set $S \subseteq V$ such that $|S| \leq k$ and $G[V \setminus S]$ has no edges?

Assuming ETH, VERTEX COVER does not admit a $2^{o(n+m)}$ -time algorithm where n is the number of vertices in the input graph and m is the number of edges. Hence, VERTEX COVER is a suitable source problem for a reduction in the ETH-framework. The construction of Yannakakis, however, reduces n -vertex instances of VERTEX COVER to $\Theta(n^2)$ -vertex instances of Π -VERTEX DELETION. This blowup in the vertex number only yields a running time lower bound of $2^{o(\sqrt{n})}$. Thus, to obtain a better lower bound, we need to “re-engineer” the original construction.

We adapt the hardness construction of Yannakakis in two ways: We take a slightly more fine-grained approach for choosing the graphs which are used as basis for the gadgets in the reduction from the VERTEX COVER instance and we reduce from a more restricted variant of VERTEX COVER. These two modifications of the reduction allow us to obtain our main result. Here, n denotes the number of vertices in the input graph and m denotes the number of its edges.

THEOREM 1.2. *Let Π be a hereditary nontrivial graph property, then:*

- (1) *If the ETH is true, then Π -VERTEX DELETION cannot be solved in $2^{o(n+\sqrt{m})}$ time,*
- (2) *If the ETH is true, then Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time if Π contains all edgeless graphs.*

¹The article by Lewis and Yannakakis explicitly attributes the construction which will form the basis of our reduction to Yannakakis.

- (3) If Π excludes some edgeless graph, then Π -VERTEX DELETION can be solved in $2^{O(\sqrt{m})} + O(n)$ time if and only if membership in Π can be recognized in $2^{O(n)}$ time.
- (4) If Π excludes some edgeless graph, then Π -VERTEX DELETION can be solved in $2^{o(m)} + O(n)$ time if and only if membership in Π can be recognized in $2^{o(m)}$ time.

The first and second statements of the theorem follow from our adaption of Yannakakis' construction, described in Section 3. The third and fourth statements follow from a simple combinatorial algorithm, described in Section 4.

In addition to our main theorem, we observe that the reduction yields a similarly tight lower bound for vertex deletion problems in planar graphs. More precisely, we show that any nontrivial variant of Π -VERTEX DELETION cannot be solved in $2^{o(\sqrt{n})}$ time on n -vertex planar graphs. This result aligns nicely with the so-called square root phenomenon [21] that many NP-hard problems on planar graphs can be solved in $2^{O(\sqrt{n})}$ time but not faster. Moreover, we also obtain tight running time lower bounds for input graphs with bounded degree and bounded degeneracy (both depending on Π) and for input graphs containing a dominating vertex. Finally, we consider CONNECTED Π -VERTEX DELETION, where the vertex set that is deleted has to induce a connected graph. We show that the bounds of Theorem 1.2 also hold for this problem variant.

Related Work. Fellows et al. [7] adapted the reduction of Yannakakis to provide NP-hardness results for the variant of Π -VERTEX DELETION where one aims to find a solution that is disjoint from and smaller than a given solution. The study considers the case that Π is closed under taking disjoint unions of graphs. These are exactly the graph properties that have only connected minimal forbidden induced subgraphs. As an introduction to their result, they describe a modified version of the original reduction which already implies Theorem 1.2 for this case. In subsequent work, Guo and Shrestha [13] investigated the complexity of the remaining case when the forbidden subgraphs have disconnected forbidden subgraphs, showing NP-hardness for most cases.

When edge deletions or additions are the allowed operations, there is no hardness result as general as the one of Lewis and Yannakakis [17]. Nevertheless, many variants of Π -EDGE DELETION are NP-hard [24]. If Π has exactly one forbidden induced subgraph, then a complexity dichotomy and tight ETH-based running time lower bounds can be achieved for edge deletion problems [1]. In similar spirit, tight running time lower bounds have been obtained for some Π -COMPLETION problems, where one may only add edges to obtain the graph property Π [6].

2 PRELIMINARIES

Unless stated otherwise, we consider undirected simple graphs $G = (V, E)$ where V is the vertex set and $E \subseteq \{\{u, v\} \mid u, v \in V \wedge u \neq v\}$ is the edge set of the graph. We denote the vertex and edge sets of a graph G also by $V(G)$ and $E(G)$, respectively. The (open) neighborhood of a vertex v is denoted by $N(v) := \{u \mid \{u, v\} \in E\}$ and the closed neighborhood by $N[v] := N(v) \cup \{v\}$. For a vertex set $S \subseteq V$, let $G[S] := (S, \{\{u, v\} \in E \mid u, v \in S\})$ denote the subgraph of G induced by S . Similarly, the open neighborhood of a vertex set $S \subseteq V$ is defined as $N(S) := \bigcup_{v \in S} N(v) \setminus S$ and its closed neighborhood is defined as $N[S] := N(S) \cup S$. For a vertex set $S \subseteq V$, let $G - S := G[V \setminus S]$ denote the induced subgraph of G obtained by deleting the vertices of S and their incident edges. Similarly, for a vertex $v \in V$ let $G - v := G - \{v\}$ denote the graph obtained by deleting this vertex and all its incident edges. A set $S \subseteq V$ is an independent set if $G[S]$ has no edges. A vertex v in a graph is a cut-vertex if $G - v$ has more connected components than G . A graph is d -degenerate if every induced subgraph contains a vertex of degree at most d .

Let $A = (a_1, a_2, \dots, a_n)$ and $B = (b_1, b_2, \dots, b_m)$ be two sequences. Then, A is lexicographically smaller than B if

- $a_i < b_i$, where i is the smallest index such that $a_i \neq b_i$, or
- $n < m$ and $a_i = b_i$ for all $i \in \{1, \dots, n\}$.

We say that a parameterized problem (I, k) has a *subexponential-time algorithm* if it can be solved in $2^{o(k)} \cdot |I|^{O(1)}$ time. In this work, the parameters under consideration are n , the number of vertices in the input graph, and m , the number of edges in the input graph. The *Exponential Time Hypothesis (ETH)* states essentially that the 3-SAT problem, which is given a boolean formula in 3-conjunctive normal form and asks whether this formula can be satisfied by assigning truth values to the variables, cannot be solved in $2^{o(n)}$ time where n is the number of variables in the input formula [14]. For a survey on lower bounds based on the Exponential Time Hypothesis, refer to [20].

We now show that, assuming the ETH, the following special case of VERTEX COVER does not admit a subexponential-time algorithm. This problem variant will be used as source problem in our reduction; here, a subcubic graph is one with maximum degree at most three.

SUBCUBIC GIRTH- d VERTEX COVER

Input: An undirected subcubic 2-degenerate graph $G = (V, E)$ such that every cycle in G has length at least d and an integer k .

Question: Does G have a vertex cover of size at most k ?

In this definition, the girth d is not part of the input. Consequently, d is considered to be a constant in the following. We show the running time lower bound by devising a simple reduction from VERTEX COVER that replaces edges by long paths on an even number of vertices.

THEOREM 2.1. *If the ETH is true, then SUBCUBIC GIRTH- d VERTEX COVER cannot be solved in $2^{o(n+m)}$ time.*

PROOF. We reduce from VERTEX COVER in subcubic graphs, which is known to admit no $2^{o(n)}$ -time algorithm (assuming ETH) [15]. Assume that d is even; otherwise, increase d by one. Given a vertex cover instance $(G = (V, E), k)$ the reduction is simply to replace each edge $\{u, v\}$ of G by a path on d vertices and to make one endpoint of this path adjacent to u and the other one adjacent to v . Equivalently, subdivide each edge d times. Let G' denote the resulting graph, and call the paths that are added during the construction the *d -paths* of G' . To complete the construction, set $k' := k + (d/2) \cdot |E(G)|$. As we will show, the two instances are equivalent. Assume for now that equivalence holds. The shortest cycle in G' clearly contains more than d vertices (in fact it contains more than $3d$ vertices). Moreover, the construction does not increase the maximum degree of G , hence G' has maximum degree three, and each original edge was replaced by a path, hence G is 2-degenerate. Finally, G' has $O(|V(G)|)$ many vertices and edges since d is a constant. Thus, any algorithm that solves (G', k') in $2^{o(|V(G')| + |E(G')|)}$ time can be used to obtain an algorithm that solves VERTEX COVER on subcubic n -vertex graphs in $2^{o(n+m)}$ time which contradicts the ETH.

