

HAMILTONIAN CYCLE PARAMETERIZED BY TREEDEPTH IN SINGLE EXPONENTIAL TIME AND POLYNOMIAL SPACE*

JESPER NEDERLOF[†], MICHAŁ PILIPCZUK[‡], CELINE M. F. SWENNENHUIS[§],
AND KAROL WĘGRZYCKI[¶]

Abstract. For many algorithmic problems on graphs of treewidth t , a standard dynamic programming approach gives algorithms with time and space complexity $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$. It turns out that when one considers the more restrictive parameter treedepth, it is often the case that a variation of this technique can be used to reduce the space complexity to polynomial, while retaining time complexity of the form $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$, where d is the treedepth. This transfer of methodology is, however, far from automatic. For instance, for problems with connectivity constraints, standard dynamic programming techniques give algorithms with time and space complexity $2^{\mathcal{O}(t \log t)} \cdot n^{\mathcal{O}(1)}$ on graphs of treewidth t , but it is not clear how to convert them into time-efficient polynomial space algorithms for graphs of low treedepth. Cygan et al. [*ACM Trans. Algorithms*, 18 (2022), 17] introduced the Cut&Count technique and showed that a certain class of problems with connectivity constraints can be solved in time and space complexity $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$. Recently, Hegerfeld and Kratsch (STACS’20) showed that, for some of those problems, the Cut&Count technique can be also applied in the setting of treedepth, and it gives algorithms with running time $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ and polynomial space usage. However, several important problems eluded such a treatment, with the most prominent examples being *Hamiltonian Cycle* and *Longest Path*. In this paper, we clarify the situation by showing that *Hamiltonian cycle*, *Hamiltonian Path*, *Long Cycle*, *Long Path*, and *Min Cycle Cover* all admit $5^d \cdot n^{\mathcal{O}(1)}$ -time and polynomial space algorithms on graphs of treedepth d . The algorithms are randomized Monte Carlo with only false negatives.

Key words. Hamiltonian cycle, connectivity, polynomial space, treedepth

MSC codes. 68Q25, 68R10, 68W20

DOI. 10.1137/22M1518943

*Received by the editors August 30, 2022; accepted for publication (in revised form) January 26, 2023; published electronically July 24, 2023. A preliminary version of this paper was presented at the 46th International Workshop on Graph-Theoretic Concepts in Computer Science 2020.

<https://doi.org/10.1137/22M1518943>

Funding: The first author was supported by the project CRACKNP that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement 853234. The second author’s work is a part of the projects BOBR and TOTAL that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreements 948057 and 677651. The third author was supported by the Netherlands Organization for Scientific Research under project 613.009.031b. The fourth author’s work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement 850979. The work leading to the results presented in this paper was initiated during the Parameterized Retreat of the algorithms group of the University of Warsaw (PARUW), held in Karpacz in February 2019. This retreat was financed by the project CUTACOMBS, which has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement 714704.

[†]Utrecht University, Utrecht, 3584 CC, The Netherlands (j.nederlof@uu.nl).

[‡]Institute of Informatics, University of Warsaw, Warsaw, 02-097, Poland (michal.pilipczuk@mimuw.edu.pl).

[§]Eindhoven University of Technology, Eindhoven, 5600 MB, The Netherlands (c.m.f.swennenhuis@tue.nl).

[¶]Saarland University and Max Planck Institute for Informatics, Saarbrücken, 66123, Germany (wegrzycki@cs.uni-saarland.de).



1. Introduction. It is widely believed that no NP-hard problem admits a polynomial time algorithm. However, actual instances of problems that we are interested in solving often admit much more structure than a general instance. This observation gave rise to the field of *parameterized complexity*, where the hardness of an instance does not depend exclusively on the input size. In the parameterized regime, we assume that each instance is equipped with an additional parameter k and the goal is to give a *fixed-parameter algorithm*: an algorithm with running time $f(k) \cdot n^{\mathcal{O}(1)}$, where f is a function independent of n . After settling that a problem admits such an algorithm, it is natural to look for one with function f as low as possible. We refer to [7, 13, 15] for an introduction to parameterized complexity.

One of the most widely used parameters is the *treewidth* t of the input graph. Usually, problems that involve only constraints of local nature admit an algorithm with running time of the form $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ [7]. For a long time, such algorithms remained out of reach for problems involving connectivity constraints, and for those only $2^{\mathcal{O}(t \log t)} \cdot n^{\mathcal{O}(1)}$ -time algorithms were known (see [7, section 7]). The breakthrough came with the Cut&Count technique, introduced by Cygan et al. [9], that allows one to design randomized Monte-Carlo algorithms with running times of the form $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ for a wide range of connectivity problems, e.g., *Hamiltonian Path*, *Connected Vertex Cover*, *Connected Dominating Set*, etc. The technique was subsequently derandomized [5, 16].

One of the main issues with standard dynamic programming algorithms is that they tend to have prohibitively large space usage. The natural goal is, therefore, to reduce the space complexity while not sacrificing much on the time complexity. However, Drucker, Nederlof, and Santhanam [14] and Pilipczuk and Wrochna [31] gave some complexity-theoretical evidence that for dynamic programming on graphs of bounded treewidth, such a reduction is probably impossible. For example, they showed that under plausible assumptions, there is no algorithm that works in time $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ and uses $2^{o(t)} \cdot n^{\mathcal{O}(1)}$ space for the *3-Coloring* or *Independent Set* problem.

Treedepth. The aforementioned issues motivate the research on a different, more restrictive parameterization, for which the reduction of space complexity would be possible. In this paper, we will consider the parameterization by *treedepth*, defined as follows.

DEFINITION 1.1. *An elimination forest of a graph G is a rooted forest F on the same vertex set as G such that for every edge uv of G , either u is an ancestor of v in F or v is an ancestor of u in F . The treedepth of G is the minimum possible depth of an elimination forest of G .*

The treedepth of a graph is never smaller than its treewidth, but it is also never larger than the treewidth times $\log n$ (see [6, 29]). In many concrete cases, the two parameters have the same advantages. For example, planar graphs have treewidth $\mathcal{O}(\sqrt{n})$, but also treedepth $\mathcal{O}(\sqrt{n})$.

It has been recently realized that on graphs of treedepth d , many algorithmic problems indeed can be solved in time $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ and using only polynomial space.¹ For the most basic problems, such as *3-Coloring* and *Independent Set*, a simple branching algorithm achieves such complexity. However, in contrast to the treewidth parameterization, for many more complex problems it is highly nontrivial, yet possible, to establish similar bounds. One technique that turns out to be useful here is the

¹Throughout the introduction, when we speak about a graph of treedepth d , we mean a graph supplied with an elimination forest of depth d . While in the case of treewidth, a tree decomposition of approximately (up to a constant factor) optimum width can be computed in time $8^t \cdot n^{\mathcal{O}(1)}$ [34, 7, 21], the existence of such an approximation algorithm for treedepth is a notorious open problem.

framework of algebraic transforms introduced by Loksthanov and Nederlof [23], who demonstrated how to reduce the space requirements of many dynamic programming algorithms to polynomial in the input size by reorganizing the computation using a suitable transform. Fürer and Yu [17] applied this framework to give $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ -time and polynomial space algorithms on graphs of treedepth d for the *Dominating Set* problem and for the problem of counting the number of perfect matchings. Pilipczuk and Wrochna [31] considered algorithms with even more restricted space requirements: they showed that *3-Coloring*, *Dominating Set*, and *Vertex Cover* admit algorithms that work in $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ time and use $\mathcal{O}(d + \log n)$ space. For *Dominating Set* they avoided the explicit use of algebraization and instead provided a more combinatorial interpretation based on what one could call *inclusion-exclusion branching* (see also [1]). Later, Pilipczuk and Siebertz [30] used color-coding to give an $2^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$ -time and polynomial space algorithm for the *Subgraph Isomorphism* problem. Recently, Belbasi and Fürer [2] presented an algorithm for counting Hamiltonian cycles in time $(4t)^d \cdot n^{\mathcal{O}(1)}$ and using polynomial space, where t is the width of a given tree decomposition and d is its (suitably defined) depth.

Treedepth and Cut&Count. Very recently, Hegerfeld and Kratsch [19] demonstrated that the Cut&Count technique can also be applied in the setting of the treedepth parameterization. Consequently, they gave randomized algorithms with running times $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ and polynomial space usage for a number of problems with connectivity constraints such as *Connected Vertex Cover*, *Connected Dominating Set*, *Feedback Vertex Set*, or *Steiner Tree*. However, Hegerfeld and Kratsch found it problematic to apply the methodology to several important problems originally considered by Cygan et al. [9] in the context of Cut&Count. Specifically, these are problems based on the selection of edges rather than vertices, such as *Hamiltonian Cycle* or *Long Cycle*. For this reason, Hegerfeld and Kratsch explicitly asked in [19] whether *Hamiltonian Cycle*, *Hamiltonian Path*, *Long Cycle*, and *Min Cycle Cover* also admit $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ -time and polynomial space algorithms on graphs of treedepth d (see Appendix B for problem definitions).²

Our contribution. In this paper, we introduce additional techniques that allow us to extend the results of [19] and to answer the abovementioned open problem of Hegerfeld and Kratsch in the affirmative. More precisely, we prove the following theorem.

