Component Order Connectivity Admits No Polynomial Kernel Parameterized by the Distance to Subdivided Comb Graphs

Jakob Greilhuber
TU Wien, Austria
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
Saarbrücken Graduate School of Computer Science, Germany

Roohani Sharma
Department of Informatics, University of Bergen, Norway

Abstract

In this paper we show that the d-Component Order Connectivity (d-COC) problem parameterized by the distance to subdivided comb graphs (`dist-to-subdivided-combs`) does not admit a polynomial kernel, unless `NP` $\subseteq$ `coNP/poly`.

The d-COC problem is a generalization of the classical Vertex Cover problem. An instance of the d-COC problem consists of an undirected graph G and a positive integer k, and the question is whether there exists a set $S \subseteq V(G)$ of size at most k, such that each connected component of $G - S$ contains at most d vertices. When $d = 1$, d-COC is the Vertex Cover problem.

Vertex Cover is a ubiquitous problem in parameterized complexity, and it admits a kernel with $O(k^2)$ edges and $O(k)$ vertices, which is tight [Dell & van Melkebeek, JACM 2014]. Our result is inspired by the work of Jansen & Bodlaender [TOCS 2013], who gave the first polynomial kernel for Vertex Cover where the parameter is provably smaller or equal to the standard parameter, solution size k. They used `fvs`, the feedback vertex set number, as the parameter. In this work, we show that unlike most other existing results or techniques for kernelization of Vertex Cover that generalize to d-COC, this is not the case when `dist-to-subdivided-combs`, which is at least as large as `fvs`, is the parameter.

Our lower bound is achieved in two stages. In the first stage we extend the result of Hols, Kratsch & Pieterse [SIDMA 2022] where they show that if a graph family $\mathcal{C}$, which is closed under taking disjoint unions, has unbounded "blocking set" size, then Vertex Cover does not admit a polynomial kernel parameterized by the size of a vertex modulator to $\mathcal{C}$, unless `NP` $\subseteq$ `coNP/poly`. We show that a similar sufficient condition for proving the non-existence of polynomial kernels also holds for d-COC. In the second stage, we show that when $\mathcal{C}$ is the family of subdivided comb graphs, contrary to Vertex Cover, where the size of minimal blocking sets of graphs in $\mathcal{C}$ is at most two [Jansen & Bodlaender, STACS 2011], the size of minimal blocking sets of graphs in $\mathcal{C}$ for the d-COC problem can be arbitrarily large. This yields the desired lower bound. In addition to this we also show that when $\mathcal{C}$ is a class of paths, then it still has blocking sets of size at most two for d-COC, indicating that polynomial kernels might be achievable when the parameter is the size of a vertex modulator to the class of disjoint unions of paths (linear forests).

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms

Keywords and phrases Component Order Connectivity, Kernelization, Structural Parameterizations, Feedback Vertex Set, Vertex Cover, Blocking Sets, Subdivided Comb Graphs

Digital Object Identifier 10.4230/LIPIcs.IPEC.2024.21

Acknowledgements The majority of the work was done at the Max Planck Institute of Informatics during a summer internship of Jakob Greilhuber. We thank the anonymous reviewers that reviewed a previous version of this paper. The presentation of this work is based on their comments. Jakob Greilhuber thanks the organizers of ALGO 2024 for their financial support.

19th International Symposium on Parameterized and Exact Computation (IPEC 2024).
Editors: Édouard Bonnet and Paweł Rzążewski; Article No. 21; pp. 21:1–21:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In this work we explore kernelization with respect to structural parameterizations for a well-studied generalization of the ubiquitous Vertex Cover problem. In the Vertex Cover problem the input is an undirected graph G and the goal is to find a minimum set of vertices of G, say S, such that $G - S$ has no edges, or equivalently, every connected component of G contains only a single vertex. Vertex Cover has garnered a lot of attention since the advent of parameterized complexity [17, 36, 6, 10, 19, 7, 13, 39, 33, 35, 28, 1, 2, 9, 16, 37, 38, 40, 8, 24, 20, 5, 26, 4] and has resulted in the inception of new tools and ideas that have proved useful for several other problems as well [6, 10, 19, 7, 31, 32].

Component Order Connectivity. One of the problems that has been at the receiving end of applications and generalizations of the tools developed for Vertex Cover (especially in the context of parameterized algorithms and kernelization), is the d-Component Order Connectivity (d-COC) problem. In this problem, d is a fixed integer, the input is an undirected graph and the goal is to find a minimum set of vertices whose deletion results in a graph in which every connected component has at most d vertices. Note that 1-COC is precisely the Vertex Cover problem.[1] Vertex Cover parameterized by the solution size k admits a kernel with $O(k^2)$ edges and $O(k)$ vertices, which is tight [14]. Starting from the Buss kernelization [6], all the way up to crown decomposition based rules [10, 19] or the (Weighted) Expansion Lemma [32], and the LP-based reduction rules using the Nemhauser-Trotter theorem [7, 31, 32]; all these techniques for getting a polynomial kernel for Vertex Cover parameterized by the solution size extend beautifully to the d-COC problem parameterized by the solution size [18, 32, 41, 21].

Amidst the kernelization results for Vertex Cover parameterized by the solution size, Jansen and Bodlaender [28] initiated the study of kernelization for Vertex Cover with a refined parameter, that is at most as large as the solution size. In particular, they studied the size of a minimum feedback vertex set of the graph, denoted by `fvs`, as the parameter, and proved that the problem admits a polynomial kernel. Observe that `fvs` could potentially be much smaller than the size of a minimum vertex cover. The work by Jansen and Bodlaender [28] led to several new and interesting insights about the problem. A next natural question is: can the result of Jansen and Bodlaender [28] be extended to the d-COC problem? In work that is independent of ours, this relevant question was negatively answered by Donkers and Jansen [15], unless $\mathtt{NP} \subseteq \mathtt{coNP/poly}$, no polynomial kernel exists for d-COC parameterized by `fvs`, for all $d \geq 2$. We study the problem with a parameter that is at least as large as `fvs`, the *distance to subdivided comb graphs* (`dist-to-subdivided-combs`). We prove that for all integers $d \geq 2$, even when parameterizing by the distance to subdivided comb graphs, d-COC does not admit a polynomial kernel, unless $\mathtt{NP} \subseteq \mathtt{coNP/poly}$. Note that, while our result is not explicitly stated as a result in [15], it may be possible to also obtain it from their proofs with some careful considerations.

For completeness, we also remark that Courcelle's theorem (see for example [12, Section 7.4] or [11]) can be used to show that the problem is `FPT`, even when parameterizing by `fvs`. Besides our work and the work by Donkers and Jansen [15] there are some other publications that study d-COC kernelization with structural parameters. For instance, Bhyravarapu et al. [3] parameterize the problem by the size of a set that is a solution to the c-COC problem, where $c \geq d$, and Jansen and Pieterse [29] study a generalization of the problem parameterized by the size of a treedepth-η modulator.

---

[1] Note that other than being a generalization of Vertex Cover, d-COC is an interesting problem in its own right. See [23].

1.1 Our results and methods

As mentioned above, we show that d-COC does not admit a polynomial kernel parameterized by the *distance to subdivided comb graphs*. A *comb graph* is a graph that consists of a central path and for each vertex of this central path, there is a unique pendant vertex adjacent to it. A *subdivided comb graph* is obtained from a comb graph by repeatedly subdividing edges. The *distance to subdivided comb graphs* is the smallest integer k such that there is a vertex set of size k whose deletion results in a disjoint union of subdivided comb graphs. Therefore, distance to subdivided comb graphs is at least as large as `fvs`. The *hairlength* of a subdivided comb graph is the maximum number of vertices of any path attached to the central path minimized over all possible choices for the central path.

▶ **Theorem 1.** *For any integer $d \geq 2$, d-Component Order Connectivity does not admit a polynomial compression when the parameter is distance to subdivided comb graphs, unless* $\mathtt{NP} \subseteq \mathtt{coNP/poly}$.

Note that showing the non-existence of a polynomial compression is stronger than showing the non-existence of a polynomial kernel (see Section 2).