To complete the proof it thus remains to show equivalence of the instances.

(G, k) is a yes-instance $\Leftrightarrow (G', k')$ is yes-instance.

(\Rightarrow) Let S be a size- k vertex cover of G . Consider the graph $G' - S$. By construction, all remaining edges are incident with vertices in d -paths. Moreover, for each d -path, the nonpath neighbor of at least one of its endpoints is contained in S since S is a vertex cover. Thus, deleting the other endpoint of the path and then every other vertex of the path gives a set of $d/2$ vertex deletions that together destroy all edges that are incident with vertices of this path. This can be done for all $|E(G)|$ many d -paths, resulting in $(d/2) \cdot |E(G)|$ vertex deletions. Consequently, $(d/2) \cdot |E(G)| + k$ vertex deletions suffice to destroy all edges in G' .




Fig. 1. A graph H with α -sequence $(3, 2, 1)$. Deleting any non-cut-vertex v gives $\alpha(H, v) = (6)$. Deleting the cut-vertex c or d gives $\alpha(H, c) = (3, 2, 1)$ and $\alpha(H, d) = (5, 1)$, respectively. The graph $J(H)$ is highlighted by a gray background.

(\Leftarrow) Observe that any vertex cover contains at least $(d/2) \cdot |E(G)|$ vertices that are on d -paths. Moreover, there is a minimum-cardinality vertex cover such that no two neighbors u and v in a d -path are deleted: Deleting u makes v a degree-one vertex which means that deleting the other neighbor of v instead of v gives a vertex cover which is at most as large as one that deletes u and v . Thus, there is a minimum-cardinality vertex cover S' that contains for each d -path exactly one of its two endpoints. Moreover, the neighbor of the other endpoint that is not in the d -path is contained in S' . Let $S := S' \cap V$ denote the vertices from V that are contained in S . For each edge $\{u, v\} \in E$, there is a d -path connecting u and v in E . By the discussion above, either u or v is in S . Thus, S is a vertex cover. Since S' contains at least $(d/2) \cdot |E(G)|$ vertices from d -paths, we have $|S| \leq k$. \square

3 LOWER BOUNDS FOR GRAPH PROPERTIES CONTAINING ALL EDGELESS GRAPHS

The main idea of the reduction of Yannakakis is the following. Reduce from VERTEX COVER by replacing edges of the VERTEX COVER instance by some graph H that is not contained arbitrarily often in any graph fulfilling Π (and thus needs to be destroyed by a vertex deletion). Then, destroying edges in the VERTEX COVER instance corresponds to destroying forbidden induced subgraphs in the Π -VERTEX DELETION instance. The main technical difficulty arises when the forbidden subgraphs for graph property Π have cut-vertices: in this case, two graphs that are used to replace two edges incident with the same vertex may form another forbidden subgraph of Π . Thus, the graph H which is used in the construction must be chosen carefully. To this end, Yannakakis introduces the notion of α -sequence for connected graphs.

Definition 3.1. Let H be a connected graph, let c be a vertex in H , and let H_1, \dots, H_ℓ denote the connected components of $H - c$ such that $|V(H_1)| \geq |V(H_2)| \geq \dots \geq |V(H_\ell)|$. Then, $\alpha(H, c) := (|V(H_1)|, |V(H_2)|, \dots, |V(H_\ell)|)$. The α -sequence of H , denoted $\alpha(H)$, is the lexicographically smallest sequence such that $\alpha(H) = \alpha(H, c)$ for some $c \in V(H)$.

An example of a graph H with its α -sequence is presented in Figure 1. The idea in the reduction of Yannakakis is to choose a forbidden subgraph that has a connected component with a lexicographically smallest α -sequence as basis for the gadget. More precisely, the edges of the vertex cover instance are replaced by the largest connected component remaining after deletion of a vertex c which produces the smallest α -sequence of H . To this end, let $c \in V(H)$ be a fixed vertex such that $\alpha(H) = \alpha(H, c)$, and let J' be a fixed largest connected component of $H - c$. Then, let $J(H) := H[V(J') \cup \{c\}]$ denote the subgraph of H consisting of this component plus c . Since H has an edge, $J(H)$ has at least two vertices. The graph $J(H)$ will be used to replace the edges of the VERTEX COVER instance.



Fig. 2. The two minimal forbidden subgraphs co-diamond (left) and co-claw (right) for Π being the complement graphs of line graphs of triangle-free graphs (see <http://www.graphclasses.org/>). Here, H_Π is the co-diamond since the lexicographically largest α -sequence of any component of the co-diamond is (1) and the co-claw has a component with α -sequence (2).

This terminology is sufficient to deal with the case that all forbidden subgraphs are connected. To deal with disconnected forbidden induced subgraphs, we introduce the notion of Γ -sequence.² Observe that the α -sequences imply a partial ordering on all connected graphs, but there maybe nonisomorphic connected graphs that have the same α -sequence. We want to avoid this, and thus the first step is to refine this partial ordering to a total ordering by breaking ties arbitrarily.

Definition 3.2. A function from the set of all connected graphs to \mathbb{N} is a γ -ordering if $\gamma(G_i) \geq \gamma(G_j)$ implies $\alpha(G_i) \geq \alpha(G_j)$ and $\gamma(G_i) = \gamma(G_j)$ if and only if G_i and G_j are isomorphic.

In the following, fix an arbitrary γ -ordering of all connected graphs. To obtain the Γ -sequence we consider the sequence of γ -values created by the connected components of a graph H .

Definition 3.3. Let H be a graph with connected components H_1, \dots, H_ℓ such that $\gamma(H_1) \geq \gamma(H_2) \geq \dots \geq \gamma(H_\ell)$. Then, the Γ -sequence of H is $\Gamma(H) := (\gamma(H_1), \gamma(H_2), \dots, \gamma(H_\ell))$.

Now, the basis for our construction will be the minimal forbidden induced subgraph of Π that has the lexicographically smallest Γ -sequence. For a graph property Π , denote this graph by H_Π , an example is given in Figure 2. We are now ready to prove the second statement of Theorem 1.2. We reduce from SUBCUBIC GIRTH-3d VERTEX COVER. The construction of the Π -VERTEX DELETION instance uses gadgets based on the graph H_Π , an example is presented in Figure 3.

Construction 1. Let H_1 be the connected component of H_Π with maximum γ -value and let d denote the number of connected components of H_Π that are isomorphic to H_1 . We reduce from SUBCUBIC GIRTH-3d VERTEX COVER. Given an instance $(G = (V, E), k)$, we build an instance of Π -VERTEX DELETION as follows. Let c denote a vertex of H_1 such that $\alpha(H_1, c) = \alpha(H_1)$, that is, c is a vertex whose deletion produces a lexicographically smallest sequence of connected component sizes. Assume without loss of generality that $|V| > |V(H_\Pi)|$ (otherwise, we can solve the SUBCUBIC GIRTH-3d VERTEX COVER instance in constant time).

Starting from an empty graph G' , construct the instance of Π -VERTEX DELETION as follows. For each connected component H_i of H_Π that is not isomorphic to H_1 , add $2 \cdot |V|$ disjoint copies of H_i to G' .³ Then, add $d - 1$ disjoint copies of H_1 to G' . Call the graph constructed so far the *base graph* of the construction.

Now, add the graph G to G' and replace each edge $\{u, v\}$ of G by the graph $J := J(H_1)$ in the following way.⁴ Since Π contains all edgeless graphs, H_1 has at least one edge. Therefore, J has at least two vertices. Accordingly, let c' be a vertex from J that is different from c . Remove the edge $\{u, v\}$ from G' , add a copy of J to G' and identify u with c and v with c' . Now, let $D :=$

²Yannakakis uses the related notion of β -sequence which does not break ties between nonisomorphic graphs that have the same α -sequence [17]. This tie-breaking, however, is necessary for our reduction. Guo and Shrestha [13] use Ω -sequences which also do not break ties between nonisomorphic graphs.

