

p-Edge/vertex-connected vertex cover

Citation for published version (APA):

Einarson, C., Gutin, G. Z., Jansen, B. M. P., Majumdar, D., & Wahlström, M. (2023). p-Edge/vertex-connected vertex cover: Parameterized and approximation algorithms. *Journal of Computer and System Sciences*, 133, 23-40. <https://doi.org/10.1016/j.jcss.2022.11.002>

Document license:
CC BY

DOI:
[10.1016/j.jcss.2022.11.002](https://doi.org/10.1016/j.jcss.2022.11.002)

Document status and date:
Published: 01/05/2023

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



p -Edge/vertex-connected vertex cover: Parameterized and approximation algorithms

Carl Einarson^a, Gregory Gutin^{a,*}, Bart M.P. Jansen^b, Diptapriyo Majumdar^c, Magnus Wahlström^a

^a Royal Holloway, University of London, Egham, United Kingdom

^b Eindhoven University of Technology, the Netherlands

^c Indraprastha Institute of Information Technology Delhi, New Delhi, India

ARTICLE INFO

Article history:

Received 17 September 2020

Received in revised form 18 October 2022

Accepted 28 November 2022

Available online 7 December 2022

Keywords:

Vertex cover

Connectivity

Parameterized algorithms

Kernels

Approximate kernels

Approximation algorithms

ABSTRACT

We introduce and study two natural generalizations of the Connected Vertex Cover (VC) problem: the p -Edge-Connected and p -Vertex-Connected VC problem (where $p \geq 2$ is a fixed integer). We obtain an $2^{\mathcal{O}(pk)}n^{\mathcal{O}(1)}$ -time algorithm for p -Edge-Connected VC and an $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$ -time algorithm for p -Vertex-Connected VC. Thus, like Connected VC, both constrained VC problems are FPT. Furthermore, like Connected VC, neither problem admits a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$, which is highly unlikely. We prove however that both problems admit time efficient polynomial sized approximate kernelization schemes. Finally, we describe a $2(p+1)$ -approximation algorithm for the p -Edge-Connected VC. The proofs for the new VC problems require more sophisticated arguments than for Connected VC. In particular, for the approximation algorithm we use Gomory-Hu trees and for the approximate kernels a result on small-size spanning p -vertex/edge-connected subgraphs of a p -vertex/edge-connected graph by Nishizeki and Poljak (1994) [30] and Nagamochi and Ibaraki (1992) [27].

© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

For a graph $G = (V, E)$, a set $C \subseteq V$ is a *vertex cover* if for every edge $uv \in E$ at least one of the vertices u, v belongs to C . The well-known classical VERTEX COVER problem is the problem of deciding whether a graph G has a vertex cover of size at most k . This problem is NP-complete and thus was studied from parameterized complexity and approximation algorithms view points. VERTEX COVER and its generalizations have been important in developing basic and advanced methods and approaches for parameterized and approximation algorithms. Thus, it was named the *Drosophila* of fixed-parameter algorithmics [29]. In particular, it is well known that VERTEX COVER is fixed-parameter tractable with respect to k , admits a kernel with $2k$ vertices, and has a trivial 2-approximation [6,28].

A well-studied variant of VERTEX COVER is CONNECTED VERTEX COVER (CVC), where the vertex cover is additionally required to be connected. This problem is FPT parameterized by the solution size k , with the fastest known algorithm running in time $\mathcal{O}^*(2^k)$ due to Cygan [7] after a sequence of improvements.¹ CVC also has a classic 2-approximation, implicit in

* Corresponding author.

E-mail addresses: einarsoncarl@gmail.com (C. Einarson), g.gutin@rhul.ac.uk (G. Gutin), b.m.p.jansen@tue.nl (B.M.P. Jansen), diptapriyo@iitd.ac.in (D. Majumdar), magnus.wahlstrom@rhul.ac.uk (M. Wahlström).

¹ The \mathcal{O}^* notation suppresses the polynomial factors.

Savage [37]. However, unlike VERTEX COVER, CVC was shown to have no polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$ [9]. Through reductions, this lower bound can be shown to imply that several other variations of VERTEX COVER, like the ones we will consider, do not admit polynomial kernels.

To get around limitations on kernelization, Lokshtanov et al. [22], in pioneering work, proposed the notion of *approximate kernels*. Although the details are somewhat technical, in essence an α -approximate kernel can be thought of as a kernel that only preserves solution optimality up to a factor of α , e.g., from a c -approximate solution to the output instance we can recover an $(\alpha \cdot c)$ -approximate solution to the input instance. Lokshtanov et al. [22] considered several problems that are known not to admit polynomial kernels (unless $\text{NP} \subseteq \text{coNP/poly}$) and analyzed the existence of α -approximate polynomial kernels. In particular, they proved that CONNECTED VERTEX COVER admits an α -approximate polynomial kernel for every fixed α .

We study two natural variations of CVC under tighter connectivity constraints, from the perspectives of FPT algorithms, approximate kernels, and approximation algorithms. Let us review the definitions.

A connected graph is called *p-edge-connected* if it remains connected whenever fewer than p edges are deleted. A connected graph is called *p-connected* if it has more than p vertices and it remains connected whenever fewer than p vertices are deleted. In this paper, we introduce the following two natural generalizations of CONNECTED VERTEX COVER, the *p-EDGE-CONNECTED VERTEX COVER* and *p-CONNECTED VERTEX COVER* problems. For both problems, p is a fixed positive integer.

p-EDGE-CONNECTED VERTEX COVER (p-EDGE-CVC)

Input: An undirected graph G , and an integer k

Parameter: k

Problem: Does G have a set of at most k vertices that is a vertex cover and induces a p -edge-connected subgraph?

p-CONNECTED VERTEX COVER (p-CVC)

Input: An undirected graph G , and an integer k

Parameter: k

Problem: Does G have a set of at most k vertices that is a vertex cover and induces a p -connected subgraph?

Our results Both *p-EDGE-CONNECTED VERTEX COVER* and *p-CONNECTED VERTEX COVER* are NP-complete; proofs for this follow from the reductions provided in Theorems 7 and 8, respectively. It is also not hard to obtain simple FPT algorithms for both *p-EDGE-CONNECTED VERTEX COVER* and *p-CONNECTED VERTEX COVER*. Unfortunately, the running time of such algorithms is $\mathcal{O}^*(2^{\mathcal{O}(k^2)})$. For the sake of completeness, we give a short proof for the existence of such an algorithm for *p-CONNECTED VERTEX COVER*. However, for *p-EDGE-CONNECTED VERTEX COVER* we can do better: our first main result is a single-exponential fixed-parameter algorithm for *p-EDGE-CONNECTED VERTEX COVER* using dynamic programming on matroids and some specific characteristics of p -edge-connected subgraphs.

Theorem 1. For every fixed $p \geq 2$, *p-EDGE-CONNECTED VERTEX COVER* can be solved in $2^{\mathcal{O}(pk)} n^{\mathcal{O}(1)}$ deterministic time and space.

Our algorithm for *p-EDGE-CVC* is as follows. First, we enumerate all minimal vertex covers of G of size at most k . The number of such vertex covers is at most 2^k , and they can be enumerated in $\mathcal{O}^*(2^k)$ time, and space (see [26]). Then, for every minimal vertex cover H of G , we use representative sets to check if it can be extended to a p -edge-connected vertex cover $S^* \supseteq H$ of size at most k . This step uses a characterization of p -edge-connected graphs due to Agrawal et al. [2].

Unfortunately, our approach to prove Theorem 1 does not work for *p-CONNECTED VERTEX COVER*.

After proving Theorem 1, we prove that there are $(1 + \varepsilon)$ -approximate polynomial kernels for both *p-Edge-CVC* and *p-CVC* for every $\varepsilon > 0$. The kernelization algorithm is the same in both cases, but the analysis differs. We provide necessary terminology and notation on approximate kernels in Section 2.3. Our results on $(1 + \varepsilon)$ -approximate kernels are as follows.

Theorem 2. For every $\varepsilon > 0$ and every fixed $p \geq 2$, the problems *p-EDGE-CONNECTED VERTEX COVER* and *p-CONNECTED VERTEX COVER* admit a $(1 + \varepsilon)$ -approximate kernel with $k + 2k^2 + \lfloor (3 + \varepsilon)k \rfloor k^{2 \lceil p / \min(\varepsilon, 1) \rceil}$ vertices.

Note that the size guarantee is mainly of interest for $\varepsilon \leq 1$, due to the term $\min(\varepsilon, 1)$. The main difficulty in obtaining a kernelization for connected variants of VERTEX COVER consists of dealing with vertices which are not needed to make a minimal vertex cover, but that may be needed to boost the connectivity of a minimal vertex cover. We give a marking procedure which selects a bounded number of vertices to preserve in the kernelized instance, and show that the role that any of the forgotten vertices plays to boost the connectivity of a vertex cover can be mimicked by a small set of marked vertices. We use existing results by Nishizeki and Poljak [30] and Nagamochi and Ibaraki [27] on sparse spanning p -vertex/edge-connected subgraphs of p -vertex/edge-connected graphs for this argument. This leads to a proof that the size of an optimal solution does not increase by more than a factor $(1 + \varepsilon)$ in the kernelized instance. From the redundancy

of the forgotten vertices for the purpose of making a small vertex cover it easily follows that any solution in the reduced graph is also a valid solution in the original, which leads to our lossy kernelizations.

Our last main result is the following:

Theorem 3. *For every fixed $p \geq 2$, p -EDGE-CONNECTED VERTEX COVER admits a polynomial-time $2(p + 1)$ -factor approximation algorithm.*

The proof of this theorem uses the notion of p -blocks which can be obtained from a Gomory-Hu tree. Unfortunately, this approach is not applicable to p -CONNECTED VERTEX COVER and the existence of a constant-factor approximation algorithm for p -CONNECTED VERTEX COVER is an open problem (recall that p is a fixed positive integer).

Finally, for the sake of completeness, we give proofs that both the problems do not admit polynomial kernels unless $\text{NP} \subseteq \text{coNP/poly}$.

Related work To the best of our knowledge, before the first version of this paper appeared in arXiv in 2020, no study in parameterized or approximation algorithms has investigated vertex-subset problems involving p -vertex/edge-connectivity constraints for $p > 2$. Li et al. [20] proved structural results for 2-EDGE-CONNECTED DOMINATING SET and 2-CONNECTED DOMINATING SET. Nutov [31] obtained an approximation algorithm with expected ratio $\mathcal{O}(\log^4 n \cdot \log \log n \cdot (\log \log \log n)^3)$ for 2-CONNECTED DOMINATING SET. Recently, Abhinav et al. [1] studied a version of the vertex cover problem, where the edge-connectivity of the vertex cover is required to be at least $n - k$, where k is the parameter. Nutov studied a p -connectivity augmentation problem [34] and domination problems with high connectivity constraints (see [32,33]).

Lokshtanov et al. [22] obtained a polynomial size approximate kernelization scheme (PSAKS, see Definition 8) for CONNECTED VERTEX COVER parameterized by the solution size. Majumdar et al. [23] considered parameters that are strictly smaller than the size of the solution and obtained a PSAKS for CONNECTED VERTEX COVER parameterized by the deletion distance of the input graph to each of the following classes of graphs: cographs, bounded treewidth graphs, and chordal graphs. Ramanujan [36] obtained a PSAKS for CONNECTED FEEDBACK VERTEX SET. Eiben et al. [12] obtained a similar result for the CONNECTED \mathcal{H} -HITTING SET problem, where \mathcal{H} is a fixed set of graphs. Eiben et al. [13] also obtained PSAKSs for CONNECTED DOMINATING SET on two classes of sparse graphs. Krithika et al. [19] designed PSAKSs for the TREE CONTRACTION, STAR CONTRACTION, OUT-TREE CONTRACTION, and CACTUS CONTRACTION problems. Dvořák et al. [11] designed a PSAKS for STEINER TREES parameterized by the number of non-terminals.

Improving a result of Lokshtanov et al. [22], Manurangsi [24] obtained a smaller size PSAKS for MAX k -VERTEX COVER, where given an edge-weighted graph G and an integer k , and the aim is to find a subset S of k vertices that maximizes the total weight of edges covered by S (an edge e is covered by S if at least one endpoint of e is in S). Lossy kernels were also obtained for problems of contraction to generalizations of trees by Agrawal et al. [3] and to classes of chordal graphs by Gunda et al. [17]. A lossy kernel was designed by Bandyapadhyay et al. [4] for SAME-SIZE CLUSTERING parameterized by the cost of clustering. Recently, Jansen and Włodarczyk [18] obtained a 510-factor approximate kernel of polynomial size for PLANAR VERTEX DELETION.

Organization The rest of the paper is organized as follows. In Section 2, we provide additional terminology and notation, and preliminaries needed for our algorithms and hardness proofs. Section 3 presents the FPT algorithms for the vertex- and edge-connectivity versions of the problem, proving Theorem 1. Section 4 presents the approximate kernelization schemes, leading to a proof of Theorem 2. The approximation algorithm of Theorem 3 is presented in Section 5. Section 6 contains the kernelization lower bounds which justify the use of approximate kernels. We conclude the paper with Section 7, where we discuss some open problems on the topic.

