Bridge-Depth Characterizes Which Structural Parameterizations of Vertex Cover Admit a Polynomial Kernel

Marin Bougeret, Bart M P Jansen, Ignasi Sau

Bridge-Depth Characterizes Which Structural Parameterizations of Vertex Cover Admit a Polynomial Kernel

Marin Bougeret
LIRMM, Université de Montpellier, CNRS, France
http://www.lirmm.fr/~bougeret/
marin.bougeret@lirmm.fr

Bart M. P. Jansen
Eindhoven University of Technology, The Netherlands
b.m.p.jansen@tue.nl

Ignasi Sau
LIRMM, Université de Montpellier, CNRS, France
ignasi.sau@lirmm.fr

---- **Abstract** ----

We study the kernelization complexity of structural parameterizations of the VERTEX COVER problem. Here, the goal is to find a polynomial-time preprocessing algorithm that can reduce any instance (G, k) of the VERTEX COVER problem to an equivalent one, whose size is polynomial in the size of a pre-determined complexity parameter of G. A long line of previous research deals with parameterizations based on the number of vertex deletions needed to reduce G to a member of a simple graph class $\mathcal{F}$, such as forests, graphs of bounded tree-depth, and graphs of maximum degree two. We set out to find the most general graph classes $\mathcal{F}$ for which VERTEX COVER parameterized by the vertex-deletion distance of the input graph to $\mathcal{F}$, admits a polynomial kernelization. We give a complete characterization of the minor-closed graph families $\mathcal{F}$ for which such a kernelization exists. We introduce a new graph parameter called *bridge-depth*, and prove that a polynomial kernelization exists if and only if $\mathcal{F}$ has bounded bridge-depth. The proof is based on an interesting connection between bridge-depth and the size of minimal blocking sets in graphs, which are vertex sets whose removal decreases the independence number.

1 Introduction

Background and motivation.   The NP-complete VERTEX COVER problem is one of the most prominent problems in the field of kernelization [3, 8, 13, 16, 26], which investigates provably efficient and effective preprocessing for parameterized problems. A *parameterized problem* is a decision problem in which a positive integer k, called the *parameter*, is associated with every instance x. A *kernelization* for a parameterized problem is a polynomial-time algorithm that reduces any parameterized instance (x, k) to an equivalent instance (x', k') of the same problem whose size is bounded by $f(k)$ for some function f, which is the *size* of the kernelization. Hence a kernelization guarantees that instances which are large compared to their parameter, can be efficiently reduced without changing their answer. Of particular interest are *polynomial kernelizations*, whose size bound f is polynomial.

An instance (G, k) of VERTEX COVER asks whether the undirected graph G has a vertex set S of size at most k that contains at least one endpoint of every edge. Using the classic Nemhauser-Trotter theorem [28], one can reduce (G, k) in polynomial time to an instance (G', k') with the same answer, such that $|V(G')| \leq 2k$. Hence when using the size of the desired solution as the parameter, VERTEX COVER has a kernelization that reduces to instances of $2k$ vertices, which can be encoded in $\mathcal{O}(k^2)$ bits. While the bitsize of this kernelization is known to be essentially optimal [10] assuming the established conjecture NP $\not\subseteq$ coNP/poly, this result does not guarantee any effect of the preprocessing for instances whose solution has size at least $|V(G)|/2$. In particular, it does not promise any size reduction when G is simply a path.

To be able to give better preprocessing guarantees, one can use *structural parameters* which take on smaller values than the size of a minimum vertex cover, a quantity henceforth called the *vertex cover number*. Such structural parameterizations can conveniently be described in terms of the vertex-deletion distance to certain graph families $\mathcal{F}$. Note that the vertex cover number $\mathrm{vc}(G)$ of G can be defined as the minimum number of vertex deletions needed to reduce G to an edgeless graph. Hence this number will always be at least as large as the *feedback vertex number* $\mathrm{fvs}(G)$ of G, which is the vertex-deletion distance of G to a forest. In 2011, it was shown that VERTEX COVER even admits a polynomial kernelization when parameterized by the feedback vertex number [21, 22]. This triggered a long line of follow-up research, which aimed to find the most general graph families $\mathcal{F}$ such that VERTEX COVER admits a polynomial kernelization when parameterized by vertex-deletion distance to $\mathcal{F}$. Polynomial kernelizations were obtained for the families $\mathcal{F}$ of graphs of maximum degree two [27], of graphs of constant *tree-depth* [5, 23], of the *pseudo-forests* where each connected component has at most one cycle [17], and for *d-quasi-forests* in which each connected component has a feedback vertex set of size at most $d \in \mathcal{O}(1)$ [18]. Note that all these target graph classes are *closed under taking minors*. Using *randomized algorithms* with a small error probability, polynomial kernelizations are also known for several parameterizations by vertex-deletion distance to graph classes that are not minor-closed, such as Kőnig graphs [25], bipartite graphs [25], and parameterizations based on the linear-programming relaxation of VERTEX COVER [19, 24]. On the negative side, it is known that VERTEX COVER parameterized by the vertex-deletion distance to a graph of treewidth two [9] does *not* have a polynomial kernel, unless NP $\subseteq$ coNP/poly. This long line of research into kernelization for structural parameterizations raises the following question:

> How can we characterize the graph families $\mathcal{F}$ for which VERTEX COVER parameterized by vertex-deletion distance to $\mathcal{F}$ admits a polynomial kernel?

We answer this question for all minor-closed families $\mathcal{F}$, by introducing a new graph parameter.

Our results. We introduce a new graph parameter that we call *bridge-depth*. It has a recursive definition similar to that of tree-depth [30] (full definitions follow in Section 3), but deals with bridges in a special way. A graph without vertices has bridge-depth zero. The bridge-depth $\mathsf{bd}(G)$ of a disconnected graph G is simply the maximum bridge-depth of its connected components. The bridge-depth of a connected nonempty graph G is defined as follows. Let G_{CB} denote the graph obtained from G by contracting each edge that is a bridge in G; the order does not matter. Then $\mathsf{bd}(G) := 1 + \min_{v \in V(G_{\mathrm{CB}})} \mathsf{bd}(G_{\mathrm{CB}} \setminus v)$. Intuitively, the bridge-depth of G is given by the depth of an elimination process [6] that reduces G to the empty graph. One step consists of contracting all bridges and removing a vertex; each of the remaining connected components is then recursively eliminated in parallel. From this definition, it is not difficult to see that $\mathsf{bd}(G)$ is at least as large as the *tree-width* of G, but never larger than the tree-depth or feedback vertex number of G. In particular, any forest has bridge-depth one.

Using the notion of bridge-depth, we characterize the minor-closed families $\mathcal{F}$ for which Vertex Cover parameterized by vertex-deletion distance to $\mathcal{F}$ admits a polynomial kernel.

▶ **Theorem 1.1.** *Let $\mathcal{F}$ be a minor-closed family of graphs, and assume* $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$. *Vertex Cover parameterized by vertex-deletion distance to $\mathcal{F}$ has a polynomial kernelization if and only if $\mathcal{F}$ has bounded bridge-depth.*

Theorem 1.1 gives a clean and unified explanation for all the minor-closed families $\mathcal{F}$ that were previously considered individually [5, 17, 18, 22, 27], and generalizes these results as far as possible. To the best of our knowledge, Theorem 1.1 captures all known (deterministic) kernelizations for structural parameterizations of Vertex Cover. (There are *randomized* kernelizations [19, 24, 25] which apply for distance to classes $\mathcal{F}$ that are *not* minor-closed, such as bipartite graphs.) For example, we capture the case of $\mathcal{F}$ being a forest [22] since forests have bridge-depth one, and the case of $\mathcal{F}$ being graphs of constant tree-depth [5, 23] since bridge-depth does not exceed tree-depth. In this sense, bridge-depth can be seen as the ultimate common generalization of feedback vertex number and tree-depth (which are incomparable parameters) in the context of polynomial kernels for Vertex Cover.

We consider it one of our main contributions to identify the graph parameter bridge-depth as the right way to capture the kernelization complexity of Vertex Cover parameterizations.

Techniques. To describe our techniques, we introduce some terminology. Let $\alpha(G)$ denote the independence number of graph G, i.e., the maximum size of a set of pairwise nonadjacent vertices. A *blocking set* in a graph G is a vertex set $Y \subseteq V(G)$ such that $\alpha(G \setminus Y) < \alpha(G)$. Hence if Y is a blocking set, then every maximum independent set in G contains a vertex from Y. Earlier kernelizations for Vertex Cover parameterized by distance to a graph class $\mathcal{F}$, starting with the work of Jansen and Bodlaender [22], all rely, either implicitly or explicitly, on having upper-bounds on the size of (inclusion-)minimal blocking sets for graphs in $\mathcal{F}$ [5, 17, 18, 22, 27]. For example, it is known that minimal blocking sets in a bipartite graph have size at most two [18, Cor. 11], while minimal blocking sets in graphs of tree-depth c have size at most 2^c [5, Lemma 1]. Similarly, all the existing superpolynomial kernelization lower bounds for parameterizations by distance to $\mathcal{F}$, rely on $\mathcal{F}$ having minimal blocking sets of arbitrarily large size. Indeed, if $\mathcal{F}$ is closed under disjoint union and has arbitrarily large blocking sets, it is easy to prove a superpolynomial lower bound (cf. [19, Thm. 1]).

Since all positive cases for kernelization are when minimal blocking sets of graphs in $\mathcal{F}$ have bounded size, while one easily obtains lower bounds when the size of minimal blocking sets of graphs in $\mathcal{F}$ is unbounded, the question rises whether a bound on the size of minimal

blocking sets is a necessary and sufficient condition for the existence of polynomial kernels. To our initial surprise, we show that for minor-closed families $\mathcal{F}$, this is indeed the case: the purely structural property of having bounded-size minimal blocking sets can always be leveraged into preprocessing algorithms.