³This is the part of the construction, where tie-breaking is necessary, as these graphs might have the same α -sequence as H_1 .

⁴This part of the construction follows the construction of Yannakakis [17].




Fig. 3. An example of the reduction (in this case from VERTEX COVER instead of SUBCUBIC GIRTH- d VERTEX COVER). Top left: the forbidden subgraph H_Π with the vertex c of H_1 ; top right: the graph G of the VERTEX COVER instance; bottom: the constructed graph G' of the Π -VERTEX DELETION instance. Vertices of G and their corresponding vertices in G' are shown in gray.

$H_1 - (V(J) \setminus \{c\})$. For each vertex $v \in V$, add a disjoint copy of D to G' and identify the vertex c of this copy with v . Call the part of G' that is obtained from G the *vc-extension* of the construction. Complete the construction of the Π -VERTEX DELETION instance by setting $k' := k$.

Using this construction, we can show our main running time lower bound.

LEMMA 3.4. *Let Π be a hereditary graph property containing all edgeless graphs. Then, if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time.*

PROOF. Let (G, k) be an instance of SUBCUBIC GIRTH- $3d$ VERTEX COVER and let (G', k) be the Π -VERTEX DELETION instance obtained from Construction 1. We first show the equivalence of the two instances, that is,

G has a vertex cover of size at most $k \Leftrightarrow G'$ has a Π -vertex deletion set of size at most k .

(\Rightarrow) Let S be a size- k vertex cover of G . Then, S is a Π -vertex deletion set in G' , that is, $G' - S$ fulfills Π . To see this, first observe that each connected component C of the *vc-extension* of $G' - S$ has γ -value lower than $\gamma(H_1)$: If C does not contain a vertex from V , then it is either a proper induced subgraph of D or of J . In both cases, C has a lexicographically lower α -sequence than H_1 and thus lower γ -value. If C contains a vertex v from V , then v is a cut-vertex in C , as deleting it disconnects the remainder of the copy of D containing v . Now, observe that the other connected components of $C - v$ have size at most $|V(J)| - 2$: These components are subgraphs of copies of J corresponding to an edge $\{u, v\}$ of G . Since S is a vertex cover, u is deleted from G , cutting the rest of the graph from this copy of J and leaving only $|V(J)| - 2$ vertices in the component. Consequently, the α -sequence of C is lexicographically smaller than the α -sequence of H_1 . This implies that the γ -value of the connected component C is smaller than $\gamma(H_1)$.

Moreover, the base part contains exactly $d - 1$ connected components with γ -value $\gamma(H_1)$. Hence, $G' - S$ fulfills Π : Any induced subgraph of $G' - S$ has at most $d - 1$ connected components with γ -value $\gamma(H_1)$. Since

- H_Π has exactly d connected components with γ -value $\gamma(H_1)$ and no connected components with larger γ -value, and
- H_Π has the lexicographically smallest Γ -sequence among all forbidden induced subgraphs of Π ,

we conclude that $G' - S$ contains no forbidden induced subgraph of Π as induced subgraph.

(\Leftarrow) Let S' be a size- k Π -vertex deletion set in G' . Since $k < 2n$, we have that the number of copies of all connected components of H_Π with γ -value smaller than $\gamma(H_1)$ is at least as large in $G' - S'$ as it is in H_Π . Hence, $G' - S'$ contains at most $d - 1$ vertex-disjoint copies of H_1 as otherwise, $G' - S'$ contains H_Π as induced subgraph. Then, create a set $S_A \subseteq V$ as follows:

- For each vertex $v \in V \cap S'$, add v to S_A ,
- for each copy D^* of D in G' such that S' contains a vertex from $V(D^*) \setminus V$, add the vertex $v \in V(D^*) \cap V$ to S_A ,
- for each copy J^* of J such that $S' \setminus V$ contains a vertex from $V(J^*) \setminus V$, add an arbitrary vertex of $V(J^*) \cap V$ to S_A .

Observe that the set S_A has size at most k . If $G - S_A$ is an edgeless graph, then G has a vertex cover of size at most k . Otherwise, consider the graph $G - S_A$ and observe that every edge $\{u, v\}$ in $G - S_A$ directly corresponds to an induced copy of H_1 in $G' - S'$: Since $\{u, v\}$ is present in $G - S_A$, we have that S' does not contain u or v , does not contain vertices from the copies of D that are attached to u or v , and does not contain a vertex from the copy of J attached to u and v . Now let q denote the size of a maximum matching in $G - S_A$ and observe that this implies that the vc-extension of $G' - S'$ has q vertex-disjoint copies of H_1 . Since $G' - S'$ contains at most $d - 1$ vertex-disjoint copies of H_1 , this implies that $q \leq d - 1$ and that S' contains at least q vertex deletions in the base graph (it needs to destroy at least q of the vertex-disjoint copies of H_1 in the base graph). Thus, $|S_A| \leq k - q$. Finally, observe that since $G - S_A$ has a maximum matching of size $q \leq d - 1$ it cannot contain a cycle as all cycles in G have length at least $3d$. Thus, $G - S_A$ has a vertex cover S_B of size at most q which together with S_A is a size- k vertex cover of G .

Since the two instances are equivalent, any algorithm deciding (G', k') decides (G, k) . Since H_Π has constant size for each fixed graph property Π , the base graph contains $O(|V(G)|)$ vertices and edges. Similarly, the vc-extension of G has $O(|V(G)|)$ copies of J (since G has $O(|V(G)|)$ edges) and $O(|V(G)|)$ copies of D and no further vertices. Consequently, G' has $O(|V(G)|)$ vertices and $O(|V(G)|)$ edges. Any algorithm deciding (G', k) in $2^{o(|V(G')|+|E(G')|)}$ time thus decides SUBCUBIC GIRTH- $3d$ VERTEX COVER in $2^{o(|V(G)|)}$ time, contradicting the ETH by Theorem 2.1. \square

Let us briefly discuss why we need to assume large girth for the VERTEX COVER instance. If Π is for example the property of being $2K_2$ -free, that is, the set of graphs not containing an induced matching of size two as induced subgraph, then given a VERTEX COVER instance (G, k) , the reduction would simply add an edge to G . Now, if G is a triangle and $k = 1$, then we have a no-instance of VERTEX COVER but the resulting Π -VERTEX DELETION instance is a yes-instance: deleting a vertex of the isolated edge is sufficient to destroy all $2K_2$ s.

Since the maximum degree is three in SUBCUBIC GIRTH- d VERTEX COVER, we obtain the following corollary on the hardness of Π -VERTEX DELETION in bounded-degree graphs.

COROLLARY 3.5. *Let Π be a hereditary nontrivial graph property containing all edgeless graphs. Let Δ be the smallest number such that there exists a γ -ordering of the forbidden subgraphs of Π such that a forbidden subgraph with lexicographically smallest Γ -sequence has maximum degree Δ . Then,*

if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time even if G has maximum degree 3Δ .

So, for example if Π is the property of being acyclic, then Π -VERTEX DELETION (this special case is known as FEEDBACK VERTEX SET) does not admit a $2^{o(n+m)}$ -time algorithm even if G has maximum degree six.

We can directly use Lemma 3.4 to obtain a lower bound for all other nontrivial graph properties. First observe that, by Ramsey's theorem, every nontrivial graph property Π contains either all edgeless graphs or all cliques: Since Π is nontrivial, the order of the graphs in Π is unbounded. Thus, for every number n , Π contains a graph which has $R(n, n)$ many vertices and therefore contains either an n -vertex clique or an edgeless graph on n vertices. Thus, if Π does not contain all edgeless graphs, then it contains all cliques. In this case, however, Lemma 3.4 implies a running time lower bound for $\bar{\Pi}$ -VERTEX DELETION, where $\bar{\Pi}$ is the graph property containing exactly the graphs that are complement graphs of a graph in Π . Now $\bar{\Pi}$ -VERTEX DELETION problem can be easily reduced to Π -VERTEX DELETION by complementing the input graph. This reduction does not change the number of vertices n in the graph and the number of edges m is $O(n^2)$, thus leading to the following lower bound for all hereditary graph properties.