We prove Theorem 1 by generalizing the notion of *blocking sets*, which was formally introduced by Hols, Kratsch & Pieterse [26] for the Vertex Cover problem,[2] to d-COC. A blocking set of a graph G, for Vertex Cover, is a subset Y of the vertices of G, such that *no* minimum vertex cover contains Y. Hence, a blocking set "blocks" itself from being extended into an optimal solution. Of special relevance are *minimal* blocking sets, which are blocking sets such that no strict subset of them is a blocking set. This concept of (minimal) blocking sets can be generalized to d-COC for any integer $d \geq 2$ in the natural way. Concretely, define a *d-blocking set* of G as a subset Y of the vertices of G such that there is no minimum solution of d-COC that contains Y.

One of the key results by Hols et al. [26] gives a sufficient condition for proving kernelization lower bounds for structural parameterizations for Vertex Cover. They show that if a graph class $\mathcal{C}$, which is closed under taking the disjoint union, contains a graph with a minimal blocking set of size α, then no kernel of size $\mathcal{O}(|X|^{\alpha-\varepsilon})$ can exist for all $\varepsilon > 0$, when parameterizing by the size of a smallest vertex modulator X to $\mathcal{C}$, unless $\mathtt{NP} \subseteq \mathtt{coNP/poly}$. Hence, when the size of minimal blocking sets of graphs in $\mathcal{C}$ is unbounded, then one can rule out a polynomial kernel altogether parameterized by $|X|$.

We extend this result by Hols et al. [26] to the d-COC problem, showing that a similar property, with minimal d-blocking sets, holds for d-COC, for all integers $d \geq 1$ (Theorem 4). We later show that even though the same condition implies kernelization lower bounds for both Vertex Cover and d-COC, the family of forests behaves differently for Vertex Cover and d-COC: no forest has a minimal 1-blocking set of size more than two (implicitly proved in [27], the conference version of [28]), whereas there are forests (in fact subdivided comb graphs) with arbitrarily large minimal d-blocking sets for every $d \geq 2$ (Lemma 5).

For the first part (as stated above), we provide a linear parameter transformation from the α-Hitting Set problem, which is the classical Hitting Set problem in which every set that needs to be hit has cardinality exactly α, and the size of the universe $\mathcal{U}$ is the parameter. For this problem, it is known that no polynomial compression with size $\mathcal{O}(|\mathcal{U}|^{\alpha-\varepsilon})$ exists for all $\varepsilon > 0$ and all integers $\alpha \geq 2$, unless $\mathtt{NP} \subseteq \mathtt{coNP/poly}$ [14, Theorem 2].

[2] The notion of blocking sets had (implicitly) appeared in a lot of other kernelization results for Vertex Cover like [28, 30, 22, 34, 5, 3].

In the parameter transformation, one creates a "universe vertex" for each element of the universe of the input α-Hitting Set instance. Then, for each set F that needs to be hit in the input instance, one creates multiple copies of a graph H of $\mathcal{C}$ that has a minimal d-blocking set of size α. Each such copy is associated with the set F, and, each vertex of the minimal d-blocking set is associated with a different element of F. Then, each vertex of the minimal d-blocking set is connected to the universe vertex corresponding to the associated element. This way, it is ensured that the set of all universe vertices is a modulator to $\mathcal{C}$, and it has the same size as the universe of the input α-Hitting Set instance. Moreover, if the bound on the size of the solution is carefully chosen, one can ensure that the d-COC instance is a yes-instance if and only if the input instances is a yes-instance. The intuition behind this is that a solution to the input instance can be easily mapped to a small d-COC solution by choosing the universe vertices of the hitting set, and by extending this selection to minimum solutions within the copies of H. On the other hand, if there is a small d-COC solution, it must be the case that the selection within the universe vertices corresponds to a small hitting set of the input instance. Otherwise, one would necessarily have to put all vertices of the minimal d-blocking set into the solution for many copies of H, which would make the solution too large.

To complete the proof of Theorem 1, for all integers $N \geq 1, d \geq 2$, we then define a blocking set graph $T_{N,d}$, which is a subdivided comb graph with hairlength at most $d+1$, such that this graph has a minimal d-blocking set of size N. Note that providing graphs with large d-blocking sets is easy, the difficulty of the construction stems from the fact that the sets must also be minimal. Combined with the previous result, this proves that when $\mathcal{C}$ is a graph class closed under taking the disjoint union that contains the blocking set graph $T_{N,d}$ for arbitrarily large N, then d-COC parameterized by the size of a smallest modulator to $\mathcal{C}$ does not have a polynomial compression, unless $\mathtt{NP} \subseteq \mathtt{coNP/poly}$.

Contrasting the fact that the class of subdivided comb graphs has minimal d-blocking sets of unbounded size when $d \geq 2$, we show that this is not the case for path graphs. We show that for paths, any minimal blocking set has size at most two. Hence, there is still hope for obtaining a polynomial kernel when $\mathcal{C}$ only contains disjoint unions of paths.

2 Preliminaries

For $n, m \in \mathbb{Z}$, by $[n, m]$ we denote the set $\{x \in \mathbb{Z} \mid n \leq x \leq m\}$. For $n, m \in \mathbb{Z}$, $k \in \mathbb{Z}^+$, we write $\{n, n+k, \dots, m\}$ for the set $\{x \in \mathbb{Z} \mid n \leq x \leq m, x \equiv n \pmod{k}\}$. In particular, $\{0, 2, \dots, 0\} = \{0\}$.

Graphs. A graph G is a pair (V, E), where V is a set and $E \subseteq \binom{V}{2}$. For a graph G, we also refer to the vertex set of G as $V(G)$ and the edge set of G as $E(G)$. In this work all considered graphs are finite, undirected, and contain no multi-edges or self-loops. For a graph G and $v \in V(G)$, the (open) neighbourhood of v is defined as $N_G(v) \coloneqq \{u \mid vu \in E(G)\}$. We may omit the subscript G if the used graph is clear from the context. Graph H is a subgraph of graph G if and only if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. For a graph G and vertex set $X \subseteq V(G)$, $G[X]$ denotes the graph induced by X, that is, $G[X] = H$, where H is the subgraph of G with $V(H) = X$ and $E(H) = \{uv \mid uv \in E(G), u \in X, v \in X\}$. For graph G and $X \subseteq V(G)$, we write $G - X$ for the graph $G[V(G) \setminus X]$. Given a graph G and edge $uv \in E(G)$, the operation of subdividing edge uv results in the graph G' with $V(G') = V(G) \cup \{w\}$ and $E(G') = (E(G) \setminus \{uv\}) \cup \{uw, wv\}$, where w is a new vertex not in $V(G)$.

Kernelization and compression. A parameterized problem Q is a subset of $\Sigma^* \times \mathbb{N}$ for a finite alphabet Σ. An instance $(I, k) \in \Sigma^* \times \mathbb{N}$ of problem Q is a yes-instance if $(I, k) \in Q$, and a no-instance otherwise. The value k of instance (I, k) is the parameter of the instance. A *kernelization algorithm* takes an instance (I, k) of a parameterized problem Q as input, and outputs an equivalent instance (I', k') of Q such that $|I'| + k' \leq f(k)$ for some computable function f. The runtime of the algorithm must be polynomial in $|I| + k$. Such an algorithm is often simply called a *kernel*, and a *polynomial kernel* if f is a polynomial function. A *polynomial compression* is defined similarly to kernelization, except that the output instance I' may be of a different unparameterized problem P. Note that if a problem has a polynomial kernel, it also has a polynomial compression, and hence, ruling out polynomial compressions also rules out polynomial kernels. We refer to the textbooks [12, 21] for more details on these concepts.

Kernelization lower bounds. In this paper, we use a well-established technique for obtaining kernelization lower bounds. It resembles classical reductions, and is fittingly called *linear parameter transformation* (see [25]).

▶ **Definition 2** (linear parameter transformation)**.** *Let P and Q be parameterized problems. A* linear parameter transformation *is an algorithm that transforms an instance (I, k) of P into an equivalent instance (I', k') of Q in time polynomial in $|I| + k$, such that $k' \in O(k)$.*

When Q has a polynomial compression of size $O(p^\alpha)$ for parameter p, and there is linear parameter transformation from P with parameter k to Q with parameter p, then by using the transformation, one obtains a polynomial compression for P of size $O(k^\alpha)$. Hence, if we know that P does not have a polynomial compression of size $O(k^\alpha)$, then neither can Q of size $O(p^\alpha)$.