For an insight into our techniques, consider an instance (G, k) of VERTEX COVER, together with a vertex set $X \subseteq V(G)$ such that $G \setminus X \in \mathcal{F}$ for some minor-closed family $\mathcal{F}$ that has bounded-size minimal blocking sets. The goal of the kernelization is then to reduce to an equivalent instance of size $|X|^{\mathcal{O}(1)}$ in polynomial time. Using ideas of the previous kernelizations [5,22], it is quite simple to reduce the *number of connected components* of $G \setminus X$ to size $|X|^{\mathcal{O}(1)}$. To obtain a polynomial kernel, the challenge is therefore to bound the size of each such component C of $G \setminus X$ to $|X|^{\mathcal{O}(1)}$, so that the overall instance size becomes polynomial in $|X|$. However, the non-existence of large minimal blocking sets does not seem to offer any handle for reducing the size of individual components of $G \setminus X$. The route to the kernelization therefore goes via the detour of bridge-depth. We prove the following relation between the sizes of minimal blocking sets and bridge-depth.

▶ **Theorem 1.2.** *Let $\mathcal{F}$ be a minor-closed family of graphs. Then $\mathcal{F}$ has bounded bridge-depth if and only if the size of minimal blocking sets of graphs in $\mathcal{F}$ is bounded.*

Using this equivalence, we can exploit the fact that all minimal blocking sets of $\mathcal{F}$ are of bounded size, through the fact that the bridge-depth of $G \setminus X \in \mathcal{F}$ is small. This means that there is a bounded-depth elimination process to reduce $G \setminus X$ to the empty graph. We use this bounded-depth process in a technical kernelization algorithm following a recursive scheme, inspired by the earlier kernelization for the parameterization by distance to bounded tree-depth [5].

Let us now discuss the ideas behind the equivalence of Theorem 1.2. We prove that the bridge-depth of graphs in a minor-closed family $\mathcal{F}$ is upper-bounded in terms of the maximum size of minimal blocking sets for graphs in $\mathcal{F}$, by exploiting the Erdős-Pósa property in an interesting way. We analyze an elementary graph structure called *necklace of length t*, which is essentially the multigraph formed by a path of t double-edges. If a simple graph $G \in \mathcal{F}$ contains a necklace of length t as a minor, then there is a minor G' of G (which therefore also belongs to $\mathcal{F}$) that has a minimal blocking set of size $\Omega(t)$. Hence to show that bridge-depth is upper-bounded in terms of the size of minimal blocking sets of graphs in $\mathcal{F}$, it suffices to show that bridge-depth is upper-bounded by the maximum length of a necklace minor of graphs in $\mathcal{F}$. Since the definition of bridge-depth allows for the contraction of all bridges in a single step, it suffices to consider bridgeless graphs. Then we argue that in a bridgeless graph, any pair of maximum-length necklace minor models intersects at a vertex (cf. Lemma 4.6). By the Erdős-Pósa property, this implies that there is a constant-size vertex set that hits all maximum-length necklace minor models, and whose removal therefore strictly decreases the maximum length of a necklace minor. If the length of necklace minor models is bounded, then after a bounded number of steps of this process (interleaved with contracting all bridges) we reduce the maximum length of necklace minor models to zero, which is equivalent to breaking all cycles of the graph. At that point, the bridge-depth is one by definition, and we have obtained the desired upper-bound on the bridge-depth in terms of the length of the longest necklace minor, and therefore blocking set size.

For the other direction of Theorem 1.2, we prove (cf. Theorem 5.4) the tight bound that a minimal blocking set in a graph G has size at most $2^{\mathsf{bd}(G)}$. We use induction to prove this statement, together with an analysis of the structure of a tree of bridges whose removal decreases the bridge-depth. The fact that bipartite graphs have minimal blocking sets of size at most two, allows for an elegant induction step.

Related work. In a recent paper, Hols, Kratsch, and Pieterse [19] also analyze the role of blocking sets in the existence of polynomial kernels for structural parameterizations of VERTEX COVER. Note that our paper is independent from, and orthogonal to [19]: we consider the setting of *deterministic* kernelization algorithms for parameterizations to *minor-closed* families $\mathcal{F}$, and obtain an exact characterization of which $\mathcal{F}$ allow for a polynomial kernelization. Hols et al. [19] consider hereditary families $\mathcal{F}$ and give kernelizations for several such parameterizations, without arriving at a complete characterization. Some of the randomized kernelizations they provide do not fit into our framework, but all the deterministic kernelizations they present are captured by Theorem 1.1. Another contribution of [19] is to prove that there is a class $\mathcal{F}$ with minimal blocking sets of size one where VERTEX COVER cannot be solved in polynomial time. In particular, there is no polynomial kernel parameterized by the distance to this family $\mathcal{F}$, and thus bounded minimal blocking set size is not sufficient to get a polynomial kernel. This implies that our minor-closed assumption of Theorem 1.1 cannot be dropped.

We refer to the survey by Fellows et al. [13] for an overview of classic results and new research lines concerning kernelization for VERTEX COVER. Additional relevant work includes the work by Kratsch [24] on a randomized polynomial kernel for a parameterization related to the difference between twice the cost of the linear-programming relaxation of VERTEX COVER and the size of a maximum matching.

Organization. Preliminaries on graphs and complexity are presented in Section 2. Section 3 introduces bridge-depth and its properties. In Section 4 we prove one direction of Theorem 1.2, showing that large bridge-depth implies the existence of large minimal blocking sets. In Section 5 we handle the other direction, proving a tight upper-bound on the size of minimal blocking sets in terms of the bridge-depth. We discuss the kernelization algorithm exploiting bridge-depth in Section 6, with the technical content being available in the full version of the article [4] due to space limitations. We conclude the article in Section 7. The proofs of the results marked with "($\star$)" have been also deferred to the full version [4].

2 Preliminaries

Graphs. We use standard graph-theoretic notation, and we refer the reader to Diestel [11] for any undefined terms. All graphs we consider are finite and undirected. Graphs are simple, unless specifically stated otherwise. A graph G has vertex set $V(G)$ and edge set $E(G)$. Given a graph G and a subset $S \subseteq V(G)$, we say that S is *connected* if $G[S]$ is connected, and we use the shorthand $G \setminus S$ to denote $G[V(G) \setminus S]$. For a single vertex $v \in V(G)$, we use $G \setminus v$ as a shorthand for $G \setminus \{v\}$. Similarly, for a set of edges $T \subseteq E(G)$ we denote by $G \setminus T$ the graph on vertex set $V(G)$ with edge set $E(G) \setminus T$. A cycle on three vertices is called a *triangle*. For a positive integer i, we denote by $[i]$ the set of all integers j such that $1 \leq j \leq i$. Given $v \in V(G)$, we denote $N_G(v) = \{u \mid \{u, v\} \in E(G)\}$, $d_G(v) = |N_G(v)|$ and given $X \subseteq V(G)$, we denote $N_G(X) = \bigcup_{v \in X} N_G(v) \setminus X$. Given $X, Y \subseteq V(G)$, we denote by $N_G^Y(X) = N_G(X) \cap Y$. We may omit the subscript G when it is clear from the context. For distinct vertices u and v of a graph G, the graph G' obtained by *identifying* u and v is defined by removing vertices u and v from G, adding a new vertex uv with $N_{G'}(uv) = N_G(\{u, v\})$, and keeping the other vertices and edges unchanged. Given two adjacent vertices u and v, we define the *contraction* of the edge $\{u, v\}$ as the identification of u and v.

Given a graph G, we denote by $\alpha(G)$ the size of a maximum independent set in G, by $\#\mathrm{cc}(G)$ the number of connected components of G, by $\mathsf{diam}(G)$ the diameter of G, and by $\Delta(G)$ the maximum degree of G. Given a graph G and a set $S \subseteq V(G)$, we say that S is a *blocking set* in G if $\alpha(G \setminus S) < \alpha(G)$. The maximum size of an inclusion-wise minimal blocking set of a graph G is denoted by $\mathsf{mbs}(G)$.

A graph H is a minor of graph G if H can be obtained from G by a sequence of edge deletions, edge contractions, and removals of isolated vertices. Let us also recall the definition of minor in the context of multigraphs. Let H be a loopless multigraph. An H-*model* M in a simple graph G is a collection $\{S_x^M \mid x \in V(H)\}$ of pairwise disjoint subsets of $V(G)$ such that $G[S_x^M]$ is connected for every $x \in V(H)$, and such that for every pair of distinct vertices x, y of H, the quantity $|\{\{u, v\} \in E(G) \mid u \in S_x^M, v \in S_y^M\}|$ is at least the number of edges in H between x and y. The vertex set $V(M)$ of M is the union of the vertex sets of the subgraphs in the collection. We say that a graph G contains a loopless multigraph H as a *minor* if G has an H-model.

For the following definitions, we refer the reader to [29] for more details and we only recall here some basic notations and facts. The *tree-depth* of a graph G, denoted by $\mathsf{td}(G)$, is defined recursively. The empty graph without vertices has tree-depth zero. The tree-depth of a disconnected graph is the maximum tree-depth of its connected components. Finally, if G is a nonempty connected graph then $\mathsf{td}(G) = 1 + \min_{v \in V(G)} \mathsf{td}(G \setminus v)$. Equivalent definitions exist in terms of the minimum height of a rooted forest whose closure is a supergraph of G. The tree-width of G is denoted $\mathsf{tw}(G)$ (cf. [2]).

Given a graph family $\mathcal{F}$, an $\mathcal{F}$-*modulator* in a graph G is a subset of vertices $X \subseteq V(G)$ such that $G \setminus X \in \mathcal{F}$. We denote by $\mathsf{dist\text{-}to\text{-}}\mathcal{F}(G)$ the size of a smallest $\mathcal{F}$-modulator in G. For a graph measure f that associates an integer with each graph, and an integer c, a *c-f-modulator* is a modulator to $\mathcal{F}_c^{\mathsf{f}} := \{G \mid \mathsf{f}(G) \leq c\}$. We denote by $c\text{-}\mathsf{f\text{-}mod}(G) := \mathsf{dist\text{-}to\text{-}}\mathcal{F}_c^{\mathsf{f}}(G)$, that is, the size of a smallest c-f-modulator of G. Typical measures f that we consider here are tree-width, tree-depth, and bridge-depth. Notice that $0\text{-}\mathsf{tw\text{-}mod}(G)$ corresponds to the minimum size of a vertex cover of G, and $1\text{-}\mathsf{tw\text{-}mod}(G)$ corresponds to the minimum size of a feedback vertex set of G. Finally, IS (resp. VC) denotes the MAXIMUM INDEPENDENT SET (resp. MINIMUM VERTEX COVER) problem.