THEOREM 1.2. *There is a randomized algorithm that, given a graph G together with its elimination forest of depth d , and number $k \in \mathbb{N}$, solves Hamiltonian Cycle, Hamiltonian Path, k -Cycle, k -Path, and Min Cycle Cover in time $5^d \cdot n^{\mathcal{O}(1)}$ using polynomial space. The algorithm has a one-sided error: it may give false negatives with probability at most $\frac{1}{2}$.*

In fact, Theorem 1.2 is an easy corollary of the following result for a generalization of the considered problems. In the *Partial Cycle Cover* problem we are given an undirected graph G and integers k and ℓ , and we ask whether in G there is a family of at most k vertex-disjoint cycles that jointly visit exactly ℓ vertices. We will prove the following theorem.

THEOREM 1.3. *There is a randomized algorithm that, given a graph G together with its elimination forest of depth d , and numbers $k, \ell \in \mathbb{N}$, solves the Partial Cycle*

²Note that when discussing the *Long Path* and the *Long Cycle* problems, we use the letter ℓ to denote the required length of a path, respectively of a cycle, instead of the letter k that is perhaps more traditionally used in this context.




FIG. 1. Construction of the graph G' from $G = C_5$, together with a simple perfect matching M that projects in π to E . Solid edges represent M , dashed edges belong to $E'_1 \setminus M$, while dotted edges comprise E'_0 . (Figure in color online.)

Cover problem for G, k, ℓ in time $5^d \cdot n^{\mathcal{O}(1)}$ and using polynomial space. The algorithm has a one-sided error: it may give false negatives with probability at most $\frac{1}{2}$.

To see that Theorem 1.3 implies Theorem 1.2, note that *Hamiltonian Cycle*, *Min Cycle Cover*, and *Long Cycle* are special cases of the *Partial Cycle Cover* (for fixed parameters k and ℓ).

To solve *Long Path*, we can simply iterate through all pairs of nonadjacent vertices s, t and apply the *Long Cycle* algorithm to the graph G with edge st added; this increases the treedepth by at most 1 and the provided elimination forest can be easily adjusted. It is easy to see that then the original graph G contains a simple path on ℓ vertices if and only if for some choice of s and t , we find a cycle of length ℓ in G augmented with the edge st . Finally, *Hamiltonian Path* is just *Long Path* applied for $\ell = |V(G)|$.

We remark that our algorithmic findings have concrete applications outside of the realm of structural parameterizations. For instance, Lokshtanov, Mnich, and Saurabh [22] gave a $2^{\mathcal{O}(\sqrt{\ell} \log^2 \ell)} \cdot n^{\mathcal{O}(1)}$ -time polynomial space algorithm for the *Long Path* problem on H -minor-free graphs, for every fixed H . In Appendix A we present how using our results one can improve the running time to $2^{\mathcal{O}(\sqrt{\ell} \log \ell)} \cdot n^{\mathcal{O}(1)}$ while keeping the polynomial space complexity.

Our techniques. Similarly to Hegerfeld and Kratsch [19] we use the Cut&Count framework, but we apply a different new view on the Count part, suited for problems based on edge selection. The main idea is that instead of counting cycle covers, as a standard application of Cut&Count would do, we count perfect matchings in an auxiliary graph, constructed by replacing every vertex with two adjacent copies; see Figure 1. The number of such perfect matchings can be related to the number of cycle covers of the original graph. However, the considered perfect matchings can be counted within the claimed complexity by either employing the previous “algebraized” dynamic programming algorithm, or the algorithm based on inclusion-exclusion branching (our presentation chooses the latter).

Applying this approach naïvely would give us a polynomial space algorithm with running time $8^d \cdot n^{\mathcal{O}(1)}$. We improve the running time to $5^d \cdot n^{\mathcal{O}(1)}$ by employing several observations about the symmetries of recursive calls of our algorithms, in a similar way as in the algorithm for $\#k$ -*Multi-Set-Cover* of Nederlof [26].

Organization of this paper. The remainder of this paper is devoted to the proof of Theorem 1.3. In section 2 we introduce the notation and present basic definitions. In section 3 we discuss the Cut&Count technique in a self-contained manner and explain the Cut part. In section 4 we reduce the Count part to counting perfect matchings in an auxiliary graph. In section 5 we give an algorithm for counting such matchings, and in section 6 we subsequently verify the correctness of the algorithm. We conclude with several open questions in section 7.

In Appendix A we present a polynomial space algorithm for *Long Path* in H -minor free graphs. In Appendix B we include the definitions of all the problems considered in this paper.

2. Preliminaries.

Notation. For a graph G , by $\text{cc}(G)$ we denote the number of connected components of G . Let F be a subset of edges of G . By $\text{cc}(F)$ we denote the number of connected components of the graph consisting of all the edges of F and vertices incident to them. For a vertex u , by $\deg_F(u)$ we mean the number of edges of F incident to u . Then F is a *matching* if $\deg_F(u) \in \{0, 1\}$ for every vertex u is a *perfect matching* if $\deg_F(u) = 1$ for every vertex u , and is a *partial cycle cover* if $\deg_F(u) \in \{0, 2\}$ for every vertex u . Note that thus we treat partial cycle covers as sets of edges.

A *cut* of a set of vertices U is just an ordered partition of U into two sets, that is, a pair (L, R) such that $L \cap R = \emptyset$ and $L \cup R = U$. A cut (L, R) of the vertex set of a graph is *consistent* with a subset of edges F if there is no edge in F with one endpoint in L and second in R .

For a function f and elements x and y , where x is not in the domain of f , by $f[x \mapsto y]$ we denote the function obtained from f by extending its domain by x and setting $f(x) = y$.

We use the $\mathcal{O}^*(\cdot)$ notation to hide factors polynomial in the input size. For convenience, throughout this paper we assume the RAM model: every integer takes a unit of space and arithmetic operations on integers have unit cost. However, it can be easily seen that all the numbers appearing during the computation have bit length bounded polynomially in the input size. Since we never specify the polynomial factors in the time or space complexity of our algorithms, without any influence on the claimed asymptotic bounds we may assume that the representation of any number takes polynomial space and arithmetic operations on the numbers take polynomial time.

Treedepth. A *rooted forest* is a directed acyclic graph T where every vertex has out-degree at most 1. The vertices of out-degree 0 in T are called the *roots*. Whenever a vertex u is reachable from a vertex v by a directed path in T , we say that u is an *ancestor* of v , and v is a *descendant* of u . Note that every vertex is its ancestor as well as its descendant. The *depth* of a rooted forest is the maximum number of vertices that can appear on a directed path in it.

We use the following notation from previous works [19, 31]. For a vertex u of a rooted forest T , we denote

$$\begin{aligned} \text{subtree}[u] &:= \{v : u \text{ is ancestor of } v\}, & \text{subtree}(u) &:= \text{subtree}[u] \setminus \{u\}, \\ \text{tail}[u] &:= \{v : v \text{ is ancestor of } u\}, & \text{tail}(u) &:= \text{tail}[u] \setminus \{u\}, \\ \text{broom}[u] &:= \text{tail}[u] \cup \text{subtree}[u]. \end{aligned}$$

Additionally, $\text{children}(u)$ denotes the set of children of u , whereas $\text{parent}(u)$ is the *parent* of u , that is, the only out-neighbor of u . If u is a root, we set $\text{parent}(u) = \perp$.

For a graph G , an *elimination forest* of G is a rooted forest T on the same vertex set as G that satisfies the following property: whenever uv is an edge in G , then in T

either u is an ancestor of v , or v is an ancestor of u . The *treedepth* of a graph is the minimum possible depth of an elimination forest of G .

Isolation lemma. The only source of randomness in our algorithm is the Isolation lemma of Mulmuley, Vazirani, and Vazirani [24]. Suppose U is a finite set and $\omega : U \rightarrow \mathbb{Z}$ is a weight function on U . We say that ω *isolates* a nonempty family of subsets $\mathcal{F} \subseteq 2^U$ if there is a unique $S \in \mathcal{F}$ such that

$$\omega(S) = \min_{X \in \mathcal{F}} \omega(X),$$

where $\omega(X) := \sum_{x \in X} \omega(x)$. Then the Isolation lemma can be stated as follows.

LEMMA 2.1 (Isolation lemma [24]). *Let U be a finite set and $\mathcal{F} \subseteq 2^U$ be a nonempty family of subsets of U . Suppose for every $u \in U$ we choose its weight $\omega(u)$ uniformly and independently at random from the set $\{1, \dots, N\}$, where $N \in \mathbb{N}$. Then ω isolates \mathcal{F} with a probability at least $1 - \frac{|U|}{N}$.*

3. The Cut part. We now proceed to the proof of Theorem 1.3. Throughout the proof we fix the input graph $G = (V, E)$, its elimination forest T of depth d , and numbers $k, \ell \in \mathbb{N}$. We may assume that G is connected, as otherwise we may apply the algorithm to each connected component separately. Thus T has to be a tree, so we will call it an *elimination tree* to avoid confusion. Also, we denote $n := |V|$.

As mentioned before, we shall apply the Cut&Count technique of Cygan et al. [9]. This technique consists of two parts: the Cut part and the Count part. The idea is that in the first part, we relax the connectivity requirements and show that it is enough to count the number of relaxed solutions together with cuts consistent with them, as this number is congruent to the number of nonrelaxed solutions modulo a power of 2. The Isolation lemma is used here to ensure that with high probability, the number of solutions does not accidentally cancel out modulo this power of 2. More precisely, having drawn a weight function at random, for each possible total weight w we count the number of solutions of total weight w . Then the Isolation lemma asserts that, with high probability, for some w there will be a unique solution of total weight w . Then comes the Count part, where the goal is to efficiently count the number of relaxed solutions together with cuts consistent with them.