3 Large minimal blocking sets as bottlenecks for polynomial kernels for d-COC

In this section, we extend the sufficient condition for proving kernelization lower bounds for Vertex Cover given by Hols et al. [26], to the d-COC problem. Concretely, we will define the notion of minimal blocking sets for d-COC, and we will show that no graph class, that is closed under taking disjoint unions and has blocking sets of unbounded size, admits a polynomial compression, unless `NP` $\subseteq$ `coNP/poly`. We begin by stating the definition of blocking sets [26, Definition 2.3] for d-COC. For a graph G, we call $S \subseteq V(G)$ a d-`coc` set of G if every connected component of $G - S$ has at most d vertices.

▶ **Definition 3** (Blocking sets)**.** *Let G be a graph and d be a positive integer. A set $Y \subseteq V(G)$ is a d-blocking set if Y is not a subset of any minimum d-`coc` set of G. A d-blocking set Y is a minimal d-blocking set if no set $Y' \subsetneq Y$ is a d-blocking set. That is, for every $y \in Y$, there exists a minimum d-`coc` set of G that contains $Y \setminus \{y\}$.*

Given a graph class $\mathcal{C}$ and a graph G, a (vertex) modulator to $\mathcal{C}$ is a set $X \subseteq V(G)$ such that $G - X$ is in $\mathcal{C}$. Hols et al. [26, Theorem 1.1] show that if $\mathcal{C}$ is a graph class that is closed under taking disjoint unions, and $\mathcal{C}$ contains a graph with a minimal 1-blocking set of size α, then, Vertex Cover parameterized by the size of a vertex modulator X to $\mathcal{C}$ does not have a kernel of size $\mathcal{O}(|X|^{\alpha-\varepsilon})$ for all $\varepsilon > 0$, unless `NP` $\subseteq$ `coNP/poly`. This powerful theorem allows one to derive kernelization lower bounds for various structural parameterizations in a convenient way.

We now extend the above result to the d-COC problem. The proof we present is inspired by that of Hols et al. [26, Theorem 1.1]. We mention the key difference between our reduction and the reduction by Hols et al. after presenting our reduction.

The proof is done via a linear parameter transformation from α-Hitting Set parameterized by the universe size, which is the problem defined below, to the d-COC problem parameterized by the size of a vertex modulator to a graph class $\mathcal{C}$ that is closed under taking the disjoint union and that contains a graph with a minimal d-blocking set of size $\alpha \geq 2$. The α-Hitting Set problem is the same as the classical Hitting Set problem, except that every set that needs to be hit has cardinality exactly α, and that the parameter is the size of the universe.

α-Hitting Set
Input: Universe $\mathcal{U}$, family of subsets $\mathcal{F} \subseteq 2^{\mathcal{U}}$ such that $|F| = \alpha$ for all $F \in \mathcal{F}$, integer $k \geq 0$
Parameter: $|\mathcal{U}|$
Question: Is there a set $H \subseteq \mathcal{U}$ such that, for all $F \in \mathcal{F}$, $H \cap F \neq \emptyset$ and $|H| \leq k$?

A lower bound for α-Hitting Set given by Dell and van Melkebeek [14, Theorem 2] shows that there is no polynomial compression of size $\mathcal{O}(|\mathcal{U}|^{\alpha-\varepsilon})$ for α-Hitting Set when $\alpha \geq 2$, for any $\varepsilon > 0$, unless $\texttt{NP} \subseteq \texttt{coNP/poly}$. We denote the problem d-COC where the parameter is the size of a modulator to $\mathcal{C}$, for some graph class $\mathcal{C}$, as d-COC/$\texttt{dist-to-}\mathcal{C}$. Note that we assume that a modulator to $\mathcal{C}$ is given together with the input.

▶ **Theorem 4.** *Let $\mathcal{C}$ be a class of graphs that is closed under disjoint union, and d be a positive integer. If there is a graph in $\mathcal{C}$ with a minimal d-blocking set of cardinality $\alpha \geq 2$, then there exists no polynomial compression for d-COC/$\texttt{dist-to-}\mathcal{C}$ of size $\mathcal{O}(|X|^{\alpha-\varepsilon})$, where X is the modulator to $\mathcal{C}$ provided with the input, for all $\varepsilon > 0$, unless* $\texttt{NP} \subseteq \texttt{coNP/poly}$.

Proof. Assume that there exists a graph H in $\mathcal{C}$ such that H has a minimal d-blocking set of size $\alpha \geq 2$. Moreover, let OPT denote the size of a minimum d-$\texttt{coc}$ set of H. We prove the polynomial compression lower bound by a linear parameter transformation from α-Hitting Set. Let $(\mathcal{U}, \mathcal{F}, k)$ be the input instance of α-Hitting Set. If $k \geq |\mathcal{U}|$, we have a trivial yes-instance at our hand, and we output a constant size yes-instance of d-COC/$\texttt{dist-to-}\mathcal{C}$. Otherwise, we will output a graph G and a positive integer k' such that G has a d-$\texttt{coc}$ set of size at most k' if and only if the input instance of α-Hitting Set is a yes-instance.

The creation of the graph G is covered next. For each $F \in \mathcal{F}$, add $k+\alpha\cdot(d-1)+1$ copies of the graph H to G, and denote these copies as $H_F^1, \ldots, H_F^{k+\alpha\cdot(d-1)+1}$. Since each copy of H has at least one minimal d-blocking set of size α, for graph H_F^i denote an arbitrary such d-blocking set as B_F^i. Next, add a vertex v_u to the graph for each $u \in \mathcal{U}$. These vertices are called the *universe vertices*, and we denote all of them as $V_{\mathcal{U}}$, that is, $V_{\mathcal{U}} \coloneqq \{v_u \mid u \in \mathcal{U}\}$. We also define $V_F \coloneqq \{v_u \mid u \in F\}$ for all $F \in \mathcal{F}$.

Now, select an arbitrary bijective function $f_{F,i} : V_F \to B_F^i$ for all $F \in \mathcal{F}$ and $i \in [1, k+\alpha\cdot(d-1)+1]$. Next, add an edge between vertex v_u and vertex $f_{F,i}(v_u)$ for all $F \in \mathcal{F}$, $u \in F$ and $i \in [1, k+\alpha\cdot(d-1)+1]$. In other words, we associate each vertex of d-blocking set B_F^i with a different universe vertex corresponding to an element in F, and add an edge between the associated vertices. This concludes the construction of the graph G.

Observe that $G - V_{\mathcal{U}}$ is a disjoint union of graphs of $\mathcal{C}$, and as $\mathcal{C}$ is closed under taking the disjoint union, $V_{\mathcal{U}}$ is a modulator to $\mathcal{C}$. Our output instance of d-COC/$\texttt{dist-to-}\mathcal{C}$ is then $(G, k', V_{\mathcal{U}})$, where $k' := k + (k+\alpha\cdot(d-1)+1)\cdot|\mathcal{F}|\cdot\mathsf{OPT}$.

First, note that the parameter did not change as $|V_{\mathcal{U}}| = |\mathcal{U}|$. Next, we show that the two instances are also equivalent.

▷ Claim. If $(\mathcal{U}, \mathcal{F}, k)$ is a yes-instance of α-Hitting Set, then $(G, k', V_{\mathcal{U}})$ is a yes-instance of d-COC/dist-to-$\mathcal{C}$.

Proof. Let $X \subseteq \mathcal{U}$ be a set such that $|X| \leq k$ and the intersection of X and each set of $\mathcal{F}$ is non-empty.

We construct a solution S for the d-COC/dist-to-$\mathcal{C}$ instance $(G, k', V_{\mathcal{U}})$ as follows. First, we construct a subset S' of the solution. For each $u \in X$, put universe vertex v_u in the set S', for each $u \in \mathcal{U} \setminus X$, put all vertices in $N_G(v_u)$ in the set S'. Begin with an empty set S. Then, for each $F \in \mathcal{F}$ and $i \in [1, k + \alpha \cdot (d-1) + 1]$, add a minimum cardinality d-coc set of the graph $H_F^i - S'$ to S. Finally, also add the vertices of S' to S.