COROLLARY 3.6. *Let Π be a hereditary graph property. Then, if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+\sqrt{m})}$ time.*

4 SUBEXPONENTIAL-TIME ALGORITHMS FOR GRAPH PROPERTIES EXCLUDING SOME EDGELESS GRAPH

The hardness results from the previous section leave open the possibility of algorithms that have subexponential running time with respect to the number of edges m of G for graph properties Π that do not contain all edgeless graphs. We now show that Π -VERTEX DELETION can be solved in $2^{o(m)} + O(n)$ time if Π can be recognized sufficiently fast and does not contain all edgeless graphs. Intuitively, the algorithm exploits the following observations: First, the number of vertices that have a high degree in G is sublinear in m . Second, the subgraph induced by the low-degree vertices in the solution has a small dominating set.

LEMMA 4.1. *Let Π be a hereditary graph property excluding some edgeless graph, then*

- Π -VERTEX DELETION can be solved in $2^{O(\sqrt{m})} + O(n)$ time if and only if membership in Π can be verified in $2^{O(n)}$ time, and
- Π -VERTEX DELETION can be solved in $2^{o(m)} + O(n)$ time if and only if membership in Π can be verified in $2^{o(m)}$ time.

PROOF. We first show an algorithm with the claimed running time for the case that membership in Π can be verified in $2^{O(n)}$ time. Assume that Π excludes the edgeless graph on d vertices. Given an input (G, k) , the first step of the algorithm is to delete all except at most $d - 1$ isolated vertices of the input graph and to decrease the size bound accordingly. These deletions are necessary by the assumption that Π does not contain the edgeless graph on d vertices. Afterwards, the graph has $O(m)$ vertices.

Now, let A denote an arbitrary but fixed maximum-cardinality set such that $G[A] \in \Pi$. Let V_h denote the vertices in G that have degree at least $2\sqrt{m}$ and observe that $|V_h| \leq \sqrt{m}$. For each $A_h \subseteq V_h$, branch into a case that assumes $A_h = A \cap V_h$. In one of these branches, the assumption is correct.

Now consider the vertices in $V_\ell := V \setminus V_h$ that have degree less than $2\sqrt{m}$ in G . Since $G[A \cap V_\ell] \in \Pi$ it has no independent set on d vertices. Thus, for each independent set $I_\ell \subseteq V_\ell$ of $G[V_\ell]$ such that $|I_\ell| < d$ we branch into a case that assumes that $I_\ell \subseteq A \cap V_\ell$ is a maximal independent set

of $G[A \cap V_\ell]$. In one of these branches, the assumption is correct. Now observe that since I_ℓ is a *maximal* independent set in $G[A \cap V_\ell]$ it is also a dominating set in $G[A \cap V_\ell]$. That is, $A \cap V_\ell \subseteq N[I_\ell]$. Thus, by branching for each $A_\ell \subseteq N(I_\ell)$ into the case that $A_\ell = (A \cap V_\ell) \setminus I_\ell$ we obtain one case where $A = A_h \cup I_\ell \cup A_\ell$. In each branch, we check whether $G[A_h \cup I_\ell \cup A_\ell] \in \Pi$ and whether $|A_h \cup I_\ell \cup A_\ell| \geq n - k$.

The running time bound can be obtained as follows: The number of subsets of V_h is $2^{O(\sqrt{m})}$. The number of subsets I_ℓ of size less than d of V_ℓ is $O(n^{d-1}) = n^{O(1)}$, and for each I_ℓ the number of subsets of $N(I_\ell)$ is $O(2^{(d-1)2\sqrt{m}}) = 2^{O(\sqrt{m})}$. Thus, the recognition algorithm for Π has to be invoked $2^{O(\sqrt{m})}$ times. Each time, it is invoked on a graph with $O(\sqrt{m})$ vertices since $|A_h| \leq \sqrt{m}$, $|I_\ell| < d = O(1)$, and $|A_\ell| < d\sqrt{m} = O(\sqrt{m})$. By the premise, this algorithm takes $2^{O(\sqrt{m})}$ time, resulting in the claimed overall running time.

For the only if part of the first claim, we observe the following. Since $m = O(n^2)$, any $2^{O(\sqrt{m})}$ -time algorithm for Π -VERTEX DELETION directly gives a $2^{O(n)}$ -time algorithm for the recognition problem, since the special case $k = 0$ is the recognition problem for Π .

If membership in Π can be verified in $2^{o(m)}$ time, then the algorithm only differs in the invoked recognition algorithm. This algorithm for Π has to be invoked $2^{O(\sqrt{m})} = 2^{o(m)}$ times, each time on a graph with at most m edges. By the premise, this algorithm takes $2^{o(m)}$ time, resulting in the claimed overall running time. The only if part of the statement follows directly from the fact that solving the special case $k = 0$ gives an algorithm for the recognition problem. \square

Since the time complexity of Π -VERTEX DELETION for the nontrivial graph properties excluding a fixed edgeless graph is settled to be $2^{O(\sqrt{m})}$ or more, it is now motivated to decrease the constants in the exponents of the running time bound. In the following, we consider the important special case of graph properties where membership can be decided in polynomial time. The generic algorithm described above has running time $2^{\sqrt{m}} \cdot n^{d-1} \cdot 2^{(d-1)2\sqrt{m}} \cdot n^{O(1)}$ for these problems. By setting the degree threshold for including a vertex in the first or in the second branching to $\sqrt{2m}/(d-1)$, this can be improved to a running time of $2^{2\sqrt{2(d-1)m}} n^{d+O(1)}$. It is more interesting, however, to determine whether the factor d before m can be removed. For a subclass of the polynomial-time decidable graph properties excluding an edgeless graph of order d , we obtain such a running time bound. The algorithm achieving the running time follows the same idea as a known $2^{O(\sqrt{m})}$ -time algorithm for CLIQUE [9].

THEOREM 4.2. *Let Π be a graph property such that*

- *there is a constant d such that for all $n > d$ every n -vertex graph G with property Π has minimum degree $n - d$, and*
- *for an n -vertex graph, membership in Π can be verified in $n^{O(1)}$ time.*

Then, Π -VERTEX DELETION can be solved in $2^{\sqrt{2m}} \cdot n^{d+O(1)}$ time.

PROOF. Let A denote an arbitrary but fixed maximum-cardinality set such that $G[A] \in \Pi$. First, consider the case that A contains a vertex v that has degree at most $\sqrt{2m}$ in G . Then, for each $A' \subseteq N(v)$, branch into a case assuming that $A' = A \cap N(v)$. In each of these cases, branch for each subset A_d of size less than d of $V \setminus N[v]$ into a case assuming $A_d = A \setminus N[v]$. Now if $|A' \cup A_d \cup \{v\}| \geq n - k$, then check whether $G[A' \cup A_d \cup \{v\}]$ fulfills Π using the algorithm promised by the premise of the lemma. If this is the case, then return “yes”. The overall number of branches is $O(2^{\sqrt{2m}} \cdot n^{d-1})$, resulting in a running time of $2^{\sqrt{2m}} \cdot n^{d+O(1)}$ for this part of the algorithm. If A contains v , then since $|A \setminus N[v]| < d$, in one of the created branches both assumptions are correct. Hence, if all of the branches for all v with degree at most $\sqrt{2m}$ return “no”, then A contains only vertices of degree

at least $\sqrt{2m} + 1$. Thus, all other vertices may be removed from the graph while decreasing the size bound k accordingly.

In the remaining graph, $n < \sqrt{2m}$. Thus a brute-force algorithm testing for each vertex subset A of size at least $n - k$ whether $G[A]$ fulfills Π has running time $2^{\sqrt{2m}} \cdot n^{O(1)}$. \square

5 CONSEQUENCES FOR SOME RESTRICTED DOMAINS OF INPUT GRAPHS

We now present running time lower bounds for Π -VERTEX DELETION when the input graph is restricted to belong to a certain graph class. Such restrictions are motivated by at least two aspects: From an application viewpoint, one should provide hardness results for realistic types of input instances, for example for sparse graphs. From a complexity-theoretic viewpoint, hardness results for restricted inputs may facilitate further reductions.