We refer the reader to [9] for a more elaborate discussion of the Cut&Count technique, while now we apply it to the particular case of *Partial Cycle Cover*. A relaxed solution is just a partial cycle cover consisting of ℓ edges. Then a solution is a relaxed solution that spans at most k cycles. Formally, the sets of *solutions* (\mathcal{S}) and *relaxed solutions* (\mathcal{R}) are defined as follows:

$$\begin{aligned} \mathcal{R} &:= \{F \subseteq E : |F| = \ell \text{ and } \deg_F(u) \in \{0, 2\} \text{ for every } u \in V\}; \\ \mathcal{S} &:= \{F \in \mathcal{R} : \text{cc}(F) \leq k\}. \end{aligned}$$

Suppose now that the input graph G is supplied with a weight function on the edges $\omega : E \rightarrow \mathbb{Z}$. Then we can stratify the families above using the total weight. That is, for every $w \in \mathbb{Z}$ we define

$$\mathcal{R}_w := \{F \in \mathcal{R} : \omega(F) = w\} \quad \text{and} \quad \mathcal{S}_w := \{F \in \mathcal{S} : \omega(F) = w\}.$$

Now, let

$$\mathcal{C}_w := \{(F, (L, R)) : F \in \mathcal{R}_w \text{ and } (L, R) \text{ is a cut of } V \text{ consistent with } F\}.$$

The following observation is the key idea in the Cut&Count technique.

LEMMA 3.1. *For every $w \in \mathbb{Z}$, we have*

$$|\mathcal{C}_w| \equiv \sum_{F \in \mathcal{S}_w} 2^{n-\ell+\text{cc}(F)} \pmod{2^{n-\ell+k+1}}.$$

Proof. Observe that for each $F \in \mathcal{R}_w$ there are exactly $2^{n-\ell+\text{cc}(F)}$ cuts of V consistent with it, because each of the $\text{cc}(F)$ cycles spanned by F can be on either side of the cut, and similarly each of $n - \ell$ vertices not incident to the edges of F can be on either side of the cut. Hence $|\mathcal{C}_w| = \sum_{F \in \mathcal{R}_w} 2^{n-\ell+\text{cc}(F)}$. However, for every $F \in \mathcal{R}_w \setminus \mathcal{S}_w$ the term $2^{n-\ell+\text{cc}(F)}$ is divisible by $2^{n-\ell+k+1}$ since $\text{cc}(F) \geq k + 1$, and thus $\sum_{F \in \mathcal{R}_w} 2^{\text{cc}(F)} \equiv \sum_{F \in \mathcal{S}_w} 2^{\text{cc}(F)} \pmod{2^{n-\ell+k+1}}$. \square

In the next sections, we will present the Count part of the technique, which boils down to proving the following lemma.

LEMMA 3.2. *Given $w \in \mathbb{Z}$ and a weight function $\omega : E \rightarrow \{1, \dots, N\}$, where $N = \mathcal{O}^*(1)$, the number $|\mathcal{C}_w|$ can be computed in time $\mathcal{O}^*(5^d)$ and space $\mathcal{O}^*(1)$.*

We now show how to combine Lemma 3.1 with Lemma 3.2 to prove Theorem 1.3.

Proof of Theorem 1.3 assuming Lemma 3.2. Let $N = 2|E|$. The algorithm proceeds as follows. First, for every edge $e \in E$, sample its weight $\omega(e)$ uniformly and independently at random from the set $\{1, \dots, N\}$. Next, for each $w \in \{1, \dots, N|E|\}$ compute the number $|\mathcal{C}_w|$ in time $\mathcal{O}^*(5^d)$ and space $\mathcal{O}^*(1)$ using the algorithm of Lemma 3.2. If for some w the number $|\mathcal{C}_w|$ is not divisible by $2^{n-\ell+k+1}$, then output that there exists a solution. Otherwise, output that there is no solution.

It is clear that the algorithm runs in time $\mathcal{O}^*(5^d)$ and uses $\mathcal{O}^*(1)$ space, so it remains to argue the correctness. On one hand, observe that if $\mathcal{S} = \emptyset$, then $\mathcal{S}_w = \emptyset$ for all $w \in \mathbb{Z}$, hence by Lemma 3.1 all the computed numbers $|\mathcal{C}_w|$ will be indeed divisible by $2^{n-\ell+k+1}$. Therefore, there are no false positives. On the other hand, if $\mathcal{S} \neq \emptyset$, then the Isolation lemma implies that with probability at least $\frac{|E|}{N} = \frac{1}{2}$ there exists $w \in \mathbb{Z}$ such that $|\mathcal{S}_w| = 1$. Note that it must hold that $w \in \{1, \dots, N|E|\}$. Denoting $\mathcal{S}_w = \{F\}$, by Lemma 3.1 we have $|\mathcal{C}_w| \equiv 2^{n-\ell+\text{cc}(F)} \pmod{2^{n-\ell+k+1}}$. As $\text{cc}(F) \leq k$, the number $|\mathcal{C}_w|$ is then not divisible by $2^{n-\ell+k+1}$ and the algorithm correctly reports the positive outcome. \square

Hence, it remains to prove Lemma 3.2.

4. From cycle covers to matchings. For the proof of Lemma 3.2, instead of counting the number of suitable partial cycle covers, we find it more convenient to count the number of perfect matchings in an auxiliary graph. Note, that this concept is natural when using the *inclusion-exclusion branching* technique: we are interested in counting cycle covers but we only know how to count perfect matchings. Therefore, we simply work with auxiliary graphs in which perfect matchings correspond to cycle-covers in the original graph. A similar auxiliary graph arises in the algorithm for $\#k$ -Multi-Set-Cover [26].

We define a graph G' as follows. The vertex set V' of G' is

$$V' := \{u^0, u^1 : u \in V\}.$$

That is, we put two copies of each vertex of G into the vertex set of G' . The edge set E' of G' is the union of the following two sets:

$$\begin{aligned} E'_0 &:= \{u^0 u^1 : u \in V\}, \\ E'_1 &:= \{u^0 v^0, u^0 v^1, u^1 v^0, u^1 v^1 : uv \in E\}. \end{aligned}$$

In other words, for every vertex $u \in V$ we put an edge in E'_0 connecting the two copies of u in V' , while for every edge $uv \in E$ we put four different edges in E'_1 , each connecting a copy of u with a copy of v in V' . See Figure 1 for a visualization of the construction of G' .

Let $\pi : E'_1 \rightarrow E$ be the natural projection from E'_1 to E : for each $uv \in E$ and $s, t \in \{0, 1\}$, we set $\pi(u^s v^t) = uv$. We extend the mapping π to all subsets $F \subseteq E'$ by setting $\pi(F) := \pi(F \cap E'_1)$. We also extend the weight function ω to the edges of E' by putting $\omega(e) = 0$ for each $e \in E'_0$ and $\omega(e) = \omega(\pi(e))$ for each $e \in E'_1$.

A set of edges F in G' shall be called *simple* if for every $e \in E$, we have

$$|F \cap \pi^{-1}(e)| \leq 1.$$

For now, we mainly focus on simple perfect matchings in G' . We observe that they are in correspondence with partial cycle covers in G , as explained next.

LEMMA 4.1. *For every simple perfect matching M in G' , the set $\pi(M)$ is a partial cycle cover in G of size $|M \cap E'_1|$. Moreover, for every partial cycle cover F in G , there are exactly $2^{|F|}$ simple perfect matchings M in G' for which $F = \pi(M)$.*

Proof. For the first assertion, observe that if for some $u \in V$, the matching M matches u^0 with some vertex of the form v^t for $v \neq u$, then u^1 has to be matched by M with some vertex of the form $\bar{v}^{\bar{t}}$, where $\bar{v} \neq u$ and $\bar{v} \neq v$; the latter inequality follows from the simplicity of M . Then $\deg_{\pi(M)}(u) = 2$. On the other hand, if M matches u^0 with u^1 , then $\deg_{\pi(M)}(u) = 0$. Thus $\deg_{\pi(M)}(u) \in \{0, 2\}$ for every vertex $u \in V$, so $\pi(M)$ is a partial cycle cover in G . The fact that $|\pi(M)| = |M \cap E'_1|$ follows directly from the construction and the simplicity of M .

For the second assertion, consider any partial cycle cover F in G . Let U be the set of vertices incident to the edges of F ; then $|U| = |F|$. Let a *binding* be any function $f : U \rightarrow F$ such that for every $u \in U$, $f(u)$ is one of the two edges of F incident to u . Observe that if for a binding f we define

$$M(f) = \{v^0 v^1 : v \in V \setminus U\} \cup \{u^{[f(u)=uv]} v^{[f(v)=uv]} : uv \in F\}$$

(where $[\varphi] = 1$ if condition φ holds and $[\varphi] = 0$ otherwise), then $M(f)$ is a simple perfect matching in G' satisfying $\pi(M(f)) = F$. Clearly, matchings $M(f)$ obtained for different bindings f are pairwise different. Moreover, it is easy to see that every simple perfect matching M satisfying $\pi(M) = F$ is of the form $M = M(f)$ for some binding f . Indeed, for every $u \in U$ we just set $f(u) = \pi(e)$, where $e \in E'_1$ is the edge of M that is incident to u^1 . Since the total number of different bindings is $2^{|F|}$, we conclude that there are exactly $2^{|F|}$ simple perfect matchings M in G' satisfying $\pi(M) = F$. \square

Lemma 4.1 motivates introducing the following analogues of the sets \mathcal{C}_w . For $w \in \mathbb{Z}$, we define

$$\mathcal{M}_w := \{(M, (L, R)) : M \text{ is a simple perfect matching in } G', \omega(M) = w, \\ |M \cap E'_1| = \ell, \text{ and } (L, R) \text{ is a cut of } V \text{ consistent with } \pi(M)\}.$$

Since for every simple perfect matching M in G' we have $\omega(M) = \omega(\pi(M))$, from Lemma 4.1 we immediately obtain the following corollary.