The set S contains at most k vertices of $V_{\mathcal{U}}$. Moreover, S is a d-coc set, as each vertex of $V_{\mathcal{U}}$ is either in S or all its neighbours are in S, and moreover, we choose vertices corresponding to a solution within each copy of H. All that remains is arguing that for an arbitrary $F \in \mathcal{F}$ and $i \in [1, k + \alpha \cdot (d-1) + 1]$, S selects at most OPT vertices of H_F^i. By the fact that $X \cap F$ is non-empty, we see that $S' \cap V(H_F^i) \subsetneq B_F^i$. Moreover, the graph $H_F^i - S'$ is not connected to any vertex of $V_{\mathcal{U}}$, nor of another copy of the graph H. By the fact that B_F^i is a minimal d-blocking set, and that S' does not contain all vertices of it, we see that S selects OPT vertices of H_F^i. Hence, $|S| \leq k + (k + \alpha \cdot (d-1) + 1) \cdot |\mathcal{F}| \cdot \mathsf{OPT} = k'$, and we indeed have a yes-instance of d-COC/dist-to-$\mathcal{C}$. ◁

▷ Claim. If $(G, k', V_{\mathcal{U}})$ is a yes-instance of d-COC/dist-to-$\mathcal{C}$, then $(\mathcal{U}, \mathcal{F}, k)$ is a yes-instance of α-Hitting Set.

Proof. Let S be a d-coc set of G of size at most k'. Define $X \coloneqq \{u \in \mathcal{U} \mid v_u \in V_{\mathcal{U}} \cap S\}$, that is, X is a subset of the universe corresponding to the universe vertices selected by S. Next, we show that X has a non-empty intersection with all $F \in \mathcal{F}$, and $|X| \leq k$.

For the size-bound, we know that any d-coc set of G must select at least OPT vertices of each copy of the graph H. Since there are exactly $(k + \alpha \cdot (d-1) + 1) \cdot |\mathcal{F}|$ copies of H in G, at most k budget is left for the selection of vertices of $V_{\mathcal{U}}$.

Now, assume that for some $F \in \mathcal{F}$, we have $F \cap X = \emptyset$. Consider an arbitrary $u \in F$. We know that vertex v_u is not in S. However, v_u is adjacent to exactly one vertex of graph H_F^i for all $i \in [1, k + \alpha \cdot (d-1) + 1]$. Given that v_u cannot be in a connected component of size more than d in $G - S$, we know that for at most $d-1$ graphs of $H_F^1, \ldots, H_F^{k+\alpha \cdot (d-1)+1}$ it can be the case that the neighbour of v_u in the graph is not in S. Now, since this is true for all α vertices in V_F, we see that overall, for at most $\alpha \cdot (d-1)$ of the graphs $H_F^1, \ldots, H_F^{k+\alpha \cdot (d-1)+1}$ their respective minimal d-blocking set is not fully contained in S. Hence, there are at least $k+1$ copies of H for which a d-blocking set is contained in S. This in turn implies that the solution S selects at least OPT + 1 vertices of $k+1$ copies of H, and we know that at least OPT vertices of any other copy must be selected. We see that the size of S is now at least $(k + \alpha \cdot (d-1) + 1) \cdot |\mathcal{F}| \cdot \mathsf{OPT} + k + 1$, which contradicts the upper bound on the size of S. ◁

By the claims above, the input and output instances are equivalent. Moreover, the parameter did not change, and the runtime of the reduction is polynomial in the input size. Hence, it is a linear parameter transformation, and the lower bound of α-Hitting Set is transferred to d-COC/dist-to-$\mathcal{C}$. ◀

It should be noted that the main difference between our construction and the one by Hols et al. [26] is that we use more copies of the graph H to account for the fact that d is no longer fixed to be 1, like it was for the Vertex Cover problem.

4 Subdivided comb graphs with arbitrarily large minimal d-blocking sets

In this section we prove our main result Theorem 1, that is, we study the d-COC problem when the parameter is the distance to subdivided comb graphs. We denote this problem by d-COC/`dist-to-subdivided-combs`. We define the d-COC/`dist-to-subdivided-combs` problem formally below, and restate Theorem 1.

> d-COC/`dist-to-subdivided-combs`
> **Input:** Graph G, $M \subseteq V(G)$ such that $G - M$ is a disjoint union of subdivided comb graphs, integer $k \geq 0$
> **Parameter:** $|M|$
> **Question:** Does there exist a d-`coc` set S of G of size at most k?

▶ **Theorem 1.** *For any integer $d \geq 2$, d-Component Order Connectivity does not admit a polynomial compression when the parameter is distance to subdivided comb graphs, unless* NP $\subseteq$ coNP/poly.

The proof of Theorem 1 is obtained by showing that the class of disjoint unions of subdivided comb graphs has unbounded minimal d-blocking set size for all integers $d \geq 2$, and hence, by Theorem 4, no polynomial compression exists for this problem. Formally, we prove Lemma 5, that together with Theorem 4 proves Theorem 1.

▶ **Lemma 5.** *For any integers $N \geq 1$ and $d \geq 2$, if $\mathcal{C}$ contains all subdivided comb graphs with hairlength at most $d+1$, then there exists a graph in $\mathcal{C}$ that has a minimal d-blocking set of size N.*

The rest of the section is devoted to the proof of Lemma 5.

The blocking set graph. For all integers $N \geq 1$, $d \geq 2$, we define a subdivided comb graph $T_{N,d}$ with hairlength at most $d+1$, which (as we will show) has a minimal d-blocking set of size N. We call these graphs blocking set graphs.

▶ **Definition 6** (blocking set graph $T_{N,d}$). *For any integers $N \geq 1, d \geq 2$ we define $T_{N,d}$ to be the graph constructed as follows.*

- *It contains a path $P := p_0, p_1, \ldots, p_{(2N+1)\cdot d}$ of length $(2N+1)\cdot d + 1$ which is called the* spine.
- *For each $i \in \{0, 2, \ldots, 2N-2\}$, there is a path $Q_{i\cdot d+1} := q^1_{i\cdot d+1}, q^2_{i\cdot d+1}, \ldots, q^{d+1}_{i\cdot d+1}$, called the $(i\cdot d+1)$-th leg. The vertices $q^1_{i\cdot d+1}$ and $q^2_{i\cdot d+1}$ are called the* first *and* second *vertex of the $(i\cdot d+1)$-th leg, respectively. Note that there are N legs in $T_{N,d}$.*
- *For each $i \in \{0, 2, \ldots, 2N-2\}$, add an edge between $q^1_{i\cdot d+1}$ and $p_{i\cdot d+1}$, that is between the first vertex of the $(i\cdot d+1)$-th leg and the vertex of the spine with index $i\cdot d+1$.*
- *For each $i \in \{1, 3, \ldots, 2N-1\}$, add a pendant vertex $p'_{i\cdot d}$ to the vertex $p_{i\cdot d}$.*
- *For each $i \in \{2, 4, \ldots, 2N\}$, add a pendant vertex $p'_{i\cdot d-1}$ to the vertex $p_{i\cdot d-1}$.*

Refer to Figure 1 for a sketch of the graph $T_{N,d}$. Observe that $T_{N,d}$ is a subdivided comb graph with hairlength at most $d+1$. Next, we analyse some important properties of these graphs.

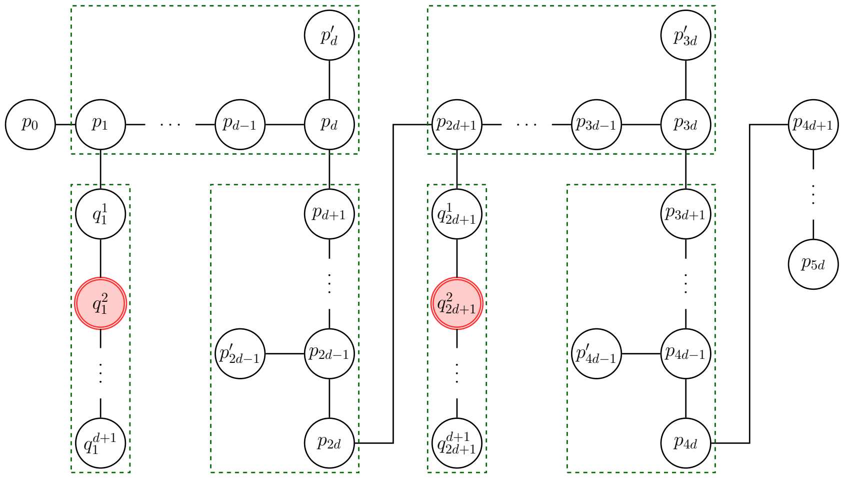

Figure 1 A depiction of the graph $T_{N,d}$ when $N = 2$. The green dashed rectangles mark $3N$ vertex-disjoint connected subgraphs containing $d + 1$ vertices each as in Lemma 7. The vertices in $Y_{N,d}$, as in Definition 9, are drawn as red double circles.