2. Preliminaries

2.1. Sets and graph theory

For $r \in \mathbb{N}$, we use $[r]$ to denote the set $\{1, 2, \dots, r\}$. Let U be a set of elements, and \mathcal{F} a family of subsets of U . (A family may have multiple copies of the same subset.) Then \mathcal{F} is said to be a *laminar family*, if for every $X, Y \in \mathcal{F}$, either $X \subseteq Y$, or $Y \subseteq X$, or $X \cap Y = \emptyset$.

For a graph G and subset $S \subseteq V(G)$, $G[S]$ denotes the subgraph of G induced by S . Given a connected graph $G = (V, E)$, a set of vertices $S \subseteq V(G)$ is called a *separator* of G if $G - S$ is a disconnected graph. For two disjoint vertex sets A and B of G , an (A, B) -cut is a set F of edges such that $G - F$ is disconnected and no connected component of $G - F$ contains both a vertex of A and a vertex of B . Given a connected graph $G = (V, E)$, for a proper subset A of $V(G)$, let $\bar{A} = V(G) \setminus A$. Then, the (A, \bar{A}) -cut is the set of edges with one endpoint in A and the other endpoint in \bar{A} . The *size* of an (A, \bar{A}) -cut is the number of edges in this set, denoted by $|E(A, \bar{A})|$. A *cut* is a set of edges equal to the (A, \bar{A}) -cut for some proper subset A of $V(G)$.

Recall that graph G is p -connected if it has more than p vertices and $G - X$ is connected for every $X \subseteq V(G)$ with $|X| < p$. A graph G is p -edge-connected if it has at least two vertices and $G - Y$ is connected for every $Y \subseteq E(G)$ with $|Y| < p$. By Menger's theorem, (i) a graph is p -connected if and only if there are at least p internally vertex-disjoint paths

between every pair of vertices, and (ii) a graph is p -edge-connected if and only if there are at least p edge-disjoint paths between every pair of vertices.

The following proposition will be useful.

Proposition 1. [27,30] *Let G be a p -vertex/edge-connected graph for some $p \geq 1$. Then, there exists a polynomial-time algorithm that computes a p -vertex/edge-connected spanning subgraph H of G such that H has at most $p|V(G)|$ edges.*

2.2. Parameterized algorithms and kernels

A parameterized problem Π is a subset of $\Sigma^* \times \mathbb{N}$ for some finite alphabet Σ . An instance of a parameterized problem is a pair (x, k) where k is called the *parameter* and x is the *input*.

Definition 1 (Fixed-parameter tractability). A parameterized problem $\Pi \subseteq \Sigma^* \times \mathbb{N}$ is said to be *fixed-parameter tractable* (or *FPT*) if there exists an algorithm for solving the problem Π that on input (x, k) , runs in $f(k)|x|^c$ time, where $f : \mathbb{N} \rightarrow \mathbb{N}$ is a computable function and c is a constant.

Definition 2 (Kernelization). Let $\Pi \subseteq \Sigma^* \times \mathbb{N}$ be a parameterized problem. A *kernelization algorithm*, or in short, a *kernelization*, for Π is an algorithm with the following property. For any given $(x, k) \in \Sigma^* \times \mathbb{N}$, it outputs in time polynomial in $|x| + k$ a string $x' \in \Sigma^*$ and an integer $k' \in \mathbb{N}$ such that

$$((x, k) \in \Pi \Leftrightarrow (x', k') \in \Pi) \text{ and } |x'|, k' \leq h(k),$$

where h is an arbitrary computable function. If \mathcal{A} is a kernelization for Π , then for every instance (x, k) of Π , the result of running \mathcal{A} on the input (x, k) is called the *kernel* of (x, k) (under \mathcal{A}). The function h is referred to as the size of the kernel. If h is a polynomial function, then we say that the kernel is polynomial.

It is well-known [8] that a decidable parameterized problem is FPT if and only if it has a kernel. Kernelization can be viewed as a theoretical foundation of computational preprocessing and thus we are interested in investigating when a parameterized problem admits a kernel of small size. We are especially interested in *polynomial kernels* for which the bound $h(k)$ is a polynomial and thus in classifying which parameterized problems admit polynomial kernels or not. Over the last decade, the area of kernelization has developed a large number of approaches to design polynomial kernels as well as a number of tools for proving lower bounds based on assumptions from complexity theory. The lower bounds rule out the existence of polynomial kernels for many parameterized problems. We refer the reader to the textbooks [8,10] for an introduction to the field of parameterized algorithms and to [15] for an introduction to kernelization.

One way to refute the existence of polynomial kernel for a parameterized problem Π_2 unless $\text{NP} \subseteq \text{coNP/poly}$ is to obtain a special reduction from another parameterized problem Π_1 known not to have a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$.

Definition 3 (Polynomial parameter transformation). Let $\Pi_1, \Pi_2 \subseteq \Sigma^* \times \mathbb{N}$ be two parameterized problems. An algorithm \mathcal{A} is called a *polynomial parameter transformation* if it takes an instance (x, k) of Π_1 and outputs an instance (x', k') of Π_2 such that

- \mathcal{A} runs in polynomial time,
- $(x, k) \in \Pi_1$ if and only if $(x', k') \in \Pi_2$, and
- k' is polynomial in k .

Proposition 2. [5] *Suppose that Π_1 and Π_2 be two parameterized problems such that as nonparameterized problems, Π_1 is NP-complete and Π_2 is in NP. If there is a polynomial parameter transformation from Π_1 to Π_2 and Π_1 has no polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$, then Π_2 has no polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$.*

2.3. Parameterized optimization problems and approximate kernels

Informally, an α -approximate kernelization is a polynomial-time algorithm that, given an instance (I, k) of a parameterized problem, outputs an instance (I', k') (called an α -approximate kernel) such that $|I'| + k' \leq g(k)$ for some computable function g , and any c -approximate solution to the instance (I', k') can be turned into a $(c\alpha)$ -factor approximate solution of (I, k) in polynomial time. We will mainly follow Lokshtanov et al. [22] to formally introduce approximate kernelization. (Several other recent papers studied approximate kernelization, see e.g. [12,13,22,36].)

Definition 4. A *parameterized optimization* (maximization or minimization) problem is a computable function $\Pi : \Sigma^* \times \mathbb{N} \times \Sigma^* \rightarrow \mathbb{R} \cup \{\pm\infty\}$.

The instances of a parameterized optimization problem Π are pairs $(x, k) \in \Sigma^* \times \mathbb{N}$, and a solution to (x, k) is simply a string $s \in \Sigma^*$ such that $|s| \leq |x| + k$. The value of the solution s is $\Pi(x, k, s)$. The problems we deal with in this paper are all minimization problems. Therefore, we state the definitions only in terms of minimization problems (the definitions of maximization problems are analogous). Consider p -EDGE-CONNECTED VERTEX COVER parameterized by solution size. This is a minimization problem where the optimization function $pECVC : \Sigma^* \times \mathbb{N} \times \Sigma^* \rightarrow \mathbb{R} \cup \{\infty\}$ is defined as follows.

$$pECVC(G, k, S) = \begin{cases} \infty & \text{if } S \text{ is not a } p\text{-edge-connected vertex cover of } G, \\ \min\{|S|, k + 1\} & \text{otherwise.} \end{cases}$$

For p -CONNECTED VERTEX COVER, the optimization function $pVVC(G, k, S)$ is defined similarly by changing the first condition in the natural way.

Definition 5. For a parameterized minimization problem Π , the optimum value of an instance $(x, k) \in \Sigma^* \times \mathbb{N}$ is $\text{OPT}_\Pi(x, k) = \min_{s \in \Sigma^*, |s| \leq |x| + k} \Pi(x, k, s)$.

Naturally for the case of p -EDGE-CONNECTED VERTEX COVER, we denote

$$\text{OPT}_{\text{ecvc}}(G, k) = \min_{S \subseteq V(G)} pECVC(G, k, S).$$

The optimal objective value OPT_{ecvc} of p -CONNECTED VERTEX COVER is defined analogously. Throughout the paper, we drop the subscript when this does not lead to confusion.

Definition 6. Let $\alpha \geq 1$ be a real number and Π a parameterized minimization problem. An α -approximate polynomial time preprocessing algorithm \mathcal{A} for Π is a pair of polynomial-time algorithms as follows. The first one is called the reduction algorithm that computes a map $\mathcal{R}_\mathcal{A} : \Sigma^* \times \mathbb{N} \rightarrow \Sigma^* \times \mathbb{N}$. Given an input instance (x, k) of Π , the reduction algorithm outputs another instance $(x', k') = \mathcal{R}_\mathcal{A}(x, k)$.

The second algorithm is called the solution lifting algorithm. This algorithm takes an input instance (x, k) , the output instance (x', k') , and a solution s' to the output instance (x', k') . The solution lifting algorithm works in time polynomial in $|x|, k, |x'|, k'$, and $|s'|$, and outputs a solution s to (x, k) such that the following holds.

$$\frac{\Pi(x, k, s)}{\text{OPT}(x, k)} \leq \alpha \cdot \frac{\Pi(x', k', s')}{\text{OPT}(x', k')}.$$

The size of the reduction algorithm \mathcal{A} is a function $\text{size}_\mathcal{A} : \mathbb{N} \rightarrow \mathbb{N}$ defined as $\text{size}_\mathcal{A}(k) = \sup\{|x'| + k' : (x', k') = \mathcal{R}_\mathcal{A}(x, k), x \in \Sigma^*\}$.

Definition 7. Let $\alpha \geq 1$ be a real number. An α -approximate kernelization (or an α -approximate kernel) for a parameterized optimization problem Π , is an α -approximate polynomial time preprocessing algorithm \mathcal{A} for Π such that $\text{size}_\mathcal{A}$ is upper bounded by a computable function $g : \mathbb{N} \rightarrow \mathbb{N}$.

By definition, an α -approximate kernel for a parameterized optimization problem Π takes input (x, k) and outputs an instance (x', k') of Π in polynomial time and satisfies the other properties as given by Definition 6. We say that (x', k') is an α -approximate polynomial kernel if g is a polynomial function (in other words if $\text{size}_\mathcal{A}$ is bounded by a polynomial function).

Definition 8. A polynomial size approximate kernelization scheme (PSAKS) for a parameterized optimization problem Π is a family of α -approximate polynomial kernelization algorithms, one such algorithm for every $\alpha > 1$.

Definition 9. A PSAKS is said to be time efficient if both the reduction algorithm and the solution lifting algorithm run in time $f(\alpha)|x|^c$ for some function f and a constant c independent of $|x|, k$, and α .

2.4. Matroids

We now present the preliminaries on matroids needed to obtain an FPT algorithm for p -EDGE-CONNECTED VERTEX COVER.

Definition 10 (Matroid). Let U be a universe and $\mathcal{I} \subseteq 2^U$. Then, (U, \mathcal{I}) is said to be a matroid if the following conditions are satisfied.

1. $\emptyset \in \mathcal{I}$,
2. if $A \in \mathcal{I}$, then for all $A' \subseteq A$, $A' \in \mathcal{I}$, and
3. if there exist $A, B \in \mathcal{I}$ with $|A| < |B|$, then there exists $x \in B \setminus A$ such that $A \cup \{x\} \in \mathcal{I}$.

A set $A \in \mathcal{I}$ is called an *independent set*.

Note that all maximal independent sets of a matroid M are of the same size, called the *rank of M* and denoted by $\text{rank}(M)$. A maximal independent set is called a *basis* of M .

We present some useful standard constructions. Let U be a universe with n elements and $\mathcal{I} = \{A \subseteq U \mid |A| \leq r\}$. Then, (U, \mathcal{I}) is a matroid called a *uniform matroid*. Next, let $G = (V, E)$ be an undirected graph and let $\mathcal{I} = \{F \subseteq E(G) \mid G' = (V, F) \text{ is a forest}\}$. Then, (E, \mathcal{I}) is a matroid called a *graphic matroid*. Finally, let U be partitioned as $U = U_1 \cup \dots \cup U_r$ and let $\mathcal{I} = \{A \subseteq U \mid |A \cap U_i| \leq 1 \forall i \in [r]\}$. Then, (U, \mathcal{I}) is a matroid called a *partition matroid*.

A matroid M is said to be *representable over a field \mathbb{F}* if there is a matrix A over \mathbb{F} and a bijection $f: U \rightarrow \text{col}(A)$, where $\text{col}(A)$ is the set of columns of A , such that $B \subseteq U$ is an independent set in M if and only if $\{f(b) \mid b \in B\}$ is linearly independent over \mathbb{F} . Clearly the rank of M is the rank of the matrix A . A matroid representable over a field \mathbb{F} is called a *linear matroid* over \mathbb{F} . A graphic matroid and partition matroid can be represented over any field, while a uniform matroid (U, \mathcal{I}) with $|U| = n$, can be represented over any field \mathbb{F}_q with $q > n$. Furthermore, all these representations can be constructed in deterministic polynomial time. See Oxley [35] for details; see also [8,25] for expositions of the issues closer to our needs.