Parameterized complexity.   A *parameterized problem* is a language $L \subseteq \Sigma^* \times \mathbb{N}$, for some finite alphabet Σ. For an instance $(x, k) \in \Sigma^* \times \mathbb{N}$, the value k is called the *parameter*. For a computable function $g \colon \mathbb{N} \to \mathbb{N}$, a *kernelization algorithm* (or simply a *kernel*) for a parameterized problem L of *size g* is an algorithm A that given any instance (x, k) of L, runs in polynomial time and returns an instance (x', k') such that $(x, k) \in L \Leftrightarrow (x', k') \in L$ with $|x'|, k' \leq g(k)$. Consult [8, 12, 14, 16, 31] for background on parameterized complexity.

3   An introduction to bridge-depth

Let G be a graph. An edge $e \in E(G)$ is a *bridge* if its removal increases the number of connected components of G. We define G_{CB} as the simple graph obtained from G by contracting all bridges of G (the order does not matter.) Observe that, as contracting an edge cannot create a new bridge, G_{CB} has no bridges, implying that $(G_{\mathrm{CB}})_{\mathrm{CB}} = G_{\mathrm{CB}}$. Given a subgraph T of a graph G, we say that T is a *tree of bridges* if T is a tree and, for every $e \in E(T)$, e is a bridge in G. Note that a single vertex is, by definition, a tree of bridges. Note also that with any vertex $v \in V(G_{\mathrm{CB}})$ we can associate, in a bijective way, an inclusion-wise maximal tree of bridges T_v of G. The set $\{T_v \mid v \in V(G_{\mathrm{CB}})\}$ is a minor model of G_{CB} in G (a

G_{CB}-model, from now on). For any $u, v \in V(G_{\mathrm{CB}})$ such that $\{u, v\} \in E(G_{\mathrm{CB}})$, there is exactly one edge $\{u', v'\} \in E(G)$ with $u' \in T_u$ and $v' \in T_v$. The latter claim can be easily verified by supposing that there are two such edges, implying that some edge in T_u or T_v is involved in a cycle, which contradicts the fact that all the edges in T_u and T_v are bridges.

▶ **Definition 3.1.** *The* bridge-depth $\mathsf{bd}(G)$ *of a graph G is recursively defined as follows:*
- *If G is the empty graph without any vertices, then $\mathsf{bd}(G) = 0$.*
- *If G has $\ell > 1$ connected components $\{G_i \mid i \in [\ell]\}$, then $\mathsf{bd}(G) = \max_{i \in [\ell]} \mathsf{bd}(G_i)$.*
- *If G is connected, then $\mathsf{bd}(G) = 1 + \min_{v \in V(G_{\mathrm{CB}})} \mathsf{bd}(G_{\mathrm{CB}} \setminus v)$.*

Informally, bd behaves like tree-depth except that at each step of the recursive definition we are allowed to delete trees of bridges instead of just single vertices, as proved in Item 4 of the following proposition. The following properties of bridge-depth follow from the definitions in an elementary way, often exploiting the fact that if e is a bridge in G, then e is also a bridge in any minor of G that still contains e.

▶ **Proposition 3.2.** *For any graph G the following claims hold:*

1. *$\mathsf{bd}(G) = 1$ if and only if G is a forest with at least one vertex.*
2. *$\mathsf{bd}(G_{\mathrm{CB}}) = \mathsf{bd}(G)$.*
3. *The parameter bd is minor-closed: if G' is a minor of G then $\mathsf{bd}(G') \leq \mathsf{bd}(G)$.*
4. *If G is connected, then $\mathsf{bd}(G) = 1 + \min_T \mathsf{bd}(G \setminus V(T))$, where the minimum is taken over all trees of bridges T of G.*
5. *For any $X \subseteq V(G)$, we have $\mathsf{bd}(G) \leq |X| + \mathsf{bd}(G \setminus X)$.*
6. *$\mathsf{tw}(G) \leq \mathsf{bd}(G)$.*

Proof. The first item follows easily from the definition, while the second one uses that $(G_{\mathrm{CB}})_{\mathrm{CB}} = G_{\mathrm{CB}}$.

Proof of 3: We prove the claim by induction on $|V(G)| + |E(G)|$. Suppose that G has multiple connected components $\{G_i \mid i \in [\#\mathrm{CC}(G)]\}$, and let $\{G'_i \mid i \in [\#\mathrm{CC}(G')]\}$ be the connected components of the minor G' of G. Then each connected component G'_j is a minor of some component G_i of G on fewer than $|V(G)|$ vertices, which gives $\mathsf{bd}(G'_j) \leq \mathsf{bd}(G_i)$ by induction. Hence we have $\mathsf{bd}(G') = \max_{j \in [\#\mathrm{CC}(G')]} \mathsf{bd}(G'_j) \leq \max_{i \in [\#\mathrm{CC}(G)]} \mathsf{bd}(G_i) = \mathsf{bd}(G)$.

We now deal with the case that G is connected. In general, if some graph G^* is a minor of G, then G^* is a minor of a graph G' obtained from G by removing an edge, contracting an edge, or removing an isolated vertex. Since G is assumed to be connected, the third case cannot occur here. Then by induction, we have $\mathsf{bd}(G^*) \leq \mathsf{bd}(G')$, so it suffices to prove that $\mathsf{bd}(G') \leq \mathsf{bd}(G)$ for any graph G' obtained by removing or contracting an edge. Let us first prove that if G'_{CB} is a minor of G_{CB}, then $\mathsf{bd}(G') \leq \mathsf{bd}(G)$. Indeed, let $v^* \in V(G_{\mathrm{CB}})$ such that $\mathsf{bd}(G) = 1 + \mathsf{bd}(G_{\mathrm{CB}} \setminus v^*)$, and consider an arbitrary component G'_i of G'. Note that $(G'_i)_{\mathrm{CB}}$ is a component of $(G')_{\mathrm{CB}}$, and therefore a minor of G_{CB} by hypothesis. If $(G'_i)_{\mathrm{CB}}$ is a minor of the graph $G_{\mathrm{CB}} \setminus v^*$, then by induction and Item 2 we have $\mathsf{bd}(G'_i) = \mathsf{bd}((G'_i)_{\mathrm{CB}}) \leq \mathsf{bd}(G_{\mathrm{CB}} \setminus v^*) < \mathsf{bd}(G)$. Otherwise, any minor model $\{S_x \mid x \in V((G'_i)_{\mathrm{CB}})\}$ of $(G'_i)_{\mathrm{CB}}$ in G_{CB} contains a branch set S_{x^*} with $v^* \in S_{x^*}$. But then $\mathsf{bd}((G'_i)_{\mathrm{CB}}) \leq 1 + \mathsf{bd}((G'_i)_{\mathrm{CB}} \setminus x^*)$ by definition, and $(G'_i)_{\mathrm{CB}} \setminus x^*$ is a minor of $G_{\mathrm{CB}} \setminus v^*$, and therefore has bridge-depth at most $\mathsf{bd}(G) - 1$, so that $\mathsf{bd}((G'_i)_{\mathrm{CB}}) \leq \mathsf{bd}(G)$. Hence for each component G'_i of G' we have $\mathsf{bd}(G'_i) = \mathsf{bd}((G'_i)_{\mathrm{CB}}) \leq \mathsf{bd}(G)$, implying $\mathsf{bd}(G') \leq \mathsf{bd}(G)$.

Thus, it only remains to prove that G'_{CB} is a minor of G_{CB}. Let us first assume that G' is obtained from G by removing an edge e. Let $\{T_v \mid v \in V(G_{\mathrm{CB}})\}$ be the G_{CB}-model in G given by the trees of bridges. If e is not a bridge, then e is an edge between T_u and T_v for some vertices $u, v \in V(G_{\mathrm{CB}})$. To obtain G'_{CB} as a minor, we start from G_{CB}, remove edge $\{u, v\}$, and for any edge e' between $T_{u'}$ and $T_{v'}$ (for any $u', v' \in V(G_{\mathrm{CB}})$) that has become a bridge in G' because of the removal of e, we contract $\{u', v'\}$. This implies that G'_{CB} is a minor of G_{CB}. Otherwise, if e is a bridge, then there exists $u \in V(G_{\mathrm{CB}})$ such that $e \in E(T_u)$, and G' has two connected components G'_1 and G'_2. To obtain $(G'_i)_{\mathrm{CB}}$ as a minor, for $i \in [2]$, we start from G_{CB} and remove any vertex v such that $T_v \cap V(G'_i) = \emptyset$ (notice that u appears both in $(G'_1)_{\mathrm{CB}}$ and $(G'_2)_{\mathrm{CB}}$). Thus, both $(G'_1)_{\mathrm{CB}}$ and $(G'_2)_{\mathrm{CB}}$ are minors of G_{CB}, hence G'_{CB} as well. The case where G' is obtained from G by contracting an edge e can be proved using similar but simpler arguments. Indeed, if e is a bridge in G, then we have that $G'_{\mathrm{CB}} = G_{\mathrm{CB}}$, and if it is not, if suffices to contract in G_{CB} the edge $\{u, v\}$ with $u, v \in V(G_{\mathrm{CB}})$ such that e is an edge between T_u and T_v.