COROLLARY 4.2. *For every $w \in \mathbb{Z}$, we have $|\mathcal{M}_w| = 2^\ell \cdot |\mathcal{C}_w|$.*

Therefore, to prove Lemma 3.2 it suffices to apply the algorithm provided by the following lemma and divide the outcome by 2^ℓ .

LEMMA 4.3. *Given $w \in \mathbb{Z}$ and a weight function $\omega : E \rightarrow \{1, \dots, N\}$, where $N = \mathcal{O}^*(1)$, the number $|\mathcal{M}_w|$ can be computed in time $\mathcal{O}^*(5^d)$ and space $\mathcal{O}^*(1)$.*

We are left with proving Lemma 4.3.

5. The Count part. In this section, we execute the Count part of the technique and prove Lemma 4.3. Let us first discuss the intuition behind the approach.

The basic idea is that we will compute the number $|\mathcal{M}_w|$ using bottom-up dynamic programming over the given elimination tree T . To achieve polynomial space complexity, this dynamic programming will be cast as a standard recursion, but for this to work, we need the recurrence equations governing the dynamic programming to have a specific form. In essence, whenever we compute an entry of the dynamic programming table at some vertex u , the value should be obtained as a simple aggregation of single entries from the tables of the children of u . The most straightforward approach to computing $|\mathcal{M}_w|$ would be to count partial perfect matchings and to remember, in the states corresponding to u , subsets of $\text{tail}[u]$ consisting of vertices matched to $\text{subtree}(u)$. This would yield a dynamic programming algorithm that is *not* of the form required for the space complexity reduction. However, we show that by counting different objects than partial perfect matchings and using the inclusion-exclusion principle at every computation step, we can reorganize the computation so that the space reduction is possible.

We remark that even though at the end of the day our algorithm relies only on basic ideas such as branching and inclusion-exclusion, there is a deeper intuition behind the definitions of the computed values. In fact, from the right angle, our algorithm can be seen as an application of the technique of *saving space by algebraization*, introduced by Lokshtanov and Nederlof [23], which boils down to applying the Fourier transform on the lattice of subsets to turn subset convolutions into pointwise products. We refer the reader to [19, 17, 2, 31] for other applications of this technique in the context of treedepth-based algorithms.

Partial objects. We start with defining the partial objects that will be counted by the algorithm.

For every vertex $u \in V$, let us order the children of u arbitrarily. Thus, every nonleaf vertex u has a unique leftmost (first in the order) child. For every $u \in V$, let $\text{left}(u)$ be the leftmost leaf descendant of u , that is, the leaf obtained by starting at u and iteratively moving to the leftmost child of the current vertex until a leaf is found. This induces a mapping $\overline{\text{left}}(\cdot)$ on the edges of E' as follows: for an edge $e = u^0 u^1 \in E'_0$, we put $\overline{\text{left}}(e) := \text{left}(u)$, while for an edge of the form $e = u^s v^t \in E'_1$, where $u \neq v$ and $s, t \in \{0, 1\}$, we let $\overline{\text{left}}(e) := \text{left}(v)$, where v is the descendant of u in T . Now, for every $u \in V$ we define the *sheaf* of u as follows:

$$\text{sheaf}[u] := \bigcup_{v \text{ is a leaf in } \text{subtree}[u]} \overline{\text{left}}^{-1}(v) \subseteq E'.$$

See Figure 2 for a schematic presentation of $\text{sheaf}[u]$. The definition of $\text{sheaf}[u]$ is technical because we need it to satisfy two properties. The first property is that $\text{sheaf}[u]$ contains all edges incident to $\text{subtree}[u]$. The second property is the next observation that directly follows from the definition.

Observation 5.1. For every vertex u that is not a leaf in T , the family $\{\text{sheaf}[v] : v \in \text{children}(u)\}$ is a partition of $\text{sheaf}[u]$.

We now move to the description of partial objects. In the following, we will use the convention that if $Z \subseteq V$, then we write $Z' := \{z^0, z^1 : z \in Z\} \subseteq V'$ for the set of




FIG. 2. Schematic definitions of $\overline{\text{left}}$ and sheaf . The left panel presents the definition of the mapping $\overline{\text{left}}(uv)$. Since v is a descendant of u , $\overline{\text{left}}(uv)$ is the leftmost descendant of v . The right panel presents the definition of $\text{sheaf}[u]$. Any edge with the lower endpoint in the yellow highlighted region is a part of $\text{sheaf}[u]$ (the red edge is not part of the $\text{sheaf}[u]$). (Figure in color online.)

copies of vertices of Z in G' . We first introduce the following notion that facilitates the definition of partial objects.

DEFINITION 5.2. Suppose X and Y are two disjoint subsets of V . Further, suppose $S \subseteq E'$ is a set of edges whose all endpoints are in $X' \cup Y'$. For a function $f : Y \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$, we shall say that a pair $(F, (L, R))$, where $F \subseteq S$ and (L, R) is a cut of $X \cup Y$, is compatible with f if the following properties hold:

- F is simple and consistent with the cut (L, R) ;
- $f^{-1}(\{1_L, 2_L\}) \subseteq L$ and $f^{-1}(\{1_R, 2_R\}) \subseteq R$;
- for every $y \in Y$ with $f(y) = 0$, no edge of F is adjacent to y^0 or to y^1 ;
- for every $y \in Y$ with $f(y) \in \{1_L, 1_R\}$, no edge of F is adjacent to y^1 ; and
- every vertex of X' is incident to some edge of F .

In essence, to define partial objects constructed for a vertex u , in the definition above we take Y to be the tail of u , X to be the subtree of u , and S to be the sheaf of u . Then with every function $f : Y \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$ we can associate all partial objects that are compatible with it. However, it makes a difference whether we include or exclude the vertex u from Y ; that is, whether we consider $Y = \text{tail}[u]$ and $X = \text{subtree}(u)$, or $Y = \text{tail}(u)$ and $X = \text{subtree}[u]$. Therefore, we distinguish *inclusive* and *exclusive* partial objects.

DEFINITION 5.3. For every $u \in V$ and every function $f : \text{tail}[u] \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$, we define the set of inclusive partial objects for u and f , denoted $\mathcal{M}[u, f]$, as the set of all pairs $(F, (L, R))$ that are compatible with f , where $X = \text{subtree}(u)$, $Y = \text{tail}[u]$, and $S = \text{sheaf}[u]$.

DEFINITION 5.4. For every $u \in V$ and every function $f : \text{tail}(u) \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$, the set of exclusive partial objects for u and f , denoted $\mathcal{M}(u, f)$, is the set of all pairs $(F, (L, R))$ that are compatible with f , where $X = \text{subtree}[u]$, $Y = \text{tail}(u)$, and $S = \text{sheaf}[u]$.

Observe that both in the inclusive and in the exclusive case we have $X \cup Y = \text{broom}[u]$, so (L, R) is a cut of $\text{broom}[u]$, and $F \subseteq S = \text{sheaf}[u]$. Note also that we *do not* require

F to be a matching in G' . For convenience, by $\text{Func}[u]$ and $\text{Func}(u)$ we shall denote the sets of all functions from $\text{tail}[u]$, respectively, $\text{tail}(u)$, to $\{0, 1_L, 1_R, 2_L, 2_R\}$.

Finally, we stratify the sets $\mathcal{M}[u, f]$ by defining, for all $a, b, c \in \mathbb{N}$, $\mathcal{M}_{a,b,c}[u, f]$ as the set of all the pairs $(F, (L, R)) \in \mathcal{M}[u, f]$ such that $\omega(F) = a$, $|F| = b$, and $|F \cap E'_1| = c$. Sets $\mathcal{M}_{a,b,c}(u, f)$ are defined analogously. The following lemma follows easily from the definitions.

LEMMA 5.5. *If r is the root of T , then $\mathcal{M}_w = \mathcal{M}_{w,n,\ell}(r, \emptyset)$.*

Proof. Note that $\text{sheaf}[r] = E'$ and $\text{subtree}[r] = V$. We observe that if $F \subseteq E'$ is such that $|F| = n$ and every vertex of V' is incident to at least one edge of F , then F has to be a perfect matching in G' . Then the remaining requirements expressed in the definition of $\mathcal{M}_{w,n,\ell}(r, \emptyset)$ exactly correspond to the restrictions on matchings considered in the definition of \mathcal{M}_w . \square

Thus, our goal is to compute all the cardinalities of the sets $\mathcal{M}_{a,b,c}[u, f]$ and $\mathcal{M}_{a,b,c}(u, f)$ for all relevant choices of a, b, c, f, u .