Lower bound on the size of d-coc sets of $T_{N,d}$. We begin by proving a lower bound on the size of d-coc sets of $T_{N,d}$ by providing a set of vertex-disjoint connected subgraphs of $T_{N,d}$, each containing $d + 1$ vertices.

▶ **Lemma 7.** *Fix any integers $N \geq 1, d \geq 2$, and define*
- $A := \{T_{N,d}[Z] \mid Z = \{q^1_{i\cdot d+1}, \ldots, q^{d+1}_{i\cdot d+1}\}, i \in \{0, 2 \ldots, 2N - 2\}\}$,
- $B := \{T_{N,d}[Z] \mid Z = \{p_{i\cdot d+1}, \ldots, p_{(i+1)\cdot d}, p'_{(i+1)\cdot d}\}, i \in \{0, 2 \ldots, 2N - 2\}\}$, *and*
- $C := \{T_{N,d}[Z] \mid Z = \{p_{i\cdot d+1}, \ldots, p_{(i+1)\cdot d}, p'_{(i+1)\cdot d-1}\}, i \in \{1, 3 \ldots, 2N - 1\}\}$.

Then, $\mathcal{D} := A \cup B \cup C$ is a collection of $3N$ vertex-disjoint connected subgraphs of $T_{N,d}$, each containing $d + 1$ vertices.

Proof. This directly follows from the description of $T_{N,d}$ in Definition 6. The described subgraphs are marked in Figure 1 for the case $N = 2$. ◀

As any d-coc set of $T_{N,d}$ must contain at least one vertex of each graph of $\mathcal{D}$ we obtain the following lower bound.

▶ **Corollary 8.** *For any integers $N \geq 1, d \geq 2$, any d-coc set S of $T_{N,d}$ fulfils $|S| \geq 3N$.*

Next, for each considered graph $T_{N,d}$, we define a set $Y_{N,d}$, which we will eventually prove to be a minimal d-blocking set.

▶ **Definition 9.** *Let N, d be integers with $d \geq 2$ and $N \geq 1$. Then, we define the vertex set $Y_{N,d} \subseteq V(T_{N,d})$ as $Y_{N,d} := \{q^2_{i\cdot d+1} \mid i \in \{0, 2, \ldots, 2N-2\}\}$, that is, $Y_{N,d}$ is the set consisting of the second vertices of all legs of $T_{N,d}$.*

The vertices in the set $Y_{N,d}$ are marked in Figure 1 for the case $N = 2$.

No minimum solution contains all of $Y_{N,d}$. Next, we prove that any d-coc set of $T_{N,d}$ that contains all vertices of $Y_{N,d}$ has size at least $3N + 1$. As we will later see that a minimum solution has size exactly $3N$, this will end up proving that $Y_{N,d}$ is a d-blocking set.

▶ **Lemma 10.** *For any integers $N \geq 1, d \geq 2$ consider the blocking set graph $T_{N,d}$. Then, any d-coc set S of $T_{N,d}$ such that $Y_{N,d} \subseteq S$ has cardinality at least $3N + 1$.*

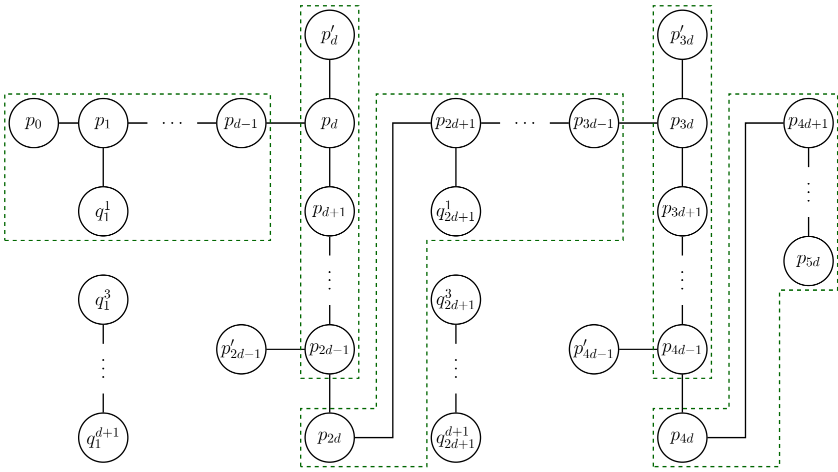

Figure 2 A depiction of the graph $T' = T_{N,d} - Y_{N,d}$, as in the proof of Lemma 10, when $N = 2$. The green dashed rectangles indicate the subgraphs of T' that lead to the lower bound of $2N + 1$.

Proof. Let $N \geq 1, d \geq 2$ be integers and set $T := T_{N,d}$ and $Y := Y_{N,d}$. We prove that any d-coc set that contains all vertices of Y has size at least $3N + 1$.

Set $T' := T - Y$. We prove the lower bound on the size of S by providing $2N + 1$ vertex disjoint connected subgraphs of T' of size $d + 1$ each. We define

$$C_i := \begin{cases} \{p_{i\cdot d}, \dots, p_{(i+1)\cdot d-1}, q^1_{i\cdot d+1}\} & \text{if } i \in \{0, 2, \dots, 2N-2\}, \\ \{p_{i\cdot d}, \dots, p_{(i+1)\cdot d-1}, p'_{i\cdot d}\} & \text{if } i \in \{1, 3, \dots, 2N-1\}, \\ \{p_{2N\cdot d}, \dots, p_{(2N+1)\cdot d}\} & \text{if } i = 2N, \end{cases}$$

and $G_i := T'[C_i]$, for all $i \in [0, 2N]$. Figure 2 depicts these subgraphs for the case $N = 2$. Observe that $C_i \cap C_j = \emptyset$ for $i, j \in [0, 2N], i \neq j$, and that G_i is a connected graph for all $i \in [0, 2N]$. Furthermore, $|C_i| = d + 1$ for all $i \in [0, 2N]$, proving that any d-coc set of T' has size at least $2N + 1$. Together with the fact that $|Y| = N$, this shows that any d-coc set S of T such that $Y \subseteq S$ fulfils $|S| \geq 3N + 1$. ◀

How much of $Y_{N,d}$ can be in a minimum solution? Now, we show that even though selecting vertices of $Y_{N,d}$ is intuitively bad for the solution size, for each strict subset Y' of $Y_{N,d}$ there is a d-coc set of $T_{N,d}$ of size $3N$ that contains all vertices of Y'. This shows that a minimum d-coc set of $T_{N,d}$ has size exactly $3N$ (when combined with Corollary 8), and moreover, that $Y_{N,d}$ is a minimal d-blocking set when considering Lemma 10.

▶ **Lemma 11.** *For any integers $N \geq 1, d \geq 2$ consider the blocking set graph $T_{N,d}$. For any $Y' \subsetneq Y_{N,d}$, there exists a d-coc set S of $T_{N,d} - Y'$ such that $|S| = 3N - |Y'|$.*

Proof. Let $N \geq 1, d \geq 2$ be integers, set $T := T_{N,d}$, and let P be the spine of T, and $Y := Y_{N,d}$. Let Y' be an arbitrary strict subset of Y, that is $Y' \subsetneq Y$.

We show that there exists a d-coc set S of $T' := T - Y'$ of size $3N - |Y'|$. Set

$$I := \{i \cdot d + 1 \mid q^2_{i\cdot d+1} \notin Y', i \in \{0, 2, \dots 2N-2\}\} \quad \text{and} \quad t := \min(I).$$

That is, I is the set of subscript indices of the legs for which the second vertex of the leg is not in Y', and t the minimum element of the index set I. Since $Y \neq Y'$, set I is non-empty.

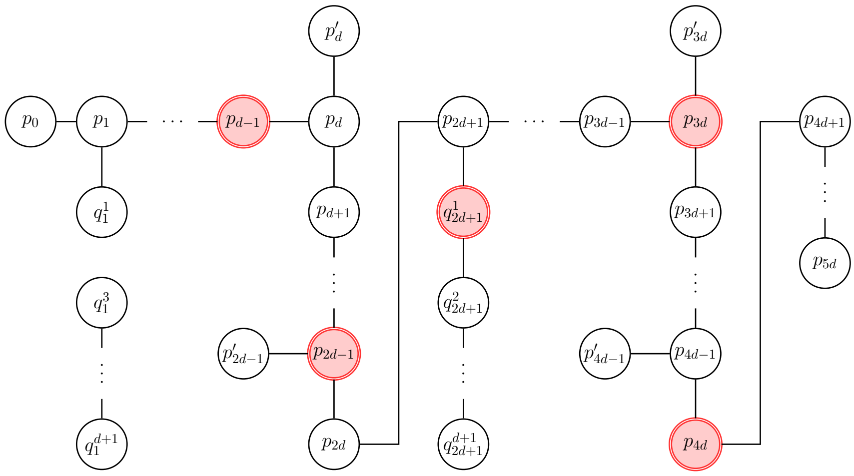

Figure 3 A depiction of $T_{N,d} - Y'$ when $N = 2$ and $Y' = \{q_1^2\}$. The vertices drawn as red double circles form a d-coc set of size $3N - |Y'|$. Note that I, t from the proof of Lemma 11 fulfil $I = \{2d + 1\}$ and $t = 2d + 1$.