Proof of 4: Let $v \in V(G_{\mathrm{CB}})$ and T_v be its associated tree of bridges in G. Observe first that we may have $(G \setminus V(T_v))_{\mathrm{CB}} \neq G_{\mathrm{CB}} \setminus v$. Indeed, if for example we consider G composed of two vertex-disjoint triangles $\{a, b, c\}$, $\{a', b', c'\}$ and an edge $e = \{a, a'\}$, and if we consider $T_v = \{e\}$, then $G_{\mathrm{CB}} \setminus v$ is composed of two disjoint edges, whereas $(G - V(T_v))_{\mathrm{CB}}$ is composed of two isolated vertices. However, it is easy to verify that $(G \setminus V(T_v))_{\mathrm{CB}} = (G_{\mathrm{CB}} \setminus v)_{\mathrm{CB}}$. Let us now prove that $\min_T \mathsf{bd}(G \setminus V(T)) \leq \min_{v \in V(G_{\mathrm{CB}})} \mathsf{bd}(G_{\mathrm{CB}} \setminus v)$. Let v^* be a vertex minimizing $\mathsf{bd}(G_{\mathrm{CB}} \setminus v)$. We have $\min_T \mathsf{bd}(G \setminus V(T)) \leq \mathsf{bd}(G \setminus V(T_{v^*})) = \mathsf{bd}((G \setminus V(T_{v^*})_{\mathrm{CB}})$ using Item 2 in the last equality, and $\mathsf{bd}((G \setminus V(T_{v^*})_{\mathrm{CB}}) = \mathsf{bd}((G_{\mathrm{CB}} \setminus v^*)_{\mathrm{CB}}) = \mathsf{bd}(G_{\mathrm{CB}} \setminus v^*)$ using again Item 2.

For the other inequality, let T^0 be a tree of bridges that minimizes $\mathsf{bd}(G \setminus V(T))$. If T^0 is not inclusion-wise maximal, let T^* be any inclusion-wise maximal tree of bridges containing T^0. Note that as $G \setminus V(T^*)$ is a subgraph of $G \setminus V(T^0)$, by Item 3 we get that $\mathsf{bd}(G \setminus V(T^*)) \leq \mathsf{bd}(G \setminus V(T^0))$, implying that T^* also minimizes $\mathsf{bd}(G \setminus V(T))$. Let $v^* \in V(G_{\mathrm{CB}})$ such that $T_{v^*} = T^*$. We have $\min_{v \in V(G_{\mathrm{CB}})} \mathsf{bd}(G_{\mathrm{CB}} \setminus v) \leq \mathsf{bd}(G_{\mathrm{CB}} \setminus v^*) = \mathsf{bd}((G \setminus T_{v^*})_{\mathrm{CB}}) = \mathsf{bd}(G \setminus T_{v^*})$.

Proof of 5: We use induction on $|X|$, the base case $X = \emptyset$ being trivial. For the induction step, pick an arbitrary $v \in X$, let $X' := X \setminus \{v\}$, and $G' := G \setminus X'$. By induction we have $\mathsf{bd}(G) \leq |X'| + \mathsf{bd}(G')$. Let G'_i be the connected component of G' containing v. Using v as a singleton tree of bridges in G'_i, Item 4 shows that $\mathsf{bd}(G'_i) \leq 1 + \mathsf{bd}(G'_i \setminus v) \leq 1 + \mathsf{bd}(G' \setminus v)$. Since all other components G'_j of G' also occur as components of $G' \setminus v$, it follows that $\mathsf{bd}(G'_j) \leq \mathsf{bd}(G' \setminus v)$, implying $\mathsf{bd}(G') \leq 1 + \mathsf{bd}(G' \setminus v) = 1 + \mathsf{bd}(G \setminus X)$ since $G' \setminus v = G \setminus X$. Hence $\mathsf{bd}(G) \leq |X'| + 1 + \mathsf{bd}(G \setminus X)$.

Proof of 6: We use induction on $|V(G)|$; the base case follows directly from the definitions. It is well-known (cf. [2, Lemma 6]) that the tree-width of G is the maximum tree-width of its biconnected components. Hence it suffices to prove that for an arbitrary biconnected component G' of G, we have $\mathsf{tw}(G') \leq \mathsf{bd}(G')$. If G' consists of a single edge, then $\mathsf{tw}(G') = \mathsf{bd}(G') = 1$. Otherwise, G' is a connected bridgeless graph. This implies $(G')_{\mathrm{CB}} = G'$, so by Definition 3.1 there is a vertex $v \in V(G')$ such that $\mathsf{bd}(G') = 1 + \mathsf{bd}(G' \setminus v)$. Since G' is a minor of G, we have $\mathsf{bd}(G') \leq \mathsf{bd}(G)$ by Item 3. By induction, the tree-width of $G' \setminus v$ is at most $\mathsf{bd}(G' \setminus v) \leq \mathsf{bd}(G) - 1$. Adding vertex v to all bags of a tree decomposition of this width, gives a valid tree decomposition of G' of width at most $\mathsf{bd}(G' \setminus v) + 1 \leq \mathsf{bd}(G')$. Hence $\mathsf{tw}(G') \leq \mathsf{bd}(G')$ for all biconnected components of G. ◀

A $(c+1) \times (c+1)$-grid is a planar graph of tree-width exactly $c + 1$ [2, Cor. 89], which implies by Item 6 of Proposition 3.2 that its bridge-depth is larger than c. This gives the following consequence of Proposition 3.2, which will be useful when invoking algorithmic meta-theorems.

▶ **Observation 3.3.** *For each $c \in \mathbb{N}$, the graphs of bridge-depth at most c form a minor-closed family that excludes a planar graph. By the Graph Minor Theorem [35], there is a finite set of forbidden minors $\mathcal{H}_c$ such that $\mathsf{bd}(G) \leq c$ if and only if G excludes all graphs of $\mathcal{H}_c$ as a minor. The set $\mathcal{H}_c$ contains a planar graph, since some planar graphs have bridge-depth $> c$.*

Observation 3.3, together with known results on minor testing, imply the following.

▶ **Proposition 3.4** (Follows from [1, Thm. 7.1]). *For each constant $c \in \mathbb{N}$, there is a linear-time algorithm to test whether the bridge-depth of a given graph G is at most c.*

Fomin et al. [15, Thm. 1.3] gave a generic approximation algorithm for finding a small vertex set that hits forbidden minors from a finite forbidden set containing a planar graph. By Observation 3.3, deleting vertices to obtain a graph of bounded bridge-depth fits into their framework.

▶ **Proposition 3.5** (Follows from [15, Thm. 1.3]). *For each fixed $c \in \mathbb{N}$ there is a polynomial-time algorithm that, given a graph G, outputs a set $X \subseteq V(G)$ such that $\mathsf{bd}(G \setminus X) \leq c$ and $|X| \leq \mathcal{O}(|X_{\mathrm{opt}}| \log^{2/3} |X_{\mathrm{opt}}|)$, where $|X_{\mathrm{opt}}|$ is the minimum size of such a set.*

The following concept will be crucial to facilitate a recursive approach for reducing graphs of bounded bridge-depth.

▶ **Definition 3.6.** *A lowering tree T of a graph G is a tree of bridges (possibly consisting of a single vertex and no bridges) such that $\mathsf{bd}(G \setminus V(T)) = \mathsf{bd}(G) - 1$.*

Item 4 of Proposition 3.2 implies that any connected graph G has a lowering tree.

▶ **Proposition 3.7.** *For each fixed $c \in \mathbb{N}$ there is an algorithm that, given a connected graph G on n vertices of bridge-depth c, computes a lowering tree in $\mathcal{O}(n^2)$ time.*

Proof. Given G, we compute its decomposition into biconnected components, which can be done in linear time taking into account that having bounded bridge-depth implies a linear number of edges [20]. From this decomposition, it is straightforward to identify the inclusion-maximal trees of bridges in G. For each tree of bridges T in G, we can test whether $\mathsf{bd}(G \setminus V(T)) < c = \mathsf{bd}(G)$ in linear time using Proposition 3.4, and we output T if this is the case. By Proposition 3.2, such a tree T exists. Since G is decomposed into at most n trees of bridges, and we need a linear-time computation for each T, this results in an $\mathcal{O}(n^2)$-time algorithm. ◀

4 Bounded minimal blocking sets imply bounded bridge-depth

The goal of this section is to prove one direction of Theorem 1.2, showing that if $\mathcal{F}$ has bounded-size minimal blocking sets, then $\mathcal{F}$ has bounded bridge-depth. As explained in Section 1, we prove this via the intermediate structure of *necklace minors* and show that the bridge-depth of a graph G can be upper-bounded in terms of the longest necklace contained in it as a minor.

This result can be seen as an analog to the fact that the tree-depth of a graph can be bounded in terms of the length of the longest simple path it contains (as a subgraph or as a minor, which is equivalent for paths). A classical proof of this fact (see [29]) is to consider a depth-first search tree of G, bounding the tree-depth of G by the depth of this tree. However, it does not seem immediate to find a similar bound for bridge-depth.

We therefore follow another approach, inspired by the following alternative proof that the tree-depth is upper-bounded by the length of the longest path (which gives a worse bound). Observe that in a connected graph G, any two longest paths intersect at a vertex. (If they did not, one could combine them to make an even longer path.) Given a connected graph G whose longest path has t vertices, we can bound its tree-depth by $f(t) := \sum_{i=1}^{t} i$ as follows. Let P be a longest path in G. Then the longest path in $G \setminus V(P)$ has strictly fewer than t vertices, and by induction the tree-depth of $G \setminus V(P)$ is at most $f(t-1)$. From the definition of tree-depth, it follows that the tree-depth of G is at most $|V(P)| = t$ larger than that of $G \setminus V(P)$, so the tree-depth of G is at most $f(t)$.

In the case of bridge-depth, where paths are replaced with necklaces contained as minors, we cannot afford to remove the entire set of vertices of the corresponding model of a longest necklace, as the size of this set cannot be bounded in terms of the length t of the necklace. To overcome this problem, we will prove in Lemma 4.6, similarly to the case of paths, that there cannot be two vertex-disjoint longest necklaces. Then we resort to the Erdős-Pósa property, which gives us a set of vertices of size $f(t)$ whose removal decreases the maximum length of a longest necklace. We now formalize these ideas.