Encoding accumulators in formal variables. Similarly as in [19, 31], we encode the different choices of $a, b, c \in \mathbb{N}$ as degrees of formal variables α, β, γ , so that all the relevant values $|\mathcal{M}_{a,b,c}[u, f]|$ can be stored as coefficients of one polynomial from $\mathbb{Z}[\alpha, \beta, \gamma]$, and similarly for the values $|\mathcal{M}_{a,b,c}(u, f)|$. Formally, for each $u \in V$ and each $f \in \text{Func}[u]$ we define the polynomial $P[u, f] \in \mathbb{Z}[\alpha, \beta, \gamma]$ as

$$P[u, f] = \sum_{a,b,c \in \mathbb{N}} |\mathcal{M}_{a,b,c}[u, f]| \cdot \alpha^a \beta^b \gamma^c,$$

and for each $f \in \text{Func}(u)$ we define the polynomial $P(u, f) \in \mathbb{Z}[\alpha, \beta, \gamma]$ as

$$P(u, f) = \sum_{a,b,c \in \mathbb{N}} |\mathcal{M}_{a,b,c}(u, f)| \cdot \alpha^a \beta^b \gamma^c.$$

Observe that since we assume that the weight function ω only assigns weights in $\{1, \dots, N\}$, in the formula above the numbers $|\mathcal{M}_{a,b,c}[u, f]|$ and $|\mathcal{M}_{a,b,c}(u, f)|$ can be nonzero only for $a \leq N \cdot |E|$, $b \leq |E|$, and $c \leq |E|$. Thus, $P[u, f]$ and $P(u, f)$ are indeed polynomials.

As argued above, each polynomial $P[u, f]$ and $P(u, f)$ has a total degree at most $(N + 2) \cdot |E|$, hence it is a sum of a polynomial (in n) number of monomials. Therefore, we may represent each of the polynomials $P[u, f]$ and $P(u, f)$ by just storing a polynomial-size table of the coefficients of the monomials. Thus, the representation of each polynomial $P[u, f]$ and $P(u, f)$ takes polynomial space, and arithmetic operations on them can be performed in polynomial time.

The computation. The idea now is that each polynomial $P(u, f)$ can be computed using polynomials $P[u, f']$ for different extensions f' of f , while each polynomial $P[u, f]$ for a nonleaf u can be computed using polynomials $P(v, f)$ for v ranging over the children of u . Moreover, polynomials $P[u, f]$ can be computed in polynomial time whenever u is a leaf. These statements are encapsulated in the following lemmas, whose proofs are postponed to section 6.

LEMMA 5.6. *For every $u \in V$ and $f \in \text{Func}(u)$, we have*

$$\begin{aligned} P(u, f) &= P[u, f[u \mapsto 2_L]] + P[u, f[u \mapsto 2_R]] - 2 \cdot P[u, f[u \mapsto 1_L]] \\ &\quad - 2 \cdot P[u, f[u \mapsto 1_R]] + P[u, f[u \mapsto 0]]. \end{aligned}$$

Algorithm 5.1 Procedure `computeExclusive`(u, f).

Input : vertex $u \in V$, function $f \in \text{Func}(u)$ **Output**: $P(u, f)$ **foreach** $s \in \{0, 1_L, 1_R, 2_L, 2_R\}$ **do**| $P_s := \text{computeInclusive}(u, f[u \mapsto s])$ **end**return $P_{2_L} + P_{2_R} - 2 \cdot P_{1_L} - 2 \cdot P_{1_R} + P_0$

Algorithm 5.2 Procedure `computeInclusive`(u, f).

Input : vertex $u \in V$, function $f \in \text{Func}[u]$ **Output**: $P[u, f]$ **if** u is a leaf **then**| Compute P using Lemma 5.8 for u and f **else**| $P \leftarrow 1$ | **foreach** $v \in \text{children}(u)$ **do**| | $P := P \cdot \text{computeExclusive}(v, f)$ | **end****end****return** P

LEMMA 5.7. For every $u \in V$ which is not a leaf in T , and every $f \in \text{Func}[u]$, we have

$$P[u, f] = \prod_{v \in \text{children}(u)} P(v, f).$$

LEMMA 5.8. For every $u \in V$ which is a leaf in T , and every $f \in \text{Func}[u]$, the polynomial $P[u, f]$ can be computed in polynomial time.

The proofs of Lemmas 5.6, 5.7, and 5.8 are given in the next section, but they immediately suggest a recursive method for the computation of polynomials $P[u, f]$ and $P(u, f)$. We now show that this idea can be used to finish the proof of Lemma 4.3.

Proof of Lemma 4.3 assuming Lemmas 5.6, 5.7, and 5.8. We give two mutually recursive procedures—`computeExclusive`(u, f) and `computeInclusive`(u, f)—that compute polynomials $P(u, f)$ and $P[u, f]$, respectively. These procedures are presented above using pseudocode as in Algorithms 5.1 and 5.2. In summary, to compute $P(u, f)$ we recursively compute the values $P[u, f[u \mapsto s]]$ for all $s \in \{0, 1_L, 1_R, 2_L, 2_R\}$, using procedure `computeInclusive`, and then apply the formula provided by Lemma 5.6. To compute $P[u, f]$ we either use the base case provided by Lemma 5.8 when u is a leaf or otherwise we recursively compute the values $P(v, f)$ for all $v \in \text{children}(u)$, using procedure `computeExclusive`, and multiply them. Lemmas 5.6, 5.7, and 5.8 assert that the presented procedures correctly compute the polynomials $P(u, f)$ and $P[u, f]$ for all $u \in V$ and relevant functions f . Hence, by Lemma 5.5, in order to compute $|\mathcal{M}_w|$ it suffices to run the procedure `computeExclusive`(r, \emptyset), where r is the root of T , and return the coefficient in the obtained polynomial that stands by the monomial $\alpha^w \beta^n \gamma^\ell$. Thus, it is clear that the presented algorithm is correct.

It remains to argue that the algorithm runs in time $\mathcal{O}^*(5^d)$ and uses polynomial space. As for the space complexity, observe that at each point, the algorithm maintains a recursion stack of depth at most d and the internal data of each recursive call on the stack takes polynomial space. As for the time complexity, the crucial observation is that throughout the recursion, for every pair (u, f) where $u \in V$ and $f \in \text{Func}(u)$ we call the procedure `computeExclusive` (u, f) exactly once: within the procedure `computeInclusive` (u', f) where u' is the parent of u , or at the very beginning if $u = r$. Similarly, for every pair (u, f) where $u \in V$ and $f \in \text{Func}[u]$ we call the procedure `computeInclusive` (u, f) also exactly once: within the procedure `computeExclusive` (u, f') , where f' is the restriction of f to $\text{tail}(u)$. Thus, the total number of recursive calls executed throughout the algorithm is bounded by the number of pairs (u, f) as above, which is at most $2n \cdot 5^d$. As internal operations within each recursive call take polynomial time, we conclude that the total time complexity is $\mathcal{O}^*(5^d)$. \square

6. Verification of recursive formulas. In this section, we provide the proofs of Lemmas 5.6, 5.7, and 5.8.

Proof of Lemma 5.6. Notice that since the postulated equality involves only sums of polynomials, it suffices to verify the equality of coefficients of each monomial $\alpha^a \beta^b \gamma^c$. In other words, we need to prove that for all $a, b, c \in \mathbb{N}$ we have

$$(6.1) \quad |\mathcal{N}(u, f)| = |\mathcal{N}[u, f[u \mapsto 2_L]]| + |\mathcal{N}[u, f[u \mapsto 2_R]]| \\ - 2|\mathcal{N}[u, f[u \mapsto 1_L]]| - 2|\mathcal{N}[u, f[u \mapsto 1_R]]| + |\mathcal{N}[u, f[u \mapsto 0]]|,$$

where we write \mathcal{N} for $\mathcal{M}_{a,b,c}$ for brevity.

First, let us define

$$\mathcal{N}^L(u, f) := \{(F, (L, R)) \in \mathcal{N}(u, f) : u \in L\}, \\ \mathcal{N}^R(u, f) := \{(F, (L, R)) \in \mathcal{N}(u, f) : u \in R\}.$$

Clearly, $(\mathcal{N}^L(u, f), \mathcal{N}^R(u, f))$ is a partition of $\mathcal{N}(u, f)$. Also, let

$$\mathcal{N}^L[u, f[u \mapsto 0]] := \{(F, (L, R)) \in \mathcal{N}[u, f[u \mapsto 0]] : u \in L\}, \\ \mathcal{N}^R[u, f[u \mapsto 0]] := \{(F, (L, R)) \in \mathcal{N}[u, f[u \mapsto 0]] : u \in R\}.$$

Again, $(\mathcal{N}^L[u, f[u \mapsto 0]], \mathcal{N}^R[u, f[u \mapsto 0]])$ is a partition of $\mathcal{N}[u, f[u \mapsto 0]]$. Thus,

$$(6.2) \quad |\mathcal{N}(u, f)| = |\mathcal{N}^L(u, f)| + |\mathcal{N}^R(u, f)|, \\ |\mathcal{N}[u, f[u \mapsto 0]]| = |\mathcal{N}^L[u, f[u \mapsto 0]]| + |\mathcal{N}^R[u, f[u \mapsto 0]]|.$$

We now observe that by the inclusion-exclusion principle, we have

$$(6.3) \quad |\mathcal{N}^L(u, f)| = |\mathcal{N}[u, f[u \mapsto 2_L]]| - 2|\mathcal{N}[u, f[u \mapsto 1_L]]| + |\mathcal{N}^L[u, f[u \mapsto 0]]|.$$