We can then construct a d-coc set S of T' of size $3N - |Y'|$, by defining

$$S :=\{q_i^1 \mid i \in I\} \cup \{p_{i \cdot d-1} \mid i \in [1, 2N], i \cdot d < t\} \cup \{p_{i \cdot d} \mid i \in [1, 2N], i \cdot d \geq t\}.$$

Clearly, $|S| = |I| + 2N = 3N - |Y'|$. We depict solution S for the set $Y' = \{q_1^2\}$ and the case $N = 2$ in Figure 3.

Next, we argue why every connected component of $\hat{T} := T_{N,d} - (Y' \cup S)$ has size at most d. In the following, whenever we talk about connected components, we mean connected components of $\hat{T}$. Either the first or second vertex of every leg is in $Y' \cup S$. Thus, at most one vertex of any leg can be in a connected component together with vertices of the spine, and all other vertices of the leg are either in $Y' \cup S$, or in a connected component of size at most d together with other vertices of the same leg.

The only remaining concern is that there could be connected components of size more than d that contain vertices of the spine. Before arguing why that is not the case, we need to make some further observations.

There are two different types of pendants in the graph $T_{N,d}$, the pendants of the *first type* are $\{p'_{i \cdot d} \mid i \in \{1, 3, \ldots, 2N - 1\}\}$, and the pendants of the *second type* are the vertices $\{p'_{i \cdot d-1} \mid i \in \{2, 4, \ldots, 2N\}\}$. As either the first or the second vertex of each leg is part of $Y' \cup S$, we know that the only vertices of the legs that could be in connected components together with vertices of the spine are the first vertices. Let $X := \{q^1_{i \cdot d+1} \mid i \in \{0, 2, \ldots, 2N - 2\}\}$ be the set of first vertices of the legs of $T_{N,d}$. The pendants of the first type, pendants of the second type, and vertices in X are all adjacent to the vertex of the spine that has the same (subscript) index they have. We additionally have that $t = j \cdot d + 1$ for some $j \in \{0, 2, \ldots, 2N - 2\}$, as t is the index of some leg.

Our solution selects either vertex $p_{i \cdot d-1}$ or vertex $p_{i \cdot d}$ for any $i \in [1, 2N]$. Moreover, the vertex $p_{2N \cdot d}$ is guaranteed to be in S. As an immediate consequence, no connected component can contain more than d vertices of the spine. We will now show that all pendants and vertices of X are separated in $\hat{T}$. That is, any path in T' between two distinct vertices that are either pendants or in X contains a vertex that is in the set S. We will call such a path broken up, as it does not exist in $\hat{T}$.

Consider a vertex v with subscript index i that is a pendant of the first type, a pendant of the second type, or in X. Observe that, when the path from v to the closest pendant of the first type with an index lower than i, and the path from v to the closest pendant of the first type with an index higher than i are broken up by vertices of the spine, so are all other paths from v to (different) pendants of the first type. Analogous properties hold for pendants of the second type and vertices in X. This limits for which paths we need to show that they are broken up.

Pendants of the first type. Consider an arbitrary pendant $p'_{i \cdot d}$ of the first type, where $i \in \{1, 3, \ldots, 2N-1\}$. If $p_{i \cdot d} \in S$, the pendant is cut off from the rest of the graph in $\hat{T}$.

Otherwise, $p_{i \cdot d - 1} \in S$ and $i \cdot d < t = j \cdot d + 1$. The path from this pendant to any other pendant of the first type contains at least $2d + 1$ vertices of the spine, and is thus broken up by the solution.

The closest pendant of the second type with a larger index is $p'_{(i+1) \cdot d - 1}$. Now, we know that i is in the set $\{1, 3, \ldots, 2N-1\}$, whereas j is in the set $\{0, 2, \ldots, 2N-2\}$, which yields $i \neq j$. Since we have $i \cdot d < j \cdot d + 1$, it follows that $i < j$. This implies $(i+1) \cdot d \leq j \cdot d$, and we obtain $(i+1) \cdot d < j \cdot d + 1$ overall. Hence, we also select vertex $p_{(i+1) \cdot d - 1}$ into the solution, which isolates the vertex of the second type. The closest pendant of the second type with a smaller index is $p'_{(i-1) \cdot d - 1}$.[3] We have $(i \cdot d) - ((i-1) \cdot d - 1) = d + 1$, and thus, the path to this vertex contains at least $d + 2$ vertices of the spine, and is broken up.

Finally, we must argue that also paths to vertices in X are broken up. The closest vertex of X with a higher index is vertex $q^1_{(i+1) \cdot d + 1}$. As $((i+1) \cdot d + 1) - (i \cdot d) = d + 1$, the path to this vertex is broken up. The closest vertex of X with a lower index is vertex $q^1_{(i-1) \cdot d + 1}$. Vertex $p_{i \cdot d - 1} \in S$ breaks up the path.

Pendants of the second type. Consider an arbitrary pendant $p'_{i \cdot d - 1}$ of the second type, where $i \in \{2, 4, \ldots, 2N\}$. Similarly to before, any path to another pendant of the second type contains too many vertices of the spine, and is broken up.

We have already argued that all paths between pendants of the second type and pendants of the first type are broken up.

It remains to argue that also paths to vertices of X are broken up. The closest vertex of X with a lower index is vertex $q^1_{(i-2) \cdot d + 1}$, which is sufficiently far away. The closest vertex of X with a higher index is $q^1_{i \cdot d + 1}$. As we select either $p_{i \cdot d}$ or $p_{i \cdot d + 1}$, the path from $p'_{i \cdot d - 1}$ to $q^1_{i \cdot d + 1}$ is broken up.

Vertices in X. The only remaining danger is that vertices of X could be too close to each other, however, a quick glance at the definition yields that the path between two distinct vertices of X contains at least $2d + 1$ vertices of the spine, and is broken up.

Connected components of $\hat{T}$. We proceed to argue about the connected components that mainly consist of vertices of the spine. We select vertices $p_{i \cdot d - 1}$ into the solution as long as $i \cdot d < j \cdot d + 1$. The last vertex that is selected due to this process is vertex $p_{j \cdot d - 1}$. We first consider the connected components of vertices with index $j \cdot d - 2$ or less. As we select every d-th vertex of the spine, each considered connected component contains at most $d - 1$ vertices of the spine. By our arguments about the paths between the different types of additional vertices, we see each component can contain at most one additional vertex, and hence, it has size at most d.

[3] In this and similar cases we ignore the possibility that the vertex may not exist for the sake of brevity.

Thus far, we have covered all connected components with vertices of index $j \cdot d - 2$ or less. As the spine vertex with index $j \cdot d - 1$ is selected, the next important vertex is $p_{j \cdot d}$. Observe that vertex $p_{(j+1) \cdot d}$ is in S. Hence, vertex $p_{j \cdot d}$ is in the same connected component as $\{p_{j \cdot d}, p_{j \cdot d+1}, \ldots, p_{(j+1) \cdot d-1}\}$. As that are already d vertices, it is important that the connected components contains no further vertices. By definition of the integer t, the first vertex of the leg with index t is selected, and thus not in a connected component. Moreover, there exist no pendant vertices that have an index in $[j \cdot d, (j+1) \cdot d - 1]$. Another way to see that no pendant is part of the connected component is the following. We already established that any path between a vertex in X and a pendant is broken up. As that argument did not consider whether the vertex in X is itself selected or not, it still holds. Vertex $p_{j \cdot d+1}$ is part of the connected component and a neighbour of $q^1_{j \cdot d+1}$. Any path from $q^1_{j \cdot d+1}$ to a pendant would necessarily go through vertex $p_{j \cdot d+1}$, and as any such path is broken up by the solution, no pendant can be in the same connected component as $p_{j \cdot d+1}$. It follows that the connected component is also not too large.