▶ **Definition 4.1.** *For $t \in \mathbb{N}$, the* necklace of length t*, denoted by N_t, is the multigraph having $t + 1$ vertices $\{v_i \mid i \in [t + 1]\}$ and two parallel edges between v_i and v_{i+1} for $i \in [t]$.*

▶ **Observation 4.2.** *A simple graph G contains N_t as a minor if and only if G contains $t + 1$ vertex-disjoint sets $S_i \subseteq V(G)$ such that each S_i is connected and, for $i \in [t]$, there are at least two edges between S_i and S_{i+1}.*

▶ **Definition 4.3.** *The* necklace-minor length *of a graph G, denoted by $\mathsf{nm}(G)$, is the largest length of a necklace contained in G as a minor, or zero if G contains no such minor.*

We need to introduce the Erdős-Pósa property for packing and covering minor models. Let $\mathcal{F}$ be a finite collection of simple graphs. An $\mathcal{F}$-*model* is an H-model for some $H \in \mathcal{F}$. Two $\mathcal{F}$-models M_1 and M_2 are *disjoint* if $V(M_1) \cap V(M_2) = \emptyset$. Let $\nu_{\mathcal{F}}(G)$ be the maximum cardinality of a packing of pairwise disjoint $\mathcal{F}$-models in G, and let $\tau_{\mathcal{F}}(G)$ be the minimum size of a subset $X \subseteq V(G)$ such that $G \setminus X$ has no $\mathcal{F}$-model. Clearly, $\nu_{\mathcal{F}}(G) \leq \tau_{\mathcal{F}}(G)$. We say that the *Erdős-Pósa property* holds for $\mathcal{F}$-models if there exists a bounding function $f : \mathbb{N} \to \mathbb{N}$ such that, for every graph G, $\tau_{\mathcal{F}}(G) \leq f(\nu_{\mathcal{F}}(G))$.

In the case where $\mathcal{F} = \{H\}$ contains a single connected graph H, Robertson and Seymour [34] proved the following result.

▶ **Theorem 4.4** (Robertson and Seymour [34])**.** *Let H be a connected graph. The Erdős-Pósa property holds for H-models if and only if H is planar.*

It is worth mentioning that a tight bounding function when H is planar has been recently obtained by van Batenburg et al. [36]. Theorem 4.4 easily implies the following corollary.

▶ **Corollary 4.5.** *For every $t \geq 1$, the Erdős-Pósa property holds for N_t-models.*

Proof. For $t \geq 1$, let $\mathcal{F}_t$ be the set containing all minor-minimal *simple* graphs that contain the necklace N_t as a minor. By definition, a simple graph G contains an N_t-model if and only if it contains an $\mathcal{F}_t$-model. Clearly, all the graphs in $\mathcal{F}_t$ are connected and planar, and

it is easy to see that $|\mathcal{F}_t|$ is bounded by a function of t. For each $F \in \mathcal{F}_t$, by Theorem 4.4 there is a function f_F such that if G does not contain k vertex-disjoint models of F, then all the F-models of G can be hit by at most $f_F(k)$ vertices. This implies that if G does not contain k models of any graph in $\mathcal{F}_t$, then the union of all hitting sets has size bounded by $\sum_{F \in \mathcal{F}_t} f_F(k)$, and since $\mathcal{F}_t$ is finite this is a valid bounding function for N_t-models. ◄

We denote by f_{N_t} the bounding function for N_t-models given by Corollary 4.5. In a connected bridgeless graph, each pair of maximum-length necklace models intersect at a vertex:

▶ **Lemma 4.6.** *If G is a connected bridgeless simple graph with $\mathsf{nm}(G) = t$, then $\nu_{N_t}(G) = 1$.*

Proof. Suppose for contradiction that G contains two disjoint models M^1 and M^2 of N_t. For $i \in [t+1]$ and $\ell \in [2]$, let S_i^ℓ be the vertex set of M^ℓ given by Observation 4.2. Note that these $2t + 2$ subsets of vertices of G are pairwise disjoint, and that for any $i \in [t]$, there are at least two edges between S_i^ℓ and S_{i+1}^ℓ. Since G is bridgeless and connected, it is 2-edge-connected and by Menger's theorem [11, § 3.3] G contains two edge-disjoint paths between any pair of vertices. Pick two arbitrary vertices $x_1 \in M^1, x_2 \in M^2$, and let P^1, P^2 be two edge-disjoint paths between them. Consider the subpath Q^ℓ of P^ℓ between the last vertex of M^1 that is visited, until the first vertex of M^2. Let $Q^\ell = (v_1^\ell, \ldots, v_{q_\ell}^\ell)$ where $v_1^\ell \in M^1$ and $v_{q_\ell}^\ell \in M^2$. Let a_ℓ such that $v_1^\ell \in S_{a_\ell}^1$ and b_ℓ such that $v_{q_\ell}^\ell \in S_{b_\ell}^2$.

Let us first show that if t is odd, then we can use Q^1 to find an $N_{t'}$-model M' for some $t' > t$ by "gluing" M^1 and M^2, leading to a contradiction. Let $S = S_{a_1}^1 \cup V(Q^1) \cup S_{b_1}^2$. If $a_1 > \frac{t+1}{2}$ define $A = \{S_1^1, \ldots, S_{a_1-1}^1\}$, and otherwise define $A = \{S_{a_1+1}^1, \ldots, S_{t+1}^1\}$. Similarly, if $b_1 > \frac{t+1}{2}$ define $B = \{S_1^2, \ldots, S_{b_1-1}^2\}$, and otherwise define $B = \{S_{b_1+1}^2, \ldots, S_{t+1}^2\}$. Note that the sets A, S, B are pairwise disjoint. Since t is odd, it can be easily checked that $M' = A \cup \{S\} \cup B$ is an $N_{t'}$-model in G for some $t' > t$; see Figure 1(a) for an illustration.

(a)   (b)
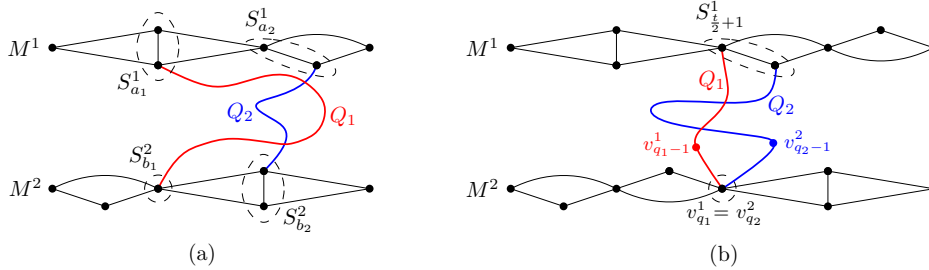
▶ **Figure 1** (a) Example with $t = 3$ and $a_1 = b_1 = 2$. (b) Example with $t = 4$.

Let us now consider the case where t is even. Note first that if there exists $\ell \in [2]$ such that $a_\ell \neq \frac{t}{2} + 1$ or $b_\ell \neq \frac{t}{2} + 1$, then we can use Q^ℓ to find an $N_{t'}$-model for some $t' > t$ as in the previous case. Hence, it only remains to consider the case where $a_1 = b_1 = a_2 = b_2 = \frac{t}{2} + 1$, meaning that Q^1 and Q^2 are two edge-disjoint paths, both between $S_{\frac{t}{2}+1}^1$ and $S_{\frac{t}{2}+1}^2$. Let $A = \{S_1^1, \ldots, S_{\frac{t}{2}}^1\}$, $B = \{S_1^2, \ldots, S_{\frac{t}{2}}^2\}$, and $S = S_{a_1}^1 \cup (V(Q^1) \setminus \{v_{q_1}^1\}) \cup (V(Q^2) \setminus \{v_{q_2}^2\})$. We claim that $M' = A \cup \{S, S_{\frac{t}{2}+1}^2\} \cup B$ is an N_{t+1}-model. Indeed, note in particular there are two edges between S and $S_{\frac{t}{2}+1}^2$ as we cannot have $v_{q_1-1}^1 = v_{q_2-1}^2$ and $v_{q_1}^1 = v_{q_2}^2$ because Q^1 and Q^2 are edge-disjoint and G is a simple graph; see Figure 1(b) for an illustration. ◄

By combining Corollary 4.5 with Lemma 4.6 we easily get the following corollary.

▶ **Corollary 4.7.** *Let G be a connected bridgeless graph and $t = \mathsf{nm}(G)$. Then G contains a set of vertices X with $|X| \leq f_{N_t}(1)$ such that $\mathsf{nm}(G \setminus X) < \mathsf{nm}(G)$, where $f_{N_t} : \mathbb{N} \to \mathbb{N}$ is the bounding function given by Corollary 4.5.*

Proof. By Lemma 4.6, it follows that $\nu_{N_t}(G) = 1$, and therefore by Corollary 4.5

$$\tau_{N_t}(G) \ \leq \ f_{N_t}(\nu_{N_t}(G)) \ = \ f_{N_t}(1).$$

Thus, there exists a set $X \subseteq V(G)$ with $|X| \leq f_{N_t}(1)$ such that $G \setminus X$ has no N_t-model, implying that $\mathsf{nm}(G \setminus X) < t$.   ◄

We are finally in position to prove the following theorem.

▶ **Theorem 4.8.** *There is a function $f : \mathbb{N} \to \mathbb{N}$ such that $\mathsf{bd}(G) \leq f(\mathsf{nm}(G))$ for all graphs G.*

Proof. We prove the statement by induction on $\mathsf{nm}(G)$, for the function f defined by $f(t) := 1 + \sum_{i=1}^{t} f_{N_i}(1)$. If $\mathsf{nm}(G) = 0$, then G is a forest, and by definition of bridge-depth we get $\mathsf{bd}(G) = 1 = f(0)$. Suppose now that $\mathsf{nm}(G) = t$ with $t > 0$.