Indeed, in partial objects $(F, (L, R))$ counted on the left-hand side, both vertices u^0 and u^1 have to be incident to an edge of F . On the right-hand side we count it by first only allowing both u^0 and u^1 to be incident to edges of F , then subtracting terms corresponding to disallowing this either for u^0 or for u^1 , and finally adding a correction term where both u^0 and u^1 are disallowed to be incident to edges of F . Note here that by symmetry (i.e., by swapping u^0 with u^1 and keeping remaining vertices

unchanged), the two subtracted terms are equal, and equal to $|\mathcal{N}[u, f[u \mapsto 1_L]]|$; hence the factor 2. Analogously we argue that

$$(6.4) \quad |\mathcal{N}^R(u, f)| = |\mathcal{N}[u, f[u \mapsto 2_R]]| - 2|\mathcal{N}[u, f[u \mapsto 1_R]]| + |\mathcal{N}^R[u, f[u \mapsto 0]]|.$$

Now (6.1) follows by adding (6.3) and (6.4) and using (6.2). \square

Proof of Lemma 5.7. We define a function (here \prod denotes the Cartesian product)

$$\xi : \mathcal{M}[u, f] \rightarrow \prod_{v \in \text{children}(u)} \mathcal{M}(v, f)$$

as follows: for $(F, (L, R)) \in \mathcal{M}[u, f]$, set

$$\xi(F, (L, R)) := ((F \cap \text{sheaf}[v], (L \cap \text{broom}[v], R \cap \text{broom}[v])) : v \in \text{children}(u)).$$

Note that $\text{sheaf}[v]$ contains all edges incident to $\text{subtree}[v]$. Hence, it is straightforward to verify from the definitions that for each $(F, (L, R)) \in \mathcal{M}[u, f]$ and $v \in \text{children}(u)$, the pair $(F \cap \text{sheaf}[v], (L \cap \text{broom}[v], R \cap \text{broom}[v]))$ belongs to $\mathcal{M}(v, f)$. Hence we may indeed set the co-domain of ξ to be $\prod_{v \in \text{children}(u)} \mathcal{M}(v, f)$. Also, it is clear that ξ is injective.

We now observe that ξ is also surjective. Indeed, since $\{\text{sheaf}[v] : v \in \text{children}(u)\}$ is a partition of $\text{sheaf}[u]$ by Observation 5.1 and $(\cup_{v \in \text{children}(u)} L_v, \cup_{v \in \text{children}(u)} R_v)$ is compatible with f since each (L_v, R_v) is compatible with f and $\{\text{subtree}[v] : v \in \text{children}(u)\}$ is a partition of $\text{subtree}(u)$, it is again straightforward to verify from definitions that the following assertion holds: If for each $v \in \text{children}(u)$ we have some $(F_v, (L_v, R_v)) \in \mathcal{M}(v, f)$, then setting

$$F := \bigcup_{v \in \text{children}(u)} F_v, \quad L := \bigcup_{v \in \text{children}(u)} L_v, \quad R := \bigcup_{v \in \text{children}(u)} R_v$$

yields a pair $(F, (L, R))$ that belongs to $\mathcal{M}[u, f]$ and satisfies

$$\xi(F, (L, R)) = ((F_v, (L_v, R_v)) : v \in \text{children}(u)).$$

This implies that ξ is a bijection between $\mathcal{M}[u, f]$ and $\prod_{v \in \text{children}(u)} \mathcal{M}(v, f)$.

Finally, we observe that by Observation 5.1 and the fact that ξ is a bijection, we have

$$\begin{aligned} P[u, f] &= \sum_{(F, (L, R)) \in \mathcal{M}[u, f]} \alpha^{\omega(F)} \beta^{|F|} \gamma^{|F \cap E'_1|} \\ &= \sum_{(F, (L, R)) \in \mathcal{M}[u, f]} \prod_{v \in \text{children}(u)} \alpha^{\omega(F \cap \text{sheaf}[v])} \beta^{|F \cap \text{sheaf}[v]|} \gamma^{|F \cap \text{sheaf}[v] \cap E'_1|} \\ &= \prod_{v \in \text{children}(u)} \sum_{(F_v, (L_v, R_v)) \in \mathcal{M}(v, f)} \alpha^{\omega(F_v)} \beta^{|F_v|} \gamma^{|F_v \cap E'_1|} \\ &= \prod_{v \in \text{children}(u)} P(v, f). \end{aligned}$$

This concludes the proof. \square

Proof of Lemma 5.8. Let $Z := \pi(\text{sheaf}[u])$. We claim that

$$(6.5) \quad P[u, f] = \prod_{xy \in Z} Q[xy, f] \cdot \prod_{x \in \text{tail}[u]} R[x, f],$$

where

$$Q[xy, f] = \begin{cases} 1 + ij \cdot \alpha^{\omega(xy)} \beta \gamma & \text{if } (f(x), f(y)) = (i_L, j_L) \text{ with } i, j \in \{1, 2\} \text{ or} \\ & (f(x), f(y)) = (i_R, j_R) \text{ with } i, j \in \{1, 2\}, \\ 1 & \text{otherwise,} \end{cases}$$

and

$$R[x, f] = \begin{cases} 1 + \beta & \text{if } f(x) \in \{2_L, 2_R\}, \\ 2 & \text{if } f(x) = 0, \\ 1 & \text{otherwise.} \end{cases}$$

Note that this will conclude the proof, as the product (6.5) can be expanded into a sum of monomials in variables α, β, γ in polynomial time.

To argue the correctness of formula (6.5), let us recall that the coefficients of $P[u, f]$ should count the pairs $(F, (L, R))$ that are compatible with f according to Definition 5.2, where we put $X = \emptyset$, $Y = \text{tail}[u]$, and $S = \text{sheaf}[u]$, so that the coefficient by $\alpha^a \beta^b \gamma^c$ is the number of such pairs with $\omega(F) = a$, $|F| = b$, and $|F \cap E'_1| = c$. We now show that each such pair $(F, (L, R))$ can be described by independent choices made for each $xy \in Z$ and each $x \in \text{tail}[u]$, which respectively correspond to the factors in formula (6.5).

For every edge $xy \in Z$, within $\pi^{-1}(xy) = \{x^0 y^0, x^1 y^0, x^0 y^1, x^1 y^1\}$ there may be at most one edge in F (because F needs to be simple), and this can happen only when $(f(x), f(y)) = (i_L, j_L)$ or $(f(x), f(y)) = (i_R, j_R)$ for some $i, j \in \{1, 2\}$. In this case, the number of possibilities for choosing the edge from F is ij . This explains the formula for $Q[xy, f]$: the summand 1 corresponds to the option of not choosing any edge from $\pi^{-1}(xy)$, while the summand $ij \cdot \alpha^{\omega(xy)} \beta \gamma$ corresponds to the option of choosing any one of the edges from $\pi^{-1}(xy)$. Note that the degrees by α, β, γ respectively correspond to the contribution of this edge to $\omega(F), |F|, |F \cap E'_1|$.

Next, for every vertex $x \in \text{tail}[u]$, the edge $x^0 x^1$ may be added to F only when we have $f(x) \in \{2_L, 2_R\}$, and in this case it contributes only to $|F|$, since $\omega(x^0 x^1) = 0$ and $x^0 x^1 \notin E'_1$. Further, x has to belong to L if $f(x) \in \{1_L, 2_L\}$, and x has to belong to R if $f(x) \in \{1_R, 2_R\}$, but if $f(x) = 0$, then we may include x either in L or in R . This explains the formula for $R[x, f]$: the summand β corresponds to the option of including $x^0 x^1$ in F , while the 2 corresponds to the two options of including x either in L or in R .

Since the choices made for different edges $xy \in Z$ and for different vertices $x \in \text{tail}[u]$ do not restrict each other, formula (6.5) for $P[u, f]$ follows. \square

7. Conclusion and further research. In this paper, we answered the open question of Hegerfeld and Kratsch [19] by presenting an $\mathcal{O}^*(5^d)$ -time and polynomial space algorithm for *Hamiltonian Path*, *Hamiltonian Cycle*, *Longest Path*, and *Longest Cycle Min Cycle Cover*, where d is the depth of a provided elimination forest of the input graph. However, there are still multiple open problems around time and space efficient algorithms on graphs of bounded treedepth. Here, we list a selection.

Approximation of treedepth. Recall that the treewidth of a graph can be approximated up to a constant factor in fixed-parameter time. For instance, the classic algorithm of Robertson and Seymour [34] (see also [7]) takes on input a graph G and integer t , works in time $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ and in polynomial space, and either concludes

that the treewidth of G is larger than t , or finds a tree decomposition of G of width at most $4t+4$. This means that for the purpose of designing $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ -time algorithms on graphs of treewidth t , we may assume that a tree decomposition of approximately optimum width is given, as it can be always computed from the input graph within the required complexity bounds. Unfortunately, no such approximation algorithm is known for the treedepth. Namely, it is known that the treedepth can be computed exactly in time $2^{\mathcal{O}(d^2)} \cdot n$ and polynomial space [33, 25] and approximated up to factor $\mathcal{O}(t \log^{3/2} t)$ in polynomial time [10], where d and t are the values of the treedepth and the treewidth of the input graph, respectively. A piece of the theory that seems particularly missing is a constant-factor approximation algorithm for treedepth running in time $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$; polynomial space usage would be also desired.

Faster algorithms. The bases of the exponent of the running times of the algorithms given by Hegerfeld and Kratsch [19] for the treedepth parameterization match the ones obtained by Cygan et al. [9] for the treewidth parameterization. In the case of our results, the situation is different: while *Hamiltonian Cycle* can be solved in time $4^t \cdot n^{\mathcal{O}(1)}$ in graphs of treewidth t [9] and in time $(2 + \sqrt{2})^p \cdot n^{\mathcal{O}(1)}$ in graphs of pathwidth p [8], we needed to increase the base of the exponent to 5 in order to achieve polynomial space complexity for the treedepth parameterization. As the treedepth of a graph is never smaller than its pathwidth, it is natural to ask whether there is an $(2 + \sqrt{2})^d \cdot n^{\mathcal{O}(1)}$ -time polynomial-space algorithm for *Hamiltonian Cycle* on graphs of treedepth d . In fact, reducing the base 5 to any $c < 5$ would be interesting.