Given two matroids $M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$, the *direct sum* $M = M_1 \oplus M_2$ is the matroid $M = (E, \mathcal{I})$ where $E = E_1 \uplus E_2$ is the disjoint union of E_1 and E_2 and $I \subseteq E$ is independent in M if and only if $I \cap E_1 \in \mathcal{I}_1$ and $I \cap E_2 \in \mathcal{I}_2$. If M_1 and M_2 are represented by matrices A_1 and A_2 , respectively, over a common field \mathbb{F} , then a representation of M can be produced as

$$A = \begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix}.$$

Clearly, this can be generalized to the direct sum of an arbitrary number of matroids M_i with representations A_i over a common field.

Given a matroid $M = (U, \mathcal{I})$, the *truncation* of M to rank r is the matroid $M' = (U, \mathcal{I}')$ where a set $A \subseteq U$ is independent in M' if and only if $A \in \mathcal{I}$ and $|A| \leq r$. Given a representation of M , a representation of a truncation of M can be computed relatively easily in randomized polynomial time [25], but can also be computed in deterministic polynomial time through more involved methods [21]. For more information on matroids, see Oxley [35].

2.5. Existence of highly connected vertex covers

We will use the following useful lemma whose vertex-connectivity part is proved in [38] (Lemma 4.2.2). We were unable to find a proof of the edge-connectivity part of the lemma in the literature and thus provide a short proof here.

Lemma 1. *If G is a p -vertex-connected (p -edge-connected, respectively) graph and G' is obtained from G by adding a new vertex y adjacent with at least p vertices of G , then G' is p -vertex-connected (p -edge-connected, respectively).*

Proof. Let C be a cut in G' . If C is the $(y, V(G))$ -cut then by construction $|C| \geq p$. Otherwise, C is an (A, B) -cut for some A and B that both contain a vertex of $V(G)$, such that $A \uplus B = V(G')$. Depending on whether $y \in A$ or $y \in B$, $(A \setminus \{y\}, B)$ or $(A, B \setminus \{y\})$ is a cut in G , respectively. As G is p -edge-connected, $(A \setminus \{y\}, B)$ or $(A, B \setminus \{y\})$ has at least p edges. This means that the cut C itself has at least p edges. \square

The next lemma will help us to identify whether a graph $G = (V, E)$ has a p -vertex/edge-connected vertex cover or not.

Lemma 2. *Let $G = (V, E)$ be a graph, L the set of vertices of G with degree at most $p - 1$ and $S = V(G) \setminus L$. Then G has a p -vertex/edge-connected vertex cover if and only if S is a p -vertex/edge-connected vertex cover of G .*

Proof. The backward direction (\Leftarrow) of the proof is trivial.

We will prove the forward direction (\Rightarrow). Let S^* be a p -vertex/edge-connected vertex cover of G . If $S^* = V(G) \setminus L$, then we are done. Hence, we may assume that $S^* \neq V(G) \setminus L$. Observe that $S^* \cap L = \emptyset$ as otherwise for every $u \in S^* \cap L$ deleting the vertices adjacent to u (the edges incident to u) will make $G[S^*]$ disconnected. (For the vertex connectivity variant, we rely here on the fact that $|S^*| \geq p + 1$ by definition of p -connectivity, so that at least one vertex besides u remains in S^* when deleting its neighbors.) Let $A = V(G) \setminus (S^* \cup L)$ and note that A is an independent set as S^* is a vertex cover. For each vertex a of the independent set A , which has degree at least p , all its neighbors belong to the vertex cover S^* . Hence, by Lemma 1 adding A to S^* will result in a p -vertex/edge-connected vertex cover of G . So, $S = S^* \cup A$ is a p -vertex/edge-connected vertex cover of G , which completes the proof. \square

3. FPT algorithms

In this section, we obtain FPT algorithms for the two VC generalizations. While the algorithm for p -CONNECTED VERTEX COVER described in Section 3.1 is quite simple and runs in time $\mathcal{O}^*(2^{\mathcal{O}(k^2)})$, the one for p -EDGE-CONNECTED VERTEX COVER obtained in Section 3.2 is based on matroid techniques and is of complexity $\mathcal{O}^*(2^{\mathcal{O}(pk)})$.

3.1. FPT algorithm for p -CONNECTED VERTEX COVER

Theorem 4. For every fixed $p \geq 1$, p -CONNECTED VERTEX COVER can be solved in $\mathcal{O}^*(2^{\mathcal{O}(k^2)})$ deterministic time and polynomial space.

Proof. Suppose that (G, k) is the input instance of p -CONNECTED VERTEX COVER. Since isolated vertices can be removed without changing the answer to the problem, we may assume without loss of generality that G does not have any isolated vertices. We compute a set H of vertices that have to be in any p -connected vertex cover of size at most k , as follows. We put a vertex u into H if the degree of this vertex is at least $k + 1$ (so any vertex cover avoiding u is too large), or it is a neighbor to a vertex v such that $\deg_G(v) \leq p - 1$ (so any vertex cover containing v is not p -connected). If $|H| > k$ then (G, k) is a no-instance, so we may assume that $|H| \leq k$. Now, we partition the vertices of $G - H$ into two parts. We put a vertex u into I if $N_G(u) \subseteq H$. Otherwise, we put u into R . Observe that I is an independent set. If $G[R]$ has more than k^2 edges, then (G, k) is a no-instance since any vertex in R has degree at most k , so that k vertices cover at most k^2 edges of this induced subgraph. Thus in the remainder we assume $G[R]$ has at most k^2 edges. This implies $|R| \leq 2k^2$ since each vertex $v \in R$ has an edge to another vertex in R , as by definition $v \in R$ has a neighbor u outside H and $u \notin I$ since u has neighbor $v \notin H$, hence $u \in R$.

We say that two vertices $u \neq v$ are *false twins* in G if $N_G(u) = N_G(v)$. We will argue that when a vertex $v \in I$ has more than $k + 1$ false twins, we can safely remove v without changing the answer: graph $G - v$ has a p -connected vertex cover of size at most k if and only if G has one. For the first direction, note that any p -connected vertex cover S of $G - v$ with $|S| \leq k$ must contain all vertices of H . For a vertex $u \in H$, either u has degree larger than k in G , and the same holds in $G - v$ since the vertex v we remove has $k + 1$ false twins remaining in $G - v$; or $N_G(u) = N_{G-v}(u)$ and there is a vertex $w \in N_G(u)$ of degree less than p . Hence S contains all of H , which by $v \in I$ implies that S contains $N_G(v)$ so that S is a vertex cover of G . Since $(G - v)[S] = G[S]$ this shows S is a p -connected vertex cover of S .

For the converse direction, suppose S is a p -connected vertex cover of G of size at most k , which implies $S \supseteq H$. If $v \notin S$, then since S is also a vertex cover in $G - v$ and $G[S] = (G - v)[S]$, this shows $G - v$ has a p -connected vertex cover of size at most k . Now suppose $v \in S$ and let $u_1, \dots, u_{k+1} \in I$ be $k + 1$ false twins of v in G . Since $|S| \leq k$, there exists some $u_i \notin S$; define $S' := (S \setminus \{v\}) \cup \{u_i\}$. The graph $G[S']$ is isomorphic to $G[S]$ since v and u_i are false twins, hence $G[S'] = (G - v)[S']$ is p -connected. Since all edges covered by v were already covered by $H \subseteq S \cap S'$, the set S' is a p -connected vertex cover of G of size at most k that does not contain v , and therefore is also a solution in $G - v$.

The argumentation above shows that a vertex $v \in I$ which has more than $k + 1$ false twins can safely be removed. After exhaustively applying this operation, for any $A \subseteq H$, there are at most $k + 1$ vertices in I whose neighborhood equals A . Hence, $|I| \leq 2^k(k + 1)$. It follows that after exhaustive reduction, the graph is partitioned in $H \cup I \cup R$ with $|H| + |R| + |I| \leq k + 2k^2 + 2^k(k + 1)$. We now consider all subsets of $H \cup I \cup R$ with at most k vertices, of which there are at most $|H \cup I \cup R|^k \leq (k + 2k^2 + 2^k(k + 1))^k \leq 2^{\mathcal{O}(k^2)}$. For each such subset S , we test whether it forms a p -connected vertex cover in polynomial time via Menger’s theorem. After trying all candidates for S , we either find a p -connected vertex cover or conclude that the answer is no in time $\mathcal{O}^*(2^{\mathcal{O}(k^2)})$. Since the graph reduction can be computed in polynomial time and space, while the iteration over vertex sets of size k and the verification of a potential solution can easily be done in polynomial space, the space usage of the algorithm is polynomial in the size of the input. \square

3.2. Single-exponential algorithm for p -edge-connected vertex cover

In this subsection, we provide a single-exponential algorithm for p -EDGE-CONNECTED VERTEX COVER using dynamic programming and the method of *representative sets*. This method was introduced to FPT algorithms by Marx [25]; Fomin et al. [14] presented additional applications and a faster method of computing such sets. We give the definitions below.

Definition 11. Let $M = (E, \mathcal{I})$ be a matroid and $X, Y \subseteq E$. We say that X *extends* Y in M if $X \cap Y = \emptyset$ and $X \cup Y \in \mathcal{I}$. Furthermore, let \mathcal{S} be a family of subsets of E . A subfamily $\hat{\mathcal{S}} \subseteq \mathcal{S}$ is q -*representative* for \mathcal{S} if the following holds: for every set $Y \subseteq E$ with $|Y| \leq q$, there is a set $X \in \mathcal{S}$ that extends Y if and only if there is a set $\hat{X} \in \hat{\mathcal{S}}$ that extends Y .

Theorem 5 (Fomin et al. [14]). Let $M = (E, \mathcal{I})$ be a linear matroid of rank $p + q = k$ over some field \mathbb{F} and let $\mathcal{S} = \{S_1, \dots, S_t\}$ be a family of independent sets in M , each of cardinality p . A q -representative subset $\hat{\mathcal{S}} \subseteq \mathcal{S}$ of size $|\hat{\mathcal{S}}| \leq \binom{p+q}{p}$ can be computed in $\mathcal{O}(k^{\mathcal{O}(1)} \binom{p+q}{p}^{\omega-1} t)$ field operations. Here, $\omega < 2.37$ is the matrix multiplication exponent.

Our algorithm is based on the following result, due to Agrawal et al. [2]. An *out-branching* of a digraph is a spanning subgraph which is a rooted tree, where every arc is oriented away from the root; or equivalently, a tree where the root has no incoming arc and every other vertex has exactly one incoming arc.

Lemma 3 (Agrawal et al. [2]). *Let $G = (V, E)$ be an undirected graph and let $v_r \in V$. Define a digraph $D_G = (V, A_E)$ by adding the arcs $(u, v), (v, u)$ to A_E for every edge $uv \in E$. Then G is p -edge-connected if and only if D_G has p pairwise arc-disjoint out-branchings rooted in v_r .*

We will use representative sets to facilitate a fast dynamic-programming algorithm. We begin by observing how out-branchings are realized using matroid tools. Let $G = (V, E)$ and $D_G = (V, A_E)$ be as defined above and let $v_r \in V$. The *out-partition matroid* (with root v_r) for a digraph D with $v_r \in V$ is the partition matroid with ground set A_E where arcs are partitioned according to their heads and where arcs (u, v_r) are dependent. It means that an arc set F is independent in the out-partition matroid if and only if v_r has in-degree 0 in F and every other vertex has in-degree at most one in F . The *graphic matroid on ground set A_E* is the graphic matroid for G , where every arc (u, v) represents its underlying edge uv and where anti-parallel arcs $(u, v), (v, u)$ represent distinct copies of uv . Note that $\{(u, v), (v, u)\}$ is thus a dependent set. The following proposition then holds true from the characteristics of out-partition matroids and graphic matroids.

Proposition 3. *F is the arc set of an out-branching rooted in v_r if and only if $|F| = |V(G)| - 1$ and F is independent in both the out-partition matroid for D_G with root v_r and the graphic matroid for G on ground set A_E .*

We extend this to construct a matroid that can be used to verify the condition of Lemma 3.

Lemma 4. *Let $v_r \in V(G)$ be a fixed vertex. Let M be the direct sum of $2p + 1$ matroids M_i as follows. Matroids $M_1, M_3, \dots, M_{2p-1}$ are copies of the graphic matroid of G on ground set A_E . Matroids M_2, M_4, \dots, M_{2p} are copies of the out-partition matroid for D_G with root v_r . Matroid M_{2p+1} is the uniform matroid over A_E with rank $p(k - 1)$. Let $F \subseteq A_E$. The following are equivalent.*

1. F is the arc set of p pairwise arc-disjoint out-branchings rooted in v_r in $D_G[S]$ for some $S \in \binom{V(G)}{k}$ with $v_r \in S$.
2. $|V(F)| = k, |F| = p(k - 1), v_r \in V(F)$, and there is an independent set I in M where every arc $a \in F$ occurs in I precisely in its copies in matroids M_{2i-1}, M_{2i} and M_{2p+1} for some $i \in [p]$.

Furthermore, a representation of M can be constructed in deterministic polynomial time.