5.1 Planar Graphs

PLANAR Π -VERTEX DELETION

Input: An undirected planar graph $G = (V, E)$ and an integer k .

Question: Is there a set $S \subseteq V$ such that $|S| \leq k$ and $G[V \setminus S]$ is contained in Π ?

In PLANAR Π -VERTEX DELETION, we are only interested in properties Π that contain all edgeless graphs. This is due to the fact that any property Π that excludes edgeless graphs of order d for some fixed d cannot contain planar graphs whose vertex number is bigger than the Ramsey number $R(5, d)$: Every graph with this many vertices contains by definition either a clique of size five or an independent set of size d . In the first case the graph is nonplanar; in the second case the graph is not contained in Π because Π is hereditary and excludes edgeless graphs on d vertices.

For graph properties Π containing all edgeless graphs and excluding at least one planar graph, we can apply our modification of Yannakakis' reduction when we reduce from PLANAR VERTEX COVER instead: First, as noted for example in [5, Theorem 14.9] it is known that PLANAR VERTEX COVER cannot be solved in $2^{o(\sqrt{n})}$ time (assuming the ETH).⁵ Second, the reduction behind Theorem 2.1 produces a planar graph if it has a planar graph as input. Third, if Π is hereditary and excludes some planar graph, then it has a nonempty family of planar forbidden induced subgraphs \mathcal{F} . Finally, in a reduction producing a planar graph, we may ignore all nonplanar forbidden induced subgraphs. Thus, we can simply carry out the reduction behind Lemma 3.4 starting from PLANAR VERTEX COVER and picking the gadget graph H_Π only among the planar forbidden induced subgraphs of Π . The replacement of edges of the PLANAR VERTEX COVER instance by planar graphs and the attachment of planar graphs to single vertices clearly yields a planar graph, as noted by Yannakakis [17]. Moreover, if the PLANAR VERTEX COVER instance has $O(n)$ vertices, then so has the PLANAR Π -VERTEX DELETION instance. Altogether, we arrive at the following.

THEOREM 5.1. *Let Π be a hereditary graph property containing all edgeless graphs and excluding at least one planar graph. Then, if the ETH is true, PLANAR Π -VERTEX DELETION cannot be solved in $2^{o(\sqrt{n})}$ time.*

5.2 Bounded-Degeneracy Graphs

For the bounded-degeneracy case, we first observe that Corollary 3.5 already implies that for Π containing all edgeless graphs, Π -VERTEX DELETION cannot be solved in subexponential time

⁵The result follows essentially from a classic reduction [10] of VERTEX COVER to PLANAR VERTEX COVER that uses constant-size uncrossing gadgets.

even on graphs with bounded degree and thus not on graphs with bounded degeneracy. Here, we improve the bound on the degeneracy.

THEOREM 5.2. *Let Π be a hereditary graph property such that all edgeless graphs are in Π . Let δ be the smallest number such that there exists a γ -ordering of the forbidden subgraphs of Π such that a forbidden subgraph with lexicographically smallest Γ -sequence is δ -degenerate. Then, if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time even if G is $(\delta + 1)$ -degenerate.*

PROOF. We modify Construction 1 as follows. Let H_Π denote the δ -degenerate forbidden subgraph of Π that has the lexicographically smallest Γ -sequence among all forbidden induced subgraphs of Π . As before, let H_1 denote the connected component of H_Π that has the highest γ -value and fix c again to be a vertex such that $\alpha(H, c) = \gamma(H)$ and let J denote the induced subgraph containing c and an arbitrary but fixed largest connected component of $H_1 - c$. Let n_J denote the order of J and fix a degeneracy-ordering of J , that is, a sequence (j_1, \dots, j_{n_J}) such that j_i has degree at most δ in $J[\{j_i, j_{i+1}, \dots, j_{n_J}\}]$. Now let c' denote the vertex with the highest index that is different from c . More precisely, if $c \neq j_{n_J}$, then $c' = j_{n_J}$; otherwise $c' = j_{n_J-1}$.

Now perform the construction as before using the new choice of c and c' when replacing an edge of the vertex cover instance by a copy of J . All other parts of the construction remain the same. Observe that in this modified construction we only make the choice of c' specific where it was arbitrary before. This implies that the modified construction is correct and that any algorithm solving Π -VERTEX DELETION on the constructed instances (G', k) implies a $2^{o(n+m)}$ -time algorithm for SUBCUBIC GIRTH- d VERTEX COVER, violating ETH.

It remains to show that the constructed instance is $(\delta + 1)$ -degenerate. The base graph of the construction is clearly δ -degenerate as each connected component is a subgraph of H_Π .

We now describe a sequence of vertex deletions for the vc-extension such that each deleted vertex has degree at most $\delta + 1$ when it is deleted. First, observe that the graph $D := H_1 - (V(J) \setminus \{c\})$ which is added for each vertex v of the original instance is δ -degenerate. In every copy D^* of D , all vertices except v (which is the vertex identified with the cut-vertex c) have no neighbors outside of this copy. Since $D - c$ is δ -degenerate, every induced subgraph of D^* that contains v plus some other vertex $u \neq v$ contains a vertex different from v that has degree at most $\delta + 1$. Consequently, this vertex can be deleted first. This vertex deletion can be performed as long as any copy D^* of D still contains a vertex different from the vertex that was identified with c .

The remaining graph consists of copies J^* of J that have two vertices which are identified with vertices V , the vertex set of the SUBCUBIC GIRTH- $3d$ VERTEX COVER instance. Consider such a copy J^* replacing an edge $\{u, v\}$, that is, the vertices c and c' are identified with u and v . By the modification of the construction, either c or c' equals j_{n_J} . As a consequence, deleting all vertices of J^* that are different from u and v in the same order as before, gives a deletion sequence in which each deleted vertex has degree at most $\delta + 1$. After deleting all these vertices in each copy of J , what remains is either an edgeless graph (if c and c' are not adjacent in J) or exactly the graph G of the original SUBCUBIC GIRTH- $3d$ VERTEX COVER instance. This graph is 2-degenerate. Since Π contains all edgeless graphs, H_Π is not 0-degenerate and thus $2 \leq (\delta + 1)$. \square

An example application of Theorem 5.2 is the following. If all forbidden subgraphs of Π are acyclic, that is, 1-degenerate, then Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time even if G is 2-degenerate. Moreover, the bound on δ is tight: For example VERTEX COVER or CLUSTER VERTEX DELETION, where the forbidden subgraph is a path on three vertices, have only 1-degenerate forbidden subgraphs and can be solved in linear time on 1-degenerate graphs.

5.3 Graphs with a Dominating Vertex

Next, we consider instances which have a dominating vertex and thus diameter two.

PROPOSITION 5.3. *Let Π be a hereditary graph property that contains all edgeless graphs. Then, if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time even if G has a dominating vertex.*

PROOF. We extend Construction 1. Recall that (G, k) is the SUBCUBIC GIRTH- $3d$ VERTEX COVER instance and that (G', k) is the instance of Π -VERTEX DELETION obtained by Construction 1. Let H_Π denote again the forbidden induced subgraph of Π that has the lexicographically smallest Γ -sequence, let H_1 denote the connected component of H_Π that has maximum γ -value and assume that d connected components of H_Π are isomorphic to H_1 . Finally, let $\mathcal{H}^<$ denote the set of graphs whose γ -value is smaller than H_1 . Now distinguish two cases.

Case 1: There is a graph \mathcal{I} whose connected components I_1, \dots, I_q are all from $\mathcal{H}^<$ such that the graph obtained by

- *taking the disjoint union of \mathcal{I} and $d - 1$ copies of H_1 , and then*
- *adding a further vertex v and making v adjacent to all vertices of the graph*

is not contained in Π . Observe that by the choice of H_Π and $\mathcal{H}^<$, \mathcal{I} is contained in Π . Perform the construction as previously, let (G', k) denote the graph as previously constructed. Now take the disjoint union of \mathcal{I} and G' . Then, add a vertex v^* , make v^* adjacent to all vertices of this graph, and call the resulting graph G^* . We show that (G', k) and $(G^*, k + 1)$ are equivalent instances.