Derandomization. Shortly after its introduction, the Cut&Count technique for the treewidth parameterization has been derandomized. Bodlaender et al. [5] presented two approaches for doing so. The first one, called the *rank-based approach*, boils down to maintaining a small set of representative partial solutions along the dynamic programming computation, and pruning irrelevant partial solutions on the fly using Gaussian elimination. Fomin et al. [16] later reinterpreted this technique in the language of matroids and extended it. The second approach, called *determinant-based*, uses the ideas behind Kirchoff's matrix-tree theorem to deliver a formula for counting suitable spanning trees of a graph, which can be efficiently evaluated by a dynamic programming over a tree decomposition.

It seems to us that none of these approaches applies in the context of the treedepth parameterization, where we additionally require polynomial space complexity. Both rank-based and matroid-based approaches are based on maintaining a set of representative solutions, which in the worst case may have exponential size. In the determinant-based approach, when computing the formula for the number of spanning trees over a tree decomposition, the aggregation of dynamic programming tables is done using operations that are algebraically more involved, and which, in particular, are noncommutative. See the work of Włodarczyk [35] for a discussion. It is unclear whether this computation can be reorganized so that in the aggregation we use only pointwise product—which, in essence, is our current methodology from the algebraic perspective.

Recently, Nederlof et al. [28] gave an isolation scheme for *Hamiltonian Cycle* that uses $\mathcal{O}(t \log(n) + \log^2(n))$ -bits of randomness. This, unfortunately, is still prohibitively expensive to derandomize our algorithm or algorithms of Hegerfeld and Kratsch [19]. We leave it as an open problem to give a deterministic algorithm in $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)}$ time and polynomial space usage.

Other graph parameters. Hegerfeld and Kratsch [19] were not the first to employ Cut&Count on structural graph parameters beyond treewidth. Pino, Bodlaender, and

van Rooij [32] used Cut&Count and rank-based approach to get single-exponential time algorithms for connectivity problems parametrized by *branchwidth*. Recently, Cut&Count was also applied in the context of *cliquewidth* [3], and of \mathbb{Q} -*rankwidth*, *rankwidth*, and *MIM-width* [4]. All these algorithms have exponential space complexity, as they follow the standard dynamic programming approach. One may expect that maybe for the depth-bounded counterparts of cliquewidth and rankwidth—*shrubdepth* [18] and *rankdepth* [12]—time-efficient polynomial-space algorithms can be designed, similarly as for *treedepth*.

Appendix A. Faster polynomial-space algorithm for *Long Path* in H -minor free graphs. Let us fix a graph H and consider the class of H -minor-free graphs, that is, graphs that exclude H as a minor. Lokshtanov, Mnich, and Saurabh [22] gave an algorithm for the *Long Path* problem on H -minor-free graphs that runs in time $2^{\mathcal{O}(\sqrt{\ell} \log^2 \ell)} \cdot n^{\mathcal{O}(1)}$ and uses polynomial space. We now show that this result can be slightly improved using a combination of our findings and previously known techniques, such as Turing kernelization for *Longest Path* and basic bidimensionality.

THEOREM A.1. *For every fixed graph H , there is a randomized algorithm for the Long Path problem on H -minor free graphs that runs in time $2^{\mathcal{O}(\sqrt{\ell} \log \ell)} \cdot n^{\mathcal{O}(1)}$ and uses polynomial space. The algorithm has a one-sided error: it can return false negatives with probability at most $\frac{1}{2}$.*

Before we prove Theorem A.1 let us consider the algorithm for the case when $n = \ell^{\mathcal{O}(1)}$. Later, we will use this algorithm as a black-box to show the general case.

LEMMA A.2. *Let G be some H -minor free graph (for some fixed H) that has $n = \ell^{\mathcal{O}(1)}$ vertices. There is a randomized algorithm that decides whether G contains a simple path of length ℓ in time $2^{\mathcal{O}(\sqrt{\ell} \log \ell)} \cdot n^{\mathcal{O}(1)}$ and uses polynomial space. The algorithm has a one-sided error: it can only return false negatives with probability at most $\frac{1}{2}$.*

Proof. Let us consider the input instance (G, ℓ) such that G is H -minor-free and has $n = \ell^{\mathcal{O}(1)}$ vertices. We use the classic approximation algorithm for the treewidth of Robertson and Seymour [34] (see also a presentation in [7]).

THEOREM A.3 (see [34, 7]). *For a given parameter $\lambda \in \mathbb{N}$, there is an algorithm that runs in $2^{\mathcal{O}(\lambda)} n^{\mathcal{O}(1)}$ time and uses polynomial space and either*

- *finds a tree decomposition of G of width at most $4\lambda + 4$, or*
- *concludes that the treewidth of G is greater than λ .*

Now, we apply Theorem A.3 on G with parameter $\lambda := h\lceil\sqrt{\ell}\rceil$, where h is some constant that we will choose later. Observe that for this choice of λ , Theorem A.3 runs in $2^{\mathcal{O}(\sqrt{\ell})}$ time. Next, similarly to Lokshtanov, Mnich, and Saurabh [22], we use the following observation of Demaine et al. [11].

Observation A.4 (see [11, 22]). *For every fixed graph H there exists a constant h such that for all $p \in \mathbb{N}$ the following holds: In any H -minor free graph of treewidth $t > hp$ there is a simple path on at least p^2 vertices.*

Equipped with this observation, we can conclude that when Theorem A.3 returns that treewidth of G is greater than λ , then G contains a path on ℓ vertices and the algorithm can just return a positive answer. Hence, we are left with investigating the former case when a tree decomposition of width $\mathcal{O}(\sqrt{\ell})$ has been found.

In this case, we use the observation by Bodlaender et al. [6], that every tree decomposition of width w can be transformed in polynomial time into an elimination

forest of depth $\mathcal{O}(w \log n)$. Note that in our case $n = \ell^{\mathcal{O}(1)}$ and $w = \mathcal{O}(\lambda)$, hence the depth d of this elimination forest is bounded by $\mathcal{O}(\sqrt{\ell} \log \ell)$. Finally, we apply the algorithm of Theorem 1.2 for *Longest Path* on this elimination forest. This algorithm runs in time $2^{\mathcal{O}(d)} \cdot n^{\mathcal{O}(1)} = 2^{\mathcal{O}(\sqrt{\ell} \log \ell)}$ and uses polynomial space.

This concludes the description of the algorithm. The space and time complexity of the algorithm follows from the construction. The probability of success follows from Theorem 1.2. \square

Now, we prove Theorem A.1.

Proof of Theorem A.1. We use the result of Jansen, Pilipczuk, and Wrochna [20] who gave a *polynomial Turing kernel* for the *Longest Path* problem in H -topological-minor-free for every fixed H . That is, they showed how to solve *Longest Path* in H -topological-minor-free graphs by a polynomial-time algorithm that has access to an oracle solving the problem on H -topological-minor-free graphs which have $\ell^{\mathcal{O}(1)}$ vertices.

We are almost done since in Lemma A.2 we have designed such an oracle. There is, however, a technical issue in that approach because the algorithm in Lemma A.2 is randomized. We need to guarantee that every call to the Lemma A.2 succeeds (with sufficiently high probability).

Note that we can use a standard trick and reduce the error probability of every oracle call to $n^{-\mathcal{O}(1)}$ by repeating $\mathcal{O}(\log n)$ times the algorithm in Lemma A.2, so that the overall probability that any oracle call returns an incorrect answer is bounded by $\mathcal{O}(n^{-3})$. This results in an algorithm `LongestPathTwoSided`(G) that given any H -minor-free graph G returns if G has a simple path of length $\geq \ell$. Note, that the algorithm `LongestPathTwoSided`(G) may return both false-positives and false-negatives, but the error probability is $\mathcal{O}(n^{-3})$.

Next, we use `LongestPathTwoSided`(G) to design an algorithm for *Longest Path* that has one-sided error and returns false-negative with probability at most $1/2$. We do this by designing a self-reduction to the *Longest Path* problem. Namely, we construct a procedure `LongestPath`(G, U), where $U \subseteq E$ is the set of *undecided* edges. Initially, we set $U = E[G]$. If $U = \emptyset$, the procedure `LongestPath`(G, \emptyset) outputs *yes* if G is an ℓ -path (otherwise it answers *no*). If $U \neq \emptyset$, we pick any edge $e \in U$ and query the oracle for `LongestPathTwoSided`(G'), where G' is the graph G without the edge e . If the oracle returned *yes*, then we permanently remove e from graph G and the set U and continue until $U = \emptyset$ (because we know that the answer to the problem is still positive after deleting e). Otherwise, the oracle returned *no*. In this case, we remove e from U and leave it in G (since it is required to construct an ℓ -path in G) and continue until $U = \emptyset$. This concludes the description of `LongestPath`(G, U). The final algorithm simply returns an answer to `LongestPath`($G, E[G]$).