Proof. Let M and $M_i, i \in [2p + 1]$ be as described. We note that graphic matroids, partition matroids (hence out-partition matroids), and uniform matroids all have deterministic representations [35]. Specifically, graphic matroids and partition matroids are representable over any field, and uniform matroids over any sufficiently large field. Hence a representation of M can be constructed as a diagonal block matrix with $2p + 1$ blocks, where each block i is a representation of the matroid M_i over some sufficiently large field \mathbb{F} (e.g., $\mathbb{F} = \mathbb{F}_q$ for some prime $q > p(k - 1)$).

On the one hand, let F meet the conditions in the first item. Since F is spanning for $D_G[S]$ we have $|V(F)| = k$ and $v_r \in V(F)$, and furthermore $|F| = p(k - 1)$ since each out-branching is spanning. Furthermore, any arc set of an out-branching is independent in both the graphic matroid and the out-partition matroid by Proposition 3. Letting $F = F_1 \cup \dots \cup F_p$ where F_i is the arc set of an out-branching for every $i \in [p]$, we can then construct I by letting an arc $a \in F_i$ be present in I in matroids M_{2i-1}, M_{2i} and M_{2p+1} . Hence all conditions in the second item are met.

Now assume that F meets the conditions in the second item. Partition $F = F_1 \cup \dots \cup F_p$ where $F_i, i \in [p]$ contains those arcs of F that are represented in matroids M_{2i-1} and M_{2i} . By a counting argument, $|F_i| = k - 1$ for every $i \in [p]$. Furthermore, F_i is the arc set of an out-forest (since its underlying undirected edge set is acyclic and every vertex has in-degree at most 1 in F_i). But then F_i must form a spanning tree of $V(F)$ by counting, hence an out-branching of $D_G[V(F)]$ by Proposition 3. Furthermore $v_r \in V(F)$ and every arc into v_r is dependent in M_{2i} ; thus every out-branching F_i is rooted in v_r . \square

To compute a representative family, we need to modify M so that the set I being described is a basis of M , not just an independent set. This is more technical, but can be done deterministically using the operation of deterministic truncation, due to Lokshtanov et al. [21], as noted in Section 2.4. We have the following lemma that ensures our matroid is linear, can be constructed in polynomial time, and satisfies the properties that we need.

Lemma 5. *Let $M' = M_1 \oplus \dots \oplus M_{2p+1}$ be the matroid defined in Lemma 4. Let M be the truncation of M' to rank $r = 3p(k - 1)$. Then a representation of M can be computed in deterministic polynomial time. Furthermore, every set I as defined in Item 2 of Lemma 4 is a basis of M .*

Proof. By standard methods [35], for every $i \in [2p + 1]$ we can compute a matrix A_i representing the matroid M_i , where furthermore all matrices A_i are over a common field \mathbb{F} . Thereby, we can also construct a matrix A representing their direct

sum M' . A matrix A' representing the r -truncation of M over a finite field \mathbb{F}' can then be constructed in deterministic polynomial time by Lokshantov et al. [21]. For the final statement, since every set I as described is independent in M' and has $|I| = r$, every such set must be a basis of M . \square

Our algorithm for p -EDGE-CVC works as follows. First, we enumerate all minimal vertex covers of G of size at most k . The number of such minimal vertex covers is at most 2^k , and they can be enumerated in $\mathcal{O}^*(2^k)$ time and space (see [26]). Then, for every minimal vertex cover H of G , we use representative sets and the above characterization to check if it can be extended to a feasible p -edge-connected vertex cover $S^* \supseteq H$ of size at most k . In detail, consider a graph G with a vertex cover H which is not p -edge-connected, where we are looking for a set $S^* \supset H$ such that $G[S^*]$ is p -edge-connected and $|S^*| \leq k$. By iteration over k , we may assume that $|S^*| < k$ is impossible.

Let us fix $v_r \in H$. By Lemma 4, there exists such a set S^* if and only if there is an independent set I in M meeting the following conditions:

1. $H \subset V(I)$ and $|V(I)| = k$.
2. $|I| = 3p(k - 1)$.
3. Every arc which is represented in I is represented in precisely three matroids $M_{2i-1}, M_{2i}, M_{2p+1}$ in M for some $i \in [p]$.

The first condition can be reformulated as $H \subset V(I)$ and $|V(I) \setminus H| = k - |H|$.

We can construct I via dynamic programming. The dynamic program is set up via a table keeping track of $|V(I)|$ and $|I|$, and we ensure that every time we add some arc a to a set I we add it in precisely three layers, as described. Thanks to the use of representative sets, each table entry in the dynamic programming only needs to contain $2^{\mathcal{O}(pk)}$ partial solutions.

We provide the details of this scheme in the proof of the main result of this section, Theorem 1 (we restate it here).

Theorem 1. For every fixed $p \geq 2$, p -EDGE-CONNECTED VERTEX COVER can be solved in $2^{\mathcal{O}(pk)} n^{\mathcal{O}(1)}$ deterministic time and space.

Proof. As outlined above, we may assume that we have a vertex cover H that is not p -edge-connected and have already tested that there is no p -edge-connected vertex cover with at most $k - 1$ vertices. Arbitrarily order the vertices of $V(G) \setminus H$: $v_1, \dots, v_{n'}$. Construct M as in Lemma 5. We create a dynamic programming table $T[(i, j, q)]$ with entries indexed by (i, j, q) for $i \leq k, j \leq n'$ and $q \leq 3p(k - 1)$. Every table slot $T[(i, j, q)]$ is a collection of independent sets in $M = (E(M), \mathcal{I})$. Any independent set $I \in T[(i, j, q)]$ satisfies $|V(I) \setminus H| = i$, the largest-index vertex of $V(G) \setminus H$ occurring in $V(I)$ is v_j , and $|I| = q$. Furthermore, for every independent set I in the table, every arc occurring in I occurs in precisely three layers, as described by Condition 3. We may then check for a solution by checking whether any slot $T[(k - |H|, j, 3p(k - 1))]$ is non-empty.

For an arc $a \in A_E$, define $F_{a,i}$ to be the set consisting of the copies of a in M_{2i-1}, M_{2i} and M_{2p+1} .

We initialize the slots $T[(0, 0, q)]$ by a dynamic programming process within $D_G[H]$. Initialize $T[(0, 0, 0)] = \{\emptyset\}$. Enumerate the arcs of $D_G[H]$ as $a_1, \dots, a_{m'}$. Then, for every $q = 3i$, where $i \in [p(k - 1)]$, fill in the slot $T[(0, 0, q)]$ from $T[(0, 0, q - 3)]$ as follows:

1. For every $I \in T[(0, 0, q - 3)]$, every arc a_j , and every $i \in [p]$ such that $F_{a_j,i}$ extends I , add $I \cup F_{a_j,i}$ to $T[(0, 0, q)]$.
2. Reduce $T[(0, 0, q)]$ to a $(3p(k - 1) - q)$ -representative set in M by applying Theorem 5, which is justified since the rank of M is $3p(k - 1)$.

This is a polynomial number of steps, where every set $I \in T[(0, 0, q)]$ is used in a polynomially bounded number of new sets. Hence every time we apply Theorem 5 at a level q , we do so with $t \leq (k + p)^{\mathcal{O}(1)} \binom{3p(k-1)}{q}$. Thus up to polynomial factors each step takes time $\binom{3p(k-1)}{q}^\omega = 2^{\mathcal{O}(pk)}$.

For slots $T[(i, j, q)]$, we process vertices v_j one at a time. The process is slightly more complex since each vertex v_j can be incident to $\mathcal{O}(k)$ arcs in A_E , but the principle is the same. We process slots $T[(i, j, q)]$ in lexicographic order by (i, j, q) . Note that we are here processing slots in a “forward” direction, i.e., we are using the sets in the slot $T[(i, j, q)]$ to populate slots $T[(i + 1, j', q')]$, which come after (i, j, q) in lexicographic order. Before we process the sets in a slot $T[(i, j, q)]$, we reduce them to a $(3p(k - 1) - q)$ -representative set in M . Then we proceed as follows. For every independent set I in $T[(i, j, q)]$ and every j' with $j < j' \leq n'$, we combine I and $v_{j'}$ as follows. As $v_{j'} \in V(G) \setminus H$ and H is a minimal vertex cover of G , we have that $N_G(v_{j'}) \subseteq H$. Therefore, there are at most $2|H|$ arcs of A_E incident on $v_{j'}$ in D_G .

1. Let $d \leq 2|H|$ be the number of arcs of A_E incident with $v_{j'}$. Create a set F for every one out of the following options: for every arc a incident with $v_{j'}$, either add $F_{a,b}$ to F for some $b \in [p]$, or do not add any set $F_{a,b}$ to F . Note that this makes $(p + 1)^d$ different sets F in total.
2. For every such non-empty set F , and every independent set I of $T[(i, j, q)]$ such that F extends I in M , add $I \cup F$ to $T[(i + 1, j', q + |F|)]$.

To roughly bound the running time, we note that the number of slots in the table is polynomial, and for every set I in a slot, at most $(p + 1)^{2|H|}$ sets $I' = I \cup F$ are added to other slots of the table. Furthermore, after the representative set reduction, every slot contains at most $2^{3p(k-1)}$ sets. Thus every time we apply Theorem 5, we have

$$|T[(i, j, q)]| < (p + k)^{O(1)} 2^{3pk} (p + 1)^{2k} = 2^{O(pk)},$$

hence the total time usage, up to a polynomial factor, is $2^{O(pk)}$.

It remains to prove correctness. Let $I \in T[(i, j, q)]$ for some (i, j, q) and let $S = V(I) \setminus H$. We note the invariants $|S| = i$, $\max_a \{v_a \in S\} = j$ and $|I| = q$ hold by induction. Furthermore, by construction I is independent in M . With these observations, we proceed. First assume that there is a set $I \in T[(k - |H|, j, 3p(k - 1))]$ for some j . Then $|V(I) \setminus H| = k - |H|$ and $|I| = 3p(k - 1)$ by the invariants. Furthermore, every arc a represented in I occurs in precisely three copies by construction. Indeed, every time we grow a set I we do so by adding a collection of sets $F_{a,i}$ to it. Hence every arc occurs at least three times, and furthermore, since $F_{a,i}$ always contains a copy of a in M_{2p+1} , we will never add two distinct sets $F_{a,i}, F_{a,i'}$ to the same set I . Hence Lemma 4 implies that $D_G[H \cup V(I)]$ has p pairwise arc-disjoint out-branchings rooted in v_r , which by Lemma 3 implies that $G[H \cup V(I)]$ is p -edge-connected.

On the other hand, assume that $G[H \cup S]$ is p -edge-connected for some $S \subseteq V(G) \setminus H$ with $|S| = k - |H|$. By Lemma 3 there exist p pairwise edge-disjoint out-branchings in $D_G[H \cup S]$ rooted in v_r , hence by Lemma 4 there is an independent set I in M formed using a set of arcs $F \subseteq A_E$ with $|F| = p(k - 1)$ and $|I| = 3p(k - 1)$ as described. Let j be the largest index such that $v_j \in V(I)$; then I is a candidate for the table slot $T[(k - |H|, j, 3p(k - 1))]$. We prove by induction that $T[(k - |H|, j, 3p(k - 1))]$ is non-empty. Observe that I is the disjoint union of sets $F_{a,i}$. We partition I according to lexicographical order of (i, j, q) as follows. First, let $F_{a_1, i_1}, \dots, F_{a_t, i_t}$ enumerate the sets $F_{a,i}$ contained in I for which a is contained in $D_G[H]$. For $r \in [t]$, let

$$I'_r = I \setminus \bigcup_{j=1}^r F_{a_j, i_j}$$

be the subset of I which is encountered “after” F_{a_r, i_r} in the natural ordering. We show by induction that for each $r \in [t]$, the slot $T[(0, 0, 3r)]$ contains a set which extends I'_r . For $r = 0$ this holds trivially. Hence, assume the statement holds for $T[(0, 0, 3r)]$ for some $r < t$, and let $I_0 \in T[(0, 0, 3r)]$ be a set which extends I'_r . While processing $(0, 0, 3r)$, the set $F_{a_{r+1}, i_{r+1}}$ is considered in the loop, and clearly it extends I_0 since $F_{a_{r+1}, i_{r+1}} \subseteq I'_r$. Hence $T[(0, 0, 3r + 3)]$ contains the set $I_1 = I_0 \cup F_{a_{r+1}, i_{r+1}}$ before the representative set computation is performed. By assumption I_1 extends I'_{r+1} . Hence by the correctness of Theorem 5, $T[(0, 0, 3r + 3)]$ contains some set I_2 that extends I'_{r+1} , as required. Hence the claim holds up to the set $T[(0, 0, 3t)]$.