First, assume that G' has a Π -vertex deletion set S' of size at most k . The proof of Lemma 3.4 shows that G has a vertex cover of size k which then implies that we can assume without loss of generality that S' does not contain vertices of the base graph. Thus, $G' - S'$ contains $d - 1$ connected components that are isomorphic to H_1 and all other connected components of $G' - S'$ have smaller γ -value than H_1 . By the choice of H_Π and by the fact that all connected components of \mathcal{I} are from $\mathcal{H}^<$, we have that the disjoint union of $G' - S'$ and \mathcal{I} is contained in Π . Thus, $G^* - (S' \cup \{v^*\})$ is contained in Π and G^* has a Π -vertex deletion set of size at most $k + 1$.

Now assume that (G', k) is a no-instance, that is, G' has no Π -vertex deletion set of size at most k . Let S' be a minimum-cardinality Π -vertex deletion set of G' . If $|S'| \geq k + 2$, then any Π -vertex deletion set of G^* has size at least $k + 2$ and the instances are equivalent. Hence, assume $|S'| = k + 1$. We first show that $G' - S'$ contains exactly $d - 1$ disjoint copies of H_1 .

Assume towards a contradiction that $d' < d - 1$ disjoint copies of H_1 remain in $G' - S'$. Then, at least $d^* = (d - 1) - d'$ copies of H_1 in the base graph are destroyed which means that S' contains at most $k + 1 - d^*$ vertices of the vc-extension. Furthermore, observe that the vc-extension of $G' - S'$ contains at most $d^* - 1$ disjoint copies of H_1 by the assumption that $G' - S'$ contains $d' < d - 1 = d' + d^*$ disjoint copies of H_1 . We show that G has a vertex cover of size at most k .

To this end, create a set $S_A \subseteq V$ as follows:

- For each vertex $v \in V \cap S'$, add v to S_A ,
- for each copy D^* of D in G' such that S' contains a vertex from $V(D^*) \setminus V$, add the vertex $v \in V(D^*) \cap V$ to S_A , and
- for each copy J^* of J such that $S' \setminus V$ contains a vertex from $V(J^*) \setminus V$, add an arbitrary vertex of $V(J^*) \cap V$ to S_A .

The set S_A has size at most $k + 1 - d^*$. Now, consider the graph $G - S_A$ and observe that every edge $\{u, v\}$ in $G - S_A$ directly corresponds to an induced copy of H_1 in $G' - S'$: Since $\{u, v\}$ is present in $G - S_A$, we have that S' does not contain u or v , does not contain vertices from the copies of D that are attached to u or v , and does not contain a vertex from the copy of J attached to u and v . Now let q denote the size of a maximum matching in $G - S_A$ and observe that this implies that the

vc-extension of $G' - S'$ has q vertex-disjoint copies of H_1 . Hence, $q < d^* < d$. Therefore, $G - S_A$ has a maximum matching of size less than d and thus (since G has girth $3d$) a vertex cover of size $q < d^*$. Consequently, G has a vertex cover of size $k + 1 - d^* + q \leq k$. This implies that G' has a Π -vertex deletion set of size at most k , contradicting the assumption that (G', k) is a no-instance.

By the above, for every Π -vertex deletion set S' of size at most $k + 1$ of G' , the graph $G' - S'$ contains exactly $d - 1$ disjoint copies of H_1 . Consequently, $G^* - S'$ contains a disjoint union of \mathcal{I} and $d - 1$ connected components isomorphic to H_1 and a vertex v^* that is adjacent to all vertices in $G^* - S'$. Hence, every Π -vertex deletion set of G^* has size at least $k + 2$.

Case 2: otherwise. Add a vertex v^* to the graph G' obtained by the original construction and make v^* adjacent to all vertices in G' , call the resulting graph G^* . We show that (G^*, k) and (G', k) are equivalent. First, if G^* has a Π -vertex deletion set of size at most k , then so does G' since G' is an induced subgraph of G^* . Second, let S' denote a Π -vertex deletion set of size at most k of G' . As discussed in the proof of Case 1, we can assume without loss of generality that S' does not contain vertices of the base graph.

Consequently, $G' - S'$ contains a disjoint union of $d - 1$ copies of H_1 and further graphs H' with $\gamma(H') < \gamma(H_1)$. Thus, with the exception of the copies of H_1 , every other connected component of $G' - S'$ is a graph from $\mathcal{H}^<$. By the case assumption, adding a vertex to $G' - S'$ and making it adjacent to all vertices of $G' - S'$ gives a graph in Π . Hence, the isomorphic graph $G^* - S'$ is also contained in Π .

Summarizing, in both cases we can modify the construction such that the resulting graph has one dominating vertex. Observe that this increases the number of edges by at most $|V(G')| + O(1)$. Thus, the resulting graph has $O(|V(G)|)$ edges and vertices, and the running time bound follows. \square

Now for graph properties excluding some fixed edgeless graph, all we can hope for is excluding $2^{o(n+\sqrt{m})}$ -time algorithms (since the general case can be solved in $2^{o(n+m)}$ time). As in the proof of Corollary 3.6, the lower bound can be easily shown by considering the property $\bar{\Pi}$ containing all graphs that are complements of graphs in Π .

PROPOSITION 5.4. *Let Π be a hereditary graph property such that Π does not contain all edgeless graphs. Then, if the ETH is true, Π -VERTEX DELETION cannot be solved in $2^{o(n+\sqrt{m})}$ time even if G has a dominating vertex.*

PROOF. Let $\bar{\Pi}$ denote the graph property that contains all graphs that are complement graphs of graphs in Π . Consider Construction 1 when applied to show that $\bar{\Pi}$ -VERTEX DELETION cannot be solved in $2^{o(n)}$ time. Let G' denote the graph obtained by this construction. Since H_1 contains at least one edge, one may safely add an isolated vertex at the end of the construction and obtain a graph F that has a $\bar{\Pi}$ -vertex deletion set of size at least k if and only if G' does. Hence, $\bar{\Pi}$ -vertex cannot be solved in $2^{o(n+m)}$ time assuming the ETH, also if the input graph F has an isolated vertex. Now the claimed hardness result follows from the fact that the complement graph \bar{F} has a dominating vertex and the same number of overall vertices and every induced subgraph $\bar{F}[X]$ of \bar{F} is in Π if and only if $F[X]$ is in $\bar{\Pi}$. \square

6 CONNECTED Π -VERTEX DELETION

Finally, we consider a popular variant of Π -VERTEX DELETION where the set of deleted vertices has to be connected. Special cases of this problem that have been studied include CONNECTED VERTEX COVER [12], CONNECTED FEEDBACK VERTEX SET [11, 22], and MINIMUM k -PATH CONNECTED VERTEX COVER [19].

CONNECTED Π -VERTEX DELETION

Input: An undirected graph $G = (V, E)$ and an integer k .

Question: Is there a set $S \subseteq V$ such that $G[S]$ is connected, $|S| \leq k$, and $G[V \setminus S]$ is contained in Π ?

For this problem, we obtain the same running time bounds as for Π -VERTEX DELETION.

THEOREM 6.1. *Let Π be a hereditary nontrivial graph property, then the following holds.*

- (1) *If the ETH is true, then CONNECTED Π -VERTEX DELETION cannot be solved in $2^{o(n+\sqrt{m})}$ time,*
- (2) *If the ETH is true, then CONNECTED Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time if Π contains all edgeless graphs.*
- (3) *If Π excludes some edgeless graph, then CONNECTED Π -VERTEX DELETION can be solved in $2^{O(\sqrt{m})} + O(n)$ time if and only if membership in Π can be recognized in $2^{O(n)}$ time.*
- (4) *If Π excludes some edgeless graph, then CONNECTED Π -VERTEX DELETION can be solved in $2^{o(m)} + O(n)$ time if and only if membership in Π can be recognized in $2^{o(m)}$ time.*

The two positive results for Π excluding some edgeless graph can be easily obtained by adapting the algorithm described in Lemma 4.1 in such a way that solutions (vertex sets to delete) are only accepted if they induce a connected subgraph. Also, the only if part of the last two statements follows directly since it is obtained by considering the special case $k = 0$, where the solution is empty and thus trivially connected. To prove the running time lower bounds given in Theorem 6.1, we extend Construction 1. One approach could be to add a universal vertex that needs to be deleted. To obtain hardness also in the case of bounded-degree graphs (if Π contains all edgeless graphs), we instead modify the construction by adding a set of vertices that is part of every connected minimum-cardinality Π -vertex deletion set and connects the solution vertices of the original instance by forming a binary tree.