Observe that every time we make a call to `LongestPathTwoSided` one edge from the set U is removed. Initially, $|U| = |E[G]| = \mathcal{O}(n^2)$, therefore, the algorithm makes at most $\mathcal{O}(n^2)$ calls to the oracle `LongestPathTwoSided` and we incur only a polynomial time overhead. Moreover, observe that `LongestPath`(G, U) never returns *yes* when the answer to the problem is *no*. Finally, because the number of calls is $\mathcal{O}(n^2)$ and each call is correct with $\mathcal{O}(n^{-3})$ probability we conclude that the probability of error is bounded by $\mathcal{O}(n^{-1}) \leq \frac{1}{2}$, which concludes the proof. \square

Appendix B. Problems and definitions. In this appendix we include the definitions of the problems considered in this paper.

*Long Path***Input:** An undirected graph G and an integer ℓ .**Task:** Is there a simple path on ℓ vertices in G ?*Long Cycle***Input:** An undirected graph G and an integer ℓ .**Task:** Is there a simple cycle of length ℓ in G ?*Hamiltonian Path***Input:** An undirected graph G .**Task:** Is there a simple path in G that visits all the vertices?*Hamiltonian Cycle***Input:** An undirected graph G .**Task:** Is there a simple cycle in G that visits all the vertices?*Min Cycle Cover***Input:** An undirected graph G and an integer k .**Task:** Can the vertices of G be covered with at most k vertex-disjoint cycles?*Partial Cycle Cover***Input:** An undirected graph G , integers k and ℓ .**Task:** Is there a family of at most k vertex-disjoint cycles in G that jointly visit exactly ℓ vertices?

Acknowledgments. We would like to thank the anonymous reviewers for their suggestions and comments. The second author would like to thank Marcin Wrochna for some early discussions on combining treedepth and Cut&Count.

REFERENCES

- [1] M. BELBASI AND M. FÜRER, *Saving Space by Dynamic Algebraization Based on Tree Decomposition: Minimum Dominating Set*, preprint, <https://arxiv.org/abs/1711.10088>, 2017.
- [2] M. BELBASI AND M. FÜRER, *A space-efficient parameterized algorithm for the Hamiltonian Cycle problem by dynamic algebraization*, in Proceedings of the 14th International Computer Science Symposium in Russia, CSR 2019, Lecture Notes in Comput. Sci. 11532, Springer, 2019, pp. 38–49.
- [3] B. BERGOUNOUX AND M. M. KANTÉ, *Fast exact algorithms for some connectivity problems parameterized by clique-width*, Theoret. Comput. Sci., 782 (2019), pp. 30–53.
- [4] B. BERGOUNOUX AND M. M. KANTÉ, *More applications of the d -neighbor equivalence: Connectivity and acyclicity constraints*, in Proceedings of the 27th Annual European Symposium on Algorithms, ESA 2019, LIPIcs. Leibniz Int. Proc. Inform. 144, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, 17.

- [5] H. L. BODLAENDER, M. CYGAN, S. KRATSCH, AND J. NEDERLOF, *Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth*, Inform. and Comput., 243 (2015), pp. 86–111.
- [6] H. L. BODLAENDER, J. R. GILBERT, H. HAFSTEINSSON, AND T. KLOKS, *Approximating treewidth, pathwidth, frontsize, and shortest elimination tree*, J. Algorithms, 18 (1995), pp. 238–255.
- [7] M. CYGAN, F. V. FOMIN, L. KOWALIK, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Parameterized Algorithms*, Springer, Cham, 2015.
- [8] M. CYGAN, S. KRATSCH, AND J. NEDERLOF, *Fast Hamiltonicity checking via bases of perfect matchings*, J. ACM, 65 (2018), 12.
- [9] M. CYGAN, J. NEDERLOF, M. PILIPCZUK, M. PILIPCZUK, J. M. M. VAN ROOIJ, AND J. O. WOJTASZCZYK, *Solving connectivity problems parameterized by treewidth in single exponential time*, ACM Trans. Algorithms, 18 (2022), 17.
- [10] W. CZERWIŃSKI, W. NADARA, AND M. PILIPCZUK, *Improved bounds for the excluded-minor approximation of treedepth*, in Proceedings of the 27th Annual European Symposium on Algorithms, ESA 2019, LIPIcs. Leibniz Int. Proc. Inform. 144, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019 34.
- [11] E. D. DEMAINE, F. V. FOMIN, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on bounded-genus graphs and H -minor-free graphs*, J. ACM, 52 (2005), pp. 866–893.
- [12] M. DEVOS, O. KWON, AND S. OUM, *Branch-depth: Generalizing tree-depth of graphs*, Eur. J. Comb., 90 (2020), 103186.
- [13] R. G. DOWNEY AND M. R. FELLOWS, *Fundamentals of Parameterized Complexity*, Texts in Computer Science, Springer, London, 2013.
- [14] A. DRUCKER, J. NEDERLOF, AND R. SANTHANAM, *Exponential time paradigms through the polynomial time lens*, in Proceedings of the 24th Annual European Symposium on Algorithms, ESA 2016, LIPIcs. Leibniz Int. Proc. Inform. 57, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2016, 36.
- [15] J. FLUM AND M. GROHE, *Parameterized Complexity Theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2006.
- [16] F. V. FOMIN, D. LOKSHTANOV, F. PANOLAN, AND S. SAURABH, *Efficient computation of representative families with applications in parameterized and exact algorithms*, J. ACM, 63 (2016), 29.
- [17] M. FÜRER AND H. YU, *Space saving by dynamic algebraization based on tree-depth*, Theory Comput. Syst., 61 (2017), pp. 283–304.
- [18] R. GANIAN, P. HLINĚNÝ, J. NEŠETŘIL, J. OBDŘÁLEK, AND P. OSSONA DE MENDEZ, *Shrub-depth: Capturing height of dense graphs*, Log. Methods Comput. Sci., 15 (2019), 7.
- [19] F. HEGERFELD AND S. KRATSCH, *Solving connectivity problems parameterized by treedepth in single-exponential time and polynomial space*, in Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, LIPIcs. Leibniz Int. Proc. Inform. 154, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020, 29.
- [20] B. M. P. JANSEN, M. PILIPCZUK, AND M. WROCHNA, *Turing kernelization for finding long paths in graph classes excluding a topological minor*, Algorithmica, 81 (2019), pp. 3936–3967.
- [21] T. KORHONEN, *A single-exponential time 2-approximation algorithm for treewidth*, in Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, IEEE, 2021, pp. 84–192.
- [22] D. LOKSHTANOV, M. MNICH, AND S. SAURABH, *Planar k -path in subexponential time and polynomial space*, in Proceedings of the 37th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2011, Lecture Notes in Comput. Sci. 6986, Springer, 2011, pp. 262–270.
- [23] D. LOKSHTANOV AND J. NEDERLOF, *Saving space by algebraization*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, ACM, 2010, pp. 321–330.
- [24] K. MULMULEY, U. V. VAZIRANI, AND V. V. VAZIRANI, *Matching is as easy as matrix inversion*, Combinatorica, 7 (1987), pp. 105–113.
- [25] W. NADARA, M. PILIPCZUK, AND M. SMULEWICZ, *Computing treedepth in polynomial space and linear FPT time*, in Proceedings of the 30th Annual European Symposium on Algorithms, ESA 2022, LIPIcs. Leibniz Int. Proc. Inform. 244, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2022, 79.
- [26] J. NEDERLOF, *Inclusion Exclusion for Hard Problems*, Master’s thesis, Utrecht University, Utrecht, The Netherlands, 2008.

- [27] J. NEDERLOF, M. PILIPCZUK, C. M. F. SWENNENHUIS, AND K. WĘGRZYCKI, *Hamiltonian cycle parameterized by treedepth in single exponential time and polynomial space*, in Proceedings on the 46th International Workshop on Graph-Theoretic Concepts in Computer Science, WG, 2020, Leeds, UK, 2020, Lecture Notes in Comput. Sci. 12301, Springer, Cham, 2020, pp. 27–39.
- [28] J. NEDERLOF, M. PILIPCZUK, C. M. F. SWENNENHUIS, AND K. WĘGRZYCKI, *Isolation schemes for problems on decomposable graphs*, in Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, LIPIcs. Leibniz Int. Proc. Inform. 219, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2022, 50.
- [29] J. NEŠETRIL AND P. O. DE MENDEZ, *Tree-depth, subgraph coloring and homomorphism bounds*, European J. Combin., 27 (2006), pp. 1022–1041.
- [30] M. PILIPCZUK AND S. SIEBERTZ, *Polynomial bounds for centered colorings on proper minor-closed graph classes*, in Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, SIAM, 2019, pp. 1501–1520, <https://doi.org/10.1137/1.9781611975482.91>.
- [31] M. PILIPCZUK AND M. WROCHNA, *On space efficiency of algorithms working on structural decompositions of graphs*, ACM Trans. Comput. Theory, 9 (2018), 18.
- [32] W. J. A. PINO, H. L. BODLAENDER, AND J. M. M. VAN ROOIJ, *Cut and Count and representative sets on branch decompositions*, in Proceedings of the 11th International Symposium on Parameterized and Exact Computation, IPEC 2016, LIPIcs. Leibniz Int. Proc. Inform. 63, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016, 27.
- [33] F. REIDL, P. ROSSMANITH, F. SÁNCHEZ VILLAAMIL, AND S. SIKDAR, *A faster parameterized algorithm for treedepth*, in Proceedings of the 41st International Colloquium on Automata, Languages, and Programming, ICALP 2014, Lecture Notes in Comput. Sci. 8572, Springer, 2014, pp. 931–942.
- [34] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory Ser. B, 63 (1995), pp. 65–110.
- [35] M. WŁODARCZYK, *Clifford algebras meet tree decompositions*, Algorithmica, 81 (2019), pp. 497–518.