Consider the case that G is connected. Then G_{CB} is also connected and has no bridge, and thus we can apply Corollary 4.7 and get a set $X \subseteq V(G_{\mathrm{CB}})$ with $|X| \leq f_{N_t}(1)$ such that $\mathsf{nm}(G') < t$, where $G' = G_{\mathrm{CB}} \setminus X$. By Item 5 of Proposition 3.2, we get that $\mathsf{bd}(G) = \mathsf{bd}(G_{\mathrm{CB}}) \leq |X| + \mathsf{bd}(G')$. Let $G'_1, \ldots, G'_\ell$ be the connected components of G'. As $\mathsf{nm}(G') < t$, we get that $\mathsf{nm}(G'_i) < t$ for every $i \in [\ell]$. Then, by induction hypothesis it follows that, for every $i \in [\ell]$, $\mathsf{bd}(G'_i) \leq f(t-1) = 1 + \sum_{i=1}^{t-1} f_{N_i}(1)$. Thus, as $\mathsf{bd}(G') = \max_{i \in [\ell]} \mathsf{bd}(G'_i) \leq 1 + \sum_{i=1}^{t-1} f_{N_i}(1)$, we get that

$$\mathsf{bd}(G) \ \leq \ |X| + \mathsf{bd}(G') \ \leq \ f_{N_t}(1) + 1 + \sum_{i=1}^{t-1} f_{N_i}(1) \ = \ 1 + \sum_{i=1}^{t} f_{N_i}(1) \ = \ f(\mathsf{nm}(G)).$$

Finally, if G is disconnected, let $G_1, \ldots, G_\ell$ be its connected components, and note that $\mathsf{bd}(G) = \max_{i \in [\ell]} \mathsf{bd}(G_i)$. Since for every $i \in [\ell]$ it holds that $\mathsf{nm}(G_i) \leq \mathsf{nm}(G)$, and since the function f is non-decreasing, by applying the above case to each connected component of G we get that

$$\mathsf{bd}(G) \ = \ \max_{i \in [\ell]} \mathsf{bd}(G_i) \ \leq \ \max_{i \in [\ell]} f(\mathsf{nm}(G_i)) \ \leq \ \max_{i \in [\ell]} f(\mathsf{nm}(G)) \ = \ f(\mathsf{nm}(G)).$$ ◄

Now that we established a relation between bridge-depth and necklace minors, our next step is to relate necklace minors to blocking sets. For this purpose, we use the known triangle-path gadget.

▶ **Definition 4.9.** *A* triangle-path *of length t is the graph consisting of t vertex-disjoint triangles, with vertex sets $\{\{a_i, b_i, c_i\} \mid i \in [t]\}$, together with the $t-1$ edges $\{\{b_i, a_{i+1}\} \mid i \in [t-1]\}$. The* triangle-path-minor length *of a graph G, denoted by $\mathsf{tpm}(G)$, is the largest length of a triangle-path contained in G as a minor, or zero if no such minor exists.*

A slight variation of this gadget was used by Fomin and Strømme [17, Def. 6]. We observe the following (cf. [17, Obs. 3–5]).

▶ **Observation 4.10.** *Let G be a triangle-path of length $t \geq 2$. Then $\mathsf{mbs}(G) \geq t + 2$, as $\{a_1, c_1\} \cup \{b_t, c_t\} \cup \{c_i \mid i \in [2, t-1]\}$ is a minimal blocking set.*

▶ **Lemma 4.11.** *For any graph G, $\mathsf{tpm}(G) \geq \lfloor \frac{\mathsf{nm}(G)+1}{2} \rfloor$.*

Proof. Let $t = \mathsf{nm}(G)$, and let $\{S_i \mid i \in [t+1]\}$ be an N_t-model in G. Let $i \in [\lfloor \frac{t+1}{2} \rfloor]$ and let $e_1 = \{u_1, v_1\}$ and $e_2 = \{u_2, v_2\}$ be the two edges between S_{2i-1} and S_{2i}, with $u_\ell \in S_{2i-1}$ and $v_\ell \in S_{2i}$. If $u_1 \neq u_2$ then there is a partition A_1, A_2 of S_{2i-1} such that $u_i \in A_i$ and A_i is connected for $i \in [2]$, and we define $L_i = \{A_1, A_2\}$, $R_i = \{S_{2i}\}$. Otherwise, if $u_1 = u_2$,

then necessarily $v_1 \neq v_2$, and we define symmetrically $L_i = \{S_{2i-1}\}$ and $R_i = \{A_1, A_2\}$. In both cases we get that $L_i \cup R_i$ is a model of a triangle, and moreover there is an edge between a vertex in R_i and a vertex in L_{i+1} for every $i \in [\lfloor \frac{t+1}{2} \rfloor - 1]$. This implies that $\bigcup_{i \in [\lfloor \frac{t+1}{2} \rfloor]} (L_i \cup R_i)$ is a model of a triangle-path of length $\lfloor \frac{t+1}{2} \rfloor$ in G. ◄

▶ **Corollary 4.12.** *There is a function* $g \colon \mathbb{N} \to \mathbb{N}$ *such that* $\mathsf{bd}(G) \leq g(\mathsf{tpm}(G))$ *for all graphs* G.

Proof. By Lemma 4.11, we have that $\mathsf{tpm}(G) \geq \mathsf{nm}(G)/2$. By letting $g(t) := f(2t)$, where f is the function given by Theorem 4.8, we get the desired result. ◄

▶ **Corollary 4.13.** *Let* $\mathcal{F}$ *be a minor-closed family of graphs. If* $\mathcal{F}$ *has unbounded bridge-depth then it contains the family* $\mathcal{F}^{\mathsf{tp}}$ *of all triangle-paths.*

Using this corollary, we can prove one direction of Theorem 1.2.

▶ **Theorem 4.14.** *Let* $\mathcal{F}$ *be a minor-closed family of graphs of unbounded bridge-depth. Then there are graphs in* $\mathcal{F}$ *that have arbitrarily large minimal blocking sets.*

Proof. By Corollary 4.13, $\mathcal{F}$ contains all triangle-paths. Since a triangle-path of length t contains a minimal blocking set of size $t + 2$ by Observation 4.10, the theorem follows. ◄

Theorem 4.14 is phrased for graph families, rather than individual graphs. There is no function h such that $\mathsf{bd}(G) \leq h(\mathsf{mbs}(G))$ for all G: a bipartite grid graph can have arbitrarily large tree-width and therefore bridge-depth, but its minimal blocking sets have size at most two (cf. Lemma 5.2).

5 Bounded bridge-depth implies bounded-size blocking sets

In this section we prove the other direction of Theorem 1.2: minimal blocking sets in a graph G have size at most $2^{\mathsf{bd}(G)}$. We need the following consequence of Kőnig's theorem.

▶ **Lemma 5.1.** *Let* G *be a bipartite graph and let* M *be a maximum matching in* G. *Every maximum independent set of* G *contains all vertices that are not saturated by* M, *and exactly one endpoint of each edge in* M.

Proof. Consider a maximum independent set S in G. Then $\overline{S} := V(G) \setminus S$ is a minimum vertex cover of G. By Kőnig's theorem (cf. [11, Thm. 2.1.1]) we have $|\overline{S}| = |M|$. Since $\overline{S}$ is a vertex cover it contains at least one endpoint of each edge of M; since $|\overline{S}| = |M|$ it contains exactly one endpoint of each edge of M, and no other vertices of G. So the complement S contains all vertices that are not saturated by M, and exactly one endpoint of each edge in M. ◄

The next lemma shows that minimal blocking sets in a bipartite graph have at most two vertices. This was known before, see for example [18, Thm. 14]. Our self-contained proof highlights an additional property of such minimal blocking sets: the two vertices of minimal blocking sets of size two belong to opposite partite sets. This will be crucial later on.

▶ **Lemma 5.2.** *Let* G *be a bipartite graph with partite sets* A *and* B. *If* $Y \subseteq V(G)$ *is a blocking set in* G, *then there is a blocking set* $Y' \subseteq Y$ *in* G *such that one of the following holds:*
- $|Y'| = 1$, *or*
- $Y' = \{a, b\}$ *for some* $a \in A$ *and* $b \in B$.

Proof. Let M be a maximum matching in G, let $V(M)$ be the saturated vertices, and let $U := V(G) \setminus V(M)$ be the unsaturated vertices. Let $R_{A \cap U}$ be the vertices that can be reached by an M-alternating path from $A \cap U$ (which necessarily starts with a non-matching edge). Let $R_{B \cap Y}$ be the vertices that can be reached by an M-alternating path that starts with a *matching edge* from a vertex of $B \cap Y$. Note that both types of alternating paths move from A to B over non-matching edges, and move from B to A over matching edges.

We first deal with some cases in which we easily obtain a blocking set Y' as desired.

Case 1: $A \cap Y \cap R_{A \cap U} \neq \emptyset$. Let $a \in A \cap Y \cap R_{A \cap U}$. Then $a \in A$ can be reached by an M-alternating path P that starts in an unsaturated vertex in the same partite set, implying that P has even length and ends with a matching edge into a. Hence $M' := M \oplus E(P)$, where $\oplus$ denotes the symmetric difference, is a new maximum matching, and it does not saturate $a \in A \cap Y$. Lemma 5.1 applied to M' implies that all maximum independent sets of G contain a, showing that $Y' := \{a\}$ is a blocking set of size one.

Case 2: $B \cap U \cap R_{B \cap Y} \neq \emptyset$. By definition, some $u \in B \cap U$ can be reached by an M-alternating path P that starts in some vertex $b \in B \cap Y$ that belongs to the same partite set. Similarly as in the previous case, $M' := M \oplus E(P)$ is a new maximum matching that does not saturate b, so by Lemma 5.1 applied to M' we conclude that $Y' := \{b\}$ is a blocking set of size one.

Case 3: $A \cap Y \cap R_{B \cap Y} \neq \emptyset$. By definition, some $a \in A \cap Y$ is reachable by an M-alternating path P from some $b \in B \cap Y$, and P starts with a matching edge. Since it ends in the other partite set, it ends with a matching edge as well; hence both a and b are saturated. We claim that $Y' := \{a, b\}$ is a blocking set in G, as desired. Let $a = a_1, b_1, \ldots, a_k, b_k = b$ be the vertices on P, so that $\{a_i, b_i\} \in M$ for all $i \in [k]$ and $\{b_i, a_{i+1}\} \in E(G) \setminus M$ for $i \in [k-1]$. By Lemma 5.1, a maximum independent set in G contains one endpoint of each of the edges $\{a_i, b_i\} \in M$. A maximum independent set avoiding a_1 therefore has to contain b_1, preventing it from containing a_2, forcing it to contain b_2, and so on. Hence a maximum independent set avoiding a_1 contains b_k, proving that $Y' := \{a, b\} = \{a_1, b_k\}$ is a blocking set in G.