LEMMA 6.2. *Let Π be a hereditary nontrivial graph property containing all edgeless graphs. If the ETH is true, then CONNECTED Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time.*

PROOF. We reduce from SUBCUBIC GIRTH-3d VERTEX COVER by adapting Construction 1. Let $(G = (V, E), k)$ denote the instance of SUBCUBIC GIRTH-3d VERTEX COVER. To simplify the construction somewhat, assume without loss of generality that $V = \{1, \dots, n\}$ and that the number n of vertices in G is a power of 2, that is, $n = 2^\mu$ for some $\mu \in \mathbb{N}$. First, perform Construction 1. Then, set $i := \mu - 1$ and add a set $V_i := \{v_i^1, \dots, v_i^{2^i}\}$ of 2^i vertices, and make each v_i^j adjacent to j and $j + 2^i$ (recall that the output graph of Construction 1 has V as a vertex subset). Now, for each vertex v_i^j of V_i , add a copy of H_1 and identify c (a fixed vertex of H_1 such that $\alpha(H_1, c) = \alpha(H_1)$) with v_i^j .

Continue the construction for decreasing i until $i = 0$ as follows. Set $i \leftarrow i - 1$, and add a vertex set $V_i := \{v_i^1, \dots, v_i^{2^i}\}$. Then, for each v_i^j , make v_i^j adjacent to v_{i+1}^j and $v_{i+1}^{j+2^i}$, and add a copy of H_1 and identify c with v_i^j . Let $V_T := \bigcup_{i \in \{0, \dots, \mu-1\}} V_i$ and call the copies of H_1 that are added in this step the V_T -attachments. Let $G' := (V', E')$ denote the graph obtained this way. Conclude the reduction by setting $k' := k + |V_T|$. Observe that $|V'| = O(n)$ and $|E'| = O(n)$. Thus, to show the lemma, it remains to show the equivalence of the SUBCUBIC GIRTH-3d VERTEX COVER instance (G, k) and of the CONNECTED Π -VERTEX DELETION instance (G', k') .

If (G, k) is a yes-instance with a size- k vertex cover S , then $S' := V_T \cup S$ is a size- k Π -vertex deletion set: the graph $G' - S'$ has at most $d - 1$ connected components with γ -value $\gamma(H_1)$ and all other connected components have γ -value less than $\gamma(H_1)$ (recall that d is the number of connected components of H_Π that are isomorphic to H_1). Then, by the definition of H_Π , $G' - S'$ fulfills Π . Moreover, $G'[S']$ is connected: $G'[S' \setminus S]$ is a binary tree and every vertex of V has one neighbor in $S' \setminus S$.

For the converse, observe that every connected Π -vertex deletion set S' deletes at least one vertex in the vc-extension. Thus, it may only delete vertices of the vc-extension and of V_T -attachments. Consequently, only the base graph may contain induced subgraphs that are isomorphic to H_1 . Note that each V_T -attachment is isomorphic to H_1 . Since, there are $|V_T|$ vertex-disjoint V_T -attachments, at least $|V_T|$ vertices of S' are not from the vc-extension. This implies that there is a set S^* of at most $k' - |V_T| = k$ vertices in the vc-extension whose deletion destroys all induced subgraphs of the vc-extension that are isomorphic to H_1 . As in the the proof of Lemma 3.4, this implies that G has a vertex cover of size at most k . \square

The reduction behind Lemma 6.2 is an extension of Construction 1 which increases the vertex degree of every vertex by at most one and additionally adds vertices whose degree is at most $\Delta(H_\Pi) + 3$, where $\Delta(H_\Pi)$ is the maximum degree in H_Π .

COROLLARY 6.3. *Let Π be a hereditary nontrivial graph property such that all edgeless graphs are in Π . Let Δ be the smallest number such that there exists a γ -ordering of the forbidden subgraphs of Π such that a forbidden subgraph with lexicographically smallest Γ -sequence has maximum degree Δ . Then, if the ETH is true, CONNECTED Π -VERTEX DELETION cannot be solved in $2^{o(n+m)}$ time even if G has maximum degree $3\Delta + 1$.*

For properties Π that do not include all edgeless graphs, we also obtain a tight result.

LEMMA 6.4. *Let Π be a hereditary nontrivial graph property excluding some edgeless graph. If the ETH is true, then CONNECTED Π -VERTEX DELETION cannot be solved in $2^{o(n+\sqrt{m})}$ time.*

PROOF. Let $\bar{\Pi}$ denote the property that contains all graphs that are complement graphs of a graph in Π . Reduce from SUBCUBIC GIRTH-3d VERTEX COVER by first performing Construction 1 with one difference: add d (instead of $d - 1$) copies of H_1 to the base graph. Let G' denote the graph obtained by this construction. Observe that at least one vertex of one of these d copies of H_1 is contained in any $\bar{\Pi}$ -vertex deletion set of the resulting graph. Moreover, it is sufficient to delete only one vertex v in this copy of H_1 . Thus, G' has a $\bar{\Pi}$ -vertex deletion set of size at most $k + 1$ if and only if G has a vertex cover of size at most k . Now, let \bar{G}' denote the complement graph of the graph obtained by the modified Construction 1. There is a minimum-cardinality Π -vertex deletion set that contains v and no further vertex from the copy of H_1 that contains v . Hence, v is in \bar{G}' adjacent to all other vertices of this minimum-cardinality Π -vertex deletion set, which makes the subgraph induced by this vertex deletion set connected. Therefore, \bar{G}' has a connected Π -vertex deletion set of size at most $k + 1$ if and only if G has a vertex cover of size at most k . The running time bound follows from observing that \bar{G}' has $O(|V|^2)$ edges and $O(|V|)$ vertices where V is the vertex set of the SUBCUBIC GIRTH-3d VERTEX COVER instance. \square

Finally, let CONNECTED PLANAR Π -VERTEX DELETION denote the variant of CONNECTED Π -VERTEX DELETION where the input graph is restricted to be planar. We can show the same running time lower bound as for the unconstrained vertex deletion problem on planar graphs.

THEOREM 6.5. *Let Π be a hereditary graph property containing all edgeless graph and excluding at least one planar graph. Then, if the ETH is true, CONNECTED PLANAR Π -VERTEX DELETION cannot be solved in $2^{o(\sqrt{n})}$ time.*

PROOF. Reduce from PLANAR VERTEX COVER. Given an instance (G, k) of PLANAR VERTEX COVER, fix an arbitrary embedding of G . Now add one vertex into every face and make this vertex adjacent to all vertices on the boundary of the face. Call the additional vertex set V_F and let E_F denote the set of edges incident with vertices of V_F . Now, perform Construction 1 on G as before, that is, add a base graph, replace each edge of G by a copy of J and add for each vertex v of G a copy

of $H_1 - (V(J) \setminus \{c\})$ and identify c with v . Now, for each vertex $v \in V_F$ add a copy of H_1 and identify c with v . Call the resulting graph G' and set $k' := k + |V_F|$. Observe that $|V_F| = O(|V|)$. Moreover, G' is planar, since it is obtained from a planar graph by replacing edges and vertices by planar graphs. Thus, if the two instances are equivalent, then any $2^{o(\sqrt{n})}$ -time algorithm for CONNECTED PLANAR Π -VERTEX DELETION implies a $2^{o(\sqrt{n})}$ for PLANAR VERTEX COVER. To complete the proof, we thus need to show that G has a vertex cover of size k if and only if G' has a connected Π -vertex deletion set of size k' .