We can now complete the proof using the same outline for entries $T[(i, j_i, q_i)]$. Enumerate S as $S = \{v_{j_1}, \dots, v_{j_{k-|H|}}\}$ in increasing order of indices j_i and for each $i \in [k - |H|]$ let I_i be the union of sets $F_{a,i}$ of I for which a is incident with v_{j_i} . Let $I_{\geq i} = \bigcup_{j=i}^{k-|H|} I_j$. We show by induction in lexicographical order that $T[(i, j_i, q_i)]$ for some q_i contains a set which extends $I_{\geq i+1}$, for each i . As a base case, the claim holds for $T[(0, 0, 3t)]$ as has already been shown. For the inductive step, the proof is precisely as above. For every $i = 1, \dots, k - |H|$, let $I'_{i-1} \in T[(i - 1, j_{i-1}, q_{i-1})]$ be a set which extends $I_{\geq i}$. Then in particular I_i extends I'_{i-1} and is added to table $T[(i, j_i, q_i)]$ in the exhaustive enumeration loop from $T[(i - 1, j_{i-1}, q_{i-1})]$. Thus before the call to Theorem 5 there was a set in $T[(i, j_i, q_i)]$ which extends $I'_{\geq i+1}$, hence the same holds after the representative set reduction. By induction, the table slot $T[(k - |H|, j, 3p(k - 1))]$ is indeed non-empty for some j . This completes the proof of the theorem. \square

4. Approximate kernels

In this section we describe the approximate kernels for the two types of connectivity. Both cases rely on a common subroutine $\text{Mark}(G, k, \varepsilon)$, which we present below. It works for a fixed value of p . It gets as input a graph G , integer k , and a value of $\varepsilon > 0$, and works as follows.

1. Let H be the vertices in G whose degree is larger than k . Let I consist of the vertices v in $V(G) \setminus H$ with $N_G(v) \subseteq H$. Note that I is an independent set. Let $R = V(G) \setminus (H \cup I)$, and note that each vertex in R has degree at least one and at most k .
2. If $|R| > 2k^2$ or $|H| > k$ then return INFEASIBLE.
3. Otherwise, we mark a set $L \subseteq I$. Initialize $L = \emptyset$. For each set $S \in \binom{H}{\leq \lceil p/\varepsilon \rceil}$, we mark common neighbors of S in I as follows:

- (a) If $|\bigcap_{v \in S} N_G(v) \cap I| \leq (3 + \varepsilon)k$, then add all vertices of $\bigcap_{v \in S} N_G(v) \cap I$ to L .
- (b) Otherwise, let L_S consist of $\lfloor (3 + \varepsilon)k \rfloor$ arbitrary vertices from $\bigcap_{v \in S} N_G(v) \cap I$ and add L_S to L .

4. Return the graph $G' = G[H \cup R \cup L]$ with parameter value $k' = \lceil (1 + \varepsilon)k \rceil$.

Note that since G' is an induced subgraph of G , any p -vertex/edge-connected vertex cover X in G with $X \subseteq V(G')$ is also a p -vertex/edge-connected vertex cover in G' . Note that a run which outputs a graph G' results in $|L| \leq \lfloor (3 + \varepsilon)k \rfloor \cdot |(\leq 2^{\lceil p/\varepsilon \rceil})^H| \leq \lfloor (3 + \varepsilon)k \rfloor k^{2^{\lceil p/\varepsilon \rceil}}$.

The following lemma encapsulates which information Mark preserves.

Lemma 6. *If $\text{Mark}(G, k, \varepsilon)$ is executed with $0 < \varepsilon \leq 1$ on a graph G which has a p -connected (resp. p -edge-connected) vertex cover X of size at most k , then it outputs a graph $G[H \cup R \cup L]$ (rather than INFEASIBLE) and $G[H \cup R \cup L]$ has a p -connected (resp. p -edge-connected) vertex cover of size at most $(1 + \varepsilon)|X|$.*

Proof. We first argue that Mark does not output INFEASIBLE when given a graph G with a p -connected (p -edge-connected) vertex cover of size at most k , as follows. We have $|H| \leq k$ as each vertex of H has degree more than k and belongs to each vertex cover of size at most k . Similarly, we must have $|R| \leq 2k^2$: as R consists of vertices outside $I \cup H$, each vertex of R has a neighbor that does not belong to H and which is therefore in R itself. Hence $G[R]$ contains no isolated vertices and has maximum degree k . If $|R| > 2k^2$, then it has more than k^2 edges which cannot be covered using at most k of its vertices. Hence if $|R| > 2k^2$ then $G[R]$ has no vertex cover of size at most k and neither does G .

It follows that under the stated assumption, the algorithm outputs a graph $G' = G[H \cup R \cup L]$. In the remainder, we refer to a p -vertex/edge-connected vertex cover in G (depending on the problem variant considered) as a *solution*. Let X be a solution in G of size at most k . We prove that $G[H \cup R \cup L]$ has a solution of size at most $(1 + \varepsilon)|X|$.

Since each vertex of H has degree more than k , it follows that $H \subseteq X$. If $X \subseteq V(G')$ then X is also a solution in G' and the lemma follows. In the remainder, we treat the case that $X \setminus V(G') = \{x_1, \dots, x_\ell\}$ for some $1 \leq \ell \leq |X| \leq k$. As $H \cup R \cup I$ is a partition of $V(G)$, since the output graph G' contains all of H and R it follows that for each $x_i \in X \setminus V(G')$ we have $x_i \in I$, which implies $N_G(x_i) \subseteq H \subseteq X$. We will show that we can replace each vertex x_i by a small set of marked vertices to obtain an approximate solution in G' .

By Proposition 1, graph $G[X]$ has a p -vertex/edge-connected spanning subgraph F on at most $p|X|$ edges. In the following argument, the degree of vertices x_i in the subgraph F will play an important role. As I is an independent set in G , it is also an independent set in its subgraph F . It follows that in the sum $\sum_{i=1}^{\ell} \text{deg}_F(x_i)$ we never count the same edge twice, so that the sum is bounded by the total number of edges in F , which is at most $p|X|$. We record this property for further use:

$$\sum_{i \in [\ell]} \text{deg}_F(x_i) \leq p|X|. \tag{1}$$

Constructing replacement sets For the replacement, we construct a sequence of sets X'_1, \dots, X'_ℓ . Each set X'_i will be used as a replacement for the corresponding vertex x_i . The sets we construct will have the following properties:

- (a) Each set X'_i is a subset of $V(G') \cap I$ that is disjoint from $X \cup \bigcup_{j < i} X'_j$.
- (b) The vertices in each set X'_i can be ordered so that each successive pair of vertices of X'_i has at least p common neighbors in the set H .
- (c) $N_F(x_i) \subseteq \bigcup_{u \in X'_i} N_G(u)$.
- (d) $|X'_i| \leq \max(\frac{\text{deg}_F(x_i)}{\lceil p/\varepsilon \rceil}, 1)$.

In the remainder of the proof, whenever we refer to successive or consecutive vertices of X'_i , we mean with respect to the ordering whose existence is guaranteed by the second condition. We construct the sets X'_i in order of increasing i . Consider a vertex $x_i \in X \setminus V(G') \subseteq I$. Let $S = N_F(x_i)$ and note that $|S| = \text{deg}_F(x_i)$. Since F is p -vertex/edge-connected, we have $|S| \geq p$. We define a partition of S as follows. If $|S| \leq \lceil p/\varepsilon \rceil$ then we use the singleton partition of $S = S_1$. If $|S| > \lceil p/\varepsilon \rceil$ then we partition S into sets S_1, \dots, S_r of size exactly $\lceil p/\varepsilon \rceil$, except for the last set which has size at least $\lceil p/\varepsilon \rceil$ and less than $2\lceil p/\varepsilon \rceil$. Such a partition always exists. Note that $r \leq \max(\frac{\text{deg}_F(x_i)}{\lceil p/\varepsilon \rceil}, 1)$, even without rounding, where the maximum is needed to deal with the case $r = 1$. For each $i \in [r - 1]$, let T_i be an arbitrary subset of S_{i+1} of size exactly p , which exists since $|S_{i+1}| \geq \lceil p/\varepsilon \rceil \geq p$ since $\varepsilon \leq 1$. Let $T_r = \emptyset$.

Note that $|S_j \cup T_j| \leq 2\lceil p/\varepsilon \rceil$ for each $j \in [r]$: for $j < r$ we have $|S_j| \leq \lceil p/\varepsilon \rceil$ and $|T_j| \leq p \leq \lceil p/\varepsilon \rceil$ (we use $\varepsilon \leq 1$ here), while the case $j = r$ holds since $T_r = \emptyset$. Each set $S_j \cup T_j$ consists of neighbors of x_i in F and therefore in G , which shows that $x_i \in \bigcap_{v \in S_j \cup T_j} N_G(v) \cap I$. Hence x_i was eligible to be marked for the set $S_j \cup T_j$, but it was not. Hence we marked a set $L_{S_j \cup T_j} \subseteq V(G') \cap I$ of $\lfloor (3 + \varepsilon)k \rfloor$ vertices. To show that there exist sufficiently many marked vertices which are not contained in X or in a set X'_j for $j' < j$, we bound the latter as follows:

$$\left| \bigcup_{j' < j} X'_{j'} \right| \leq \sum_{1 \leq j' < j} \max\left(\frac{\text{deg}_F(x_{j'})}{\lceil p/\varepsilon \rceil}, 1\right) \tag{By the fourth condition}$$

$$\begin{aligned}
 &\leq \sum_{j' \in [\ell]} \left(\frac{\deg_F(x_{j'})}{\lceil p/\varepsilon \rceil} + 1 \right) \\
 &\leq \frac{\sum_{j' \in [\ell]} \deg_F(x_{j'})}{\lceil p/\varepsilon \rceil} + \ell \leq \frac{p|X|}{\lceil p/\varepsilon \rceil} + \ell \quad \text{By (1)} \\
 &\leq k + \varepsilon k. \quad \text{Since } \ell, |X| \leq k
 \end{aligned}$$

So the set X contains at most k vertices of $L_{S_j \cup T_j}$, while $\bigcup_{j' < j} X'_{j'}$ contains at most $k + \varepsilon k$ vertices of $L_{S_j \cup T_j}$. It follows that there are at least k vertices of $L_{S_j \cup T_j}$ which belong neither to X nor to sets X'_i we already constructed. As this holds for each of the r sets into which we partitioned S , while $r \leq \deg_F(x_i) \leq \deg_G(x_i) \leq |X| \leq k$, there exist *distinct* vertices $u_1, \dots, u_r \in V(G') \setminus (X \cup \bigcup_{j' \leq i} X'_{j'})$ such that for each $j \in [r]$ we have $u_j \in L_{S_j \cup T_j}$, which implies that $N_G(u_j) \supseteq S_j \cup T_j$. We set $X'_i = \{u_1, \dots, u_r\}$. To see that this satisfies all four conditions mentioned above, observe that we get the first by construction. The second follows from the fact that for $j \in [r - 1]$, both u_j and u_{j+1} are adjacent to $T_j \subseteq H$. Since the sets S_j partition $S = N_F(x_i)$, we satisfy the third condition. The fourth follows from the given bound on r . This completes the construction of the sets X'_i .

Using these sets we complete the proof. Let $X' = (X \setminus \{x_1, \dots, x_\ell\}) \cup \bigcup_{i \in [\ell]} X'_i$. By (1) and the fourth condition, we can infer that $|X'| \leq (|X| - \ell) + (\frac{p|X|}{\lceil p/\varepsilon \rceil} + \ell) \leq (1 + \varepsilon)|X|$. Hence X' is a vertex set in G' of the appropriate size. It remains to analyze its connectivity. We split the proof into two cases here, based on whether we are considering vertex or edge connectivity. Note that $G[X'] = G'[X']$ since G' is an induced subgraph of G with $X' \subseteq V(G')$.

Edge connectivity In the case of edge connectivity, the subgraph F of $G[X]$ we chose above is p -edge-connected. We will argue that $G[X']$ is p -edge-connected. Assume for a contradiction that $G[X']$ has a cut (A', B') of less than p edges. For each set X'_i inserted into X' , the members of X'_i can be ordered so that successive vertices have at least p common neighbors in H . As H belongs to X and $V(G')$ and therefore to X' , successive vertices of X'_i have at least p common neighbors in $G[X']$, and therefore belong to the same side of any cut of less than p edges. Hence for each set X'_i inserted into X' , we have $X'_i \subseteq A'$ or $X'_i \subseteq B'$. This allows us to transform (A', B') into a cut (A, B) of the subgraph F on vertex set X in the natural way, by replacing each set X'_i by the corresponding vertex x_i . The key observation is now that this transformation does not increase the number of edges in the cut: for each edge in the cut (A, B) , either it is an edge between two vertices of $X' \cap X$ (and therefore also an edge of the cut (A', B')), or it is an edge incident on some vertex $x_i \in I$ whose other endpoint v therefore belongs to H . But then the vertex set X'_i is in the same side of the cut in (A', B') and contains a vertex adjacent to v , as $N_F(x_i) \subseteq \bigcup_{u \in X'_i} N_G(u)$. Hence the size of cut (A, B) of F is not larger than the cut (A', B') of $G[X']$, which contradicts that F is p -edge-connected. Hence $G[X']$ is p -edge-connected.