Next, consider connected components that contain vertices with index larger than $(j+1) \cdot d$, but lower than $2N \cdot d$. We can, similarly to before, observe that we select every d-th vertex, and hence, any connected component can contain at most $d - 1$ vertices of the spine. For the same reason as earlier, no such component can contain two additional vertices, and thus, they are not too large. Observe that the vertex with the highest index that is part of the solution is $p_{2N \cdot d}$.

Finally, consider connected components that contain vertices with index larger than $2N \cdot d$. There are no legs or pendants with an index that high, and hence, there is only a single connected component that contains the vertices $\{p_{2N \cdot d+1}, \ldots, p_{(2N+1) \cdot d}\}$. ◀

Combining the results. Now, we can combine our knowledge about the graph $T_{N,d}$ to prove that $Y_{N,d}$ is a minimal d-blocking set.

▶ **Corollary 12.** *Let $N \geq 1$ and $d \geq 2$ be integers. Then, the set $Y_{N,d}$ is a minimal d-blocking set of $T_{N,d}$ of cardinality N.*

Proof. By Definition 9, the size of $Y_{N,d}$ is N. By Corollary 8, each d-coc set of $T_{N,d}$ has size at least $3N$. Then, Lemma 11 shows that a minimum d-coc set of $T_{N,d}$ has size exactly $3N$, and moreover, that for any $Y' \subsetneq Y_{N,d}$, there is a minimum d-coc set of $T_{N,d}$ that contains all vertices of Y'. Finally, Lemma 10 shows that no d-coc set that takes all of $Y_{N,d}$ is of minimum cardinality, concluding the proof. ◀

Corollary 12 implies Lemma 5.

5 Minimal d-blocking sets of paths have size at most two

In the previous section we have illustrated that the class of subdivided comb graphs contains graphs with arbitrarily large minimal d-blocking sets, for every integer $d \geq 2$. Contrasting this result, we show that all minimal d-blocking sets of paths have size at most two for every integer $d \geq 1$. This will follow from the lemma we prove next.[4]

▶ **Lemma 13.** *Let G be a graph and $X \subseteq V(G)$ be a d-blocking set of G for some integer $d \geq 1$. If there exists a subset $X' \subseteq X$ such that G_1 and G_2 are distinct connected components of $G - X'$, $V(G_1) \cap X \neq \emptyset$ and $V(G_2) \cap X \neq \emptyset$, then X is not a minimal d-blocking set.*

[4] Similar properties for (minimal) 1-blocking sets were shown in [26].

Proof. Let G be a graph and X, X' be sets with the properties described in the statement of the lemma. Moreover, let G_1 be a connected component of $G - X'$ that contains a vertex, say $x_1 \in X$, and G_2 be a different connected component of $G - X'$ that contains a vertex, say $x_2 \in X$.

Towards a contradiction, assume that X is a minimal d-blocking set. Then, there exists a minimum d-coc set S_1 of G that contains all vertices of $X \setminus \{x_1\}$, and a minimum d-coc set S_2 of G that contains all vertices of $X \setminus \{x_2\}$. Given that $x_1, x_2 \notin X'$ and $X' \subseteq X$, we see that $X' \subseteq S_1$ and $X' \subseteq S_2$. Then, it must be the case that S_i selects vertices of G_j corresponding to a minimum d-coc set of G_j for all $i, j \in \{1, 2\}$. However, we now see that the vertex set $(S_2 \setminus V(G_2)) \cup (S_1 \cap V(G_2))$ is a d-coc set of G of minimum size that contains all vertices of X, contradicting that X is a d-blocking set. ◀

An immediate corollary of this is that we obtain a bound on the size of minimal d-blocking sets of paths.

▶ **Corollary 14.** *For all integers $d \geq 1$, any minimal d-blocking set of a path has cardinality at most two.*

Proof. Let $P = v_1, v_2, \ldots, v_n$ be a path, and $X \subseteq V(P)$ a d-blocking set. If $|X| \geq 3$, then it suffices to set $X' := \{v_i\}$, where $v_i \in X$ and X contains vertices with higher and lower subscript index than v_i. Then, Lemma 13 shows that X is not a minimal d-blocking set. ◀

For completeness, we remark that there are paths for which all minimal d-blocking sets have size exactly two. For instance, consider the path on $d + 1$ vertices. There is no minimal d-blocking set of size one, however, any two vertices of the graph form a minimal d-blocking set of size two.

In contrast to the situation we had with the distance to subdivided comb graphs, Theorem 4 cannot be used to exclude polynomial kernels when the parameter is the size of a modulator to the graph class of linear forests, that is, forests in which every tree is a path. This indicates that there is still hope for a polynomial kernel in this setting.

6 Conclusion

The d-COC problem is a relevant generalization of Vertex Cover, and hence, efficient preprocessing algorithms are of high interest. For the Vertex Cover problem, many kernelization algorithms are known, and usually they can be extended to also work for d-COC. In this work we show that unlike Vertex Cover, for any integer $d \geq 2$, d-COC does not admit a kernel parameterized by the distance to subdivided comb graphs.

Several related questions gain relevance due to the new kernelization lower bound. Is a polynomial kernel possible when the parameter is even larger or at least incomparable? Reasonable possibilities would be the size of sets M such that $G - M$ consists of a disjoint union of paths, or a disjoint union of a constant number of trees, or a single subdivided comb graph, or even a single path. Answering these questions would almost certainly require very different tools from the ones we used in this paper, as the potential source of hardness would be completely different from the one that we exploited.

An interesting insight of our lower bound proof is that we show that the size of minimal d-blocking sets of families $\mathcal{C}$ still remains a bottleneck for getting polynomial kernels for d-COC parameterized by the size of a vertex modulator to $\mathcal{C}$. The contrast between Vertex Cover and d-COC is that for the former, the class $\mathcal{C}$ of forests has minimal blocking sets of size at most two (implicitly proved in [27]), whereas when $d \geq 2$, the size of minimal d-blocking sets is unbounded even for subdivided comb graphs.

There is a result by Bougeret, Jansen & Sau [4] that complements the sufficient condition defined by Hols, Kratsch & Pieterse [26]. In particular, Bougeret et al. show that if $\mathcal{C}$ is a minor-closed class, then the VERTEX COVER problem admits a polynomial kernel parameterized by the size of a vertex modulator to $\mathcal{C}$ if and only if there is a bound on the size of minimal blocking sets of graphs in $\mathcal{C}$. An interesting open question is whether we can obtain such a characterization also for d-COC. As of now, Theorem 4 only shows one side of this characterization when $\mathcal{C}$ is additionally closed under taking the disjoint union. If this characterization was also possible for d-COC, then Corollary 14 would immediately imply a polynomial kernel for d-COC parameterized by a vertex modulator to a linear forest.

References

1 Faisal N. Abu-Khzam, Rebecca L. Collins, Michael R. Fellows, Michael A. Langston, W. Henry Suters, and Christopher T. Symons. Kernelization algorithms for the vertex cover problem: Theory and experiments. In Lars Arge, Giuseppe F. Italiano, and Robert Sedgewick, editors, *Proceedings of the Sixth Workshop on Algorithm Engineering and Experiments and the First Workshop on Analytic Algorithmics and Combinatorics, New Orleans, LA, USA, January 10, 2004*, pages 62–69. SIAM, 2004.

2 R. Balasubramanian, Michael R. Fellows, and Venkatesh Raman. An improved fixed-parameter algorithm for vertex cover. *Inf. Process. Lett.*, 65(3):163–168, 1998. doi:10.1016/S0020-0190(97)00213-5.

3 Sriram Bhyravarapu, Satyabrata Jana, Saket Saurabh, and Roohani Sharma. Difference determines the degree: Structural kernelizations of component order connectivity. In Neeldhara Misra and Magnus Wahlström, editors, *18th International Symposium on Parameterized and Exact Computation, IPEC 2023, September 6-8, 2023, Amsterdam, The Netherlands*, volume 285 of *LIPIcs*, pages 5:1–5:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.IPEC.2023.5.

4 Marin Bougeret, Bart M. P. Jansen, and Ignasi Sau. Bridge-depth characterizes which minor-closed structural parameterizations of vertex cover admit a polynomial kernel. *SIAM J. Discret. Math.*, 36(4):2737–2773, 2022. doi:10.1137/21M1400766.