(\Rightarrow) If S is a size- k vertex cover of G , then $S \cup V_F$ is a connected size- k Π -vertex deletion set of G' : After the deletion of $S \cup V_F$ all connected components of G' have γ -value lower than H_1 and thus, by the definition of H_Π and the construction of the base graph, the remaining graph is contained in Π . It remains to show that $G'[S \cup V_F]$ is connected. Every vertex in S has a neighbor in V_F . Moreover, every pair of vertices $u, v \in V_F$ that were placed into faces of G whose boundaries share an edge e are connected in $G'[S \cup V_F]$ because one of the two endpoints of the edge e is contained in S and both endpoints of e are neighbors of u and v . Thus, all vertices of V_F are connected in $G'[S \cup V_F]$ since the dual graph of the planar graph G is connected.

(\Leftarrow) Let S' denote a connected Π -vertex deletion set of size at most k' in G' . Since $G'[S']$ is connected, S' cannot contain any vertices of the base graph. Consequently, by the construction of the base graph and the choice of Π , all induced subgraphs isomorphic to H_1 must be destroyed in the connected component containing the vc-extension. Thus, for each copy of H_1 that was attached to a vertex of V_F at least one vertex must be deleted. This implies that S' contains at most k vertices of the vc-extension. Observe that for each subgraph isomorphic to H_1 in the vc-extension, there are two vertices $\{u, v\}$ that are from V and deleting these two vertices cuts the other vertices from the rest of the graph. Thus, since $G'[S']$ is connected at least one of these two vertices is contained in S' . Let S denote the set of these at most k vertices. Since for each edge $\{u, v\}$ of G , the graph G' contains a such subgraph isomorphic to H_1 having u and v as cut-vertices, S is a vertex cover in G . \square

7 CONCLUSION

We have shown that, by assuming the Exponential Time Hypothesis (ETH), a classic generic NP-hardness result for Π -VERTEX DELETION can be adapted to show tight running time lower bounds for Π -VERTEX DELETION. In particular, we have shown that for hereditary and nontrivial graph properties Π , the Π -VERTEX DELETION problem cannot be solved in $2^{o(n)}$ time on n -vertex graphs if the ETH is true. If membership in Π can be determined in polynomial time, then Π -VERTEX DELETION can be trivially solved in $2^n \cdot n^{O(1)}$ time. It is desirable, however, to obtain faster exponential-time algorithms for these problems [9]. In this context it is an important open question, which we repeat here, whether or not for each graph property Π such that containment in Π can be determined in polynomial time, there is some $\epsilon > 0$, such that Π -VERTEX DELETION can be solved in $O((2 - \epsilon)^n)$ time.

ACKNOWLEDGMENTS

This work is supported by the Deutsche Forschungsgemeinschaft, DFG, project KO 3669/4-1. The main part of the work was done while the author was with the Friedrich Schiller University Jena.

REFERENCES

- [1] N. R. Aravind, R. B. Sandeep, and Naveen Sivadasan. 2017. Dichotomy Results on the Hardness of H-free Edge Modification Problems. *SIAM Journal on Discrete Mathematics* 31, 1 (2017), 542–561.
- [2] Balabhaskar Balasundaram, Sergiy Butenko, and Illya V. Hicks. 2011. Clique Relaxations in Social Network Analysis: The Maximum k -Plex Problem. *Operations Research* 59, 1 (2011), 133–142.

- [3] Sebastian Böcker and Jan Baumbach. 2013. Cluster Editing. In *Proceedings of the 9th Conference on Computability in Europe (CiE '13) (LNCS)*, Vol. 7921. Springer, Heidelberg, Germany, 33–44.
- [4] Kevin Buchin, Maïke Buchin, Jaroslaw Byrka, Martin Nöllenburg, Yoshio Okamoto, Rodrigo I. Silveira, and Alexander Wolff. 2012. Drawing (Complete) Binary Tanglegrams - Hardness, Approximation, Fixed-Parameter Tractability. *Algorithmica* 62, 1-2 (2012), 309–332.
- [5] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer, Heidelberg, Germany.
- [6] Pål Grønås Drange, Fedor V. Fomin, Michal Pilipczuk, and Yngve Villanger. 2015. Exploring the Subexponential Complexity of Completion Problems. *ACM Transactions on Computation Theory* 7, 4 (2015), 14.
- [7] Michael R. Fellows, Jiong Guo, Hannes Moser, and Rolf Niedermeier. 2011. A Complexity Dichotomy for Finding Disjoint Solutions of Vertex Deletion Problems. *ACM Transactions on Computation Theory* 2, 2 (2011), 5.
- [8] Bernold Fiedler, Atsushi Mochizuki, Gen Kurosawa, and Daisuke Saito. 2013. Dynamics and Control at Feedback Vertex Sets. I: Informative and Determining Nodes in Regulatory Networks. *Journal of Dynamics and Differential Equations* 25, 3 (2013), 563–604.
- [9] Fedor V. Fomin and Dieter Kratsch. 2010. *Exact Exponential Algorithms*. Springer, Heidelberg, Germany.
- [10] Michael R. Garey, David S. Johnson, and Larry J. Stockmeyer. 1976. Some simplified NP-complete graph problems. *Theoretical Computer Science* 1, 3 (1976), 237–267.
- [11] Alexander Grigoriev and René Sitters. 2010. Connected Feedback Vertex Set in Planar Graphs. In *Proceedings of the 35th International Workshop on Graph-Theoretic Concepts in Computer Science (WG '09) (Lecture Notes in Computer Science)*, Vol. 5911. Springer, Heidelberg, Germany, 143–153.
- [12] Jiong Guo, Rolf Niedermeier, and Sebastian Wernicke. 2007. Parameterized complexity of Vertex Cover variants. *Theory of Computing Systems* 41, 3 (2007), 501–520.
- [13] Jiong Guo and Yash Raj Shrestha. 2014. Complexity of Disjoint II-Vertex Deletion for Disconnected Forbidden Subgraphs. *Journal of Graph Algorithms and Applications* 18, 4 (2014), 603–631.
- [14] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which Problems Have Strongly Exponential Complexity? *J. Comput. System Sci.* 63, 4 (2001), 512–530.
- [15] David S. Johnson and Mario Szegedy. 1999. What are the Least Tractable Instances of max independent set?. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '99)*. ACM/SIAM, Philadelphia, PA, USA, 927–928.
- [16] Christian Komusiewicz. 2016. Multivariate Algorithmics for Finding Cohesive Subnetworks. *Algorithms* 9, 1 (2016), 21.
- [17] John M. Lewis and Mihalis Yannakakis. 1980. The node-deletion problem for hereditary properties is NP-complete. *J. Comput. System Sci.* 20, 2 (1980), 219–230.
- [18] Kun Liu and Evimaria Terzi. 2008. Towards identity anonymization on graphs. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD '08)*. ACM, New York, NY, USA, 93–106.
- [19] Xianliang Liu, Hongliang Lu, Wei Wang, and Weili Wu. 2013. PTAS for the minimum k-path connected vertex cover problem in unit disk graphs. *Journal of Global Optimization* 56, 2 (2013), 449–458.
- [20] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. 2011. Lower bounds based on the Exponential Time Hypothesis. *Bulletin of the EATCS* 105 (2011), 41–72.
- [21] Dániel Marx. 2013. The Square Root Phenomenon in Planar Graphs. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP '13) (Lecture Notes in Computer Science)*, Vol. 7966. Springer, Heidelberg, Germany, 28.
- [22] Neeldhara Misra, Geevarghese Philip, Venkatesh Raman, Saket Saurabh, and Somnath Sikdar. 2012. FPT algorithms for Connected Feedback Vertex Set. *Journal of Combinatorial Optimization* 24, 2 (2012), 131–146.
- [23] André Nichterlein. 2014. *Degree-Constrained Editing of Small-Degree Graphs*. PhD Thesis. TU Berlin.
- [24] Mihalis Yannakakis. 1981. Edge-Deletion Problems. *SIAM J. Comput.* 10, 2 (1981), 297–309.

Received February 2007; revised March 2009; accepted June 2009