Vertex connectivity In the case of vertex connectivity, the subgraph F of $G[X]$ is p -connected. We argue that $G[X']$ is also p -connected. Since $|X'| \geq |X| \geq p + 1$, it suffices to verify that $G[X']$ cannot be disconnected by removing less than p vertices. Consider a vertex set Z' of size less than p ; we will argue that $G[X'] - Z'$ is connected. Let \mathcal{I} consist of those indices $i \in [\ell]$ such that $X'_i \cap Z' \neq \emptyset$, and let $Z = (Z' \setminus \bigcup_{i \in \mathcal{I}} X'_i) \cup \{x_i \mid i \in \mathcal{I}\}$ be obtained by replacing each set X'_i intersecting Z' by the single vertex x_i . Since the sets X'_i are pairwise disjoint by the first condition, $|Z| \leq |Z'| < p$ and therefore $F - Z$ is connected. Let $X'' = X' \setminus \bigcup_{i \in \mathcal{I}} X'_i$. We shall first prove that $G[X''] - Z'$ is connected, and later show how this implies connectivity of $G[X'] - Z'$ itself.

Assume for a contradiction that $G[X''] - Z'$ is not connected. Consider a pair of vertices u, v that belong to different connected components of $G[X''] - Z'$. Each vertex of each set X'_i has at least p neighbors in H : if $|X'_i| > 1$ this follows from the second condition on the subsets, while for $|X'_i| = 1$ the third condition implies the single vertex in X'_i has at least $\deg_F(x_i)$ neighbors in H , while $\deg_F(x_i) \geq p$ due to p -connectivity of F . Hence any connected component of $G[X''] - Z'$ that contains a vertex of a set X'_i for $i \in [\ell]$, also contains a vertex of H . It follows that there are two vertices u, v of $G[X''] - Z'$ belonging to different connected components and $u, v \notin \bigcup_{i \in [\ell]} X'_i$, so that $u, v \in X \cap X'$. Since the process of turning Z' into Z only affected vertices outside $X \cap X'$, vertices u and v exist in $F - Z$ and are connected by a path P in $F - Z$ since F is p -connected. We transform P into a path connecting u and v in $G[X''] - Z'$, as follows. For each occurrence of a vertex $x_i \in X \setminus X'$ on path P , we know $x_i \notin Z'$ so $X'_i \cap Z' = \emptyset$. Each pair of successive vertices from X'_i has p common neighbors in $H \subseteq X''$, of which at most $p - 1$ belong to Z' , so each pair of successive vertices from X'_i is connected in $G[X''] - Z'$; hence all vertices of X'_i belong to the same connected component of $G[X''] - Z'$. By the third condition on the subsets X'_i , some vertex of X'_i is adjacent to the predecessor of x_i on P , and some vertex of X'_i is adjacent to the successor of x_i on P . Hence each occurrence of a vertex $x_i \in X \setminus X'$ on path P can be replaced by a path through $G[X''] - Z'$. This transforms P into a path P' connecting u and v in $G[X''] - Z'$; a contradiction. Hence $G[X''] - Z'$ is connected.

From the fact that $G[X''] - Z'$ is connected, we derive that $G[X'] - Z'$ is connected as follows. Each vertex of $X' \setminus X''$ belongs to some set X'_i for $i \in \mathcal{I}$. As observed above, each vertex of X'_i has at least p neighbors in $H \subseteq X''$, of which at most $p - 1$ belong to Z' ; hence each vertex of $X' \setminus X''$ is adjacent to a vertex of $G[X''] - Z'$. Therefore $G[X'] - Z'$ can be obtained from $G[X''] - Z'$ by inserting non-isolated vertices, which leaves the graph connected. This completes the proof that $G[X']$ is p -connected.

As the above two cases show that $G[X'] = G'[X']$ is p -vertex/edge-connected, while we already derived $|X'| \leq (1 + \varepsilon)|X|$, this proves that X' is a p -vertex/edge-connected vertex cover of G of the appropriate size. As $X' \subseteq V(G')$ and G' is an induced subgraph of G , it is also a valid solution in G' , which completes the proof. \square

Using the previous lemma we now prove the existence of approximate kernels for the two considered problems.

Theorem 2. For every $\varepsilon > 0$ and every fixed $p \geq 2$, the problems p -EDGE-CONNECTED VERTEX COVER and p -CONNECTED VERTEX COVER admit a $(1 + \varepsilon)$ -approximate kernel with $k + 2k^2 + \lfloor (3 + \varepsilon)k \rfloor k^{2\lceil p/\min(\varepsilon, 1) \rceil}$ vertices.

Proof. Fix $p \geq 2$. Before presenting the main argument, we show that to prove the theorem it suffices to prove it for $0 < \varepsilon \leq 1$. Let $s(k, \varepsilon) = k + 2k^2 + \lfloor (3 + \varepsilon)k \rfloor k^{2\lceil p/\min(\varepsilon, 1) \rceil}$ denote the guarantee on the number of vertices in a reduced instance claimed by the theorem for a certain value of ε and parameter value k . For any $k \in \mathbb{N}$ and $\varepsilon > 1$ we have $s(k, 1) \leq s(k, \varepsilon)$; we rely here on the $\min(\varepsilon, 1)$ term in the exponent, which means that the exponent stops becoming smaller when ε becomes larger than 1. Suppose that the theorem holds for $\varepsilon = 1$, which means that for each type of connectivity considered there is a polynomial-time reduction algorithm \mathcal{R}_1 reducing any instance (G, k) to an instance (G', k') on at most $s(k, 1)$ vertices, and a polynomial-time solution lifting algorithm \mathcal{L}_1 that can lift α -approximate solutions for (G', k') to $(\alpha \cdot (1 + \varepsilon)) = (\alpha \cdot (1 + 1))$ -approximate solutions in (G, k) . The algorithms \mathcal{R}_1 and \mathcal{L}_1 also form a valid $(1 + \varepsilon)$ -approximate kernel for any $\varepsilon > 1$: as just argued, the output of \mathcal{R}_1 has at most $s(k, 1) \leq s(k, \varepsilon)$ vertices, while the solution lifting algorithm produces a solution whose approximation factor is $(\alpha \cdot (1 + 1)) \leq (\alpha \cdot (1 + \varepsilon))$. Hence to prove the theorem it suffices to prove it for $0 < \varepsilon \leq 1$.

Consider $0 < \varepsilon \leq 1$. The approximate kernelization algorithm has two parts. The first part is a reduction algorithm, and the second part is a solution lifting algorithm. Let (G, k) be an input instance of p -EDGE-CONNECTED VERTEX COVER or p -CONNECTED VERTEX COVER.

Reduction algorithm First we invoke Lemma 2 to check whether G has a p -vertex/edge-connected vertex cover. If not, then the reduction algorithm outputs the instance $(2K_2, 1)$, that is, a matching of two edges with a parameter value of one.

If $|V(G)| \leq k^{2\lceil p/\varepsilon \rceil}$, then the instance is already small in terms of the parameter. To ensure the running time of the reduction algorithm is bounded by a polynomial whose degree does not depend on ε , in this case we simply output (G, k) unchanged. If $|V(G)|$ is larger, we run $\text{Mark}(G, k, \varepsilon)$. If it outputs **INFEASIBLE**, we output $(K_{2p}, 1)$, that is, a clique of size $2p$ with a parameter value of one. If Mark outputs an instance $(G' = G[H \cup R \cup L], k' = (1 + \varepsilon)k)$, we use (G', k') as the output of the reduction algorithm.

As the for-loop of Mark only happens when $|H| \leq k$, the running time of the algorithm can be bounded as $n^{\mathcal{O}(1)} \cdot k^{2\lceil p/\varepsilon \rceil}$ as it spends $n^{\mathcal{O}(1)}$ time for each subset of H of size at most $2\lceil p/\varepsilon \rceil$ while $|H| \leq k$. By our assumption on $|V(G)|$, we have $k^{2\lceil p/\varepsilon \rceil} \leq |V(G)|$ so that the running time is $n^{\mathcal{O}(1)}$ for some absolute constant not depending on p, k or ε . Hence the approximate kernelization scheme is time-efficient. If the output is not equal to the result of Mark , its size is trivially bounded as required. The output of Mark is $G[H \cup R \cup L]$, where $|H| \leq k$ and $|R| \leq 2k^2$ follow from the definition of the algorithm, while $|L| \leq \lfloor (3 + \varepsilon)k \rfloor k^{2\lceil p/\varepsilon \rceil}$ as observed below its presentation. Hence the number of vertices in the output graph is as claimed. Since p is a fixed constant, this is suitable for a polynomial-sized approximate kernelization scheme.

Solution lifting algorithm Given a solution $S' \subseteq V(G')$ for the instance (G', k') , the solution lifting algorithm proceeds as follows. If G does not have a p -edge/vertex-connected vertex cover, it outputs \emptyset as the solution. If S' is not a valid solution in G' , or S' does not contain all vertices of H , then we output the trivial p -edge/vertex-connected vertex cover of G found via Lemma 2 as the solution to (G, k) . Otherwise, we output S' as the solution S for (G, k) ; we argue below that it is a valid solution.

It remains to argue that the output S of the solution lifting algorithm is of sufficient quality. Formally, we need to establish that:

$$\frac{pCVC(G, k, S)}{\text{OPT}(G, k)} \leq (1 + \varepsilon) \frac{pCVC(G', k', S')}{\text{OPT}(G', k')} \tag{2}$$

Depending on the type of connectivity considered, $pCVC$ corresponds to either $pECVC$ or $pVCVC$, the functions defined below Definition 4 that map solutions of the considered parameterized optimization problems to their cost value. If G does not have any p -vertex/edge-connected vertex cover, then the value of each solution is $+\infty$ so each solution is optimal, which implies that (2) holds since the left-hand side becomes 1 and the right-hand side is never smaller. Similarly, if G has a p -vertex/edge-connected vertex cover, but not of size at most k , then $\text{OPT}(G, k) = k + 1$ and by definition of the function $pCVC$, each solution has cost at most $k + 1$ and is therefore optimal. Hence it remains to consider the case that G has a p -vertex/edge-connected vertex cover of size at most k .

By Lemma 6, the fact that G has a p -vertex/edge-connected vertex cover of size at most k implies that the reduction algorithm outputs a nontrivial graph G' which has a p -vertex/edge-connected vertex cover of size at most $(1 + \varepsilon)|X|$. This implies that $\text{OPT}(G, k) = |X|$ and $\text{OPT}(G', k') \leq (1 + \varepsilon)|X|$, so that $\text{OPT}(G', k') \leq (1 + \varepsilon)\text{OPT}(G, k)$. To analyze the result of the solution lifting algorithm, we consider two cases depending on the structure of the solution S' given to the algorithm.

- Suppose S' is a p -vertex/edge-connected vertex cover of G' of size at most k . We argue that S' contains all vertices of H : each vertex of H has degree more than k in G , and for each vertex $v \in I$ that is not marked by the algorithm and therefore no longer occurs in G' , we marked $\lfloor (3 + \varepsilon)k \rfloor > k$ vertices for each neighbor of v in H . Hence each vertex of H also has degree more than k in G' , which means it is contained in each vertex cover of size at most k . Hence $H \subseteq S'$. Each vertex of $V(G) \setminus V(G')$ belongs to the independent set I and has all its neighbors in H . Hence $S' \supseteq H$ covers all edges incident on vertices of $V(G) \setminus V(G')$. As $G[S'] = G'[S']$, this implies that S' is a p -vertex/edge-connected vertex cover of G , which is a valid output for the solution lifting algorithm. This satisfies Equation (2) since $\text{OPT}(G', k') \leq (1 + \varepsilon)\text{OPT}(G, k)$.
- Now suppose S' is a p -vertex/edge-connected vertex cover of G' of size more than k , which implies $p\text{CVC}(G', k', S') \geq k + 1$. Since the solution lifting algorithm outputs a valid p -vertex/edge-connected vertex cover of G whenever there is one, we have $p\text{CVC}(G, k, S) \leq \min(|S|, k + 1) \leq k + 1$. Now we derive:

$$\frac{p\text{CVC}(G, k, S)}{\text{OPT}(G, k)} \leq \frac{k + 1}{\text{OPT}(G, k)} \leq (1 + \varepsilon) \frac{k + 1}{\text{OPT}(G', k')} \leq (1 + \varepsilon) \frac{p\text{CVC}(G', k', S')}{\text{OPT}(G', k')}$$

where the middle inequality follows from $\text{OPT}(G', k') \leq (1 + \varepsilon)\text{OPT}(G, k)$.

This concludes the proof. \square

5. Constant factor approximation algorithm for p -EDGE-CONNECTED VERTEX COVER

In this section, we describe a $2(p + 1)$ -approximation algorithm for p -EDGE-CONNECTED VERTEX COVER. We begin by defining the notion of a Gomory-Hu tree.

Definition 12 (Gomory-Hu tree). Let $G = (V, E)$ be a graph, and let $c(u, v) \geq 0$ be the capacity of edge $uv \in E$, letting $c(u, v) = 0$ if $uv \notin E$. Denote the minimum capacity of an $(s-t)$ -cut by λ_{st} for each $s, t \in V(G)$. Let $T = (V_T, E_T)$ be a tree with $V_T = V(G)$, and let us denote the set of edges in the $s-t$ path in T by P_{st} for each $s, t \in V_T$. Then T is said to be a Gomory-Hu tree of G if $\lambda_{st} = \min_{e \in P_{st}} c(S_e, T_e)$ for all $s, t \in V(G)$, where

- S_e and T_e are sets of vertices of the two connected components of $T - e$ such that $s \in S_e$ and $t \in T_e$, and
- $c(S_e, T_e) = \sum_{u \in S_e} \sum_{v \in T_e} c(u, v)$ is the capacity of the cut in G .