Case 4: $B \cap U \cap R_{A \cap U} \neq \emptyset$. Then some unsaturated vertex of A can reach an unsaturated vertex of B by an M-alternating path P. But then M is not a maximum matching since $M \oplus E(P)$ is larger; a contradiction. Hence this case cannot occur.

Assume now that none of the cases above hold. We will conclude the proof of the lemma by deriving a contradiction. Let $R := R_{A \cap U} \cup R_{B \cap Y}$. The following will be useful.

▷ **Claim 5.3.** If $a \in A \cap R$ and $\{a, b\} \in E(G)$, then $b \in B \cap R$.

Proof. By definition, $a \in A \cap R$ implies a is reachable by some M-alternating path P that moves to A over matching edges and moves to B over non-matching edges, such that P starts in a vertex $v \in (A \cup U) \cup (B \cap Y)$. But then b is also reachable by such an alternating path from v: if $\{a, b\} \in M$ then, since P ends at a, edge $\{a, b\}$ must be the last edge of P, so a prefix of P is an M-alternating path reaching b; if $\{a, b\} \notin M$ then appending $\{a, b\}$ to P yields such an M-alternating path. Hence $b \in R$, and $b \in B$ follows since G is bipartite.   ◁

Now consider the following set: $S := (A \cap R) \cup (B \setminus R)$.

We will prove that S is a maximum independent set of G disjoint from Y, contradicting the assumption that Y is a blocking set. To see that S is indeed an independent set, consider any vertex from $A \cap S$, which belongs to $A \cap R$. By Claim 5.3 all neighbors of a belong

to $B \cap R$, and are therefore not contained in S. Hence S is indeed an independent set. To see that it is maximum, by Lemma 5.1 it suffices to argue it contains all of U and one endpoint of each edge in M.

To see that S contains all vertices of $A \cap U$, note that all such vertices are trivially in $R_{A \cap U}$ and therefore in R, implying their presence in $A \cap R$ and therefore in S. To see that S contains all vertices of $B \cap U$, it suffices to show that $B \cap U \cap R = \emptyset$, which follows from the fact that neither Case 2 nor Case 4 is applicable. Hence S contains all vertices of U.

To see that S contains an endpoint of each edge of M, let $\{a, b\} \in M$ be arbitrary with $a \in A$ and $b \in B$. If $b \notin R$ then clearly $b \in S$, as desired. If $b \in R$, then this is witnessed by an alternating path P that reaches b and ends with a non-matching edge. Extending P with the edge $\{a, b\} \in M$ then shows that $a \in R$, so that $a \in A \cap R$ is an endpoint of the edge contained in S.

Hence S is a maximum independent set in G. Since Case 1 and Case 3 do not apply, it follows that $A \cap R \cap Y = \emptyset$, so that $S \cap A$ contains no vertex from Y. Since all vertices of $B \cap Y$ are trivially in $R_{B \cap Y}$ and therefore in R, it follows that $B \setminus R$ contains no vertex from Y. Hence S is a maximum independent set in G disjoint from Y, contradicting the assumption that Y is a blocking set.   ◀

We will use Lemma 5.2 to power the induction step in the proof of the next theorem, which gives the desired upper-bound on the size of minimal blocking sets in terms of bridge-depth. The main idea in the induction step is as follows. For a connected graph G, we consider a tree of bridges T for which $\mathsf{bd}(G \setminus V(T)) < \mathsf{bd}(G)$. We can summarize the relevant ways in which a maximum independent set in G can be composed out of maximum independent sets for the connected components of $G \setminus E(T)$, into a weighted tree T' that is obtained from T by adding a pendant leaf to each vertex. In turn, maximum-weight independent sets in T' correspond to maximum independent sets a bipartite graph obtained from T' by replacing each vertex by a set of false twins. Applying Lemma 5.2 to this bipartite graph points to two vertices that form a blocking set. We translate this back into two components of $G \setminus E(T)$ that are sufficient for constructing a blocking set in G, and apply induction using the fact that $\mathsf{bd}(G \setminus V(T)) < \mathsf{bd}(G)$.

▶ **Theorem 5.4** (⋆)**.** *Let G be a graph and $Y_G \subseteq V(G)$ a blocking set in G. There is a blocking set $Y_G' \subseteq Y_G$ in G of size at most $2^{\mathsf{bd}(G)}$.*

Note that Theorem 5.4 and Theorem 4.14 together prove Theorem 1.2. We finish the section by showing that the upper-bound of $2^{\mathsf{bd}(G)}$ on the size of minimal blocking sets is tight.

▶ **Theorem 5.5.** *For every $c \in \mathbb{N}$, there is a graph G with $\mathsf{bd}(G) \leq c$ that contains a minimal blocking set of size 2^c.*

Proof. Recall the notion of triangle-path from Definition 4.9. For $t \geq 2$, let a *truncated triangle-path of length t* be the graph U_t obtained from a triangle-path of length t by removing vertices a_1 and b_t; see Figure 2. Analogously to Observation 4.10, we show that $Y_t := \{c_i \mid i \in [t]\}$ is a minimal blocking set in U_t. Since Y_t is an independent set of size t, while (the remainders of) the triangles in U_t partition the vertices of U_t into t cliques, it follows that $\alpha(U_t) = t$. The set Y_t is a blocking set, since $U_t \setminus Y_t$ is a path on $2(t-1)$ vertices, whose independence number is only $t-1$. Finally, it is easy to see that for any $y \in Y_t$, there is a size-t independent set in $U_t \setminus (Y_t \setminus y)$ that consists of the vertex y and, for every (remainder of a) triangle in U_t, the vertex closest to y.
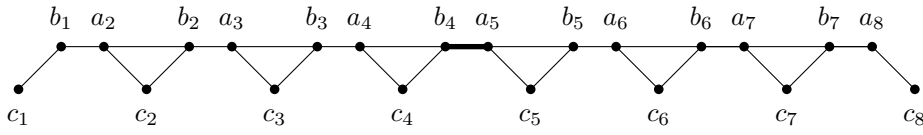
Figure 2 Truncated triangle path U_8 of length 8, illustrating Theorem 5.5. Removing the fat middle bridge and its incident vertices, leaves two connected components isomorphic to U_4.

Hence U_t has a minimal blocking set of size t, for all $t \geq 2$. To prove the theorem, it therefore suffices to show that $\mathsf{bd}(U_{2^c}) \leq c$ for all $c \in \mathbb{N}$. We prove this by induction on c. For $c = 1$, note that the graph U_2 is just the four-vertex path. Hence it is a forest, implying $\mathsf{bd}(U_2) = 1$ by Proposition 3.2. For $c > 1$, consider the graph U_{2^c}. By construction, the middle edge $e = \{b_{2^{c-1}}, a_{2^{c-1}+1}\}$ is a bridge in U_{2^c}. Let T be the tree in U_{2^c} consisting of the single bridge e. Note that removing $V(T)$ splits U_{2^c} evenly, into two connected components that are both isomorphic to $U_{2^{c-1}}$. By induction, $\mathsf{bd}(U_{2^{c-1}}) \leq c - 1$. Then Proposition 3.2 shows that $\mathsf{bd}(U_{2^c}) \leq 1 + \mathsf{bd}(U_t \setminus V(T)) = 1 + (c - 1) = c$. ◀

6   Kernelization for modulators to bounded bridge-depth

To establish the positive direction of Theorem 1.1, we develop a polynomial kernel for VERTEX COVER parameterized by the size of a modulator X whose removal leaves a graph of constant bridge-depth; an approximately optimal such set X can be computed using Proposition 3.5. As the kernelization is technical and consists of many different reduction rules, with a nontrivial size analysis, the material is deferred to the full version [4]. In this limited space, we present the high-level idea behind the kernelization and the role of bridge-depth.

Consider an instance (G, k) of VERTEX COVER with a modulator X such that $\mathsf{bd}(G \setminus X) \in \mathcal{O}(1)$. As explained in the introduction, using the fact that minimal blocking sets for the components C of $G \setminus X$ have bounded size, the number of such components can easily be bounded by $|X|^{\mathcal{O}(1)}$. To bound the size of individual components, the definition of bridge-depth ensures that in each connected component C of $G \setminus X$ there is a tree of bridges $T \subseteq E(C)$ (called a *lowering tree*) such that removing the vertex set $V(T)$ from C decreases the bridge-depth of C. By designing new problem-specific reduction rules, we shrink the tree of bridges to size polynomial in the parameter. This is where the main technical work of the kernelization step lies. It properly subsumes the earlier kernelization for the parameterization by distance to a forest, which is imported as a black box in all previous works [5, 17–19, 27]. Having bounded the number of components of $G \setminus X$, together with the size of a lowering tree of bridges in each component, we now proceed as follows: in each component C of $G \setminus X$ we move the vertices from a lowering tree of bridges into the set X. This blows up $|X|$ by a polynomial factor, but strictly decreases the bridge-depth of the graph $G \setminus X$. We then recursively kernelize the resulting instance. When the bridge-depth of $G \setminus X$ reaches zero, the graph $G \setminus X$ is empty and the kernelization is completed. Full details can be found in the full version [4].

The negative direction of Theorem 1.1 is much easier to establish. Using the fact that a minor-closed family $\mathcal{F}$ of unbounded bridge-depth contains all triangle paths, a kernelization lower bound for modulators to such $\mathcal{F}$ follows easily using known gadgets.

7   Conclusion

In this paper we introduced the graph parameter bridge-depth and used it to characterize the minor-closed graph classes $\mathcal{F}$ for which VERTEX COVER parameterized by $\mathcal{F}$-modulator has a polynomial kernel. It would be interesting to see whether the characterization can be extended to subgraph-closed or even hereditary graph classes. If a characterization exists of the hereditary graph classes whose modulators lead to a polynomial kernel, it will likely not be as clean as Theorem 1.1: it will have to deal with the fact that bipartite graphs can be arbitrarily complex in terms of width parameters, while bipartite modulators allow for a polynomial kernel. Hence such a characterization has to capture parity conditions of $\mathcal{F}$.