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 Jonathan F. Buss and Judy Goldsmith. Nondeterminism within P. *SIAM J. Comput.*, 22(3):560–572, 1993. doi:10.1137/0222038.

7 Jianer Chen, Iyad A. Kanj, and Weijia Jia. Vertex cover: Further observations and further improvements. *J. Algorithms*, 41(2):280–301, 2001. doi:10.1006/JAGM.2001.1186.

8 Jianer Chen, Iyad A. Kanj, and Ge Xia. Improved upper bounds for vertex cover. *Theor. Comput. Sci.*, 411(40-42):3736–3756, 2010. doi:10.1016/J.TCS.2010.06.026.

9 Jianer Chen, Lihua Liu, and Weijia Jia. Improvement on vertex cover for low-degree graphs. *Networks*, 35(4):253–259, 2000. doi:10.1002/1097-0037(200007)35:4<253::AID-NET3>3.0.CO;2-K.

10 Benny Chor, Mike Fellows, and David W. Juedes. Linear kernels in linear time, or how to save k colors in $O(n^2)$ steps. In Juraj Hromkovic, Manfred Nagl, and Bernhard Westfechtel, editors, *Graph-Theoretic Concepts in Computer Science, 30th International Workshop, WG 2004, Bad Honnef, Germany, June 21-23, 2004, Revised Papers*, volume 3353 of *Lecture Notes in Computer Science*, pages 257–269. Springer, 2004. doi:10.1007/978-3-540-30559-0_22.

11 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000. doi:10.1007/S002249910009.

12 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.

13 Frank K. H. A. Dehne, Michael R. Fellows, Frances A. Rosamond, and Peter Shaw. Greedy localization, iterative compression, modeled crown reductions: New FPT techniques, an improved algorithm for set splitting, and a novel 2k kernelization for vertex cover. In Rodney G. Downey, Michael R. Fellows, and Frank K. H. A. Dehne, editors, *Parameterized and Exact Computation, First International Workshop, IWPEC 2004, Bergen, Norway, September 14-17, 2004, Proceedings*, volume 3162 of *Lecture Notes in Computer Science*, pages 271–280. Springer, 2004. doi:10.1007/978-3-540-28639-4_24.

14 Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *Journal of The Acm*, 61(4):23:1–23:27, 2014. doi:10.1145/2629620.

15 Huib Donkers and Bart M. P. Jansen. A turing kernelization dichotomy for structural parameterizations of F-minor-free deletion. *J. Comput. Syst. Sci.*, 119:164–182, 2021. doi:10.1016/J.JCSS.2021.02.005.

16 Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness I: basic results. *SIAM J. Comput.*, 24(4):873–921, 1995. doi:10.1137/S0097539792228228.

17 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013. doi:10.1007/978-1-4471-5559-1.

18 Pål Grønås Drange, Markus S. Dregi, and Pim van 't Hof. On the computational complexity of vertex integrity and component order connectivity. *Algorithmica*, 76(4):1181–1202, 2016. doi:10.1007/S00453-016-0127-X.

19 Michael R. Fellows. Blow-ups, win/win's, and crown rules: Some new directions in FPT. In Hans L. Bodlaender, editor, *Graph-Theoretic Concepts in Computer Science, 29th International Workshop, WG 2003, Elspeet, The Netherlands, June 19-21, 2003, Revised Papers*, volume 2880 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2003. doi:10.1007/978-3-540-39890-5_1.

20 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.

21 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: theory of parameterized preprocessing*. Cambridge University Press, 2019. doi:10.1017/9781107415157.

22 Fedor V. Fomin and Torstein J. F. Strømme. Vertex cover structural parameterization revisited. In Pinar Heggernes, editor, *Graph-Theoretic Concepts in Computer Science - 42nd International Workshop, WG 2016, Istanbul, Turkey, June 22-24, 2016, Revised Selected Papers*, volume 9941 of *Lecture Notes in Computer Science*, pages 171–182, 2016. doi:10.1007/978-3-662-53536-3_15.

23 Daniel Gross, Monika Heinig, Lakshmi Iswara, L William Kazmierczak, Kristi Luttrell, John T Saccoman, and Charles Suffel. A survey of component order connectivity models of graph theoretic networks. *WSEAS Transactions on Mathematics*, 12(9):895–910, 2013.

24 David G. Harris and N. S. Narayanaswamy. A faster algorithm for vertex cover parameterized by solution size. In Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov, editors, *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, March 12-14, 2024, Clermont-Ferrand, France*, volume 289 of *LIPIcs*, pages 40:1–40:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.STACS.2024.40.

25 Danny Hermelin and Xi Wu. Weak compositions and their applications to polynomial lower bounds for kernelization. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 104–113. SIAM, 2012. doi:10.1137/1.9781611973099.9.

26 Eva-Maria C. Hols, Stefan Kratsch, and Astrid Pieterse. Elimination distances, blocking sets, and kernels for vertex cover. *SIAM J. Discret. Math.*, 36(3):1955–1990, 2022. doi:10.1137/20M1335285.

27 Bart M. P. Jansen and Hans L. Bodlaender. Vertex cover kernelization revisited: Upper and lower bounds for a refined parameter. In Thomas Schwentick and Christoph Dürr, editors, *28th International Symposium on Theoretical Aspects of Computer Science, STACS 2011, March 10-12, 2011, Dortmund, Germany*, volume 9 of *LIPIcs*, pages 177–188. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2011. doi:10.4230/LIPICS.STACS.2011.177.

28 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.

29 Bart M. P. Jansen and Astrid Pieterse. Polynomial kernels for hitting forbidden minors under structural parameterizations. *Theor. Comput. Sci.*, 841:124–166, 2020. doi:10.1016/J.TCS.2020.07.009.

30 Bart MP Jansen. *The power of data reduction: Kernels for fundamental graph problems*. PhD thesis, Utrecht University, 2013.

31 Samir Khuller. Algorithms column: the vertex cover problem. *SIGACT News*, 33(2):31–33, 2002. doi:10.1145/564585.564598.

32 Mithilesh Kumar and Daniel Lokshtanov. A 2lk kernel for l-component order connectivity. In Jiong Guo and Danny Hermelin, editors, *11th International Symposium on Parameterized and Exact Computation, IPEC 2016, August 24-26, 2016, Aarhus, Denmark*, volume 63 of *LIPIcs*, pages 20:1–20:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICS.IPEC.2016.20.

33 Michael Lampis. A kernel of order 2 k-c log k for vertex cover. *Inf. Process. Lett.*, 111(23-24):1089–1091, 2011. doi:10.1016/J.IPL.2011.09.003.

34 Diptapriyo Majumdar, Venkatesh Raman, and Saket Saurabh. Polynomial kernels for vertex cover parameterized by small degree modulators. *Theory of Computing Systems*, 62(8):1910–1951, 2018. doi:10.1007/S00224-018-9858-1.

35 N. S. Narayanaswamy, Venkatesh Raman, M. S. Ramanujan, and Saket Saurabh. LP can be a cure for parameterized problems. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012, February 29th - March 3rd, 2012, Paris, France*, volume 14 of *LIPIcs*, pages 338–349. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012. doi:10.4230/LIPICS.STACS.2012.338.

36 Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press, 2006. doi:10.1093/ACPROF:OSO/9780198566076.001.0001.

37 Rolf Niedermeier and Peter Rossmanith. Upper bounds for vertex cover further improved. In Christoph Meinel and Sophie Tison, editors, *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*, volume 1563 of *Lecture Notes in Computer Science*, pages 561–570. Springer, 1999. doi:10.1007/3-540-49116-3_53.

38 Rolf Niedermeier and Peter Rossmanith. On efficient fixed-parameter algorithms for weighted vertex cover. *J. Algorithms*, 47(2):63–77, 2003. doi:10.1016/S0196-6774(03)00005-1.

39 Arezou Soleimanfallah and Anders Yeo. A kernel of order 2k-c for vertex cover. *Discret. Math.*, 311(10-11):892–895, 2011. doi:10.1016/J.DISC.2011.02.014.

40 Ulrike Stege and Michael Ralph Fellows. An improved fixed parameter tractable algorithm for vertex cover. *Technical report/Departement Informatik, ETH Zürich*, 318, 1999. doi:10.3929/ethz-a-006653305.

41 Mingyu Xiao. Linear kernels for separating a graph into components of bounded size. *Journal of Computer and System Sciences*, 88:260–270, 2017. doi:10.1016/J.JCSS.2017.04.004.