The capacity of an edge uv of T is equal to λ_{uv} .

Theorem 6. [16] Every weighted graph (G, c) has a Gomory-Hu tree which can be constructed in polynomial time.

For an unweighted graph $G = (V, E)$, we can introduce weights by setting $c(u, v) = 1$ for every $uv \in E$. Let T be a Gomory-Hu tree of G . Then, by Definition 12, for every pair of vertices $u, v \in V(G)$, the size of a minimum edge cut between u and v in G is the minimum capacity of an edge cut between u and v in T . For $i \in [p]$, consider the set E_i of all the edges of total capacity at most $i - 1$ in T . Deleting E_i disconnects T into several subtrees. We call the vertex set of each such subtree an i -segment in G . Thus, a subset of vertices $S \subseteq V(G)$ is an i -segment in G if and only if for every $u, v \in S$, there are at least i edge-disjoint paths between u and v in G and S is a maximal such subset. It is obvious from the construction that the i -segments of G form a partition of the vertex set of G , and can be computed in polynomial time.

Now let $G = (V, E)$ be an undirected graph, and $X \subseteq V(G)$. For $i \in [p]$, let an i -block of X in G be a maximal subset $X' \subseteq X$ such that for every $u, v \in X'$, there are at least i edge-disjoint paths between u and v in G . We can use the Gomory-Hu tree to compute the i -blocks of X in G , as follows.

Lemma 7. Let $G = (V, E)$ be an undirected graph, and $X \subseteq V(G)$. The i -blocks of X in G are precisely the sets $X \cap S$ over all i -segments S of G .

Proof. Let T be the Gomory-Hu tree of G , and let $u, v \in X$ be distinct vertices. By the definition of a Gomory-Hu tree, $\lambda_{uv}(G) \geq i$ if and only if there is no edge e on the path from u to v in T with capacity $c(e)$ less than i . Since this is also equivalent to u and v being in the same i -segment of G , the statement follows. \square

Based on Lemma 7, we can compute the collection of i -blocks in polynomial time using a Gomory-Hu tree of G .

Lemma 8. Let G be a graph, $X \subseteq V(G)$, and p be a fixed integer. Then the collection of all i -blocks of X for all $i \in [p]$ forms a laminar family.

Algorithm 1: LaminarTree($G = (V, E), X, p$).

```

input :  $G = (V, E), X \subseteq V(G), p$ 
output: A laminar tree of  $X$  in  $G$ 
Initialize a tree  $T := \emptyset$ ;
for  $i = 1, \dots, p$  do
    Compute the set of all  $i$ -blocks of  $X$  in  $G$ ;
     $\mathcal{A}_i \leftarrow$  the set of all  $i$ -blocks of  $X$  in  $G$ ;
end
Set  $X$  as the root of  $T$ ;
Make all 1-blocks the children of  $X$ ;
for  $i = 2, \dots, p$  do
    for every  $Y \in \mathcal{A}_i$  do
        Make  $Y$  a child of  $W \in \mathcal{A}_{i-1}$  in  $T$  such that  $Y \subseteq W$ ;
    end
end
Output  $T$  as the laminar tree of  $X$  in  $G$ ;

```

Algorithm 2: Approximation algorithm for p -Edge-CVC.

```

input :  $G = (V, E)$ 
output: An approximate  $p$ -edge-connected vertex cover of  $G$ 
 $L \leftarrow \{v \in V(G) \mid \deg_G(v) < p\}$ ;
if  $V(G) \setminus L$  is not a  $p$ -edge-connected vertex cover of  $G$  then
    Output "No feasible solution";
end
Compute a maximal matching  $M$  of  $G - N_G[L]$ ;
 $X \leftarrow N_G(L) \cup V(M)$ ;
 $T \leftarrow$  LaminarTree( $G[X], X, p$ );
 $Y \leftarrow X$ ;
while  $T$  has at least two leaves do
    Let  $u \in V(G) \setminus (Y \cup L)$  s.t.  $N_G(u)$  intersects two distinct  $p$ -blocks in  $T$ ;
     $Y \leftarrow Y \cup \{u\}$ ;
     $T \leftarrow$  LaminarTree( $G[Y], X, p$ );
end
Output  $Y$  as a solution;

```

Proof. Note first that for every i , the i -segments of G form a partition of the vertex set, hence similarly, the i -blocks of X form a partition of X for every $i \in [p]$. Furthermore, let $i, j \in [p]$ with $1 \leq i < j \leq p$. It is obvious from the definition that the j -segments of G form a refinement of the i -segments of G , since they are formed from the Gomory-Hu tree by deleting an additional set of edges. Hence the j -blocks of X also form a refinement of the i -blocks of X in G by Lemma 7, and the statement follows. \square

The approximation algorithm exploits the Gomory-Hu tree, which can be used to derive the p -blocks which capture the edge-connectivity of the graph. The proof of Lemma 8 leads to Algorithm 1 which sets all i -blocks of X , $i \in [p]$, as nodes of a (laminar) tree. Algorithm 1 works as follows. For all $i \in [p]$, it computes all the i -blocks of X in G . Then, it sets X as the root of the tree, and makes all 1-blocks children of X . After that, for every 1-block X and 2-block Y , it sets Y as a child of X if $Y \subseteq X$. It repeats this process for 2-blocks, 3-blocks, up to and including p -blocks (in this order).

Lemma 9. Let G be a graph, $u \in V(G)$ and let $A, B \subseteq V(G - u)$ such that $A \cap B = \emptyset$. If the size of a minimum (A, B) -cut in $G - u$ is i and $\min\{|N(u) \cap A|, |N(u) \cap B|\} = j$ then the size of a minimum (A, B) -cut in G is at least $i + j$.

Proof. Consider a minimum (A, B) -cut (S, T) in G ; $A \subseteq S$ and $B \subseteq T$. Assume without loss of generality that $u \in S$. Suppose that the size of (S, T) is at most $i + j - 1$. Consider $(S \setminus \{u\}, T)$, which is a cut in $G - u$. Since $|N(u) \cap B| \geq j$, we have that the size of the cut $(S \setminus \{u\}, T)$ is at most $i - 1$, but this contradicts the fact that the size of a minimum (A, B) -cut in $G - u$ is i . \square

Our approximation algorithm, Algorithm 2, proceeds as follows.

- Let L denote the set of vertices with degree less than p . By Lemma 2, G has a p -edge-connected vertex cover if and only if $V(G) \setminus L$ is such a vertex cover. Hence if it is not, we output that 'G has no feasible solution'.
- Compute a maximal matching M of $G - N_G[L]$ and initialize $X = N_G(L) \cup V(M)$ (the vertices matched by M) and initialize $Y = X$ as the partial solution.
- Invoke Algorithm 1 to construct the laminar tree T of X in $G[X]$.

- As long as T has at least two leaves, the fact that G has a p -edge-connected vertex cover implies that there is a vertex $u \in V(G) \setminus (Y \cup L)$ whose neighborhood intersects two distinct p -blocks in T . We add u into the solution Y and recompute the laminar tree T of X in $G[Y]$.
- Output Y as the solution.

Our ‘recompute laminar tree of X in $G[Y]$ ’ ensures that at the end, every pair of vertices in X can be connected by p edge-disjoint paths. Finally, based on the other characteristics, we are ensured that $G[Y]$ actually becomes p -edge-connected (proof as part of Theorem 2 proof).

We are ready to prove the main result of this section, Theorem 3, which we restate for completeness.

Theorem 3. For every fixed $p \geq 2$, p -EDGE-CONNECTED VERTEX COVER admits a polynomial-time $2(p + 1)$ -factor approximation algorithm.

Proof. We will show that Algorithm 2 is a $2(p + 1)$ -approximation algorithm.

Observe that the first output “No feasible solution” is correct due to Lemma 2. Note $X = N_G(L) \cup V(M)$ is a vertex cover of G . Suppose that T has more than one leaf and there is no vertex $u \in V(G) \setminus (Y \cup L)$ such that $N_G(u)$ intersects two distinct p -blocks in T . Then even adding all vertices of $V(G) \setminus (Y \cup L)$ to Y will not make Y into a p -edge-connected vertex cover of G since there will be the same number of p -blocks of X in $G[Y]$ before and after the addition because the vertices of $V(G) \setminus (Y \cup L)$ are not vertices of M and thus form an independent set. However, this is impossible as $V(G) \setminus L$ is a p -edge-connected vertex cover of G . Thus, as long as T has more than one leaf there is a vertex $u \in V(G) \setminus (Y \cup L)$ such that $N_G(u)$ intersects two distinct p -blocks in T .

When T has just one leaf, G has only one p -block of X in $G[Y]$. This means that for every pair x, y of vertices in X there are p edge-disjoint paths in $G[Y]$ between x and y . Let $u, v \in Y \setminus X$. Since $|N_G(u)| \geq p$ and $|N_G(v)| \geq p$ by Menger’s theorem, there are p edge-disjoint paths in $G[Y]$ between $N_G(u)$ and $N_G(v)$ with distinct end-vertices and hence p edge-disjoint paths in $G[Y]$ between u and v . Similarly, we can see that there are p edge-disjoint paths in $G[Y]$ between u and any $x \in X$. Therefore, Y is a p -edge-connected vertex cover of G .

Let us analyze how T changes after $u \in V(G) \setminus (Y \cup L)$ is added to Y . First we consider T before u is added to Y . By the description of Algorithm 2, u has neighbors in two distinct p -blocks X_1, X_2 of X in $G[Y]$. Let X_0 be the least common ancestor of X_1 and X_2 in T and let X_0 be an i -block of X in $G[Y]$. Observe that X_0 has two children X_a and X_b such that $X_1 \subseteq X_a$ and $X_2 \subseteq X_b$. Note that X_a and X_b are $(i + 1)$ -blocks of X in $G[Y]$ and the minimum size of a (X_a, X_b) -cut is i (otherwise, X_0 is not an i -block of X in $G[Y]$). Now consider what happens just after u is added to Y . By Lemma 9, the size of any (X_a, X_b) -cut increases by at least one. Thus, the minimum size of a (X_a, X_b) -cut becomes $i + 1$ and so X_a and X_b become part of a new $(i + 1)$ -block of X in $G[Y]$. Thus, the number of the nodes of T decreases.

Let us now bound the approximation factor of Algorithm 2. The number of leaves in T becomes one only when there is just one node on each level of T , i.e., T has $p + 1$ vertices. Initially, T may have up to $p|X| + 1$ nodes. Thus, at most $p(|X| - 1)$ nodes will be added to X before a solution Y is obtained. Hence, $|Y| \leq |X| + p(|X| - 1) \leq (p + 1)|X|$. Since M is a maximal matching of $G - N_G(L)$, at least one endpoint of each edge of M has to be in any vertex cover of G . Thus, $\text{OPT}(G) \geq |N_G(L)| + |M|$, where $\text{OPT}(G)$ is the minimum number of vertices in a p -edge-connected vertex cover of G . Since $|X| = 2|M| + |N_G(L)|$, $|X| \leq 2\text{OPT}(G)$. Therefore, $|Y| \leq (p + 1)|X| \leq 2(p + 1)\text{OPT}(G)$. \square

6. Hardness proofs

We will first show that p -CONNECTED VERTEX COVER admits no polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$ using the fact that CONNECTED VERTEX COVER admits no polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$ [9].

Theorem 7. For every fixed $p \geq 1$, p -CONNECTED VERTEX COVER is NP-hard and does not admit a polynomial kernel parameterized by the solution size k unless $\text{NP} \subseteq \text{coNP/poly}$.

Proof. Let (G, k) be an instance of CONNECTED VERTEX COVER. We construct an instance of p -CONNECTED VERTEX COVER as follows. For every $i \in [p - 1]$, add vertices v_i, v'_i to G , with edges $\{v_i v'_i \mid i \in [p - 1]\}$ and $\{uv_i \mid i \in [p - 1], u \in V(G)\}$, obtaining a new graph G' . Observe that G' can be computed in polynomial time. Furthermore, (G, k) is a yes-instance of CONNECTED VERTEX COVER if and only if $(G', k + p - 1)$ is a yes-instance of p -CONNECTED VERTEX COVER.

This reduction runs in polynomial time and p is a constant. Hence, it satisfies the conditions of Definition 3. Thus, it is a polynomial parameter transformation from CONNECTED VERTEX COVER to p -CONNECTED VERTEX COVER. As CONNECTED VERTEX COVER does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$, due to Proposition 2, neither does p -CONNECTED VERTEX COVER unless $\text{NP} \subseteq \text{coNP/poly}$. This completes the proof of the kernel lower bound. Since CONNECTED VERTEX COVER is NP-hard, this also proves NP-hardness of the problem. \square

The reduction in the next theorem is from RED BLUE DOMINATING SET. In the problem, given a bipartite graph G with partite sets R and B , we are to decide whether there is $R' \subseteq R$ such that $|R'| \leq k$ and $N_G(R') = B$. Dom et al. [9] proved that RED BLUE DOMINATING SET parameterized by $k + |B|$ does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$.