A natural attempt to generalize our approach to deal with bipartite graphs is to consider the following parameter, which we call *bipartite-contraction-depth*: we mimic the definition of bridge-depth (cf. Definition 3.1), except that we redefine the graph G_{CB} to be the graph obtained from G by *simultaneously* contracting all edges that do *not* lie on an odd cycle. Note that bipartite-contraction-depth generalizes bridge-depth, in the sense that bridges do not lie on an odd cycle, and that the bipartite-contraction-depth of a graph with an odd cycle transversal of size k is at most $k + 1$. Having defined this parameter, we would need, in order to obtain a statement similar to Theorem 4.14, that large bipartite-contraction-depth implies the existence of structures that allow to obtain kernel lower bounds, similarly to the fact that large bridge-depth implies the existence of large triangle-paths (cf. Corollary 4.13). The appropriate structure here seems to be an *odd-cycle-path of length* t, defined as a set of t vertex-disjoint odd cycles $C_1, \ldots, C_t$, and a set of $t - 1$ vertex-disjoint paths (of any length) connecting C_i to C_{i+1} for $i \in [t - 1]$, in such a way that for every $i \in \{2, \ldots, t - 1\}$, the two attachment vertices in C_i are distinct. Now the expected property would be that large bipartite-contraction-depth forces long odd-cycle-paths. Unfortunately, this is not true. Indeed, consider the *Escher wall of size* h depicted in [33, Fig. 3]. It is proved in [33] that this graph does not contain two vertex-disjoint odd cycles, but a smallest hitting set for odd cycles has size h. Since there are no two vertex-disjoint cycles, a longest odd-cycle-path has length one. On the other hand, it can be easily verified that an Escher wall of size h has bipartite-contraction-depth $\Omega(h)$. Informally, this can be seen by noting that, initially, all edges lie on an odd cycle, hence a vertex removal is required, and that each such removal cascades in a constant number of contractions until all edges lie again on an odd cycle. Since a smallest hitting set for odd cycles of an Escher wall of size h has size h, the claimed bound follows. Therefore, summarizing this discussion, if one aims at a result similar to Theorem 1.1 that also applies to families $\mathcal{F}$ containing bipartite graphs, it seems that significant new ideas are required.

Another open research direction consists of a further algorithmic exploration of the merits of bridge-depth. We expect that several polynomial-space fixed-parameter tractable algorithms that work for graphs of bounded tree-depth [7, 32] can be extended to work with bridge-depth instead. Which other ways to enrich the recursive definition of tree-depth lead to novel algorithmic insights? As for kernelization purposes, it is plausible that bridge-depth also characterizes the existence of polynomial kernels for other problems other than VERTEX COVER, parameterized by the vertex-deletion distance of the input graph to a minor-closed graph class.

───── **References** ─────

1   Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996. `doi:10.1137/S0097539793251219`.

2   Hans L. Bodlaender. A partial *k*-arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2):1–45, 1998. `doi:10.1016/S0304-3975(97)00228-4`.

3   Hans L. Bodlaender. Kernelization, exponential lower bounds. In *Encyclopedia of Algorithms*, pages 1013–1017. Springer, 2016. `doi:10.1007/978-1-4939-2864-4_521`.

4   Marin Bougeret, Bart M. P. Jansen, and Ignasi Sau. Bridge-Depth Characterizes which Structural Parameterizations of Vertex Cover Admit a Polynomial Kernel. *CoRR*, abs/2004.12865, 2020. `arXiv:2004.12865`.

5   Marin Bougeret and Ignasi Sau. How much does a treedepth modulator help to obtain polynomial kernels beyond sparse graphs? *Algorithmica*, 81(10):4043–4068, 2019. `doi:10.1007/s00453-018-0468-8`.

6   Jannis Bulian and Anuj Dawar. Graph isomorphism parameterized by elimination distance to bounded degree. *Algorithmica*, 75(2):363–382, 2016. `doi:10.1007/s00453-015-0045-3`.

7   Li-Hsuan Chen, Felix Reidl, Peter Rossmanith, and Fernando Sánchez Villaamil. Width, depth, and space: Tradeoffs between branching and dynamic programming. *Algorithms*, 11(7):98, 2018. `doi:10.3390/a11070098`.

8   Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms.* Springer, 2015. `doi:10.1007/978-3-319-21275-3`.

9   Marek Cygan, Daniel Lokshtanov, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. On the hardness of losing width. *Theory Comput. Syst.*, 54(1):73–82, 2014. `doi:10.1007/s00224-013-9480-1`.

10 Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *J. ACM*, 61(4):23:1–23:27, 2014. `doi:10.1145/2629620`.

11 Reinhard Diestel. *Graph Theory.* Springer-Verlag, Heidelberg, 5th edition, 2016.

12 Rod G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity.* Texts in Computer Science. Springer, 2013.

13 Michael R. Fellows, Lars Jaffke, Aliz Izabella Király, Frances A. Rosamond, and Mathias Weller. What is known about vertex cover kernelization? In Hans-Joachim Böckenhauer, Dennis Komm, and Walter Unger, editors, *Adventures Between Lower Bounds and Higher Altitudes - Essays Dedicated to Juraj Hromkovič on the Occasion of His 60th Birthday*, volume 11011 of *Lecture Notes in Computer Science*, pages 330–356. Springer, 2018. `doi:10.1007/978-3-319-98355-4_19`.

14 Jörg Flum and Martin Grohe. *Parameterized Complexity Theory.* Springer-Verlag, 2006. `doi:10.1007/3-540-29953-X`.

15 Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, Geevarghese Philip, and Saket Saurabh. Hitting forbidden minors: Approximation and kernelization. *SIAM J. Discrete Math.*, 30(1):383–410, 2016. `doi:10.1137/140997889`.

16 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: Theory of Parameterized Preprocessing.* Cambridge University Press, 2019. `doi:10.1017/9781107415157`.

17 Fedor V. Fomin and Torstein J. F. Strømme. Vertex cover structural parameterization revisited. In *Proc. 42nd WG*, volume 9941 of *LNCS*, pages 171–182, 2016. `doi:10.1007/978-3-662-53536-3_15`.

18 Eva-Maria C. Hols and Stefan Kratsch. Smaller parameters for vertex cover kernelization. In *Proc. 12th IPEC*, volume 89 of *LIPIcs*, pages 20:1–20:12, 2017. `doi:10.4230/LIPIcs.IPEC.2017.20`.

19 Eva-Maria C. Hols, Stefan Kratsch, and Astrid Pieterse. Elimination distances, blocking sets, and kernels for vertex cover. In *Proc. 37nd STACS*, volume 154 of *LIPIcs*, pages 36:1–36:14, 2020. `doi:10.4230/LIPIcs.STACS.2020.36`.

20   John E. Hopcroft and Robert Endre Tarjan. Efficient algorithms for graph manipulation (algorithm 447). *Commun. ACM*, 16(6):372–378, 1973. `doi:10.1145/362248.362272`.

21   Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited: Upper and lower bounds for a refined parameter. In *Proc. 28th STACS*, volume 9 of *LIPIcs*, pages 177–188, 2011. `doi:10.4230/LIPIcs.STACS.2011.177`.

22   Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited - upper and lower bounds for a refined parameter. *Theory Comput. Syst.*, 53(2):263–299, 2013. `doi:10.1007/s00224-012-9393-4`.

23   Bart M. P. Jansen and Astrid Pieterse. Polynomial kernels for hitting forbidden minors under structural parameterizations. In *Proc. 26th ESA*, volume 112 of *LIPIcs*, pages 48:1–48:15, 2018. `doi:10.4230/LIPIcs.ESA.2018.48`.

24   Stefan Kratsch. A randomized polynomial kernelization for vertex cover with a smaller parameter. *SIAM J. Discrete Math.*, 32(3):1806–1839, 2018. `doi:10.1137/16M1104585`.

25   Stefan Kratsch and Magnus Wahlström. Representative sets and irrelevant vertices: New tools for kernelization. In *Proc. 53rd FOCS*, pages 450–459, 2012. `doi:10.1109/FOCS.2012.46`.

26   Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Kernelization - Preprocessing with a guarantee. In *The Multivariate Algorithmic Revolution and Beyond*, pages 129–161, 2012. `doi:10.1007/978-3-642-30891-8_10`.

27   Diptapriyo Majumdar, Venkatesh Raman, and Saket Saurabh. Polynomial kernels for vertex cover parameterized by small degree modulators. *Theory Comput. Syst.*, 62(8):1910–1951, 2018. `doi:10.1007/s00224-018-9858-1`.

28   George L. Nemhauser and Leslie E. Trotter Jr. Vertex packings: structural properties and algorithms. *Math. Program.*, 8:232–248, 1975. `doi:10.1007/BF01580444`.

29   Jaroslav Nesetril and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. `doi:10.1007/978-3-642-27875-4`.

30   Jaroslav Nesetril and Patrice Ossona de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *Eur. J. Comb.*, 27(6):1022–1041, 2006. `doi:10.1016/j.ejc.2005.01.010`.

31   Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006. `doi:10.1093/acprof:oso/9780198566076.001.0001`.

32   Michal Pilipczuk and Marcin Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. *ACM TOCT*, 9(4):18:1–18:36, 2018. `doi:10.1145/3154856`.

33   Dieter Rautenbach and Bruce A. Reed. The Erdös-Pósa Property for Odd Cycles in Highly Connected Graphs. *Combinatorica*, 21(2):267–278, 2001. `doi:10.1007/s004930100024`.

34   Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986. `doi:10.1016/0095-8956(86)90030-4`.

35   Neil Robertson and Paul D. Seymour. Graph minors. XX. Wagner's conjecture. *J. Comb. Theory, Ser. B*, 92(2):325–357, 2004. `doi:10.1016/j.jctb.2004.08.001`.

36   Wouter Cames van Batenburg, Tony Huynh, Gwenaël Joret, and Jean-Florent Raymond. A tight Erdös-Pósa function for planar minors. In *Proc. 30th SODA*, pages 1485–1500. SIAM, 2019. `doi:10.1137/1.9781611975482`.