Theorem 8. For every fixed $p \geq 1$, p -EDGE-CONNECTED VERTEX COVER is NP-hard and does not admit a polynomial kernel parameterized by the solution size k unless $\text{NP} \subseteq \text{coNP/poly}$.

Proof. Since p -edge- and p -vertex-connectivity are equivalent for $p = 1$, the lower bound for $p = 1$ follows from Theorem 7. In the remainder we consider an arbitrary $p \geq 2$. Let $(G = (R \uplus B, E), k + |B|)$ be an instance of RED BLUE DOMINATING SET and let $B = \{b^1, b^2, \dots, b^t\}$. Without loss of generality, we may assume that for every $v \in R$, there exists $u \in B$ such that $uv \in E(G)$ (otherwise we can just delete v). Similarly, we may assume that for every $v \in B$, there exists $u \in R$ such that $uv \in E(G)$ (otherwise, there is no feasible solution). We will also assume that $t \geq p$ and $k \geq p$, as the instance can otherwise be solved in polynomial time. Construct a new graph H from G as follows.

- Add a complete graph K_p with vertex set $A = \{a_1, \dots, a_p\}$ such that $V(G) \cap A = \emptyset$, and edges $a_i r$ for every $i \in [p]$ and $r \in R$.
- Replace every vertex b^j of B by a complete graph K_p with vertices $\{b^j_1, \dots, b^j_p\}$ such that if $rb^j \in E(G)$ then $rb^j_i \in E(H)$ for every $i \in [p]$. Thus, B is replaced by \hat{B} of size pt .
- Attach a pendant vertex to every vertex in $\hat{B} \cup A$.
- Set $k' = k + p(t + 1)$. Note that since p is a fixed constant, we have that k' is $\mathcal{O}(k + t)$.

To complete the proof, it suffices to prove that $(G, k + t)$ is a yes-instance of RED BLUE DOMINATING SET if and only if (H, k') is a yes-instance of p -EDGE-CONNECTED VERTEX COVER.

(\Leftarrow) Let S be a p -edge-connected vertex cover of H such that $|S| \leq k'$. Observe that $\hat{B} \cup A \subseteq S$ since all of the vertices in $\hat{B} \cup A$ have pendant neighbors. Since every vertex in \hat{B} has just $p - 1$ neighbors in \hat{B} and $G[S]$ is p -edge-connected, S must contain at least one neighbor in R of every vertex of \hat{B} . Since $|\hat{B} \cup A| = p(t + 1)$, $k' = |\hat{B} \cup A| + k$ and so there must be a subset R' of R of size at most k such that $N_G(R') = B$. Hence, (G, k) is a yes-instance of RED BLUE DOMINATING SET.

(\Rightarrow) Let $R^* \subseteq R$ with $|R^*| = k$ such that $N_G(R^*) = B$. Let $R^* = \{r_1, \dots, r_k\}$, recall that $k \geq p$. We claim that $A \cup R^* \cup \hat{B}$ is a p -edge-connected vertex cover of H . Clearly, it is a vertex cover, so it remains to prove that $H[A \cup R^* \cup \hat{B}]$ is p -edge-connected. Observe that by construction of H , $H[A \cup R^*]$ contains a spanning subgraph $K_{p,k}$, where $k \geq p$. Hence, $H[A \cup R^*]$ is p -edge connected. Note that adding a clique of size p and making it adjacent to one vertex in a p -edge-connected graph keeps it p -edge connected. This implies that by construction of H , $H[A \cup R^* \cup \hat{B}]$ is p -connected.

This reduction runs in polynomial time. As p is a constant, this reduction provides a polynomial parameter transformation. As RED BLUE DOMINATING SET parameterized by $k + |B|$ does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$, by Proposition 2 p -EDGE-CONNECTED VERTEX COVER does not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP/poly}$. As RED BLUE DOMINATING SET is well-known to be NP-hard (cf. [9, §4.1]), the NP-hardness result follows as well. \square

7. Conclusions

We presented time efficient polynomial sized approximate kernelization schemes (PSAKS) for both p -EDGE-CONNECTED VERTEX COVER and p -CONNECTED VERTEX COVER. We also gave a $\mathcal{O}^*(2^{\mathcal{O}(pk)})$ time algorithm for p -EDGE-CONNECTED VERTEX COVER. The approach we use in this FPT algorithm does not work for p -vertex-connectivity. Hence, an interesting open problem would be to determine whether there exists a singly exponential FPT algorithm for p -CONNECTED VERTEX COVER.

We also obtained a polynomial-time $2(p + 1)$ -factor approximation algorithm for p -EDGE-CONNECTED VERTEX COVER. Again, the main idea used in this algorithm does not work for p -CONNECTED VERTEX COVER, which appears to be the harder variant of the two. We mention without proof that it is possible to use the decomposition into biconnected components to derive a constant-factor approximation algorithm for BICONNECTED VERTEX COVER ($p = 2$) using an approach similar to p -EDGE-CONNECTED VERTEX COVER. However, it is currently unknown how to get such a result for p -CONNECTED VERTEX COVER for any arbitrary fixed $p \geq 3$. Finally, another interesting open problem would be to obtain a size efficient PSAKS for p -EDGE-CONNECTED VERTEX COVER and for p -CONNECTED VERTEX COVER, but this is also open for CONNECTED VERTEX COVER.

CRedit authorship contribution statement

The authors provided an approximately the same contribution to obtain the results, writing the original version of the paper, and revising it.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgment

We are grateful to the reviewers for numerous useful suggestions, which led to a significant improvement over our initial presentation.

References

- [1] Ankit Abhinav, Susobhan Bandopadhyay, Aritra Banik, Saket Saurabh, Parameterized algorithms for finding highly connected solution, in: *Computer Science – Theory and Applications*. CSR, in: *Lecture Notes in Computer Science (LNCS)*, vol. 13296, 2022, pp. 1–16.
- [2] Akanksha Agrawal, Pranabendu Misra, Fahad Panolan, Saket Saurabh, Fast exact algorithms for survivable network design with uniform requirements, in: *WADS*, in: *Lecture Notes in Computer Science*, vol. 10389, Springer, 2017, pp. 25–36.
- [3] Akanksha Agrawal, Saket Saurabh, Prafullkumar Tale, On the parameterized complexity of contraction to generalization of trees, *Theory Comput. Syst.* 63 (3) (2019) 587–614.
- [4] Sayan Bandyapadhyay, Fedor V. Fomin, Petr A. Golovach, Nidhi Purohit, Kirill Simonov, Lossy kernelization of same-size clustering, *CoRR*, arXiv:2107.07383 [abs], 2021.
- [5] Hans L. Bodlaender, Stéphan Thomassé, Anders Yeo, Kernel bounds for disjoint cycles and disjoint paths, *Theor. Comput. Sci.* 412 (35) (2011) 4570–4578.
- [6] Jianer Chen, Iyad A. Kanj, Weijia Jia, Vertex cover: further observations and further improvements, *J. Algorithms* 41 (2) (2001) 280–301.
- [7] Marek Cygan, Deterministic parameterized connected vertex cover, in: *Proceedings of SWAT 2012*, 2012, pp. 95–106.
- [8] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, Saket Saurabh, *Parameterized Algorithms*, Springer, 2015.
- [9] Michael Dom, Daniel Lokshtanov, Saket Saurabh, Kernelization lower bounds through colors and IDs, *ACM Trans. Algorithms* 11 (2) (2014) 13:1–13:20.
- [10] Rodney G. Downey, Michael R. Fellows, *Fundamentals of Parameterized Complexity*, *Texts in Computer Science*, Springer, 2013.
- [11] Pavel Dvorák, Andreas Emil Feldmann, Dusan Knop, Tomáš Masarík, Tomáš Toufar, Pavel Veselý, Parameterized approximation schemes for Steiner trees with small number of Steiner vertices, *SIAM J. Discrete Math.* 35 (1) (2021) 546–574.
- [12] Eduard Eiben, Danny Hermelin, M.S. Ramanujan, On approximate preprocessing for domination and hitting subgraphs with connected deletion sets, *J. Comput. Syst. Sci.* 105 (2019) 158–170.
- [13] Eduard Eiben, Mithilesh Kumar, Amer E. Mouawad, Fahad Panolan, Sebastian Siebertz, Lossy kernels for connected dominating set on sparse graphs, *SIAM J. Discrete Math.* 33 (3) (2019) 1743–1771.
- [14] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, Efficient computation of representative families with applications in parameterized and exact algorithms, *J. ACM* 63 (4) (2016) 29:1–29:60.
- [15] Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Meirav Zehavi, *Kernelization: Theory of Parameterized Preprocessing*, Cambridge University Press, 2019.
- [16] R.E. Gomory, T.C. Hu, Multi-terminal network flows, *J. Soc. Ind. Appl. Math.* 9 (1961).
- [17] Spoorthy Gunda, Pallavi Jain, Daniel Lokshtanov, Saket Saurabh, Prafullkumar Tale, On the parameterized approximability of contraction to classes of chordal graphs, *ACM Trans. Comput. Theory* 13 (4) (2021) 27:1–27:40.
- [18] Bart M.P. Jansen, Michal Włodarczyk, Lossy planarization: a constant-factor approximate kernelization for planar vertex deletion, *CoRR*, arXiv:2202.02174 [abs], 2022.
- [19] R. Krithika, Diptapriyo Majumdar, Venkatesh Raman, Revisiting connected vertex cover: FPT algorithms and lossy kernels, *Theory Comput. Syst.* 62 (8) (2018) 1690–1714.
- [20] Hengzhe Li, Yuxing Yang, Baoyindureng Wu, 2-edge connected dominating sets and 2-connected dominating sets of a graph, *J. Comb. Optim.* 31 (2) (2016) 713–724.
- [21] Daniel Lokshtanov, Pranabendu Misra, Fahad Panolan, Saket Saurabh, Deterministic truncation of linear matroids, *ACM Trans. Algorithms* 14 (2) (2018) 14:1–14:20.
- [22] Daniel Lokshtanov, Fahad Panolan, M.S. Ramanujan, Saket Saurabh, Lossy kernelization, in: *Proceedings of STOC 2017*, 2017, pp. 224–237.
- [23] Diptapriyo Majumdar, M.S. Ramanujan, Saket Saurabh, On the approximate compressibility of connected vertex cover, *Algorithmica* 82 (10) (2020) 2902–2926.
- [24] Pasin Manurangsi, A note on max k -vertex cover: faster FPT-AS, smaller approximate kernel and improved approximation, in: Jeremy T. Fineman, Michael Mitzenmacher (Eds.), *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*, Dagstuhl, Germany, in: *OpenAccess Series in Informatics (OASIS)*, vol. 69, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018, pp. 15:1–15:21.
- [25] Dániel Marx, A parameterized view on matroid optimization problems, *Theor. Comput. Sci.* 410 (44) (2009) 4471–4479.
- [26] Daniel Mölle, Stefan Richter, Peter Rossmanith, Enumerate and expand: improved algorithms for connected vertex cover and tree cover, *Theory Comput. Syst.* 43 (2) (2008) 234–253.
- [27] Hiroshi Nagamochi, Toshihide Ibaraki, A linear-time algorithm for finding a sparse k -connected spanning subgraph of a k -connected graph, *Algorithmica* 7 (1992) 583–596.
- [28] George L. Nemhauser, Leslie E. Trotter Jr., Vertex packings: structural properties and algorithms, *Math. Program.* 8 (1) (1975) 232–248.
- [29] Rolf Niedermeier, *Invitation to Fixed-Parameter Algorithms*, Oxford University Press, 2006.
- [30] Takao Nishizeki, Svatopluk Poljak, k -Connectivity and decomposition of graphs into forests, *Discrete Appl. Math.* 55 (3) (1994) 295–301.
- [31] Zeev Nutov, 2-Node-connectivity network design, in: *Proceedings of WAOA 2020*, in: *Lecture Notes in Computer Science*, vol. 12806, Springer, 2020, pp. 220–235.
- [32] Zeev Nutov, A 4-approximation for k -connected subgraphs, *J. Comput. Syst. Sci.* 123 (2022) 64–75.
- [33] Zeev Nutov, Approximating k -connected m -dominating sets, *Algorithmica* 84 (6) (2022) 1511–1525.
- [34] Zeev Nutov, Parameterized algorithms for node connectivity augmentation problems, *CoRR*, arXiv:2209.06695 [abs], 2022.
- [35] James Oxley, *Matroid Theory*, Oxford University Press, 2011.
- [36] M.S. Ramanujan, An approximate kernel for connected feedback vertex set, in: *Proceedings of ESA 2019*, 2019, pp. 77:1–77:14.
- [37] C. Savage, Depth-first search and the vertex cover problem, *Inf. Process. Lett.* 14 (5) (1982) 233–235.
- [38] Douglas B. West, *Introduction to Graph Theory*, 1st edition, Prentice Hall, 1996.